Development of Monte Carlo Models for the Optimisation of Positron Emission Particle Tracking Experiments

By

Matthew Herald

A thesis submitted to
the University of Birmingham
for the degree of
DOCTOR OF PHILOSOPHY

Chemical Engineering
School of Chemical Engineering
College of Engineering and Physical Sciences
University of Birmingham
January 2023

UNIVERSITY OF BIRMINGHAM

University of Birmingham Research Archive
e-theses repository

This unpublished thesis/dissertation is under a Creative Commons Attribution 4.0 International (CC BY 4.0) licence.

You are free to:

Share — copy and redistribute the material in any medium or format

Adapt — remix, transform, and build upon the material for any purpose, even commercially.

The licensor cannot revoke these freedoms as long as you follow the license terms.

Under the following terms:

Attribution — You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.

No additional restrictions — You may not apply legal terms or technological measures that legally restrict others from doing anything the license permits.

Notices:

You do not have to comply with the license for elements of the material in the public domain or where your use is permitted by an applicable exception or limitation.

No warranties are given. The license may not give you all of the permissions necessary for your intended use. For example, other rights such as publicity, privacy, or moral rights may limit how you use the material.

Unless otherwise stated, any material in this thesis/dissertation that is cited to a third-party source is not included in the terms of this licence. Please refer to the original source(s) for licencing conditions of any quotes, images or other material cited to a third party.

© Copyright by UNIVERSITY OF BIRMINGHAM, 2023

All Rights Reserved

ABSTRACT

Positron emission particle tracking (PEPT) is a non-invasive technique used to measure the three-dimensional position of positron-emitting tracers. PEPT is useful for studying myriad industrial and/or scientific systems which often are optically inaccessible. However, when running an experiment, often little attention is paid to optimising the tracers, detectors, algorithms, and experimental procedures. As a result, trajectories can be degraded leading to inefficient use of resources. To address this opportunity, Monte Carlo simulations are employed to model experiments and predict the tracer activity, detector geometry, and algorithm parameters that will produce the best trajectory resolution possible and even determine whether an experiment is feasible or how long the experiment should be conducted to reduce uncertainty to an acceptable degree. Importantly, this simulated work can be conducted prior to experiments. In this thesis, a general procedure for simulating PEPT experiments is described which can be applied to any PEPT experiment. The results of this work demonstrate that not only is this method able to produce realistic synthetic PEPT data, but allows, for the first time, a quantitative comparison of PEPT algorithms, the ability to optimise experiments, and to develop new PEPT methodologies using information difficult or impossible to extract from real experiments.

DEDICATION

Dedicated to Robyn for lovingly supporting me through this journey and to my parents for always encouraging me to follow my dreams.

ACKNOWLEDGMENTS

There are too many people to list who have inspired and helped me along the way, but I would not be writing this acknowledgement today without the invaluable support and guidance from my advisors, Christopher Windows-Yule and Jonathan Seville. They believed in me from the beginning of my PhD and pushed me to be more rigorous, professional, and a better person. I would also like to extend my sincere thanks to my fellow PhD students, Dominik Werner, Jack Sykes, Leonard Nicuşan, and Dawid Hampel. Without their contributions, conversations, and creativity this thesis would be far shorter and decidedly less useful. Further, I would like to mention a few others with who it has been a pleasure to work with these last years; David Parker for creating the PEPT technique and championing its use, Ben Phoenix for running the cyclotron, Samuel Manger for helping me with GATE simulations, Tichaona Dauramanzi for the productive and friendly conversations, Peter James for brewing coffee every day at 10 am and 3:30 pm sharp, and Tzany Kokalova-Wheldon for use of PEPT facilities.

LIST OF PUBLISHED WORKS AND AUTHOR CONTRIBUTIONS

Journal Papers

1. *Monte Carlo model validation of a detector system used for Positron Emission Particle Tracking*, **M. Herald**, T. Kokalova-Wheldon, C. R. K. Windows-Yule, Nuclear Instruments and Methods in Physics Research Section A, Volume 993, (2021), 10.1016/j.nima-.2021.165073

   M. Herald created the Monte Carlo model of the detector system, performed the validation experiments, conducted the analysis, and wrote the paper. T. Kokalova-Wheldon provided access to the PEPT facility as well as the time and resources to conduct the experiments. C. R. K. Windows-Yule provided supervision and editing.

2. *DEM2GATE: Combining discrete element method simulation with virtual positron emission particle tracking experiments*, **M. Herald**, J. Sykes, D. Werner, J. P. K. Seville, C. R. K. Windows-Yule, Powder Technology, Volume 401, (2022), 10.1016/j.powtec.2022.117302

   M. Herald created the Monte Carlo simulation, analysed the simulated results, and wrote the paper. J. Sykes and D. Werner assisted with data processing, created the discrete element method simulation, and contributed to editing. J. P. K Seville and C. R. K. Windows-Yule provided supervision and editing.

3. *Recent advances in positron emission particle tracking: a comparative review*, C. R. K. Windows-Yule, **M. Herald**, A. Nicuşan, C. Wiggins, G. Pratx, S. Manger, A. Odo, T.

Leadbeater, J. Pellico, R. Rosales, Reports on Progress in Physics, Volume 85, (2022), 10.1088/1361-6633/ac3c4c

C. R. K. Windows-Yule organised the international collaboration and wrote the majority of the paper. M. Herald created the simulations and framework for assessing algorithm performance, conducted the analysis of each algorithm's performance, and authored the section regarding the simulations and comparative review of algorithms. A. Nicuşan provided analysis with PEPT-ML algorithm. C. Wiggins provided analysis with the G-Means and FPI algorithms. G. Pratx provided analysis with the SBSR algorithm. S. Manger provided analysis with the Birmingham Method and PEPT-EM algorithms. A. Odo provided analysis with the SDM algorithm. T. Leadbeater assisted in authoring the PEPT detectors section. J. Pellico and R. Rosales assisted in authoring the PEPT tracers section.

4. *Autonomous Digitizer Calibration of a Monte Carlo Detector Model through Evolutionary Simulation*, **M. Herald**, A. Nicuşan, T. Kokalova-Wheldon, J. P. K. Seville, C. R. K. Windows-Yule, Nature Scientific Reports, Volume 12, (2022), 10.1038/s41598-022-24022-x

   M. Herald created the Monte Carlo model, performed the validation experiments, used the evolutionary algorithm to calibrate the model, analysed the results, and wrote the paper. A. Nicuşan developed the software to use the evolutionary algorithm. T. Kokalova-Wheldon provided access to the PEPT facility as well as the time and resources to conduct the experiments. J. P. K Seville and C. R. K. Windows-Yule provided supervision and editing.

5. *Improving the Accuracy of PEPT Algorithms through Dynamic Parameter Optimisation*, **M. Herald**, J. Sykes, D. J. Parker, T. Kokalova-Wheldon, J. P. K Seville, C. R. K. Windows-Yule, Nuclear Instruments and Methods in Physics Research Section A, Volume 1047, (2023), 10.1016/j.nima.2022.167831

M. Herald performed the static tracer experiments, conducted the Monte Carlo simulations, analysed the results, developed the dynamic parameter optimisation method, and wrote the paper. J. Sykes performed the moving tracer experiment. D. J. Parker provided consultation when using the dynamic parameter optimisation method. T. Kokalova-Wheldon provided access to the PEPT facility as well as the time and resources to conduct the experiments. J. P. K Seville and C. R. K. Windows-Yule provided supervision and editing.

6. *Monte Carlo Model of the Large Modular Array for Positron Emission Particle Tracking*, **M. Herald**, D. M. Hampel, T. Kokalova-Wheldon, J. P. K. Seville, C. R. K Windows-Yule, IEEE Access, Volume 11, (2023) 10.1109/ACCESS.2023.3255505

   M. Herald developed the Monte Carlo model, performed the validation experiment, analysed the results, calibrated the model, and wrote the paper. D. M. Hampel operated the detector and provided details of its construction. T. Kokalova-Wheldon provided access to the PEPT facility as well as the time and resources to conduct the experiments. J. P. K Seville and C. R. K. Windows-Yule provided supervision and editing.

Books

1. *Positron Emission Particle Tracking: A comprehensive guide*, C. R. K. Windows-Yule, A. Nicuşan, **M. Herald**, S. Manger, D. J. Parker, Institute of Physics Publishing, (2022), 10.1088/978-0-7503-3071-8

C. R. K. Windows-Yule was contracted to write the book and contributed the majority of the text. A. Nicuşan contributed to the PEPT algorithms chapter and created interactive Python notebooks to accompany the text. M. Herald wrote the chapter covering

the simulation of PEPT experiments. S. Manger was a co-author and contributed to the writing. D. J. Parker provided consultation and information regarding the PEPT technique.

Conference Papers

1. *Optimising Experimental Parameters for Positron Emission Particle Tracking using Monte Carlo Simulation* **M. Herald**, C. R. K. Windows-Yule, IEEE Nuclear Science Symposium and Medical Imaging Conference, 2021

   M. Herald conducted the simulated work, analysed the results, developed the method of optimising tracer activity and detector geometry, wrote the conference paper, and presented a poster at the conference. C. R. K. Windows-Yule provided supervision and editing.

2. *Positron Emission Particle Tracking of Y-90 Microsphere Clusters for Digital Twin Dosimetry* **M. Herald**, A. Nicuşan, J. P. K. Seville, C. R. K. Windows-Yule, IEEE Nuclear Science Symposium and Medical Imaging Conference, 2022

   M. Herald conducted the simulated work, analysed the results, wrote the conference paper, and presented a poster at the conference. A. Nicuşan developed the algorithm for locating clusters. J. P. K Seville and C. R. K. Windows-Yule provided supervision and editing.

Curiosity is a part of human nature

and there will always be

science for the sake of science

Carl David Anderson

Contents

List of Figures

List of Tables

Chapter One

General Introduction

1.1 Positron Emission Particle Tracking (PEPT)

Positron emission particle tracking (PEPT) is a fully three-dimensional experimental imaging technique typically used to track the motion of a positron-emitting tracer over time as it moves through a system of industrial, scientific, or engineering significance [148]. Positrons are particles with the same mass as an electron but with a positive charge [3]. When positrons and electrons collide, they annihilate each other and produce a burst of energy in the form of two back-to-back 511 keV gamma rays, shown in Equation 1.1.

$$e^- + e^+ \rightarrow 2\gamma \tag{1.1}$$

If coincident gamma rays from an annihilation are detected with a position-sensitive radiation detector, then the point where the positron was annihilated should fall along the line that connects the two rays [5]. This line is called a Line-of-Response (LoR). Since the annihilation point lies somewhere along an LoR, this is exploited by PEPT to locate a positron-emitting tracer by using multiple LoRs, typically on the order of 100, collected

during the short time relative to the tracer's velocity and finding the point where the LoRs converge [100]. An example of a typical PEPT tracer is a solid spherical particle < 1 mm in diameter which contains a small amount of a positron-emitting isotope and follows the flow patterns of the system it is placed within. The process of locating a tracer with LoRs is shown in Figure 1.1. When this process is conducted on successive samples of LoRs, the history of the tracer's position over time can be determined. This history is called the tracer's trajectory.

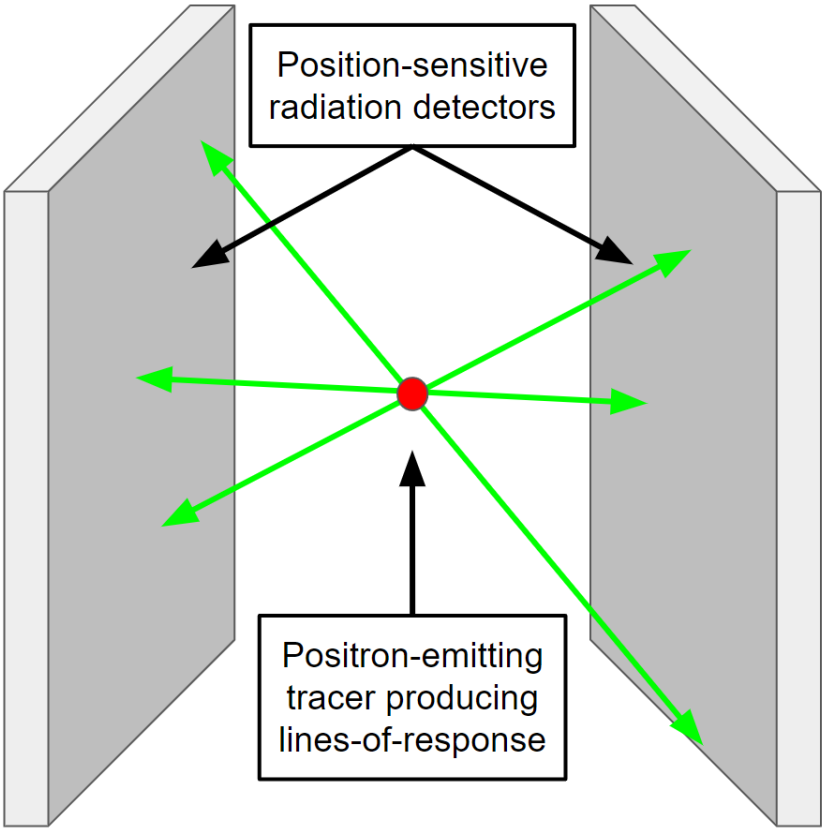

Figure 1.1: Positron-emitting tracer producing LoRs which can be used for PEPT.

Crucially for PEPT experiments, the trajectory of a single tracer can be used to understand the motion of all particles in the system if the system is at steady-state and assumed to be ergodic [143]. The main advantage of PEPT over conventional optical imaging methods is that gamma rays can pass through an amount of material without being significantly

attenuated, providing the ability to image through vessel walls and opaque media [100]. The basic work-flow of a PEPT experiment is the following:

1. Produce a discrete tracer which is labelled with a positron-emitting isotope and place this tracer in a system where the tracer is identical to the other particles of the system or has a Stokes Number « 1 (Section 1.4.1).

2. Record coincident 511 keV gamma rays which are generated by positrons annihilating within a short distance from the tracer by using a position-sensitive radiation detector (Section 1.4.2).

3. Apply a PEPT algorithm on the LoRs to reconstruct the tracer's position over time (Section 1.4.3).

4. Post-process the tracer trajectory to extract information about the system, such as time-averaged velocity fields (Section 1.4.4).

The PEPT technique was first developed at Rutherford Appleton Laboratories and the University of Birmingham in the late 1980s and early 1990s. This technique grew out of early research using a positron camera, initially developed for medical applications, to instead image the distribution and flow of positron-emitting liquid inside industrial systems [43]. From these early experiments, it was realised that more information about the system could be collected if a discrete tracer particle was used and this method could help researchers understand flow dynamics in chemical processes [7]. Following this, an algorithm was developed which could use samples LoRs collected from the positron camera, calculate the tracer position during a short time slice, and then later analyse the history of the tracer's location over multiple time slices [100]. This algorithm, termed the Birmingham Method, described in Section 1.4.3, was tested in a series of experiments with stationary and moving tracers demonstrating its effectiveness [100].

In the years since its development, PEPT has been used to study a wide range of different granular and liquid systems. Some of the most commonly imaged systems include rotating drums, fluidised beds, and pipe-flows [87, 146, 137]. Further, experiments in more complex systems such as coffee roasters, washing machines, and radio-labelled yeast and cancer cells have also been conducted [119, 58, 70, 60].

1.2 PEPT vs PET

Throughout this Chapter, PEPT will be explained in detail: from the physics of positron emission to the types of information that can be extracted from a measurement. However, it is worth pointing out the distinguishing features of the PEPT technique from the far more commonly encountered positron emission tomographic (PET) imaging [75]. PET is a medical imaging technique that uses positron-emitting material to visualise and measure changes in metabolic processes; including measuring blow flow, identifying tumours, or studying brain activity. The major difference from PEPT is that PET is used to quantitatively reconstruct the amount and distribution of a positron-emitting material. This is very different from the information reconstructed by PEPT, since PEPT only returns the discrete position of a tracer [100]. To get a quantitative distribution of the positron-emitting material, two-dimensional or three-dimensional images are produced from the data and treated with corrections for attenuation or random coincidences. These are used to reduce the noise in the images. In PEPT, no such corrections are used, but filtering techniques are applied to help distinguish whether or not data originated from the tracer. Additionally, there are nomenclature differences between the two fields. For example, a PEPT LoR is simply the line connecting two gamma-rays within a coincidence window. In PET, the meaning of an LoR is subtly different. A PET LoR is any possible connection between two areas of the detector system, typically two individual crystals. Further, the detection rate of coincidences for PET is com-

monly in counts-per-second (cps), while PEPT prefers to report this count-rate in terms of frequency (Hz). In this work, the PEPT nomenclature is used.

1.3 Background and Motivation

The purpose of this thesis is to present a general method of simulating positron emission particle tracking (PEPT) and apply these simulations to optimise the technique itself, experiments, and develop new applications. This is an important and significant endeavour because it provides, for the first time, a way to study the technique in isolation; methods for optimising real experiments in terms of the tracers, detectors, and algorithms; and synthetic data can be generated to explore new uses for PEPT. In this way, these simulations of PEPT are an optimisation strategy, a complement to pure experimental work, and a way of gaining deeper insight into the PEPT technique.

The need for simulating PEPT grew out of the realisation that the time and resources of the Positron Imaging Centre (PIC) could be used more effectively if a method of predicting the results of an experiment were developed. Given that experimental geometries and materials of construction are often fixed in engineering applications, this is immensely useful. Over the course of years of running PEPT experiments, some work inevitably produced results which were unusable for a variety of reasons: either the tracer contained too much or too little activity, the detector geometry was not optimised for the experimental system, or the PEPT algorithms were not used with the right settings. Before running experiments, it is not possible to determine what the effect of experimental parameters or settings for PEPT algorithms will have on the quality of data recorded. Therefore, it would be beneficial to produce a simulated model of the experiment and be able to test a range of conditions.

In the following sections, the fundamentals of the PEPT technique, the simulation

software used, the development of a method for simulating PEPT, and the myriad applications for this new approach will be discussed in depth.

1.4 Fundamentals of PEPT

1.4.1 PEPT Tracers

The ideal PEPT tracer has all the same properties as other media inside the experimental system, but also emits positrons at a high enough rate such that it can be located frequently [99, 97]. There are two ways to label materials with positrons: direct or indirect activation [96]. Direct activation is when some atoms of the tracer material are transmuted into positron-emitting isotopes after being placed in a beam of high-energy protons, deuterons, or helium nuclei [32]. Linear accelerators and cyclotrons are capable of directly activating materials containing oxygen, iron, zinc, and others [32]. Alternatively, an indirect activation can be achieved. One method of indirect activation is to irradiate a solution of water enriched with oxygen-18 using protons, which produces fluorine-18. This solution can then be used to coat a tracer [32]. Common types of positron-emitting isotopes used for PEPT, the production routes, and properties are listed in Table 1.1.

Fluorine-18 is one of the most commonly used positron-emitting isotopes in PEPT because of its useful half-life and ability to be readily produced [32]. With a half-life of 109.7 minutes, this is typically long enough to be used in a PEPT experiment over the course of approximately an hour and then set aside to decay to background levels of activity. This is useful because it limits the dose absorbed by personnel and means that the equipment can be used again without lingering activity contaminating future experiments. This is also useful when experiments are conducted with equipment that will be used outside of the lab, such as

Table 1.1: Positron-emitting isotopes used in PEPT, reactions for production, and half-lives.

Isotopes	**Production**	**Half-life**	β_{max} **(keV)**
Carbon-11	$C_6^{12} + p_1^1 \rightarrow C_6^{11} + p_1^1 + n_0^1$ $N_7^{14} + p_1^1 \rightarrow C_6^{11} + He_2^4$	20.3 mins	961
Oxygen-15	$N_7^{14} + d_1^2 \rightarrow O_8^{15} + n_0^1$ $O_8^{16} + p_1^1 \rightarrow O_8^{15} + p_1^1 + n_0^1$ $N_7^{15} + p_1^1 \rightarrow O_8^{15} + n_0^1$	2.0 mins	1735
Fluorine-18	$O_8^{18} + p_1^1 \rightarrow F_9^{18} + n_0^1$ $O_8^{16} + He_2^3 \rightarrow F_9^{18} + p_1^1$	109.7 mins	634
Sodium-22	$Mg_{12}^{24} + d_1^2 \rightarrow Na_{11}^{22} + He_2^4$	2.6 years	544
Cobalt-55	$Fe_{26}^{54} + d_1^2 \rightarrow Co_{27}^{55} + n_0^1$	17.5 hours	1498
Copper-61	$Zn_{30}^{64} + p_1^1 \rightarrow Cu_{29}^{61} + He_2^4$	3.3 hours	1220
Gallium-66	$Zn_{30}^{66} + p_1^1 \rightarrow Ga_{31}^{66} + n_0^1$	9.5 hours	4153
Gallium-68	$Ge_{32}^{68} + e_{-1}^0 \rightarrow Ga_{31}^{68}$	67.7 mins	1899

industrial food processing systems where contamination with radiation is highly undesirable [119].

Positrons are generated on a continuous beta spectrum in which the shape is asymmetrical, shifted towards low energies, and has a well-defined maximum energy [114]. This is because when a positron emitter decays, a daughter isotope, a positron, and a neutrino are produced. These three particles share the energy of the decay, demonstrated in Equation 1.2. Sometimes all the energy is carried by the positron and sometimes it is carried by the neutrino. More often the energy is shared unequally between them.

$$F_9^{18} \rightarrow O_8^{18} + e^+ + \nu \tag{1.2}$$

Positron energy is important because the energy determines how far, on average, the positrons can travel before interacting with a nearby electron. For PEPT and PET, the positron range affects the spatial resolution of the measurement, with a larger range leading to lower spatial resolution [88]. The range is also influenced by the materials which the positron is travelling through. In air, positrons can travel distances on the order of a metre before annihilating with an electron whereas, in water, which is much denser than air, the positrons only travel on the order of a millimetre [21]. In Figure 1.2 the distribution of the positron range in water for carbon-11, oxygen-15, and fluorine-18 is shown.

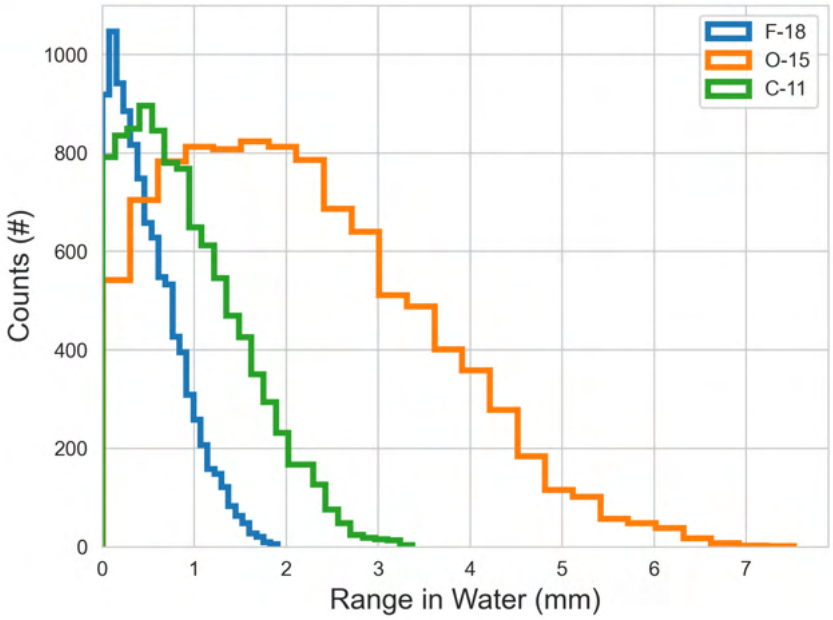

Figure 1.2: The positron range distribution in water for carbon-11, oxygen-15, and fluorine-18 [21].

Fluorine-18 can also be readily produced by a cyclotron such as the MC40 cyclotron located at the University of Birmingham or cyclotrons located at medical facilities [98]. The irradiation of oxygen-18 enriched water only produces other short-lived isotopes which quickly decay, leaving only a solution of fluorine-18 and water. This solution can be absorbed

into anionic exchange resin tracers and coated onto other tracers and left to dry or sealed using an ultraviolet-curing epoxy resin [148]. However, if a tracer is used that does not have the same properties as other materials in the system being studied using PEPT, then the trajectory of the tracer may not be representative. To address this, surface coatings with specific properties can be used. These include hydrophobic or hydrophilic coatings or ones that match the cohesive, frictional, density, or the coefficient of restitution [17].

1.4.2 PEPT Detectors and Systems

Once a tracer has been made, the coincident gamma rays must be recorded for use in a PEPT algorithm. The 511 keV gamma rays produced by positron annihilation are highly penetrating and capable of passing through several millimetres of steel without significant attenuation [129]. This is a major advantage for PEPT since highly penetrating gamma rays make studying flow inside opaque media possible. Typically, PEPT relies on the same positron-sensitive radiation detectors as PET imaging [97]. However, in addition, custom-built detector systems for PEPT have also been produced [71, 39].

The general goal of radiation detection is to convert the interaction of radiation with a detector element into a usable signal. 511 keV gamma rays interact with matter through two main mechanisms: photo-electric absorption and Compton scattering. In photoelectric absorption the full energy of the gamma-ray, $h\nu_0$, is absorbed into an electron, raising the energy of the absorbing electron, E_e, and liberating it from the shell of an atom [29]. This atom later releases a characteristic x-ray as it captures an electron from an outer shell to fill the vacancy in the inner shell. The Equation for photoelectric absorption is shown in Equation 1.3.

$$E_e = h\nu_0 - E_b \tag{1.3}$$

Compton scattering is a form of inelastic scattering that occurs when a gamma ray with initial wavelength, λ, interacts with an electron, but only imparts some of its energy [19]. The ray then changes wavelength, λ', and direction, θ, shown in Equation 1.4. Compton scattering is the dominant mode of scattering for 511 keV gamma rays and thus is important in PEPT [129, 88]. The scattered photon typically still has a significant amount of energy and is detected along with other rays. However, because the direction of one or both gamma rays has changed, the LoR created by these two rays no longer intersects the point of positron-emission, corrupting the LoR.

$$\lambda' - \lambda = \frac{h}{m_e c}(1 - \cos\theta) \tag{1.4}$$

To detect gamma-rays for PEPT, scintillation detectors are commonly used [148]. Only certain types of materials are scintillating and these can be broadly classed into inorganic and organic. Inorganic scintillators are typically crystalline and have relatively high densities which improve their intrinsic efficiency, which is the likelihood of an incident gamma-ray interacting with the crystal. These include sodium-iodide (NaI), bismuth germanate (BGO), and lutetium oxyorthosilicate (LSO) [85]. Organic scintillators contain carbon and can be liquid, crystal, or plastic. These are typically less dense, have poorer energy resolution, and typically are cheaper than inorganic scintillators [106]. Examples of organic scintillators are p-terphenyl, anthracene, and polyethylene naphthalate [106].

The basic design for a position-sensitive detector which can be used for PEPT can be broken down into three main components: (1) scintillation crystals which have a large photon cross-section, (2) photo-multiplier tubes to amplify the scintillation light and localise

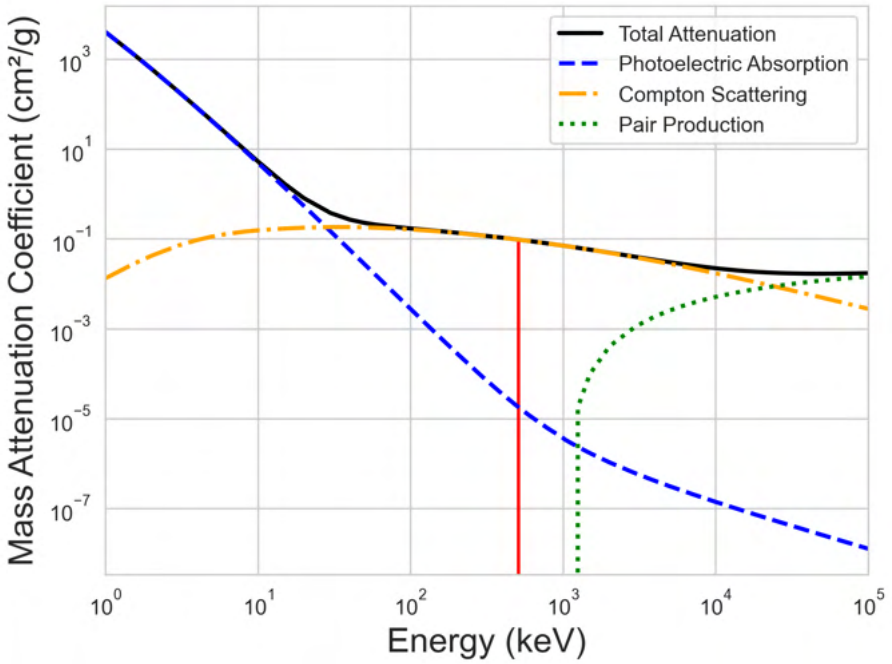

Figure 1.3: The total, photoelectric absorption, incoherent (Compton scattering), and pair-production photon mass attenuation coefficients for water [129].

where the interaction occurred, and (3) fast electronics to quickly process the data and identify coincidences between detector elements to form LoRs [75]. Once scintillation light is produced in the crystal, it must be converted into an electrical signal. This is accomplished using a photo-multiplier tube (PMT) [122]. Light travels through the crystal and enters the PMT where it interacts with the photo-cathode, which releases an electron when the low-energy photon is absorbed. The electron then interacts with the dynode (an electrode operated in a vacuum) charged with a high voltage which releases a large number of secondary electrons when struck by the primary electron, effectively amplifying the original signal [54].

As these detectors process particle interactions, a finite amount of time is needed to trigger the detector, record the event, and return to a state which can accept new events. During this time, called the dead-time, no interactions can be processed and the information is simply lost. There are two main models of dead-time: paralysable and non-paralysable

[134]. In the paralysable model, shown in Equation 1.5, when the detector is in dead-time, any subsequent interactions with the detector fully reset the dead-time. When a detector is placed in an environment with sufficiently high count-rates, no interactions can be recorded since the detector is always in dead-time. The measured count-rate, m, can be found from the expected count-rate, n, using an exponential function and a constant, τ, to represent the dead-time.

$$m = n \exp\left(-n\tau_p\right) \tag{1.5}$$

In the non-paralysable model, shown in Equation 1.6, the detector dead-time is triggered and cannot be reset by new events.

$$n = m(1 + m\tau_{np}) \tag{1.6}$$

This behaviour applies to detecting single events in sections of the detector system, such as a block detector. From these recorded interactions, an additional set of electronics searches for coincidences by looking for two events which were detected within a short time from one another. This time window called the coincidence window (CW) is dependent on the detector but typically in the range of a few nano-seconds to a few hundred pico-seconds [67]. The total coincidence-rate, shown in Equation 1.7, is comprised of true coincidences, scattered coincidences, and random coincidences, Ideally, the CW should be set as low as possible to record only valid coincidences while minimising the chance of recording a random coincidence.

$$m_{Total} = m_{True} + m_{Scattered} + m_{Random} + m_{Background} \tag{1.7}$$

True coincidences are formed when both annihilation photons are detected without being attenuated. Scattered coincidences are formed when one or both of the photons is attenuated, such as undergoing Compton scattering before reaching the detector. Random coincidences are generated when two unrelated events are detected within the coincidence. This could happen when two photons from separate annihilations are recorded within the CW. The amount of random coincidences increases linearly with the CW length. Additionally, the rate at which random events are detected in coincidence, R, is proportional to the detection rate between two independent sections of the detector, S, and the CW 2τ, shown in Equation 1.8 [125]. At high activities, random coincidences can comprise a large fraction of the detected coincidences. Additionally, there are also background counts which come from radiation outside of the FOV (such as cosimc rays) or from noise in the electronics.

$$m_{Random} = 2\tau S_1 S_2 \tag{1.8}$$

The systems used for PEPT come in a variety of shapes and sizes, but the most common geometries are rings and opposed parallel planes. Rings provide the best geometric efficiency, but dual-headed geometries are better for their flexibility in accommodating experiments. Some examples of these devices are the ADAC Forte, the Large Modular Array (LaMA), and SuperPEPT detector systems at the University of Birmingham [103, 72, 39]. In addition, the Siemens Inveon, the Siemens ECAT EXACT HR+ and Siemens Truepoint, and the Siemens ECAT EXACT3D and ADAC Vertex are more examples of PEPT systems [139, 73, 50, 14, 136, 9].

1.4.3 PEPT Algorithms

There are a number of PEPT algorithms, all of which use different approaches to locate a tracer. The first PEPT algorithm was the Birmingham Method [100]. This is still the most widely used PEPT algorithm. However, in the nearly 30 years since its inception, other algorithms have been developed which take advantage of machine learning (PEPT-ML) or Bayesian statistics through expectation maximisation (PEPT-EM) [91, 82]. A brief description of these and other algorithms can be found in Section 3.2. In this section, the Birmingham Method PEPT algorithm, which is the original PEPT algorithm and most widely used, will be discussed along with general considerations for PEPT algorithms.

The Birmingham Method's core assumption is that in a sample of LoRs, there will be a fraction of LoRs which do not lead to the tracer. If these corrupted LoRs are removed from the sample, then the remaining LoRs can be used by finding the point in three-dimensional space which minimises the distance to the LoRs [100]. This is achieved by first locating an initial minimum distance point (MDP) using all LoRs in the sample. Then, individual LoRs are iteratively removed until only a user-specified fraction remains. The lines which are removed are the lines furthest from the MDP. The squared minimum distance from a given point, (x, y, z) to a line can be found using Equation 1.9, where r^2 is the scalar length of the LoR and x_{12} is the distance between the two ends of a given LoR in the x-axis and x_2 is the x-coordinate of the LoR. The total distance, D, between a point and the LoRs in a sample can be found using Equation 1.10. The Birmingham Method is illustrated in Figure 1.4.

$$\delta^2 = \frac{[(x-x_2)z_{12} - (z-z_2)x_{12}]^2}{r^2} + \frac{[(y-y_2)x_{12} - (x-x_2)y_{12}]^2}{r^2} + \frac{[(z-z_2)y_{12} - (y-y_2)z_{12}]^2}{r^2} \quad (1.9)$$

$$D = \sqrt{\frac{1}{N}\sum_{i=0}^{N}\delta_i} \quad (1.10)$$

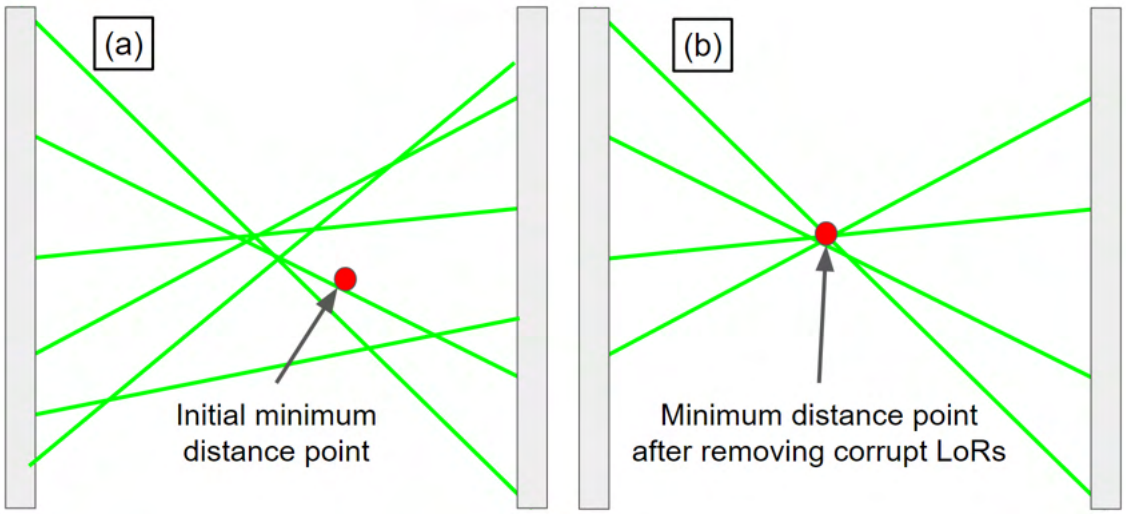

Figure 1.4: The Birmingham Method: (a) The MDP is calculated using all of the LoRs in a sample. (b) The MDP is recalculated after removing the corrupted LoRs.

The point which minimises D is thus the global MDP. There are two parameters that must be set by the user when using the Birmingham Methods, N, the number of LoRs in each initial sample, and, f, the fraction of remaining LoRs used to locate the tracer. Determining the optimal value of these parameters is non-trivial and will be discussed in more detail in Chapter 6. However, general rules can be applied to PEPT data sets. First, the sample size of LoRs should be kept reasonably small in comparison with the tracer velocity. Ideally, during the sample, the tracer should only move a small distance, roughly

comparable to the position uncertainty of the measurement, such that Equation 1.11 holds, where w is an experimentally measured PEPT resolution of the detector system, and v is the tracer velocity. If N is substituted for the LoR acquisition rate, R, and multiplied by the time between the first and last LoR, Δt, then by rearranging the terms, Δt can be found in Equation 1.12

$$v\Delta t = \frac{w}{\sqrt{fN}} \tag{1.11}$$

$$\Delta t = \left[\frac{w^2}{v^2 f R}\right]^{\frac{1}{3}} \tag{1.12}$$

$$N = R\Delta t = f^{\frac{-1}{3}} \left[\frac{Rw}{v}\right]^{\frac{2}{3}} \tag{1.13}$$

However, in general, the tracer velocity is unknown, and f is dependent on a complex relationship between the tracer activity, its position, and scattering of gamma rays [100]. Due to this, it is often easier to conservatively estimate the parameters and apply them uniformly throughout the data set, even if they are non-optimal [100, 18].

There are two things that are important in regards to the performance of PEPT algorithms: (1) locating the tracer or tracers accurately, and (2) locating them often [149]. These are often described as spatial and temporal resolution. Sometimes a combined metric of spatiotemporal resolution is used to convey the performance of an algorithm [149]. Spatial resolution is measured in terms of the distance that the PEPT located tracer is, on average, away from the tracer's actual location, measured through Equation 1.14. However, the real position of the tracer cannot be known absolutely in real experiments. One way this is avoided is to use a stationary tracer, find its mean position with a PEPT algorithm, and then compute the standard deviation of the instantaneous positions. This is often termed the

uncertainty and calculated using Equations 1.15 and 1.16, where $\chi = x, y, z$. It is important to understand which metric is being used when comparing different algorithms.

$$\bar{\varepsilon}_{spatial} = \frac{1}{N}\sum_{i=1}^{N}\sqrt{(x_{real}-x_{exp})^2+(y_{real}-y_{exp})^2+(z_{real}-z_{exp})^2} \tag{1.14}$$

$$\sigma_{\chi} = \sqrt{\frac{\Sigma(\chi-\bar{\chi})^2}{N}} \tag{1.15}$$

$$\sigma_{spatial} = \sqrt{\sigma_x^2+\sigma_y^2+\sigma_z^2} \tag{1.16}$$

Temporal resolution is a measure of how often the tracer can be detected. This is limited by the acquisition rate of LoRs and how many LoRs are needed per sample to locate the tracer. A typical data rate for the ADAC Forte PEPT detector system is approximately 20,000 LoRs per second, or 20 kHz [103, 44]. When 100 LoRs are used per sample, this equates to a temporal resolution of around 200 locations per second or 200 Hz. Since the intervals between LoRs are variable, the sensitivity of the PEPT detector system changes as a function of the tracer position, and the activity of the tracer decays, it is easiest to calculate an average temporal resolution during an experiment [44, 149]. In this work, temporal resolution is calculated following Equations 1.17 and 1.18, where C is the number of LoRs collected during an experiment and t_{exp} is the length of the experiment in seconds.

$$\tau_{temporal} = \frac{N_{locations}}{t_{exp}} \tag{1.17}$$

$$N_{locations} = C/N_{LoRs} \tag{1.18}$$

With regard to the spatial and temporal resolution of a PEPT algorithm, this is dependent on the sample size of LoRs [149]. For stationary tracers, larger sample sizes increase the spatial resolution since there is more information about where the tracer is located, with error decreasing for static tracers proportionally to the inverse square root, i.e. the standard error. Conversely, with larger sample sizes the temporal resolution is decreased since fewer tracer locations can be produced. The trade-off between spatial and temporal resolution is well defined, and shown in Figure 1.5, and can be stated as the spatial resolution, σ, being proportional to the inverse of the square root of the number of LoRs, N_{LoRs}, in the sample, following Equation 1.19, [100, 103, 97]. Combining these two metrics is useful for characterising the performance of a PEPT algorithm since the influence of using different numbers of LoRs, and thus the effect this has on the temporal resolution can be removed. To combine the two metrics, the first important insight is that $\sigma_{spatial}$ multiplied by the $\sqrt{N_{LoRs}}$ equals a constant, shown in Equation 1.20. Then, rearranging the terms of Equation 1.17 to solve for N_{LoRs}, this can be substituted into Equation 1.20 to form Equation 1.22 where the constant is now ξ, the spatial resolution.

$$\sigma_{spatial} \propto \frac{1}{\sqrt{N_{LoRs}}} \tag{1.19}$$

$$1 \propto \sigma_{spatial}\sqrt{N_{LoRs}} \tag{1.20}$$

$$N_{LoRs} = \frac{C}{t_{exp}\tau_{temporal}} \tag{1.21}$$

$$\xi = \sigma_{spatial}\sqrt{\frac{C}{t_{exp}\tau_{temporal}}} \tag{1.22}$$

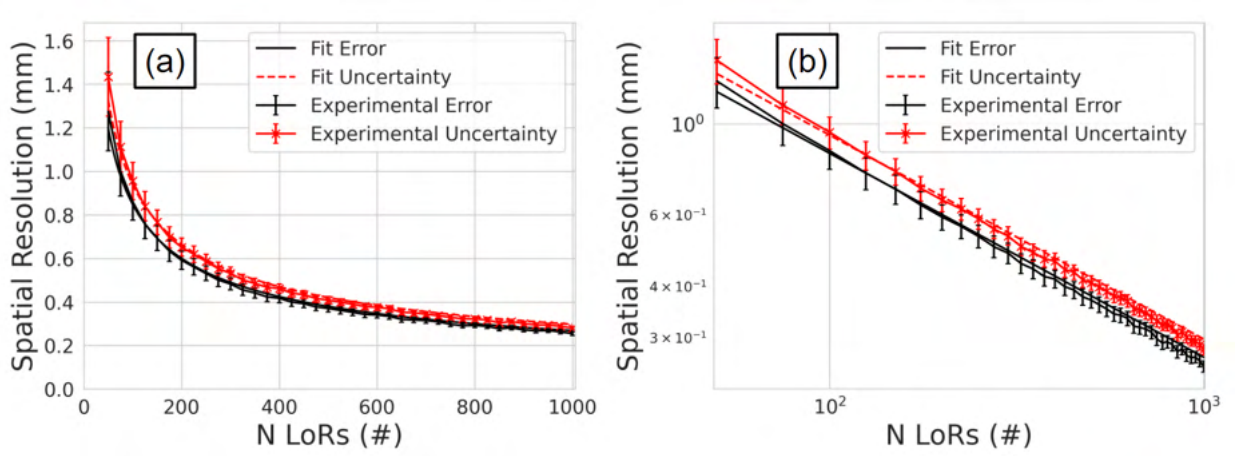

Figure 1.5: (a) The spatial resolution as a function of the number of LoRs used per location for the Birmingham Method with a stationary tracer plotted on a linear scale. (b) The spatial resolution as a function of the number of LoRs used per location for the Birmingham Method with a stationary tracer plotted on a log-log scale [148].

The outputs of a PEPT algorithm are the time, position, and optionally other quantities such as velocity or estimated spatial error [148]. These can be manipulated or further processed to study the dynamics of a system. For example, consider a tracer inside a common household blender. During the experiment, the tracer will circulate with a constantly changing velocity. After using a PEPT algorithm, the tracer trajectory can be visualised either as a time-series plot with each position component plotted against time or as a fully three-dimensional plot. Examples of a time-series plot and a three-dimensional plot are shown in Figures 1.6 and 1.7, respectively. From the plots, there is evidence of periodic motion of the tracer, suggesting the presence of circulatory flow in the system, as would be expected for a blender.

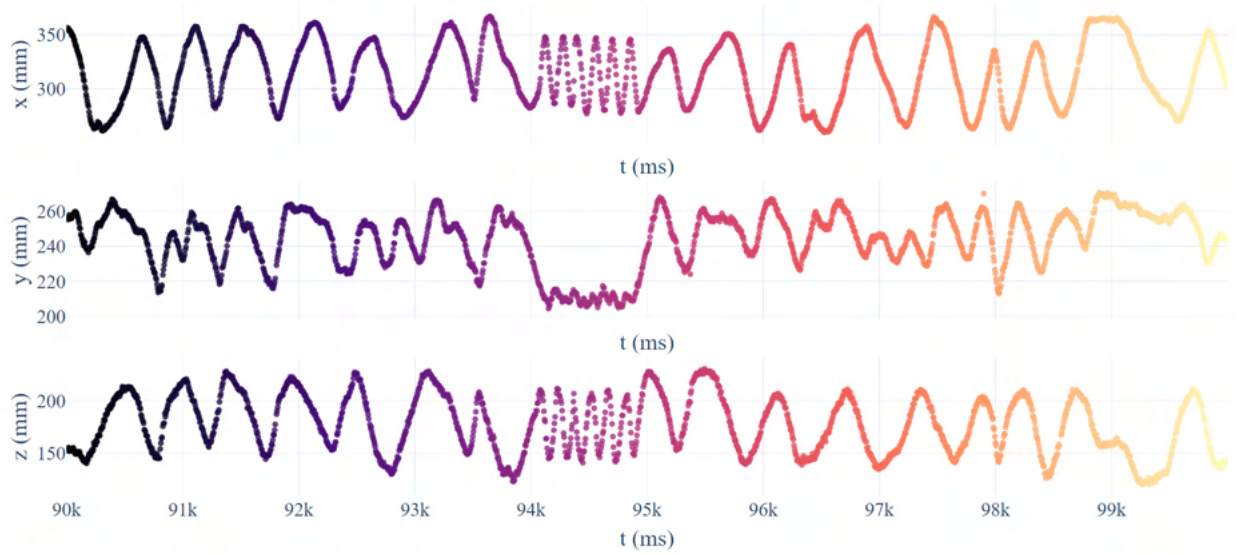

Figure 1.6: An example of a time-series plot showing a tracer moving through a blender coloured by time.

1.4.4 Post-Processing PEPT Trajectories

Once a PEPT experiment is finished and tracer trajectories produced, the next step is to post-process the trajectories in order to extract useful information about the experimental system [149]. A great depth and diversity of information can be extracted from experiments. Trajectories themselves are fully three-dimensional, time-dependent, Lagrangian information [147]. However, most of the useful information produced by PEPT experiments is the time-averaged behaviour of tracers, whether it be one, two, or three-dimensional, including the occupancy of tracers, velocity and acceleration fields, granular temperature, and many other types of measurements [148, 46].

Perhaps the easiest information extracted from a PEPT experiment is the occupancy [143, 148]. In other words, the fraction of the total time that the tracer spends in a particular volume. To perform an occupancy measurement, the volume of the system is divided into cells which define areas of the system of interest [148]. From the trajectory, it can be observed

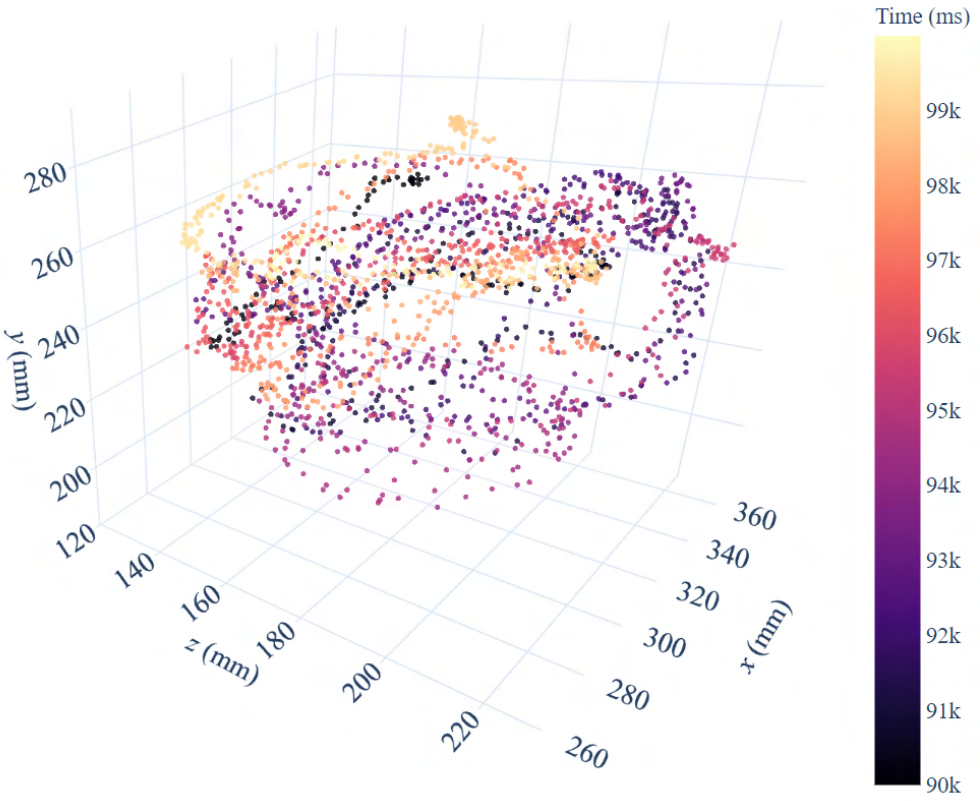

Figure 1.7: An example of a three-dimensional plot showing a tracer moving through a blender coloured by time.

when the tracer enters and exits a volume. This metric of occupancy can be calculated using Equation 1.23, where O_i is the occupancy in a particular volume, $N_{locations}$ is the number of times the tracer was detected in that volume, and V_i is the volume [148].

$$O_i = \frac{N_{locations}}{V_i} \tag{1.23}$$

Simply counting the number of times a tracer is detected within an area of the system is a crude way to calculate the occupancy and is not a valid approach to the problem. This is because the time intervals between successive locations in a PEPT tracer's trajectory are not always the same size due to the variable detection rate of LoRs. This depends on the position

of the tracer, the sensitivity of the PEPT detector system, the amount of attenuation that occurs, and the radioactive half-life of the positron-emitting isotope [44]. In order to improve the occupancy, the time that the tracer spends in each volume must be taken into account on a location-by-location basis using Equation 1.24.

$$O_i = \frac{1}{T_{exp}} \sum_{i=0}^{N_{locations}} \frac{N_{locations_i} \Delta t_i}{V_i} \tag{1.24}$$

Applying this occupancy method to the tracer in a blender, mentioned in Section 1.4.3, the one-dimensional vertical occupancy profile can be produced and is shown in Figure 1.8. By looking at the history of the trajectory over a relatively long time, statistics on the occupancy can be collected. The amount of time needed to run an experiment varies by the size of the volume and the frequency that the tracer enters and exits the volume [144]. A discussion on the time needed to run a particular experiment can be found in Chapter 5.

Moving beyond spatial distributions, another type of information that is useful to extract from a PEPT experiment is the velocity field [148]. The crudest, yet simplest, way to calculate instantaneous velocity is by using the central finite difference method. This is done by calculating the distance between points on either side of the tracer position and then dividing by the time to travel that distance. The number of points, or steps, on either side, can be adjusted depending on how long it is desired to average over, but typically one or two steps on either side are enough [148]. The central finite difference method is presented in Equation 1.25 for individual velocity components. The velocity components can also be combined to produce an instantaneous velocity magnitude through Equation 1.26.

$$v_{x_i} = \frac{1}{2} \left[\frac{(x_i - x_{i-1})}{(t_i - t_{i-1})} + \frac{(x_{i+1} - x_i)}{(t_{i+1} - t_i)} \right] \tag{1.25}$$

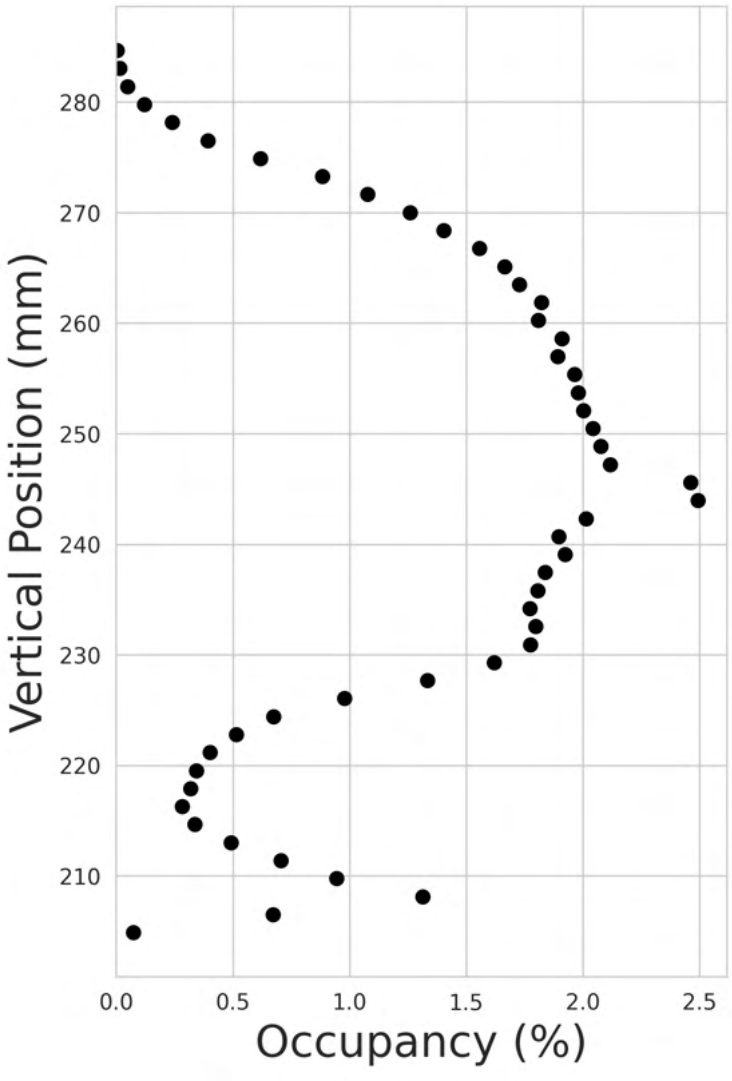

Figure 1.8: An example of a one-dimensional occupancy plot for a tracer in a blender.

$$v_{mag_i} = \sqrt{v_{x_i}^2 + v_{y_i}^2 + v_{z_i}^2} \tag{1.26}$$

By applying this method to the blender example previously used, a velocity magnitude field in the x-axis and z-axis can be produced, also known as a depth-averaged velocity field. To do this, the velocity field in each direction need to be computed, and then combined in a separate step. Similar to the occupancy plot, the system needs to be broken down into different volume elements such that the instantaneous velocities in these elements can be

averaged. In this example, a grid of 30 x 30 cells is used with a cell size of approximately 3.5 mm. The instantaneous velocity magnitudes in each cell are averaged following Equation 1.27 and a velocity magnitude field is produced in Figure 1.9. As expected, the velocities are highest in the centre of the system since this is where the rotating blades are located.

$$v_i = \frac{1}{N_{locations}} \sum_{i=0}^{N_{locations}} v_{mag_i} \tag{1.27}$$

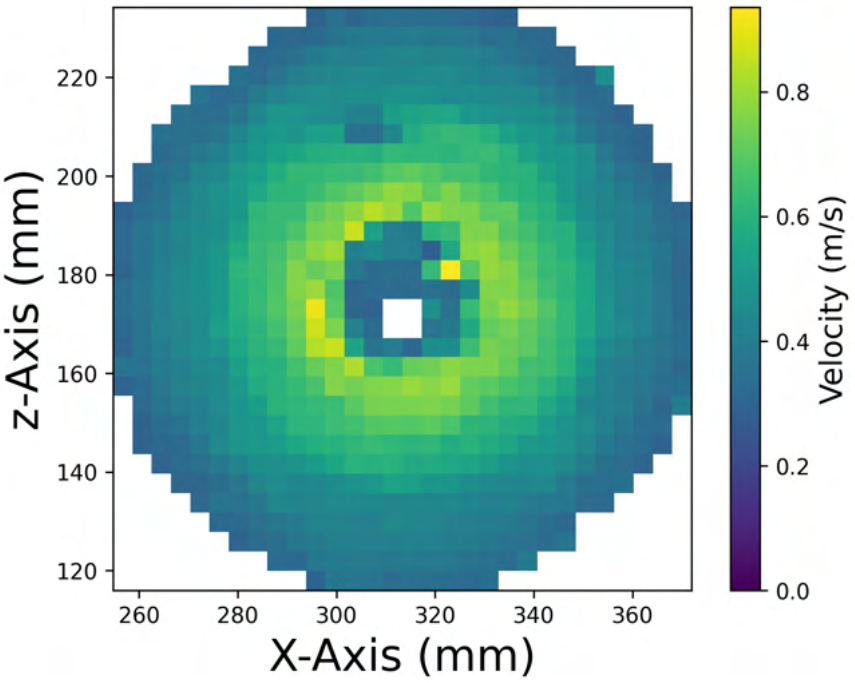

Figure 1.9: An example of a top-down view of a two-dimensional velocity field for a tracer in a blender.

1.5 Simulating PEPT

This section will cover the types of simulation and specific simulation software which will be used in later Chapters. Two main simulation approaches are used (1) The Monte Carlo

(MC) method for simulating radioactive sources and the radiation transport, then (2) the discrete element method (DEM) to simulate trajectories of granular media.

1.5.1 The Monte Carlo Method

MC is a general method of modelling stochastic events and generating probability distributions from a large number of random events [86]. MC methods are typically used in problems such as radiation transport when it is too complex to use deterministic or analytical methods [61]. For example, it is relatively easy to determine the intensity of a perfectly parallel beam of mono-energetic photons passing through an infinitely wide but finitely thick plate of uniform material using the Beer-Lambert Law, shown in Equation 1.28 [10]. However, it is much more complex to model radiation transport inside a nuclear reactor. For these situations, and indeed, PEPT experiments, to determine the probability of particles interacting with an object like a radiation detector it is necessary to use MC methods [112].

$$I = I_0 e^{-\mu \delta x} \tag{1.28}$$

In essence, to use MC to model a radiation transport problem, first a source with a spatial distribution (point-like, linear, planar, or volumetric) is defined which emits particles at specified energies and angular distributions. Then a single particle is generated at a random position within the source and assigned random energy, momentum, and angle from the distributions prescribed. Next, the particle is simulated to move a small increment relative to the spatial and temporal scale of the particular system and physics. At the end of the step, the chance of the particle interacting is calculated using properties which depend on the particular materials the particle passed through during the step. Next, from the previously calculated probability of interaction, it is determined whether or not an interaction actually

occurs. If an interaction occurs, the particle's energy and momentum are recalculated. If an interaction does not occur, the particle's energy and momentum stay the same. These steps are repeated for the single particle until it either is absorbed, exits the simulated system, or reaches an energy threshold beyond which the particle is no longer considered. As new particles are then created and their interaction history is recorded, the probability distribution of interactions becomes well-defined. One benefit of MC simulations is that they can be made arbitrarily accurate by increasing the number of particles simulated [76]. The number of particles that need to be simulated in order to develop good statistics about the interactions occurring in a particular system varies depending on the system size, the number of radioactive sources and the type of radiation to name a few factors [76].

1.5.2 Geant4 Application for Tomographic Emission (GATE)

In this work, GATE v9.1 is used to run MC simulations of radiation transport and also, crucially, to emulate the response of the PEPT detector to the interactions occurring within the scintillation crystal [56]. GATE is built on top of the Geant4 toolkit which is a general MC radiation transport code [1]. Geant4 is primarily used in applications of high energy physics, nuclear experiments, accelerator, and space physics. This software is written in C++ and is open-source. The novel developments of GATE are the introduction of a user-friendly scripting language which makes building a detector system much easier and also introduces expanded functionality to use the recorded interactions of particles to mimic the detector pulse-processing to generate a realistic response [62].

GATE was created by the medical physics community to address the need for an easy-to-use MC simulation software which was flexible enough to model the complexities of medical detector systems and also have the user support and geometry visualisation utilities needed to ensure the simulations were accurate [55]. While other MC simulation software

exists, each has its own drawbacks which ultimately result in them not being detailed or flexible enough to allow realistic simulations [31, 121, 115]. GATE has been used to model a wide range of different detector systems such as PET systems like the Phillips Vereos and General Electric Discovery Single Photon Emission Computed Tomography (SPECT) systems like the Philips AXIS and High-Resolution Animal SPECT Imaging System, and also multi-modal imaging platforms like the Siemens Inveon pre-clinical Scanner [113, 118, 123, 131, 74]. To validate GATE models of these systems, typically the National Electrical Manufacturers Association (NEMA) standards for characterising system performance are used [90].

As a side effect of the research that has gone into producing MC simulation software, the tools have become powerful and flexible enough to be used for research not originally considered. The software GATE v9.1 is used in this work to model both detector systems designed for medical imaging which have now been adapted for use in PEPT experiments, like the ADAC Forte and Siemens Inveon [44, 74], as well as new detector systems which were designed for PEPT, such as the LaMA which is composed of modular boxes of detector which can be arranged around an experiment [39, 97]. In addition to the detectors, experimental systems studied using PEPT, such as blenders, rotating drums, and mills can also be recreated using GATE and PEPT tracers can be modelled and movement representative of their behaviour in these systems prescribed [46, 47].

1.6 Building a GATE Model of a PEPT Detector System

Detectors and experimental systems used for PEPT are complex and often have intricate geometric components, several types of radiation transport physics to consider, and multi-stage pulse-processing chains which all must be captured to a high degree of fidelity to

ensure the synthetic data produced by the model are as close to reality as possible [56, 62, 36]. However, by using the tools provided by GATE it is possible to create a digital twin of a detector system and experiment [115]. In this Section, the process of developing a Monte Carlo detector and models of the experiments and PEPT tracers using GATE v9.1, examples of existing detector models, and considerations when calibrating a GATE model to match experimental results will be discussed. There are six main steps to setting up a GATE simulation: (1) defining the geometries for the detector and experiment, (2) defining the source, (3) setting-up the physical processes of radiation interactions with matter, (4) describing the detector pulse-processing chain, (5) assigning the data output format, (6) and prescribing the acquisition settings.

1.6.1 Geometry

The PEPT technique is applicable to a wide range of systems and thus there is a wide variation in the size of the systems studied, types of tracers used, and behaviour of the tracers [148]. Therefore, methods for replicating the complex geometry of these systems, the different types of tracers used, and the PEPT tracer trajectories are needed. Commonly used PEPT detectors can be arranged in rings, parallel planes (dual-headed), or modular, shown in Figure 1.10 [97]. Similarly, PEPT experiments can be as simple as a cylindrical tank filled with water or as complex as a dynamic flowing granular system in which the volume is a constantly evolving mix of voids and dense regions [18, 87]. Illustrations of these systems are shown in Figure 1.11.

The geometry creation commands for GATE use the same core functions as Geant4, but GATE makes them easier to use through the GATE macro language and introduces several convenience functions like geometry repeaters [56]. The list of built-in analytical volumes includes boxes, spheres, cylinders, cones, ellipsoids, hexagonal prisms, and wedges

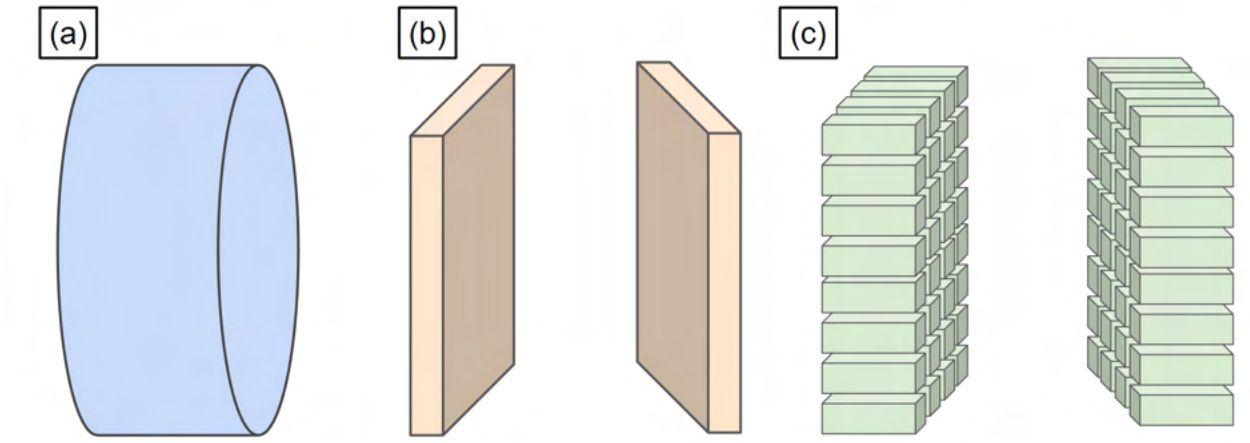

Figure 1.10: Illustrations of a (a) cylindrical ring, (b) parallel planes, and (c) modular PEPT detector system.

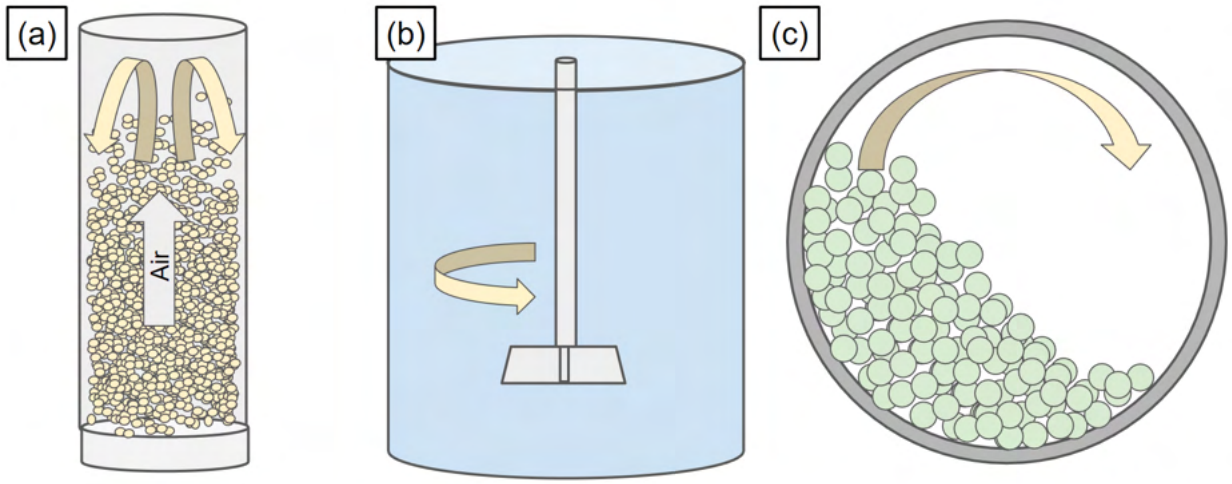

Figure 1.11: An example of a (a) fluidised bed, (b) stirred tank, and (c) rotating drum.

[56]. To use an analytical geometry, the type of geometry must be selected and values for the parameters of the geometry prescribed. For instance, in order to create a sphere, only the radius needs to be defined. However, partial spheres can be made by describing the radius along with the angular span that the partial sphere inhabits. Similarly, boxes are created by defining the side lengths in the x, y, and z directions. Examples of some analytical geometries are shown in Figure 1.12.

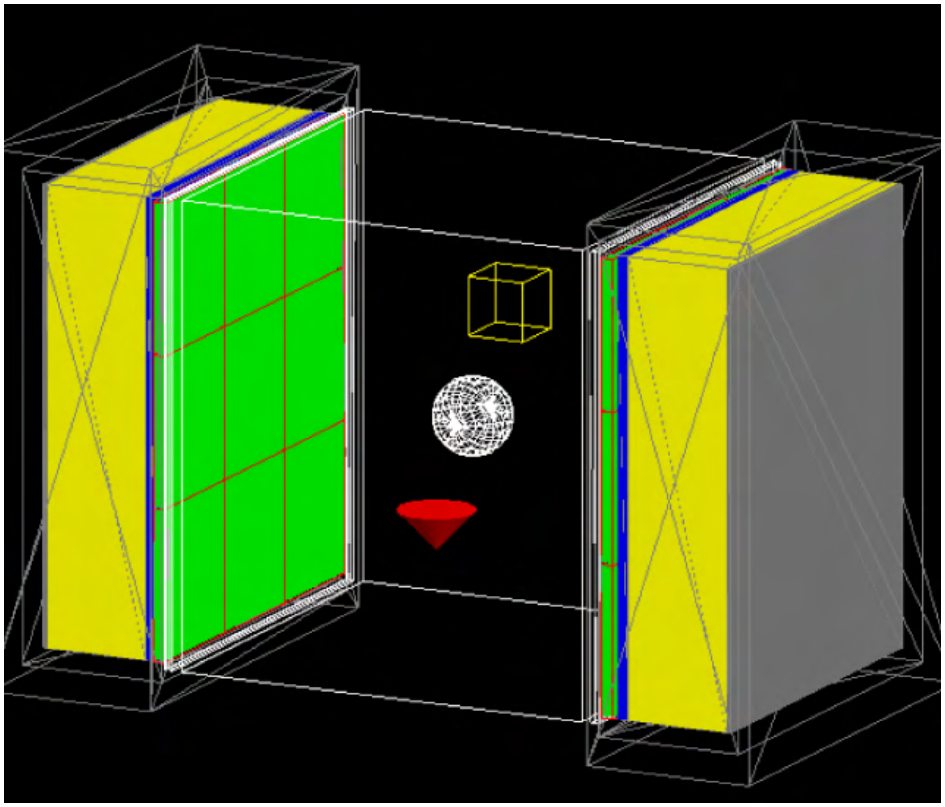

Figure 1.12: A box, sphere, and cone in the field of view of the ADAC Forte.

Additionally, and usefully, there is also an option to import a volume using a tessellated list of vertices, for example, a stereolithography file. Tessellated volumes are created by using groups of three or more vertices to form a surface, then the surfaces are joined to form a coherent volume. Each surface can be a different size and joined to other surfaces which share two or more common vertices. This means that essentially any shape can be created by a tessellated volume and used in a GATE simulation. However, one drawback of a tessellated volume is that only one material can be used per volume. An example of a tessellated geometry is shown in Figure 1.13.

Further, volumes can also be imported as voxelised volumes. Voxels are three-dimensional pixels, usually cube-shaped, and each voxel can be comprised of a different material. This allows for complex geometries with varied material properties to be imported

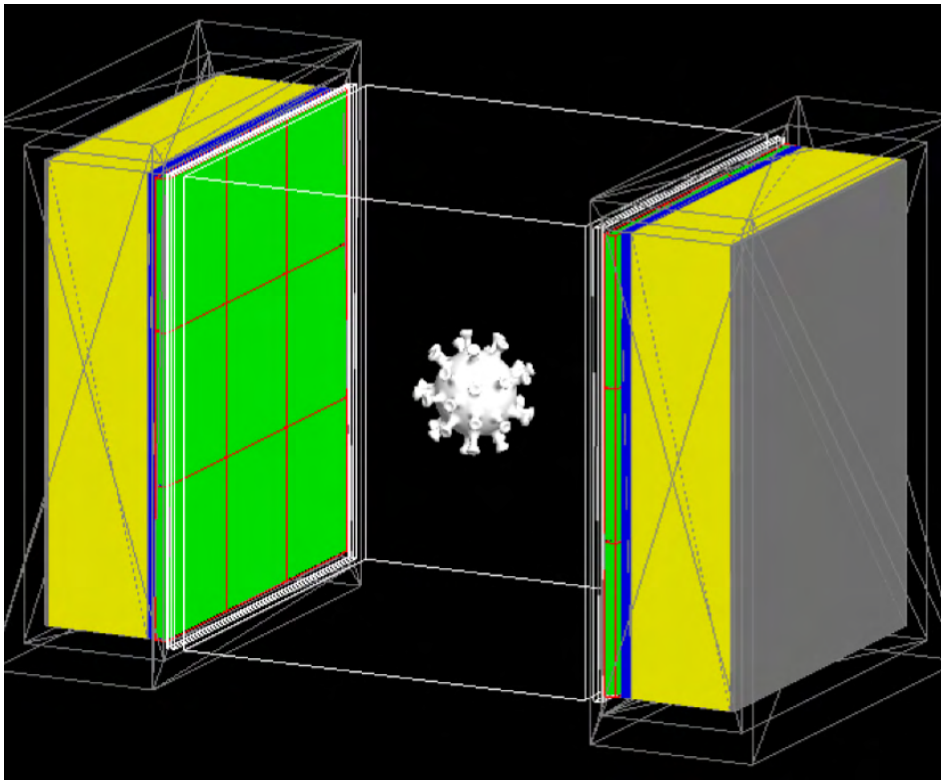

Figure 1.13: A tessellated volume of the coronavirus in the field of view of the ADAC Forte.

as a single volume. The use of voxelised geometries requires a three-dimensional array with the voxel values and an additional file called a 'range translator' which maps the voxel values to GATE materials. Originally, voxelised geometries were developed so that realistic representations of human or animal anatomies could be used in GATE simulations. However, voxelised geometries are also useful for capturing the dynamics of flowing granular systems if the voxel volume is time-averaged or updated as the GATE simulation progresses. An example of a voxelised volume is shown in Figure 1.14.

The first geometry created in a simulation is the 'world', which defines the volume in which particles can be tracked. The world is a box centred at the origin. For all particles, tracking stops when they escape the world. The world can be any size but should be large enough to include all volumes of the simulation while not being excessively large, with excessively large being defined as being more than an order of magnitude of space in the model

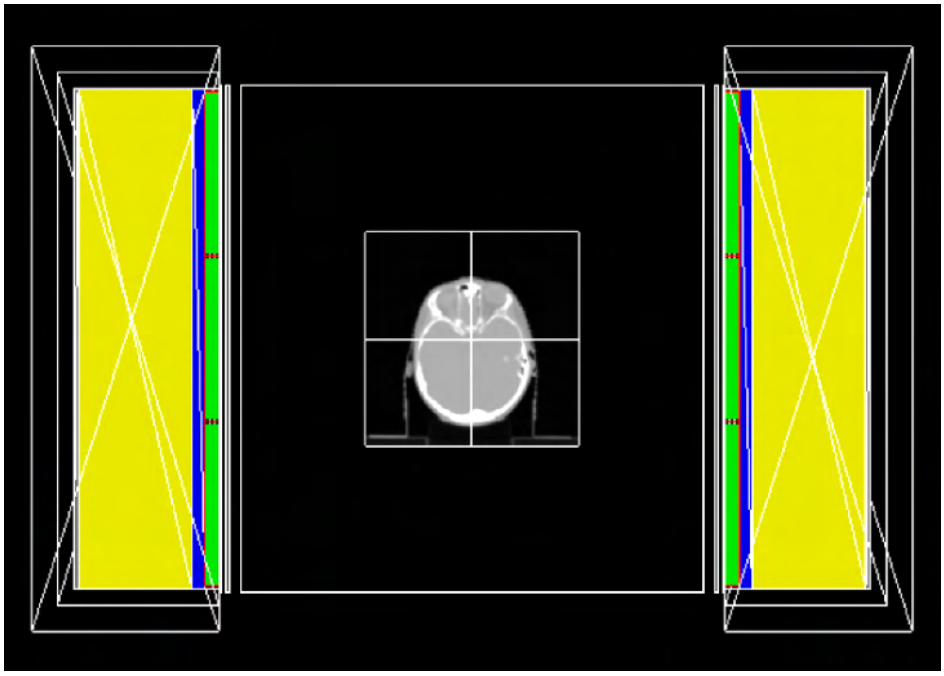

Figure 1.14: A slice of voxelised human head in the field of view of the ADAC Forte.

not containing any volumes of interest [56]. Excessively large world volumes will not affect the results of the simulations, but will slow them down since the paths of particles must be simulated until they reach the edge of the volume, adding additional computation. To define a material, the GATE materials database is used which includes a list of already defined materials common for simulation, such as sodium iodide (NaI), a common scintillation crystal, but can also be updated with custom-made materials [89]. In general, the material name must be defined and prescribed a density, and then the components of the material can be defined with their cross-sections for interacting with particles calculated from a reference table [129].

Once the world is created, to start building a detector, one of the detector architectures must be selected [56]. These architectures describe how many 'levels' are expected in the geometry and what type of data will ultimately be produced at the end of the simulation. The levels of geometry are templates used to help modellers set up the detector geometry

and include common detector architectures such as cylindrical PET systems or CT systems, but also allow for custom architectures [75]. In the cylindrical PET system, there are four levels, with the largest system unit, termed a 'rsector': rsectors, modules, submodules, and crystals. When considering a typical cylindrical PET system, the rsector corresponds to a group of block detectors mounted to the ring in a large panel [57]. The module would be an individual block detector in this panel. The submodule would be considered the scintillation crystals which are optically coupled to PMTs [36]. Finally, each individual crystal element would be the crystal level. Using this architecture makes it relatively simple to use one of GATE's repeater functions to populate the ring with rsectors using a ring repeater, fill in each rsector with block detector modules using an array repeater, and fill in each submodule with individual crystal elements. In this way, only one rsector and its sub-layers need to be fully described and can be copied to create a full detector system geometry. Additionally, with the same geometry commands, an experimental system, such as a mill, rotating drum, or fluidised bed can be replicated [47, 46, 148].

1.6.2 Sources and Tracer Motion

With the detector and experimental geometries defined, sources can now be added to the simulation which emulates PEPT tracers. In GATE, a source is defined as any volume which emits particles. The types of sources used to model PEPT experiments are positron-emitting sources and back-to-back gamma-ray sources [56]. In addition to the source definition, movement of the source volume can also be emulated by updating the source definition each time step. For each time step, all volumes are static but by making the time step small in relation to the source velocity, the appearance of movement is produced similar to how each frame of a movie is static, but at high enough frame rates changes between frames appear smooth.

To model a source, the same geometric commands used in Section 1.6.1 can be used in addition to new commands to prescribe the type of particles emitted, their energy, angular distribution, activity, and any movement of the source. The simplest source which is similar to a PEPT tracer is a point source which emits back-to-back 511 keV gamma rays [148]. To do this the particle type is set to a 'gamma' source. Next, the emission type is prescribed as 'backtoback' as this is already included in GATE and by default produces two particles opposed by 180°. However, positron annihilation does not produce two gamma rays which are exactly 180°, there is always some accolinearity caused by the conservation of momentum when electron and positron are not at rest [88]. To account for this, acollinearity can be prescribed which captures this real behaviour. Typically, the accolinearity in positron annihilations observed in PEPT experiments is approximately $\pm 0.5°$[149]. Next, the energy of the gamma rays is set to 511 keV to match that produced from positron annihilation and the angular distribution is defined as to isotropic to have equal probability of emission for all angles. Further, the shape of the source can be either a point, line, plane (circle, annulus, ellipsoid, square, or rectangle), or volumetric source (sphere, ellipsoid, cylinder, or parallelepiped) [56]. For sources other than a point source, each emission will occur at a random position within the source volume. After this, the activity of the source can be set and the half-life of the isotopes used prescribed, which will determine how many emissions are produced per unit of time and the decay in activity as the simulation progresses. The actual number of emissions per unit of time varies based on a Poisson distribution [56]. However, the mean is the source activity. Finally, the source can be placed at any user-specified position within the world volume.

While back-to-back gamma point sources are simple approximations, real PEPT tracers have a non-negligible volume and positron range [148]. Thus, defining a volumetric source which emits positrons is a more accurate source since it takes into account the real tracer size and positron range. In this case, a 1 mm diameter spherical source will be modelled which

has an energy spectrum equivalent to that of fluorine-18. To create this source, most of the commands are the same, but instead of setting the particle to gamma, a positron source is selected and the type of source is set to a volumetric source instead of a point source. Fluorine-18 emits positrons with a beta energy spectrum characteristic of fluorine-18 and is defined by a maximum energy [88]. GATE already contains the beta energy spectrum for fluorine-18 which is used in this example, but a user-defined spectrum could also be added which contains a list of particle energies and their probability of emission.

It is often the case that PEPT tracers move during the experiment. Movement can be described as either simple analytical trajectories such as when the tracer is attached to a rotating object, creating periodic motion, or described by a more complex trajectory, as would be expected of real tracers inside flowing granular systems [47, 46]. In simple analytical cases, GATE provides a way to prescribe this motion using either constant velocities, constant rotation about an axis, or sinusoidal motion along one axis. To translate a source with constant velocity only the velocity vector needs to be defined. For rotation, the angular velocity and axis of rotation need to be defined. For sinusoidal motion, the amplitude, frequency, and initial phase need to be defined.

Except in rare instances, tracer motion does not follow simple analytical trajectories. For these cases, GATE has the ability to read user-input files describing a trajectory as a series of times, positions, and rotations using the 'Generic Move' command. In this way, any arbitrary tracer trajectory can be simulated [46]. Importantly, it is the user's responsibility to ensure that the time steps of the file and the time steps of the simulation match up such that each new position in the file corresponds to a unique time step. Otherwise, if the simulation time step is larger than the trajectory file time steps, some positions in the file will not be used since GATE only applies the position which is closest in time to the simulated time step. To generate this trajectory file, either user-generated code, a trajectory extracted from a real PEPT experiment, or a trajectory file produced from a DEM simulation could

be used. Using DEM to produce a trajectory file to be read by GATE ensures a trajectory which behaves as close as possible to a real tracer [46]. This will be explored more in-depth in Chapter 5, where it will be shown how a DEM simulation of a rotating drum can be combined with a simulated PEPT experiment to investigate the limitations of the PEPT technique and better understand the trajectory data produced by a PEPT algorithm in comparison to the prescribed trajectory. An example of trajectory extracted from a DEM simulation to be used with GATE is shown in Figure 1.15.

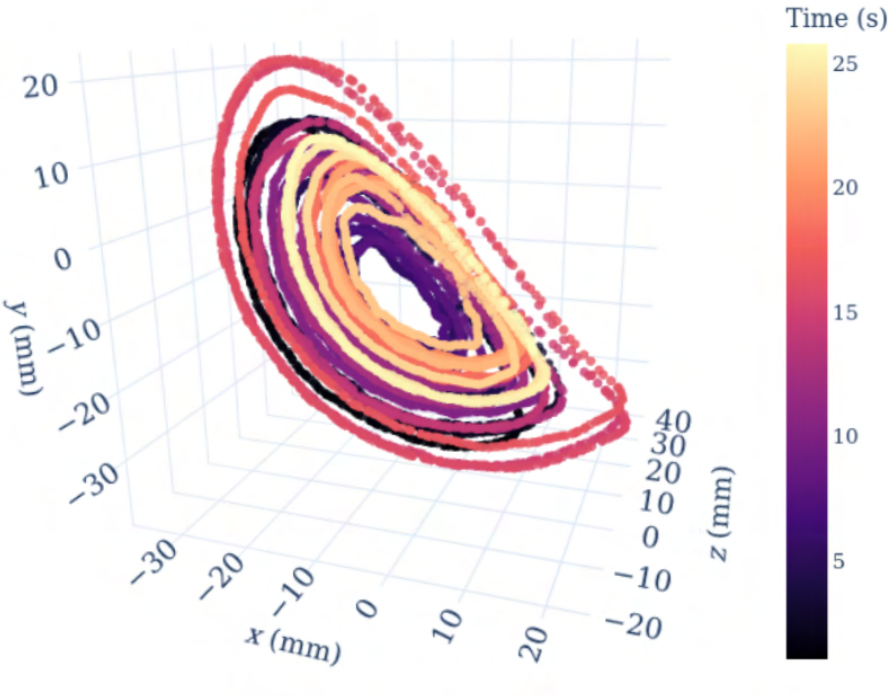

Figure 1.15: An example of a realistic tracer trajectory extracted from a DEM simulation which can be used to move a source in GATE [46].

1.6.3 Physical Processes and Particle Interactions

Once the geometry is designed, the next step is to set up the physics. For a radiation transport problem, such as would be encountered when simulating a positron-emitting source, the

annihilation of positrons with electrons, and the transport of gamma rays through material several physics processes need to be considered which are electromagnetic in nature. GATE also has the capability to simulate hadronic process which would be needed if proton or neutron sources were used, as opposed to leptons such as positrons and electrons [1]. However, hadrons are not a factor in any PEPT experiments to date and will not be discussed here. To simulate the electromagnetic interaction of particles with matter, one of the most important things to consider is the mean free path of a process, λ, which can be given in terms of the total cross section for interacting through any of the possible processes, $\sigma(Z_i, E)$, as a function of the particle energy, E, in Equation 1.29, where n_i is the number density of the atoms, and Z_i is the number of protons in the atom.

$$\lambda(E) = \left(\sum_i [n_i \sigma(Z_i, E)] \right)^{-1} \tag{1.29}$$

The mean free paths for each material are tabulated by GATE at the beginning of the simulation using cross-section libraries so that the step size used to progress a particle through a material is much smaller than the mean free path. For photons, the cross-section library used is the Evaluated Photon Data Library, 1997 [24]. For electrons and positrons, the cross-section library used is the Evaluated Electrons Data Library [108]. Physics processes are imported through GATE's built-in physics lists. These are used to ensure that all the relevant physics processes are added to the simulation. However, the ability to import processes manually is also included for debugging purposes. For most work, including this work, the standard electromagnetic process is sufficient, which is effective between 1 keV and 100 TeV. Since the highest energies of gamma rays are 511 keV and the highest positron energies are in the MeV range, this model is appropriate. The electromagnetic processes this covers, but is not limited to, are the photoelectric effect, Compton scattering, Rayleigh scattering, bremsstrahlung, positron and electron annihilation, and single and multiple Coulomb scattering [1, 65].

1.6.4 Pulse-Processing

Once the physics processes have been set up and the simulation initialised, the next step is to define the pulse-processing chain. This step determines how particles that interact with the detector are ultimately recorded by applying filters and flow logic to the stream of data [56, 62, 36]. The pulse-processing chain in GATE is referred to as the 'digitizer' since it describes how particle interactions are converted into a digital signal which can be manipulated [62]. The digitizer is an important aspect of GATE simulations because it is responsible for converting the history of particle interactions into an output analogous to what would be produced by a real detector system [115]. There are a number of aspects of real detector systems which must be factored into the digitizer model. Generally, these can be considered as three separate categories: (1) limitations in the resolution of the detector response, (2) the readout scheme of the detector, and (3) the trigger logic for pulse-processing.

When a particle is recorded as interacting with the detector in a GATE simulation, the time, position and energy of the particle are known absolutely, but in real systems, there always exists some uncertainty. This uncertainty results in a degradation of the temporal, spatial, or energy resolution and is caused by a variety of factors such as the timing counter frequency, crystal size, and crystal light output [88, 20]. However, in GATE, the root causes of uncertainty do not need to be modelled explicitly, but can rather be treated as the lump sum of many causes and the time, spatial, and energy resolutions calibrated to match experimental results [62]. One way to achieve this is by Gaussian blurring of the particle's information. A Gaussian function, in this sense, blur the information from its simulated value by an amount within the probability density function as defined by a characteristic parameter σ, shown in Equation 1.30. A Gaussian blur is an ideal way to mimic the behaviour of real systems since this describes the probability of recording a value that is the sum of many independent processes, such as measurement errors [79].

$$g(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp - \left(\frac{x-\mu}{2\sigma} \right) \tag{1.30}$$

Additionally, the digitizer needs to capture the readout scheme of the detector which refers to how hits are grouped into pulses. Consider a block detector which has many crystal elements which are connected to four PMTs. The role of PMTs is to convert the scintillation light into a usable signal and locate which crystal an interaction occurred inside. If only one scintillation crystal produces light, the job of the PMT is relatively straightforward in figuring out where the interaction occurred using an Anger logic positioning algorithm [5]. However, if two interactions in two separate crystals produce light at approximately the same time, the detector will only produce one signal with the combined light output used to assign the most likely crystal interaction through Anger logic [5]. Thus, in GATE simulations, it must be defined how to group pulses within crystals that are associated with each other and to define a time limit in which it is possible for two signals to 'pile-up' on one another. The readout policy is defined by using the level depth of the volumes, such that, as previously mentioned when describing the detector system architecture, the readout is applied only within groups of crystals which are a part of the same block detector or similar structure. Further, the time limit in which signals can pile-up on one another is set by defining a pile-up time, typically on the order of nanoseconds [115].

GATE also allows users to describe how pulses are treated by trigger logic. In the types of detector systems used for PEPT, only a certain subset of pulses contain useful information. Thus not every pulse should be recorded. Moreover, when a pulse triggers the detector electronics, there is a period during which the segment of the detector recording the pulse must be reset so that it is ready to record a new pulse [134]. Further, to form LoRs out of the recorded pulses, they must also meet certain criteria. Of the pulses generated by the readout, it is known that the gamma rays produced by positron annihilation have an energy of 511 keV. If pulses have higher energies, they either are multiple rays piled-up on

one another or are generated by some other type of decay process [134]. In either case, these are not useful for PEPT and should be discarded. To do this, two filters are used: (1) energy thresholds and (2) energy windows.

Energy thresholds can be thought of as a coarse filter in which if the light output of the scintillation crystal indicates that the energy of the pulse may be below the minimum energy threshold or above the maximum energy threshold, the detector does nothing and does not record the pulse [62]. On the other hand, if the pulse falls within these thresholds, a detector response is triggered and the pulse's time, position, and energy are recorded. However, if the recorded pulse energy falls outside of a tighter energy window centred on 511 keV, the recorded pulse is not considered for forming an LoR [62]. The values for the energy thresholds are detector specific. Generally, 100 keV and 1 MeV are typical values for lower and upper thresholds, respectively [115]. However, the energy windows are a function of the energy resolution of the detector [89]. The strongest factor in the energy resolution of the detector is the scintillation crystal material. Crystals which are composed of BGO typically have relatively poor energy resolution with a full-width half maximum of 50% at 511 keV; NaI crystals are approximately 15%; and LSO crystals are approximately 5% [89]. For the purposes of PEPT, all of the true LoRs are useful, even if there is a large fraction of corrupted LoRs detected [103]. This means that in practice, the energy windows can be set wide, typically 50% of the 511 keV photo-peak [103]. This is approximately $\pm$127 keV, or a lower energy of 380 keV and an upper energy of 640 keV.

If these criteria are met, the detector records the pulse. At this stage dead-time becomes important. Dead-time is the time spent by the detector during and after detection in which no new pulses can be recorded in that part of the detector [134]. Dead-time modelling is crucial for replicating the count-losses experienced by real detectors, and thus is an important part of the GATE model [62]. There are two main models for dead-time included in GATE, as mentioned in Section 1.4.2: (1) paralysable and (2) non-paralysable

[134]. The model names refer to the behaviour of the detector when new events interact in the scintillation crystal while dead-time is being experienced. In the paralysable model, shown in Equation 1.5, the dead-time is reset with each new interaction. In the non-paralysable model, shown in Equation 1.6, the dead-time is fixed after which a new pulse can be recorded. Of the original events that interacted with the crystal, only the ones which make it through this stage are passed on to further pulse-processing stages to form coincidences [56].

To form a coincidence in GATE, a coincidence processor must be defined. First, a CW must be set such that only two pulses which are detected within a short amount of time can be paired. As shown in Equation 1.8, the rate of random coincidences is proportional to the length of the CW, thus it should be as short as possible, while still being long enough to capture the majority of true coincidences [125]. CWs are a function of the detector electronics, typically their length is similar to the time resolution, and are usually a few nanoseconds to hundreds of pico-seconds in modern PET systems [103, 113]. To form a coincidence, the two singles must also come from separate regions of the detector system. It would be pointless to form a coincidence between two adjacent block detectors, as this coincidence is most likely false. To fix this, geometric constraints are set to ensure the two events are far enough away from each other to potentially be a true coincidence. In GATE, this is done by setting a minimum rsector difference, which excludes LoRs from being formed under this threshold. Moreover, in real experiments, it is often the case that more than two events meet the criteria to become a coincidence. Multiple coincidences can also be formed when more than two events are recorded in the CW. When this happens, a multiple coincidence policy is needed to describe how the real system and GATE model will treat these cases. With GATE there are several options, but the most common policy is to form a coincidence with the pair of events that have the energy closest to 511 keV.

In parallel to the coincidences formed during this stage of the digitizer, which are termed ‘prompt’ coincidences, an additional set of coincidences are formed by using a CW

that is the same length as the original window, but offset in time by an amount such that no coincidences formed in this 'delayed' line could have come from positron annihilation. In real detector systems, this is done to provide an estimate of the random coincidences which could be subtracted from tomographic PET images but does not benefit PEPT measurements [8]. The result is two sets of coincidences that must be merged to create 'multiplexed coincidences' containing both sets. During this merger, an additional, optional, non-paralysable dead-time model can be added which mimics the time needed to write each coincidence to file and a data buffer can be added if the detector system uses a buffer to temporarily cache information about the LoRs to prevent data from being lost while other data is being written [62]. An example of a digitizer model is shown in Figure 1.16.

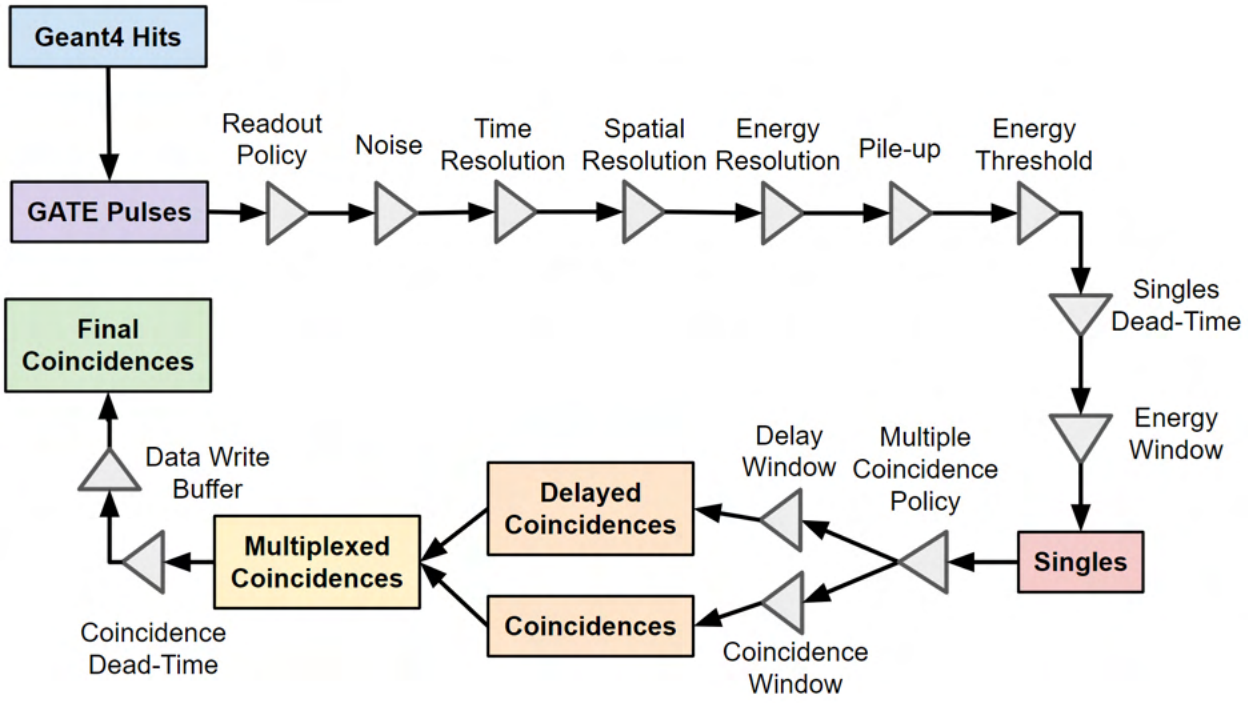

Figure 1.16: An example of the typical digitizer model used in PEPT detectors [45].

1.6.5 Data Output Format and Acquisition Settings

The data output and acquisition settings are key to reproducing realistic synthetic data. These are responsible for ensuring that the end result of the simulation is comparable to that of a real detector and that the size of the time step is small enough to resolve the motion of the tracer. These settings include the type of information which is recorded, data storage considerations, the run time of the simulation, and the size of the time step.

The data output by a real detector system, like the ADAC Forte, are LoRs in a time-sorted list [103]. Each LoR is a line in the list and the columns contain a time step and the x, y, and z positions for both ends of the LoR. An example of a list containing three LoRs is shown in Table 1.2 and illustrated in Figure 1.17. However, in GATE, much more information can be generated alongside this. When comparing experiments to simulations, the equivalent data is always given to a PEPT algorithm, but GATE can also record the particle IDs for each event, the number of scatterings that occurred before the particle was detected, which volumes the scatterings occurred in, and the ID of volume in which the particles were created. This is useful information which allows users to better understand the experimental data. For example, when examining the LoRs, the true, scattered, and random fraction contributions can be determined directly without the use of the image processing techniques which are needed in real experiments where this information is not known [47]. Additionally, the volume IDs where a particle was emitted can be useful when using multiple tracers in a simulated PEPT experiment so as to verify that a PEPT algorithm is appropriately using the correct LoRs to locate each tracer without using LoRs from other tracers in the FOV.

As for the acquisition settings, these are the parameters which prescribe how long to run the simulation and how short the time step should be. The activity of the particle is calculated at the beginning of the time step, thus the time step should not be long relative

Table 1.2: An example of a list of LoRs.

t_1	x_1	y_1	z_1	x_2	y_2	z_2
1	100	100	-250	-100	-100	250
2	0	-100	-250	0	100	250
3	50	0	-250	-50	0	250

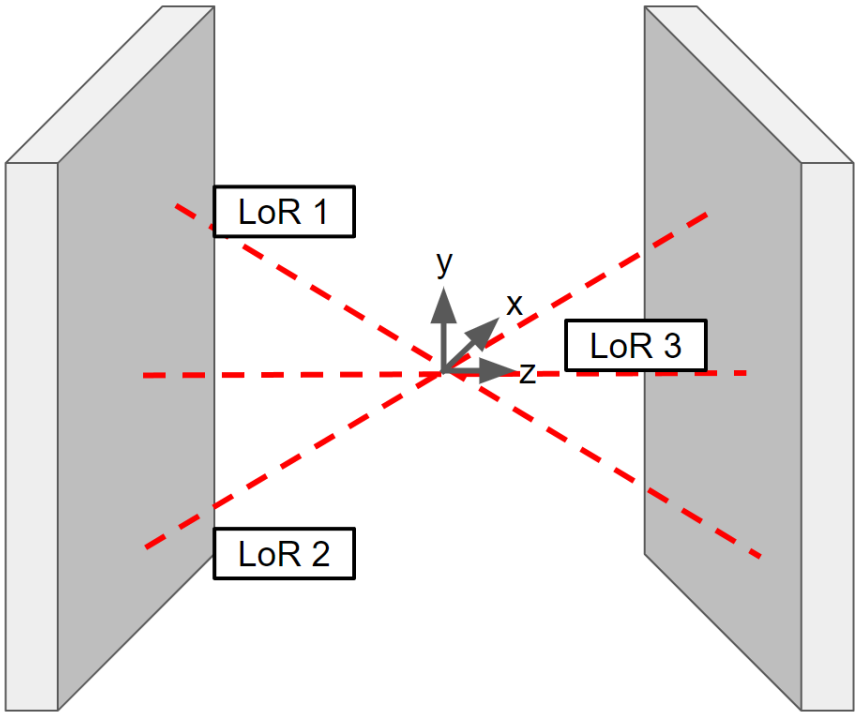

Figure 1.17: The LoRs from Table 1.2 illustrated.

to the decay constant of the isotope used in the experiment [56]. A good rule of thumb is to keep the change in activity less than 1%, for fluorine-18 this is a maximum time step of 96 seconds and for carbon-11 this is 18 seconds. Additionally, if the tracer is moving, it is important to ensure that the time step is small in comparison to the tracer velocity. This ensures that the tracer motion appears smooth since during each time step all volumes are static. Ideally, the tracer should not move more than a tracer diameter per time step

[148, 97]. This means that when using a tracer with a diameter of 1 mm that is moving a maximum of 1000 mm/s, the time slice should not exceed 1 ms. Alternatively, to run times and time slices, the total number of particles emitted can be prescribed. This method is independent with respect to time and is typically only useful when a handful of particles are wanted for visualisation purposes.

1.7 Existing Models of PEPT Detectors Systems

In the years since GATE was first released, myriad detector models have been developed [115]. These were typically designed for medical applications [115]. However, these same systems can be used for PEPT experiments [97]. For PEPT experiments at Cape Town, South Africa the Siemens ECAT Exact HR+ and at Stanford University and the University of Tennessee the Siemens Inveon Pre-Clinical scanner have been used for PEPT experiments [18, 73]. For these detector systems, there also exist GATE models which can be used to better understand the detector performance and optimise experiments [57, 74].

1.7.1 Siemens ECAT Exact HR+

The ECAT Exact HR+ is a typical whole-body ring PET scanner like those typically found in hospitals [57]. Like other whole-body ring scanners, the ECAT offers high sensitivity and spatial resolution. The design of this system is four rings, with each ring having 72 block detectors, and each block detector having an 8×8 array of BGO crystals measuring $4.0 \times 4.1 \times 30$ mm^3, resulting in an 82.7 cm diameter ring and an axial FOV of 15.5 cm [57]. The FOV of the ECAT is useful for accommodating large PEPT experiments but is fairly limited in the axial direction. The measured energy resolution of the detector is approximately 25%. In practice, an energy window of 350–650 keV is used. In order to match the experimental

results for the count-rate response and function of activity a paralysable singles dead-time of 5 μs per block is added. Additionally, random spatial blurring is added within each crystal element. While this digitizer model is not sophisticated, it still manages to capture the general response of the detector. In Figure 1.18, the detector geometry is visualised.

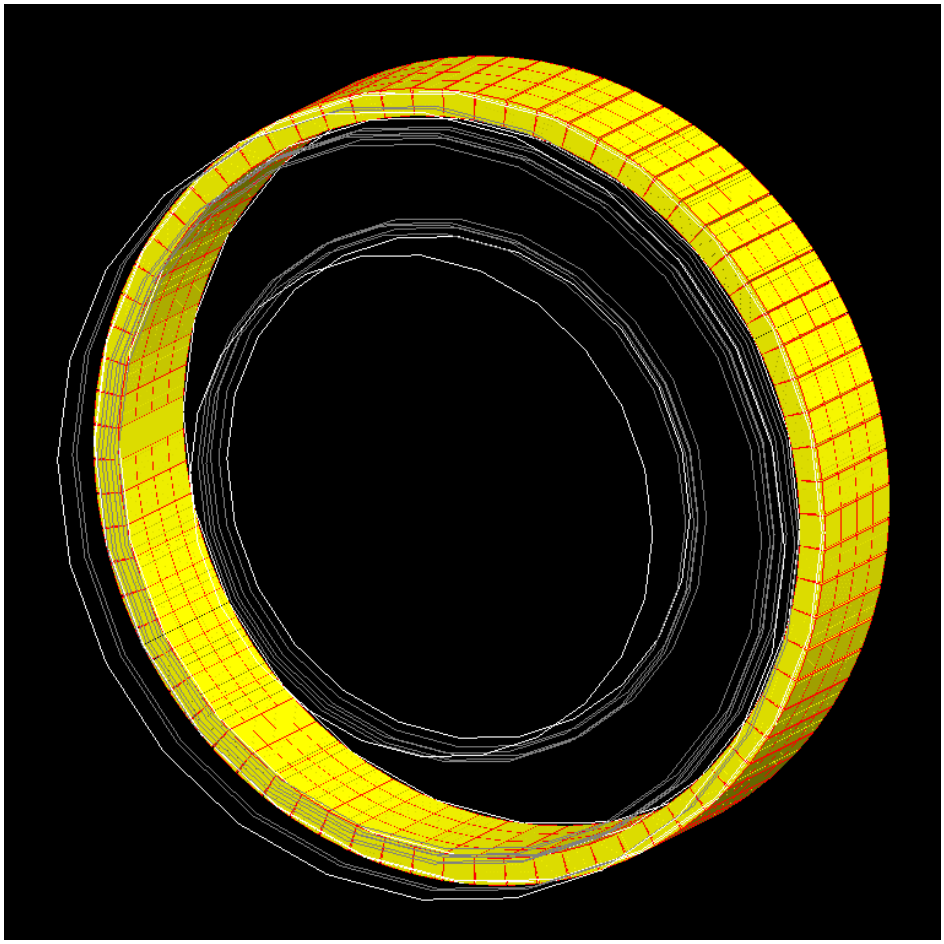

Figure 1.18: The GATE model of the Siemens ECAT Exact HR+ [57].

1.7.2 Siemens Inveon Pre-Clinical Scanner

The Siemens Inveon is a pre-clinical, small animal scanner which was designed such that new imaging treatments and novel PET imaging agents could be tested [81]. The main design difference between pre-clinical scanners and clinical whole-body PET scanners is a smaller

ring diameter and axial FOV. However, because of their smaller size, pre-clinical scanners have better sensitivity and spatial resolution than larger scanners [74]. This is an important advantage when using these devices for PEPT experiments because it means lower activity tracers can be used and higher spatial accuracy achieved [88]. The Inveon has been used for PEPT studies at Stanford University and the University of Tennessee for single-cell tracking and turbulent pipe flow measurements, respectively [140, 73]. These are two types of PEPT experiments that would be very difficult to conduct using larger scanners due to the low amount of activity that can be labelled onto a single cell and the temporal and spatial resolution needed to resolve turbulent fluctuations [70, 137].

To facilitate pre-clinical imaging research, a GATE model of the Inveon was created and made freely available [74]. This model consists of 16 detector modules, with each module composed of four detector blocks. Each block has a 20 x 20 array of 1.59 mm $\times$ 1.59 mm $\times$ 10.00 mm LSO crystals. The overall system possesses 25,600 crystals in a 161 mm diameter ring with a 127 mm axial length. To accurately replicate the contribution from scattered LoRs, a back compartment is included in the model which adds the glass light guide, PMTs, electronics, and a ring of lead shielding to reduce background counts. The digitizer model closely follows a typical PET digitizer and prescribes an energy resolution of 14.6%, energy windows of 250-750 keV, a paralysable singles dead-time of 7 μs per block, and a coincidence window of 3.432 ns. In Figure 1.19 the detector geometry is visualised.

1.8 Calibrating PEPT Detector System Models

Once a GATE model is created, the digitizer model must be calibrated so that the synthetic data is realistic [56, 62, 36]. Before using a GATE model, it must be tested against real experimental data and shown to replicate the experiments to a reasonable degree. The cali-

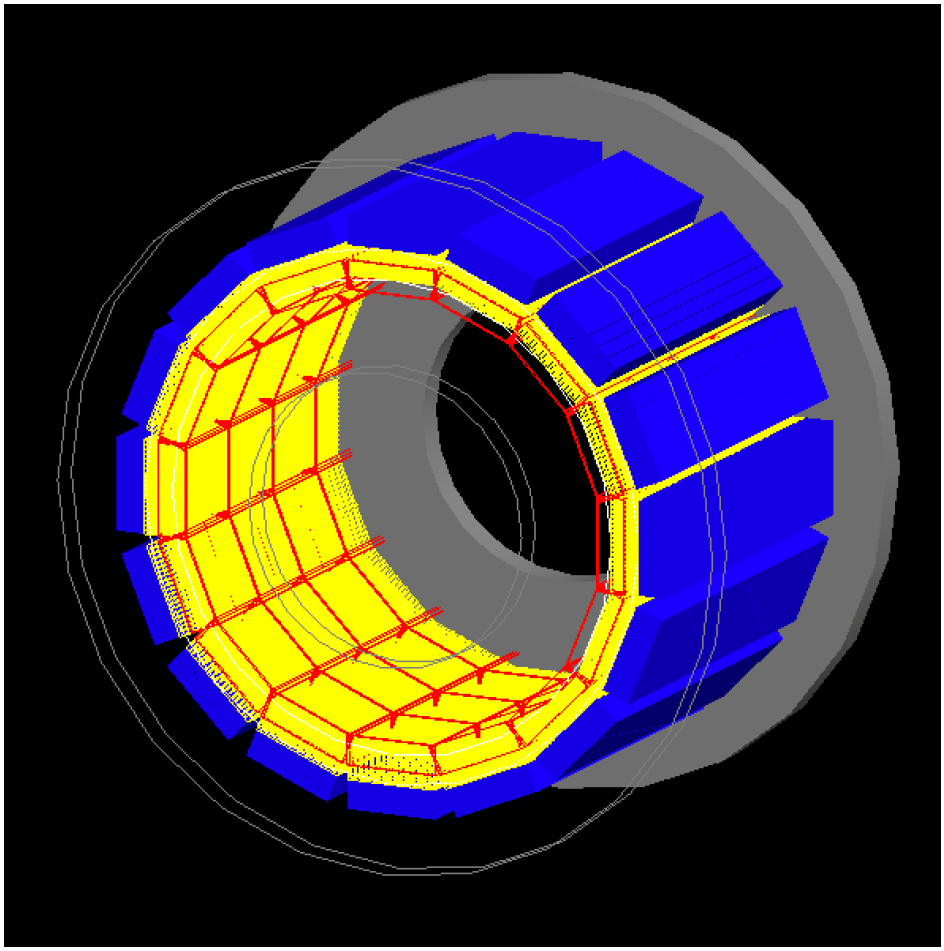

Figure 1.19: The GATE model of the Siemens Inveon Pre-Clinical Scanner [74].

bration of GATE models typically follows a format similar to that laid out in the performance characterisation experiments described by NEMA [90]. These experiments are designed to test the spatial resolution and count-rate response using a series of standard experiments to compare the performance of various detector systems [90]. By following experiments similar to these standards, the relevant metrics for detector performance can be calculated in a reliably reproducible way and used to compare simulated results to that of real experiments [90]. The most important consideration when comparing simulated results to real experiments is to ensure the conditions of the real experiments are accurately replicated [148, 44].

1.8.1 Spatial Resolution

Spatial resolution is a measure of how accurately a point-like source can be resolved with a PEPT detector system. This is different from the spatial resolution observed with PEPT algorithms and instead refers to how blurred a point source appears when a two-dimensional image is produced by binning the detected LoRs on a plane which intersects the source location [88]. The spatial resolution is reported as the full-width half-maximum (FWHM) of the projection image [90]. To test the spatial resolution, point-like sources are placed at specified locations within the detector FOV then the FWHM is measured from the projections.

To conduct a spatial resolution measurement, a small, low-activity source, not exceeding 1 mm in diameter, is placed in specified locations in the FOV. Then, 1,000,000 LoRs are recorded and a planar grid with a cell size of approximately 1 mm is created [90]. This grid is significantly smaller than the spatial resolution and the number of LoRs passing through each cell is counted. Following this, the grid is summed along one axis and the FWHM of the one-dimensional projection is calculated, as shown in Figure 1.20. This process is repeated for the other NEMA-specified positions and the individual FWHM measurements are combined using Equations 1.31-1.34 to produce the spatial resolution. In ring detectors, the axial direction is defined as the vector through the centre of the system parallel to the detector faces. In planar detectors like the ADAC Forte, the axial direction is defined as the vector perpendicular to the detector faces. Additionally, the measurements taken at 1/4th of the FOV refer to locations halfway between the edge of the FOV and the centre of the FOV.

$$Res_{Trans_{Center}} = (Res_{y_{x=0,y=0,z=0}} + Res_{x_{x=0,y=0,z=0}} + Res_{y_{x=1/4,y=0,z=0}} + Res_{x_{x=1/4,y=0,z=0}})/4 \quad (1.31)$$

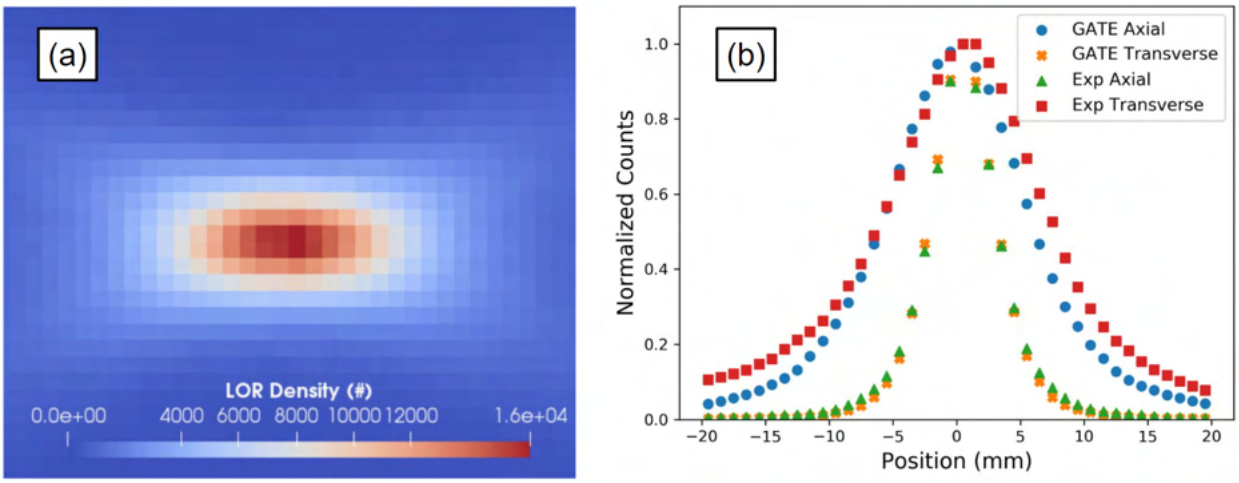

Figure 1.20: A spatial resolution measurement of a point source. Note that Exp Axial and Exp Transverse should be reversed due to an error in the original manuscript. [44].

$$Res_{Trans_{1/4}} = (Res_{y_{x=0,y=1/4,z=1/4}} + Res_{x_{x=0,y=1/4,z=1/4}} + Res_{y_{x=0,y=1/4,z=0}} + Res_{x_{x=0,y=1/4,z=0}})/4 \quad (1.32)$$

$$Res_{Axial_{Center}} = (Res_{z_{x=0,y=0,z=0}} + Res_{z_{x=1/4,y=0,z=0}})/2 \quad (1.33)$$

$$Res_{Axial_{1/4}} = (Res_{z_{x=0,y=0,z=1/4}} + Res_{z_{x=1/4,y=0,z=1/4}})/2 \quad (1.34)$$

1.8.2 Count-Rate Response

The count-rate response of the detector refers to the count-rate curves for total, true, scattered, and random LoRs as a function of source activity. As higher activity sources are used, the effects of dead-time and pile-up induce count-losses and the higher frequency of interactions causes the random LoRs to increase exponentially [134, 125]. The NEMA standards

836 describe an experiment which uses a specialised high-density polyethylene (HDPE) cylinder
837 in which a source is inserted as the basis for calculating count-rate response as a function
838 of activity [90]. This source decays over time and by imaging over several half-lives, the
839 count-rate response as a function of the source activity can be produced. As reported in
840 other work, similar experiments using non-specialised equipment can be used in place of this
841 with little effect on the overall results if an accurate model of the equipment is replicated
842 in GATE [59, 68]. In this work, non-standard phantoms and a slightly modified NEMA
843 protocol are used but is still an applicable approach to validating detector models because
844 the simulated and experimental data are treated the same.

845 To set up a count-rate response experiment, initially, the equipment is filled with a
846 high-activity source. Then as the source decays a projection image is produced at regular
847 intervals and the total, true, scattered, and random count-rates are extracted as a function
848 of the source activity [90]. This is achieved by applying the NEMA protocol to projection
849 images of the source [90]. A demonstration of the workflow for extracting count rates from
850 the acquisition is shown in Figure 1.21. First, samples of a minimum of 500,000 lines-of-
851 response (LoRs) are summed on a two-dimensional grid with a 1 mm cell size. The grid is
852 then collapsed into a line profile of the pixel intensities. All points within ±20 mm of the
853 maximum pixel are summed, and then the background is subtracted by averaging both ends
854 of the ± 20 mm and multiplying by the size of the window. The counts that remain are
855 true LoRs [90, 44]. The total counts are the sum of all LoRs passing through the grid and
856 the scattered plus random coincidence count rate is the total counts subtracted by the true
857 counts. The extracted coincidence count rates are shown in Figure 1.21.

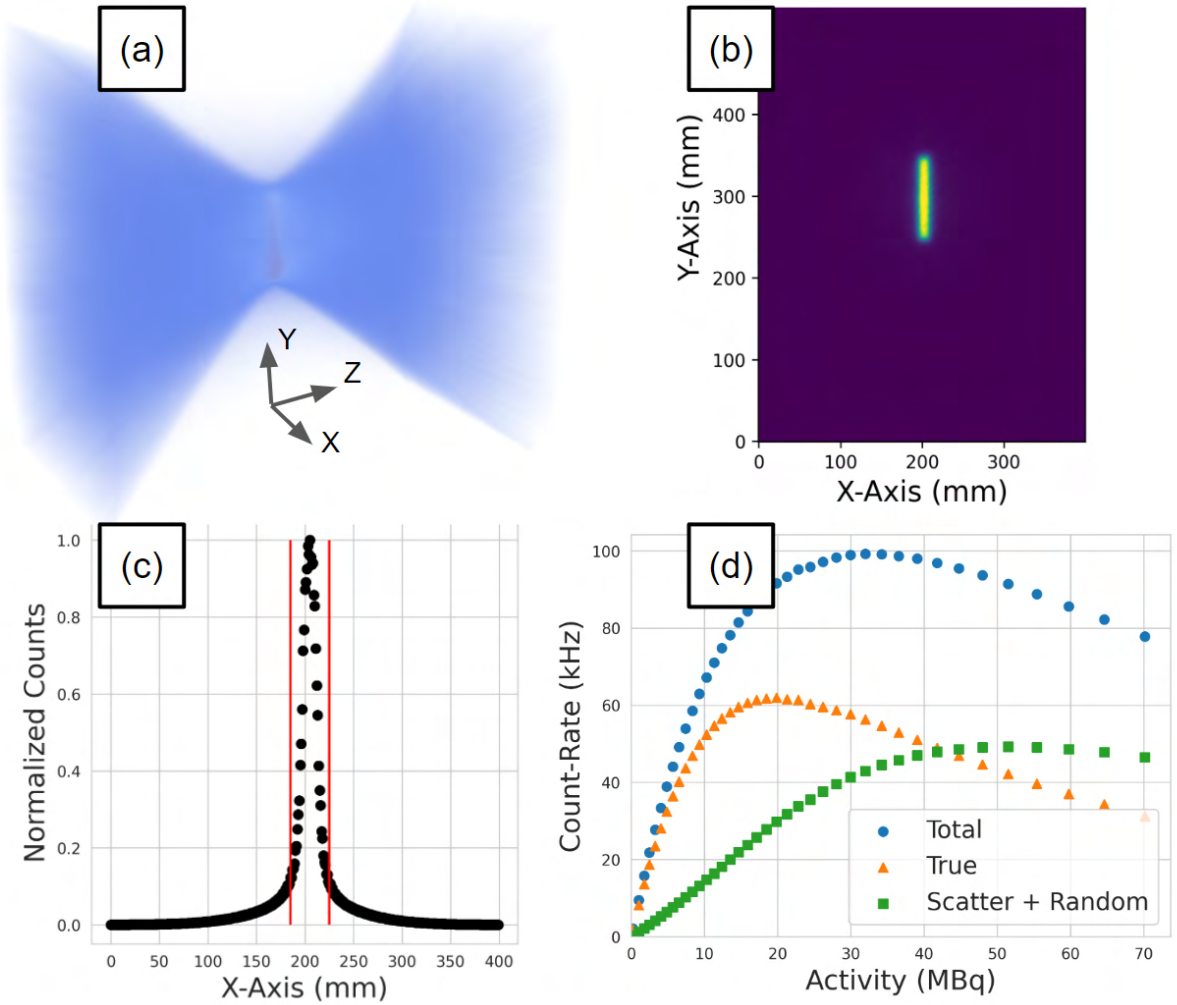

Figure 1.21: A demonstration of the protocol for extracting count rates from a sample of LoRs: (a) a sample of LoRs collected during the experiment is converted into voxels, (b) the slice containing the maximum number of LoRs is extracted, (c) the slice is collapsed into a line profile and the counts in the central 40 mm strip are summed and background counts subtracted to yield a total, true, and scattered + random count-rate (d). Steps a-c are repeated for multiple samples to generate the count-rate response as a function of activity [45].

1.8.3 Calibration Methods

The calibration of GATE models refers to tuning the parameter values of the digitizer such that a simulation and real experiment are in reasonable agreement with one another. To calibrate a model, first, the performance characterisation experiments for spatial resolution and count-rate response are conducted both for the real detector system and for the GATE model and then compared [90, 44, 45]. Often these initial results do not match perfectly, with the match improving by tuning the parameters of the digitizer model with a calibration method.

Calibrating GATE models of PEPT detector systems is a multi-parameter and multi-objective optimisation problem since each spatial resolution measurement and each of the count-rate response curves is in themselves an objective and there are numerous free-parameters in the digitizer model [45]. The typical method of calibration is to use the known properties of the detector provided by the manufacturer or by fitting models to count-rate experiments [36, 113, 127, 132]. Many steps in the digitizer model correspond directly to measurable properties of the detector, such as the energy resolution, dead-time, or time resolution [113]. Values for these properties are often provided by the manufacturer and serve as a reliable starting point, but manual tuning is still needed to match the simulated and experimental response of the detector due to variation between each detector [44, 127]. This is done by adjusting parameter values, re-simulating the characterisation experiment, and assessing the agreement with the simulation iteratively until an optimal match between the simulated and real performance characteristics is achieved [44, 45].

Additionally, model-fitting approaches to tuning the digitizer parameter values can be used [36, 118]. For instance, since the spatial resolution is predominately determined by the spatial blurring prescribed in the digitizer model only this parameter needs to be fitted [62]. In order to calibrate the GATE model's spatial resolution, a range of spatial

blurring values can be tested and the optimal value selected by minimising the differences to the experimentally observed spatial resolution [44]. However, when using a model-fitting approach with the count-rate response, manual tuning can be avoided, but other challenges arise. For example, when fitting a dead-time model to the count rates, this relies on having both the singles and coincidence count rates, which may be available, and also involves fitting simplified models to the detector response, which may not capture the complexity of a real system [134].

1.9 Discussion

In the sections above, the background for the work presented in this thesis has been presented which includes a general introduction to PEPT, information on the simulation software used, a description of methods to simulate PEPT experiments, and considerations for creating and calibrating GATE models of PEPT detector systems. PEPT is a powerful experimental technique, but through complementing experiments with GATE simulations, it will be shown that the technique can be more fully understood, its use optimised, and the technique extended to infer new types of information from a system.

PEPT detector systems are complex machines which have intricate components, multi-stage pulse-processing, and non-linear responses to source activity [56]. However, the ability to create faithful recreations of PEPT detector systems, such as those described in Section 1.7, has created an opportunity to be able to simulate the PEPT technique itself [44]. Therefore, to support the work in this thesis and to simulate PEPT with systems that are currently being used in real experiments at the PIC or other PEPT facilities, in Chapter 2, two new PEPT detector system models created using GATE are presented. These are the ADAC Forte and the LaMA [44, 72, 97]. To calibrate the models, performance characteri-

sation experiments are conducted and then replicated in simulation, with a novel method of calibrating the models using an evolutionary algorithm used [45].

In Chapter 3, GATE simulations are used to benchmark the performance of various PEPT algorithms in terms of both their spatial and temporal resolution by using a common simulated data set where the position of a PEPT tracer is prescribed [149]. This work is a first-of-its-kind effort to bring together PEPT researchers from across the globe to examine the relative strengths and weaknesses of the wide range of PEPT algorithms available. The performance of individual algorithms has been assessed in the past, but these were conducted in isolation, using different systems, and varying methods of performance characterisation [100, 73, 139, 138, 91, 94, 82]. Additionally, the use of a common data set and knowing the exact position of the GATE-simulated PEPT tracer provides a rigorous and quantitative measure of performance.

In Chapters 4, 5, and 6 novel methods of using GATE simulations of PEPT experiments to optimise aspects of real experiments are demonstrated. These methods include the optimisation of the tracer activity and the head separation of the ADAC Forte PEPT detector system, estimating the optimal length of an experiment to collect adequate statistics of tracer behaviour, correcting experimentally measured quantities like the granular temperature, and establishing parameters of the Birmingham Method PEPT algorithm which maximise the spatial and temporal resolution of trajectories [46, 47]. While there exist methods of optimising PEPT experiments using real data, the time and resources needed to do this are impractical. GATE simulations are thus presented as a useful new tool to conduct this type of analysis which can be conducted before equipment is moved to a PEPT facility and which can extract more information than could be generated experimentally.

Chapter 7 will discuss the development of a novel PEPT methodology to determine the size of a PEPT tracer as it loses material to the surrounding system through attrition,

dissolution, or other processes. This is an application for PEPT which has previously been unable to do, representing a useful extension of the technique. This is demonstrated through simulated work, where a tracer loses material and activity to the surrounding medium then the f parameter of the Birmingham Method is used to infer how much activity has been lost, with this information used to infer the tracer's size. Importantly, this novel measurement of tracer size can be generated alongside the time and position of a tracer. If this could be used in real experiments where the loss of particle material occurs, then not only could the loss of material over time be analysed, but also the areas where it occurs could be identified. The results of this work are an important first step towards realising this new capability and demonstrating how GATE simulations can be used to verify new PEPT methodologies without expending vast amounts of time and resources on physical experimentation.

Finally, in Chapter 8, the sum of the work presented in this thesis will be tied together, the future outlook for simulating PEPT experiments discussed, and conclusions presented.

Chapter Two

Development of New Detector Models and a Novel Calibration Method

2.1 Introduction

The simulation of PEPT is dependent on the availability of a validated GATE model of the detector system. This ensures that realistic synthetic LoRs can be generated and comparable data produced from real experiments. This Chapter seeks to introduce two new GATE models of PEPT detector systems which have been created, calibrated, and validated. These two detector models are for the ADAC Forte and Large Modular Array (LaMA), systems which have been, and still are, used at the University of Birmingham's Positron Imaging Centre (PIC) for PEPT experiments [103, 102]. These detector systems are similar to those previously mentioned, but with key differences that make them particularly suited for PEPT. The main difference is that both the ADAC Forte and LaMA systems can change their geometry. The Forte is a dual-headed gamma camera and each head of the detector can be rotated around the gantry ring and the separation of the heads can be adjusted [103]. This results in a flexible system which can change its geometry. Additionally, the LaMA is

a modular system composed of up to 48 'boxes' which contain four block detectors each [97, 48]. These boxes can be placed in any configuration around a PEPT experiment, allowing for unique geometries. A feature of these flexible geometries as will be demonstrated in later Chapters is that the geometry can be optimised for individual experiments. A demonstration of these geometry features for each of the two PEPT detector systems is shown in Figure 2.1.

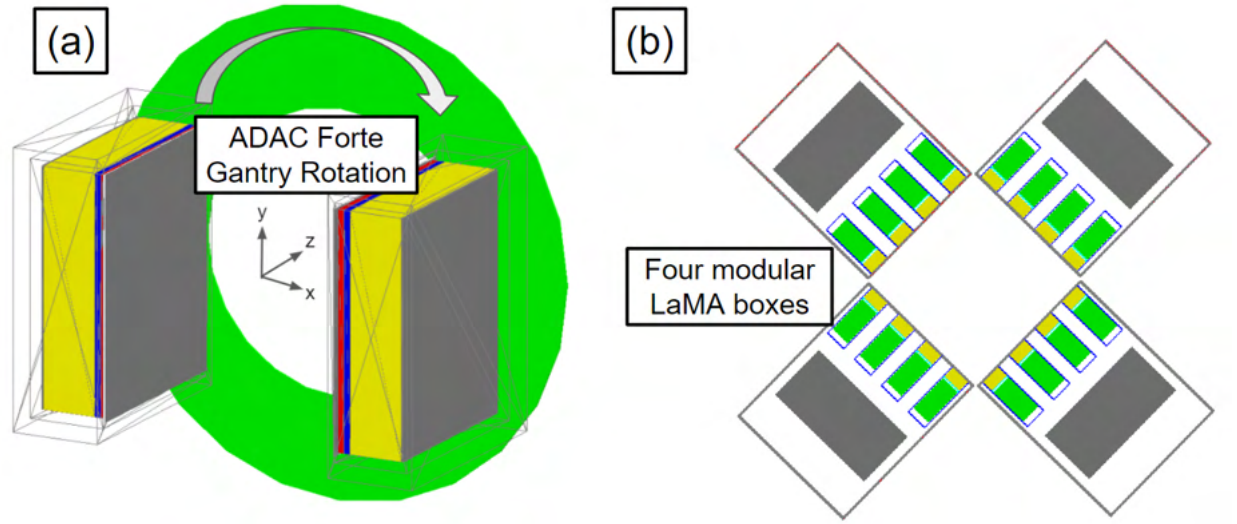

Figure 2.1: (a) Demonstration of the rotation of the ADAC Forte detector head around the gantry ring (the heads can also move independently), (b) four modular boxes of the Large Modular Array.

Due to this high degree of flexibility, they are well suited for PEPT because the size and geometry of experimental systems can vary greatly [148]. However, this poses a challenge for GATE models in that the models must be able to be rapidly changed in order to match the specific configuration of the detector for a given experiment and still produce realistic simulations without each version of the GATE model needing to be calibrated and validated, which would limit their usefulness.

To do this, a new method of GATE model digitizer calibration has been developed

[45]. In this method, described in Section 2.2, first, the performance characteristics are measured, and then a candidate set of digitizer parameters are applied to the model. The performance characteristics from the simulation are measured and compared to the real characteristics. Next, an evolutionary algorithm is used to generate a new set of candidate parameter values and compare their ability to reproduce the real performance characteristics [51]. After several iterations, the algorithm learns how to best replicate the real performance characteristics in an autonomous way, eventually reaching a calibration that matches the experiment as best as possible. Moreover, this is achieved efficiently and without bias by avoiding manual tuning.

The following Sections will cover the development of the novel digitizer calibration method which is then applied to calibrate two PEPT detector system models in Sections 2.3.1 and 2.3.2.

2.2 Model Calibration using Evolutionary Algorithms

In the previous Chapter, existing GATE models of detector systems and methods of achieving calibration have been covered. However, these methods can take a considerable amount of time to produce a calibration, could be seen as arbitrary, and are not guaranteed to produce the optimal calibration [115]. Here, a new calibration procedure is proposed which leverages recent advances in high-level problem-independent algorithmic frameworks to perform an efficient optimisation of a GATE model's digitizer [51]. When trying to calibrate any type of simulation's free parameters, it is often useful to test a range of conditions and assess how the tested parameter values replicate the measurement. In the simplest case with only one free parameter, the value that minimises the error to the measurement can be easily found and visualised by plotting candidate parameter values and the error as a two-dimensional

scatter plot. This can also be extended to two free parameters by plotting the error as a third dimension on the plot. Beyond three dimensions, the relationship between the parameters becomes non-intuitive in large part because they cannot be easily visualised. For these problems, an optimisation algorithm is needed to efficiently test a range of parameter values, determine the difference between simulation and reality, and then converge to the optimal values. However, in simulations and experiments, there often exist noisy measurements, thus a function defining the difference between experiment and simulation will be non-smooth and potentially have many false local minima [51]. This means that gradient-based optimisers are ill-suited for calibrating simulations [135].

In these difficult optimisations, evolutionary algorithms excel [51]. Evolutionary algorithms are a type of bio-inspired computing which mimics natural selection. For example, in a population where individuals have a randomised set of genes and selective pressure is exerted, only the individuals which have genes that enable them to survive will reproduce. Due to this, the next generation of individuals will be more adapted to selective pressure [27]. Similarly, when an evolutionary algorithm is applied to a model with free parameters which can be tuned, the parameter values act as genes, a model with a specific set of parameters is an individual, and a group of individual simulations is a generation [51]. For each generation, a cost function determines an individual's fitness which acts as a selective pressure for the next generation. A flow diagram of how an evolutionary algorithm can be applied to digitizer calibration is shown in Figure 2.2.

Applying the same logic to GATE simulations, from a large population of individual GATE simulations, each with a variation in the parameter values for aspects of the digitizer, some sets of candidate solutions will more realistically replicate the response of the detector. The most successful candidates can then be used to improve the next generation. After multiple generations, the parameter values should converge to the optimal values which best replicate the real performance characteristics [45]. This method is expected to offer

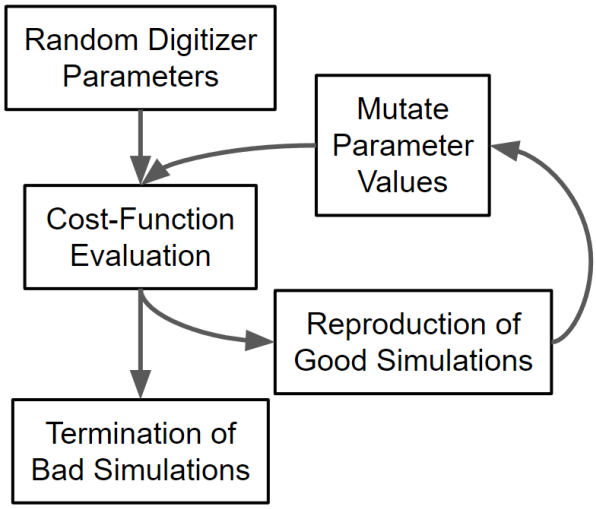

Figure 2.2: The flow diagram for ACCES is applied to optimising free parameters in a digitizer of a GATE detector model [45].

improvements over previous calibration procedures since the optimal parameters are chosen by directly comparing the performance of the optimised digitizer to experiments and removes the need to perform manual calibration, which eliminates user bias.

The evolutionary algorithm used for calibrating a GATE model's digitizer is the Covariance Matrix Adaptation Evolutionary Strategy (CMA-ES), which is a stochastic optimiser for non-linear non-convex numerical optimisation [41, 42]. While there exist several types of evolutionary algorithms, the CMA-ES algorithm is used since it performed well in a comparative review of optimisation algorithms and there is a well-documented Python implementation CMA v3.0.3 [111, 42]. With this approach, candidate solutions are generated following a multivariate normal distribution for each free parameter. A particular advantage of this is that the underlying optimisation function does not need to have a continuous response. The addition of stochastic 'mutations' to the inputs tried to mimic the injection of new genetic material in the population allows CMA-ES to escape local minima which gradient-based optimisers are prone to falling into [135]. Due to the complex relationship

between parameters in the digitizer, it is expected that there will be many local minima, making this approach more likely to be successful than a gradient-based optimiser [62].

One problem in using an optimiser with software like GATE is that GATE was not designed to be used this way. In order to interface simulation software with the existing CMA-ES optimiser and extend the types of problems it can be used with, a Python library has been developed which is called the Autonomous Calibration and Characterisation via Evolutionary Software (ACCES) v0.2.2 [92]. The purpose of ACCES is to use a Python script defining the simulation to populate the user-defined free parameters, populate the script with candidate solutions generated by CMA-ES, then autonomously re-launch the simulation, analyse the results, and generate improved candidate solutions until a termination criterion is met [92, 42]. By supplying ACCES with a cost-function evaluation of the results of the simulation, the fitness of the candidate sets of solutions to parameters in the digitizer model can be compared [92]. In this case, the cost function is the absolute percent difference between the total, true, and scattered plus random coincidence count rates determined by the NEMA performance characterisation experiments [90].

ACCES needs only the bounds of the search parameters, and the number of individuals in a population, and stores the results after each generation, or 'epoch', so that the optimisation state can be restored at any point. The default implementation of CMA-ES requires the use of a single initial standard deviation for all parameters, assuming that all parameters have comparable value ranges and sensitivities. ACCES scales the parameter values by 40 % of each parameter's allowed range, such that parameters of vastly different scales can be optimised together. As parameter combinations are drawn from normal distributions, an initial standard deviation of 40 % naturally covers the entire parameter range.

In order to allow the use of complex, potentially thread-unsafe simulations written in

different programming languages, ACCES launches each simulation as a completely separate OS process, which is either scheduled by the kernel to be run locally on a shared-memory machine, like a laptop or using an external workload manager to launch jobs on multi-node clusters. ACCES offers improvements over other interfaces to optimisers in that it is fault-tolerant and designed for high-performance computing. In this study, ACCES automatically sets up and launches batch jobs for each parameter combination to be evaluated using GATE. To summarise, the two critical CMA-ES configuration parameters are automatically determined by the computing resources available and the possible parameter ranges, such that no manual adjustments of optimiser settings for a given problem is necessary.

In the next Sections, ACCES is applied to GATE models of the ADAC Forte and LaMA to calibrate the model in comparison with experimentally observed performance characteristics. However, an initial model of the ADAC Forte was developed before this calibration was available. In this case, both the model developed through manual calibration and the new model calibrated using ACCES will be discussed.

2.3 New Detector Models

2.3.1 ADAC Forte

Background

In 1999, the University of Birmingham acquired the ADAC Forte dual-headed gamma camera for PEPT experiments [103]. The Forte is the most extensively used PEPT detector system at the PIC, having a history spanning over 20 years [97]. The basic feature of the Forte is two large-area NaI crystal scintillators measuring 470 mm x 590 mm and 16 mm in thickness [103]. The two crystals are optically coupled to a glass light guide and 55 PMTs each which

are used to locate a gamma ray's point-of-interaction within the crystal [5]. Each crystal is contained within a 'head' which includes the electronics, aluminium casing, and lead shielding. Each head is mounted on two adjustable arms, such that the distance between each head can be adjusted from 250 mm to 800 mm of separation. This adjustment allows the geometric efficiency of the camera to be adjusted, as well as allowing equipment of varying sizes into the field-of-view (FOV). A custom data acquisition system records LoRs [103]. The system can produce a maximum of approximately 100 kHz of LoRs, which when using a sample of size of 100 LoRs per location results in a temporal resolution of PEPT tracer of approximately 1 ms [103]. This is sufficient to easily track tracers moving in excess of 1000 mm/s [97].

Manual Calibration

To begin developing the GATE model, a geometric representation of the detector is created using the commands discussed in Section 1.6.1. Each head of the detector is modelled as a rsector and copied with a ring repeater to form two heads [44]. Importantly, the distance of the rsector to the origin of the system and its rotation angle can be varied using a single parameter for each. This allows for the head separation and the rotation angle of the detector to be easily changed, reflecting the behaviour of the real system. Since the scintillation crystal covers a wide area and multiple interactions can be recorded in different regions of a single head simultaneously, the crystal layer is implemented as nine rectangular pixels, with each measuring approximately 170 mm x 126 mm [103, 44]. The pixels are not components of the real detector system but are needed to capture the previously mentioned behaviour. A back-compartment of the glass light guide, PMTs, and electronics are added to account for scattering [109]. Additionally, there is lead shielding around the back and sides of the crystals to reduce background counts. The gantry ring is modelled as an aluminium cylinder with an outer diameter of 1600 mm and an inner diameter of 700 mm with a thickness of

165 10 mm. A description of the geometry model is listed in Table 2.1. The GATE physics list
166 used for the model is the 'emstandard' list. A side-by-side view of the real ADAC Forte and
167 the GATE model of the detector is shown in Figure 2.3.

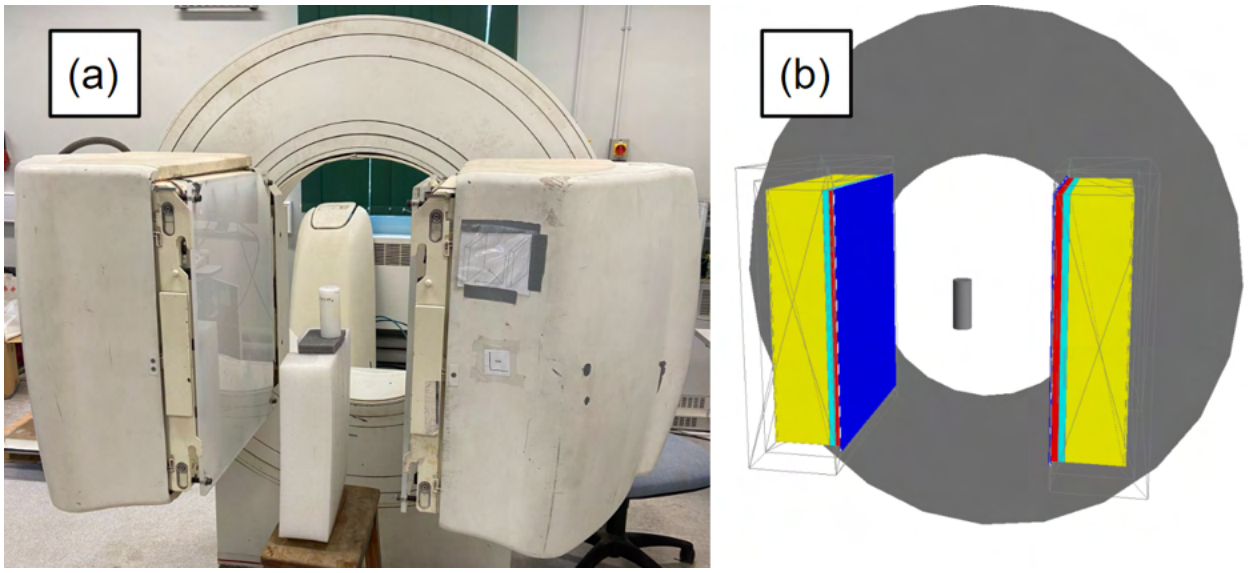

Figure 2.3: A side-by-side comparison of the (a) real ADAC Forte and (b) GATE modelled ADAC Forte with a model of the count-rate experiment in the FOV [44].

Table 2.1: ADAC Forte geometry description [44].

Geometric Model	**Values**
Detector Dimensions	590 x 470 x 16 mm^3
Useful Detector Area	510 x 380 mm^3
Dual-Head Separation	250 - 800 mm
Detector Crystal	Sodium Iodide
Number of Pixels	9
Pixel Dimensions	170 x 126 x 16 mm^3
Gantry Ring Diameter	1600 mm outer, 700 mm inner
Gantry Ring Thickness	10 mm

The characterisation experiments described in Section 1.8.1 and 1.8.2 were conducted to provide a basis to calibrate the model. These two aspects of the detector system were calibrated independently. For the spatial resolution test, the source is a 0.2 mm diameter anionic exchange resin bead. Anionic exchange beads selectively absorb fluorine-18 ions from a solution of water and fluorine-18, volumetrically activating the bead with activity [32]. After soaking in the solution for approximately 30 minutes, the resin bead activity was measured with a well counter to be approximately 27.8 MBq ±0.1. The resin bead was then placed in the FOV.

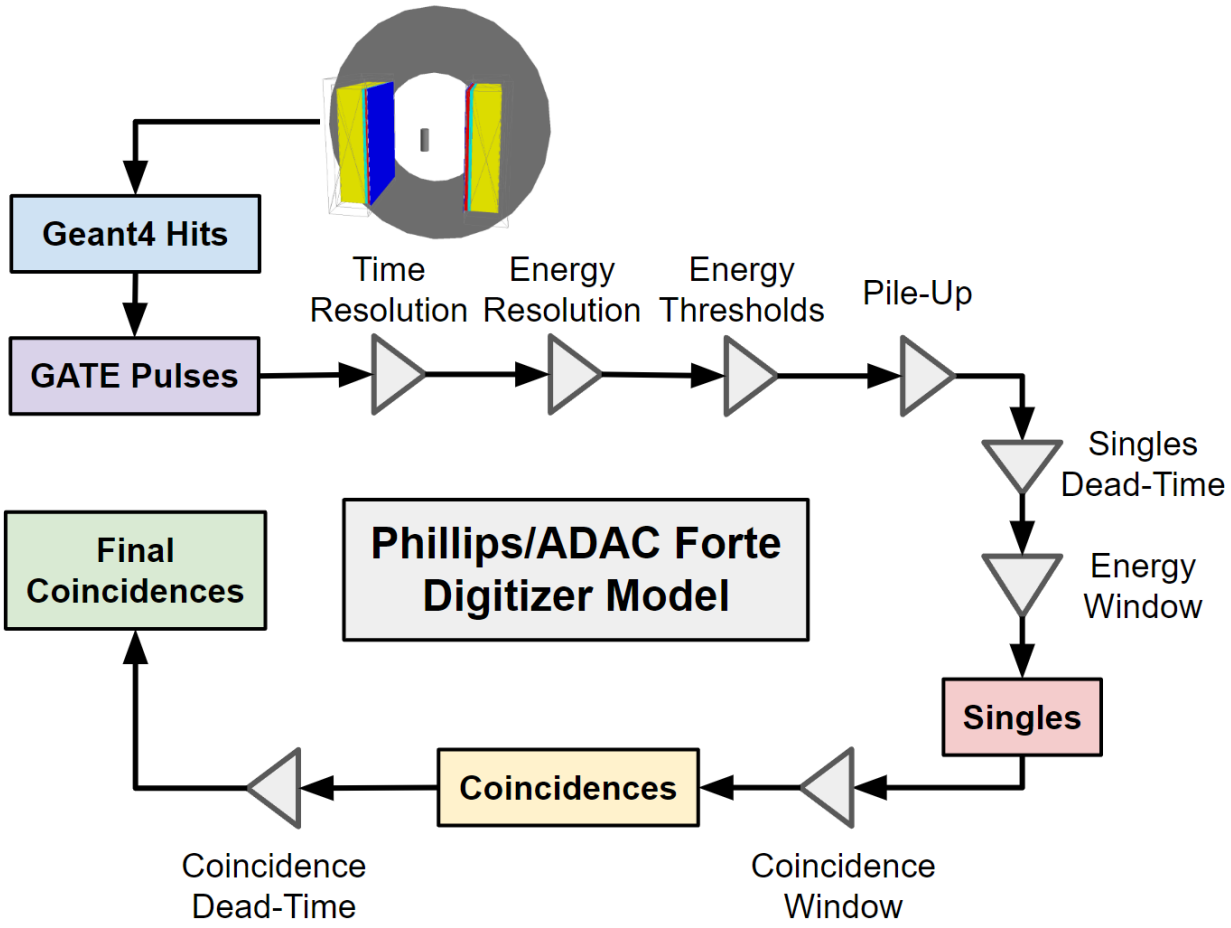

Figure 2.4: A diagram of the ADAC Forte digitizer model [45].

A GATE model of the spatial resolution source was designed to recreate the experimental geometry, with the source modelled as a 0.2 mm diameter plastic sphere, encapsulated

in a spherical plastic phantom 2 mm in diameter. This outer sphere represents the ultraviolet-activated epoxy and provides sufficient material for positrons from the source to annihilate instead of escaping into the air around the tracer while also not contributing significantly to scattering. By the end of the experiments, the source activity decayed to approximately 19 MBq. The positions used to measure the spatial resolution are shown in Table 2.2.

Table 2.2: Positions of the spatial resolution experiments in the Forte FOV [44]. The coordinate system is shown in Figure 2.1.

X-Position (mm)	**Y-Position (mm)**	**Z-Position (mm)**
0	0	0
0	127.5	0
0	0	95
150	0	0
150	127.5	0
150	0	95

To measure the count-rate response of the Forte, a cylindrical HDPE phantom was used and the NEMA methods were applied to the recorded LoRs to extract the total, true, scattered, and random count-rates as a function of the source activity [90]. This cylinder has an outside diameter of 50 mm and length of 120 mm, with an internal cylindrical cavity with a diameter of 10 mm and length of 100 mm which is filled with a solution of fluorine-18 and water [44]. Initially, the activity was approximately 100 MBq at the start of the experiment then decayed until approximately 1 MBq was left. For this experiment, the head separation was set to 600 mm. At regular intervals, the recorded LoRs are used to produce a projection image of the source and the method described in Section 1.8.2 was used to extract the count-rates. The source activity was calculated based on an exponential decay equation using the initial activity, the time elapsed since the start of the experiment, and the decay constant

for fluorine-18 [90].

Manual Results

The spatial resolution experiment conducted in the centre of the FOV of the Forte produces a central axial FWHM of 16.28 mm and transverse FWHM of 5.39 mm. The spatial resolution was also measured at 1/4th of the FOV. The 1/4th FOV axial and transverse FWHM are measured to be 19.62 mm and 5.13 mm, respectively. Compared to the experiment, the GATE simulation without spatial blurring results in an under-prediction of the FWHM in both the axial and transverse directions, indicating that spatial blurring is required. The 2.5 mm, 5 mm, and 7.5 mm spatial blurring simulations were compared, with the 5 mm blurring GATE simulation agreeing best with the experiment, as shown in Figure 2.5. In previous work, the Forte was measured to have a transverse spatial resolution of approximately 6 mm by fitting a one-dimensional Gaussian function to a point source profile in air [103]. The measurements in this work are consistent with this previously reported value. While this previous measurement is useful for comparison, it was unknown whether the ADAC Forte PEPT detector system, which has been used at the PIC for 20 years since this measurement, would still produce the same values, due to the degradation of components and the age of the system. Additionally, an added spatial blurring of 6 mm would not produce a spatial resolution since gamma-ray accolinearity is present in the simulation.

After manual calibration of the digitizer model, the results of the count-rate experiment and simulation are in good agreement, with the total, true, scattered, and random count-rates differing by less than 10% over the range of 1-100 MBq. The count-rate response as a function of activity for the real experiment and simulation are shown in Figure 2.6 and the fraction of true and corrupted LoRs as a function of activity are shown in Figure 2.7. The total and true count-rates perform particularly well, with an average error of 2.87% and

Table 2.3: Experimental and GATE simulated spatial resolution results for the Forte [44].

Spatial Resolution	**Experiment**	**0 mm Blurring**	**5 mm Blurring**
Central axial (mm)	14.680	10.965	13.294
Central transverse (mm)	5.567	3.159	5.675
1/4 FOV axial (mm)	17.217	13.368	16.904
1/4 FOV transverse (mm)	4.941	2.871	5.477

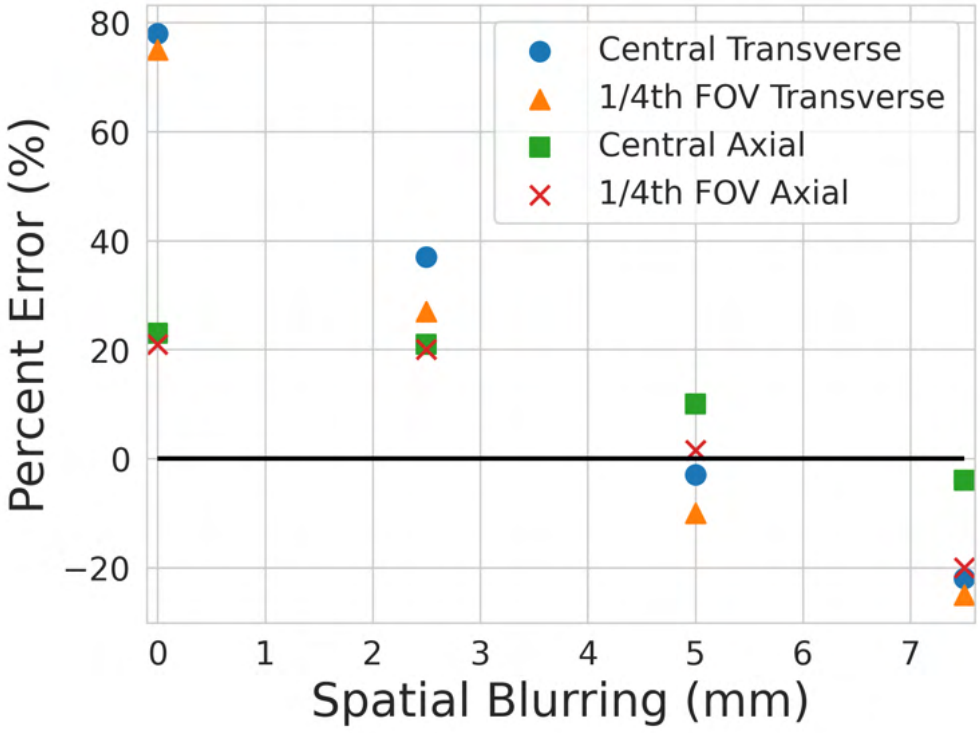

Figure 2.5: A comparison of the results of the spatial resolution test with varied amounts of spatial blurring added to the digitizer model [44].

6.04%, respectively over the range of tested activities. The peak true coincidence rate of the experiment is 50.3 kHz at 21.4 MBq compared to the simulated values of 48.5 kHz at 21.4 MBq.

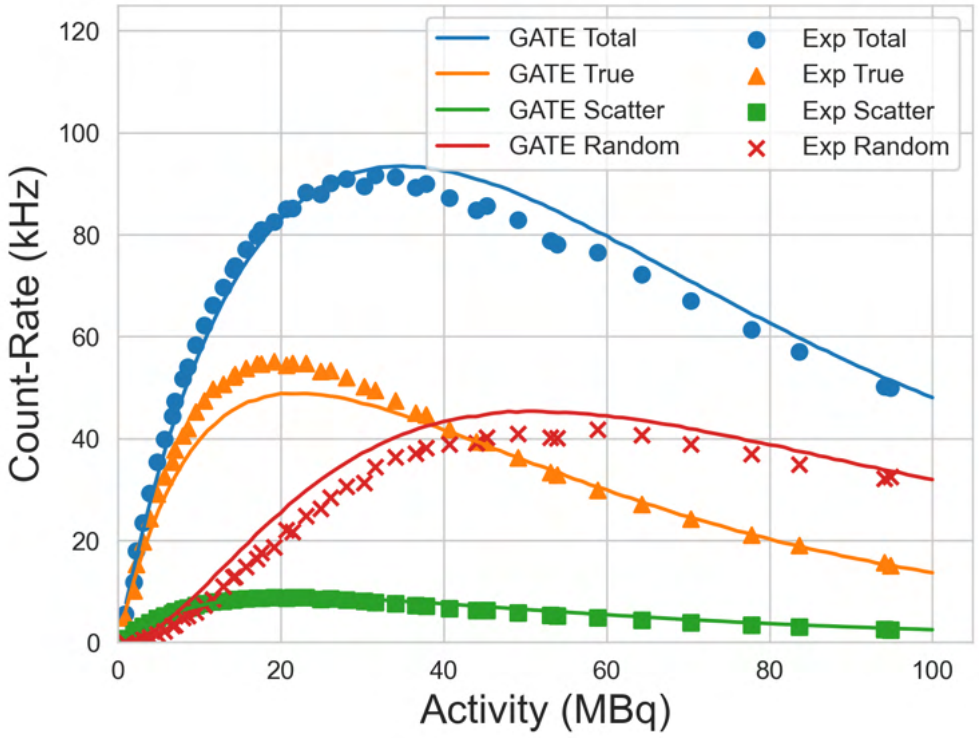

Figure 2.6: A comparison of the manually-calibrated results of the count-rate experiment for the total, true, scattered, and random count-rates as a function of source activity [44].

Table 2.4: ADAC Forte digitizer description [44].

Digitizer Model	**Values**
Coincidence Window	15 ns
Time Resolution	15 ns
Paralysable Dead-Time Per Pixel	1.15 μs
Non-Paralysable Dead-Time	1.15 μs
Energy Resolution at 511 keV	14%
Energy Window	350 - 650 keV (50% Photo-peak)
Intrinsic Efficiency	24%

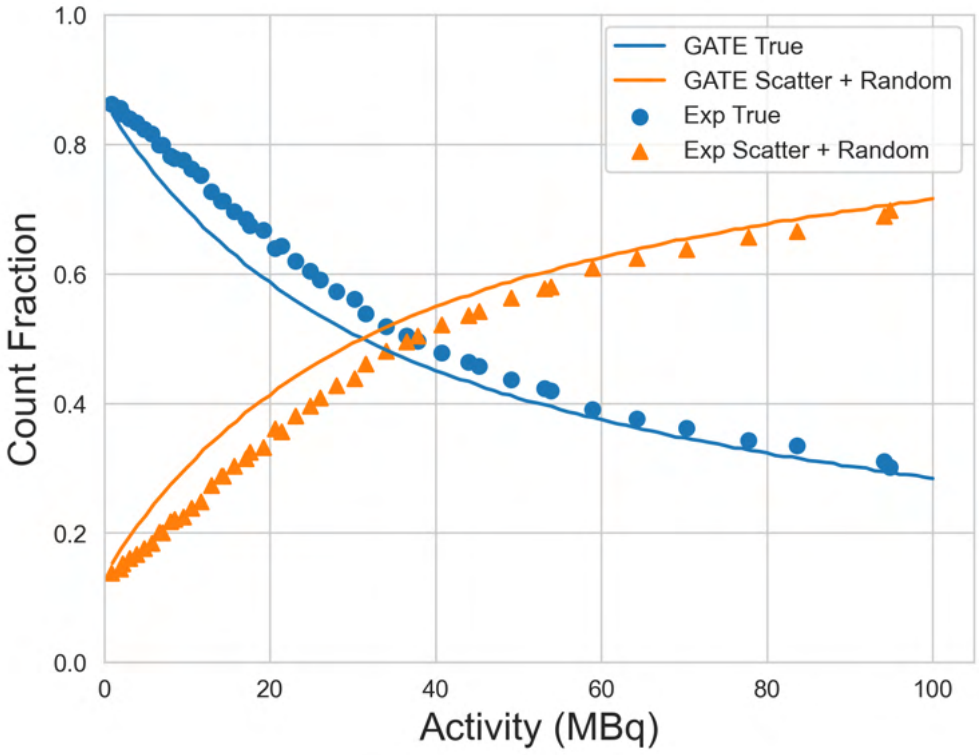

Figure 2.7: A comparison of the manually-calibrated results of the count-rate experiment is shown as the fraction of true LoRs compared with the fraction of corrupted LoRs (scattered + random) [44].

Overall, this calibrated GATE model of ADAC Forte is in good agreement with the experimentally observed detector performance characteristics. This model has been shown to generate synthetic LoRs from conditions which replicate real experiments. However, while the model has been calibrated for this head separation, the model was not tested in other configurations. In the next section, an autonomous calibration of the digitizer's free parameters in regard to the count-rate experiment will be conducted using performance characteristics measured at the minimal, median, and maximal head separations.

Autonomous Calibration

As introduced in Section 2.2, a new method of digitizer calibration has recently been developed which seeks to find optimal values for digitizer parameters using an evolutionary algorithm [45]. This method is demonstrated by simultaneously calibrating six parameters within the ADAC Forte digitizer model to match the count-rate response and the like before a manual calibration for spatial resolution was done since the spatial blurring parameter has little effect on the count-rate response [92, 41]. The parameters chosen to be optimised are the singles dead-time, coincidence dead-time, pileup, lower energy discriminator, upper energy discriminator, and time resolution. These parameters were chosen because they have not been measured directly through a characterisation experiment meaning there is uncertainty in the optimal values.

In a similar manner to the previously described count-rate response characterisation experiments, a new set of experiments are conducted which measure the coincidence count-rates of the Forte as a function of source activity at three head separations representing the minimum, median, and maximal head separations, or 250 mm, 525 mm, and 800 mm. The initial source activities for each separation were selected to test both the high-activity range where the effect of detector dead-time induces count-losses and to test the low-activity range where count-rates are linearly proportional to the source activity. In the experiments, the HDPE cylinder is filled with an initial activity, then placed in the centre of the FOV of the Forte and imaged over several half-lives until the activity is below 1 MBq. The three head separations and initial activities for each experiment are found in Table 2.5.

The experimental count-rates are then compared to the simulated count rates generated from a digitizer with a set of candidate solutions for the free parameters. For each of the three head separations, the total, true, and corrupted (scattered + random) count-rates are extracted at five different activities which cover the low-activity range where the detector

Table 2.5: Head separations and initial activities for each calibration experiment [45].

Experiment	**Head Separation (mm)**	**Initial Activity (MBq)**
Experiment 1	800	75
Experiment 2	525	60
Experiment 3	250	40

response is linearly proportional to source activity and the high-activity range when the dead-time and count-losses cause the detector response to be highly non-linear [90]. Since there are six free parameters and 45 individual data points for each parameter combination in the optimisation (three detector separations, five activities per separation, and three coincidence count rates per activity), the optimisation problem is considered to be well-constrained [45]. Further, since the count-rates have a complex relationship to the digitizer parameters, source activity, and detector configuration, this makes the count-rate response characterisation an ideal metric for comparison.

The fitness of parameter combinations was determined through a cost function which measures the percent difference between the experimentally observed and simulated count-rates for the total, true, and corrupted count-rates across all three head separations and activities. The sum for each of these percent differences is denoted as ε_R, ε_T, and ε_{SR} respectively and computed using Equation 2.1. Each type of count rate is treated as an objective to optimise and combined into a multi-objective optimisation by multiplying them together using Equation 2.2. In this case, each type of count rate is treated as equally important. Other types of cost functions could be devised in the future which treat some count-rates as more important than others.

$$\varepsilon_R = \sum 100 \frac{|R_{exp} - R_{GATE}|}{R_{exp}} \tag{2.1}$$

$$\varepsilon = \varepsilon_R \varepsilon_T \varepsilon_{SR} \tag{2.2}$$

The number of simulations per epoch is set to 150 so that a wide range of different parameter combinations can be tried and the simulations can be launched in parallel while not affecting the shared usage of the BlueBEAR high-performance computing (HPC) system [133]. In this work, each set of simulations has a candidate set of parameter solutions and is simulated using a single Intel Icelake core of the BlueBEAR HPC with 8 GB of memory allocated [133]. The maximum run time is set to 4 hours and 30 minutes, which is approximately twice as long as the expected mean run time. In the event that a set of simulations takes longer than 4 hours and 30 minutes, the job is terminated and the results are not included in the next generation of parameter solutions.

Additionally, the bounds of the parameter values are set to only explore solutions which make physical sense, excluding options like a negative dead-time or upper energy level being below the upper energy window. The bounds are also limited where needed such that the solution space is finite, yet spanning a range likely to contain the optimal value based on an estimate from a previous manual calibration in Section 2.3.1 [44]. A list of the bounds and the initial values are shown in Table 2.6.

The termination criterion for the optimisation is the standard deviation for each parameter reaching 10% of the initial standard deviation. This range is chosen such that variation in the parameter values will not significantly affect the accuracy of the model. The initial standard deviation is equal to the range of the bounds at the beginning of the optimisation and the initial scaled standard deviation is defined as unity.

Once the optimal values are identified, they are input to the digitizer model and a count-rate response is generated. These simulations are run at 2 MBq intervals starting at 1 MBq and reaching into the upper activities for each experiment. A study of the accuracy of

Table 2.6: Digitizer parameter bounds and initial parameter values [45].

Parameter	**Lower Bound**	**Upper Bound**	**Initial Values**
Singles Dead-Time (ns)	0	2	1
Coincidence Dead-Time (ns)	0	2	1
Pile-Up (ns)	0	600	300
Lower Energy Discriminator (keV)	0	360	180
Upper Energy Discriminator (keV)	640	1200	920
Time Resolution (ns)	10	20	15

extracted count-rates for the simulation with the highest separation and lowest activity (800 mm and 2 MBq) at different numbers of LoRs used to produce projection images showed that at least 10,000 events are needed to ensure that variance in the extracted count-rates is well below 10%. The results of this study are shown in Figure 2.8. The lowest count rate that would be expected in an experiment is approximately 1 kHz. As a result, it is determined that simulations should be run for 10 seconds of simulated time at each activity in order to ensure that 10,000 events are captured. In order to contextualise this proposed method of calibrating Monte Carlo detector models through evolutionary simulation to the existing methods, the ACCES-calibrated model is compared to the existing manually-calibrated model [45, 44].

Autonomous Results

In total, the ACCES optimisation took 56 epochs, 8400 cost function evaluations, and approximately 4 days to complete. At the beginning of the ACCES optimisation, the candidate solutions for the six free parameters are broad to explore the solution space. After this initial period, the values begin to converge to their optimal values, as shown in Figure 2.9. The

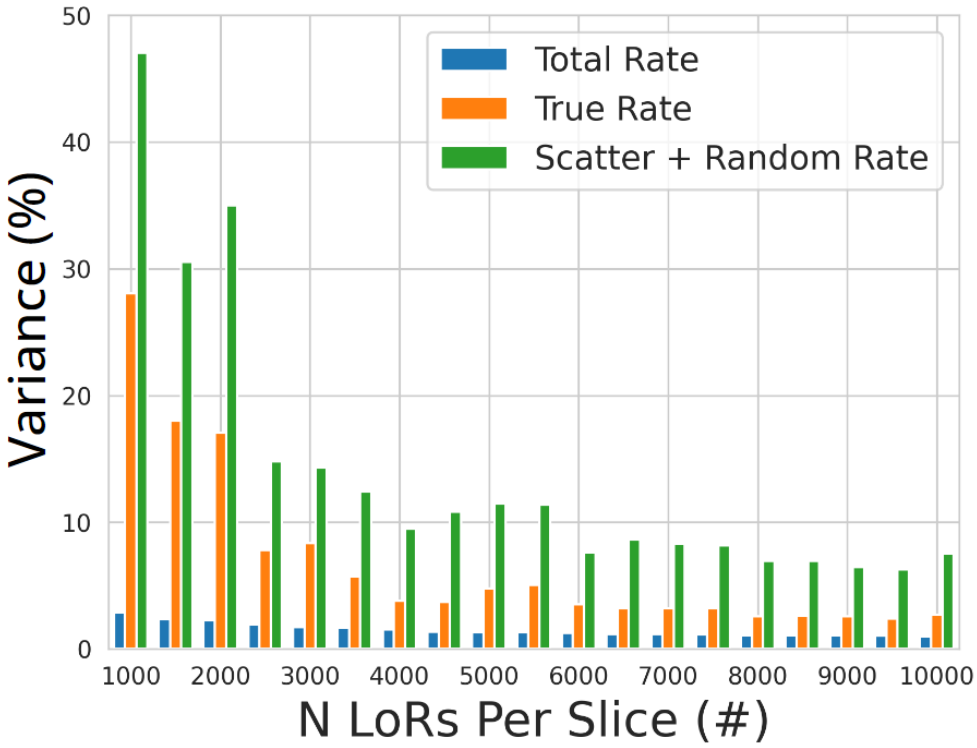

Figure 2.8: The variance of the total, true, and scattered plus random count rates for the 800 mm and 2 MBq simulation as a function of the number of LoRs used to generate a projection image [45].

scaled standard deviations are shown in Figure 2.10 to depict how the uncertainty in the optimisation decreases as the optimisation progresses. Initially, the standard deviations are wide, reflecting there is a large uncertainty in the optimal parameter values. However, as more candidate solutions are tried and evaluated with the cost function, ACCES learns how to best emulate the real detector response.

Once the optimisation reached 10% uncertainty for each parameter, the parameters were considered calibrated. The final calibrated values are presented in Table 2.7. For all parameters except for the coincidence dead-time, the optimal solutions are well within their bounds, suggesting an optimal calibration that would not change with different bounds. However, the optimal coincidence dead-time is found to be approximately 0 ns. While this

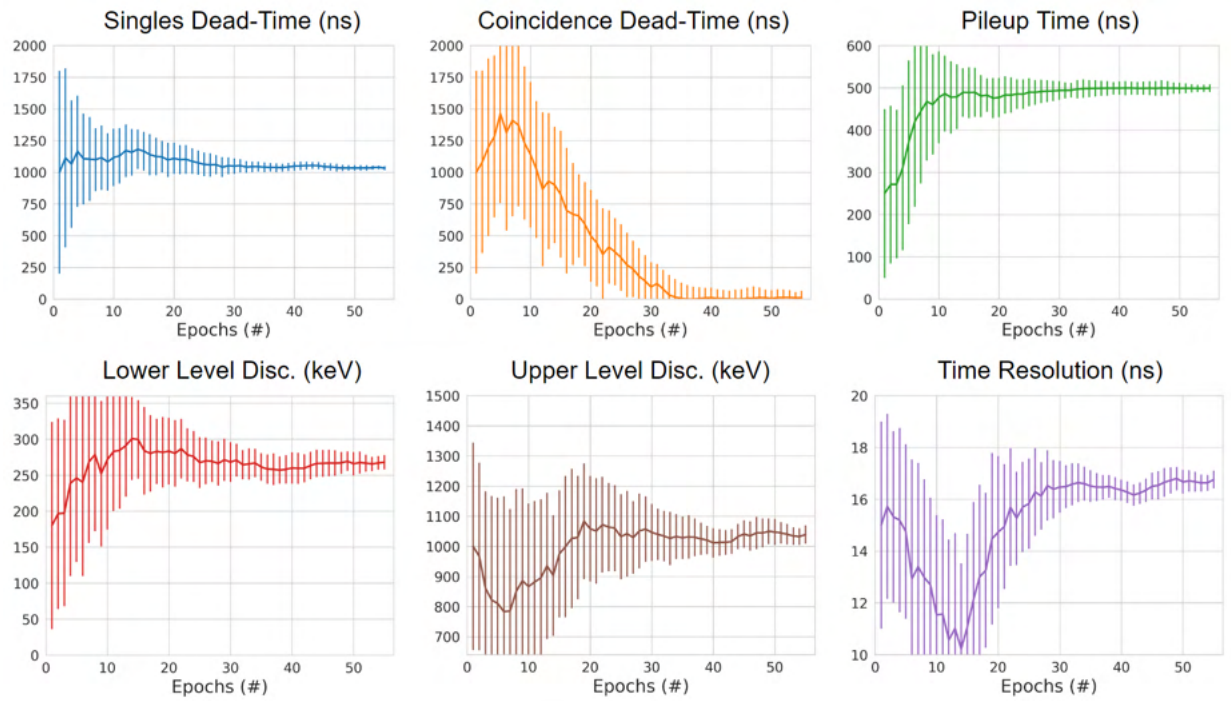

Figure 2.9: The mean parameter values for each of the six free parameters with the standard deviation of the values are plotted as error bars. After 56 epochs all parameters are below 10% standard deviation and the optimisation is completed [45].

could be due to the bounds being ill-suited to the problem, in this case, it is believed this demonstrates that coincidence dead-time is insignificant to the digitizer model. Further, support for this is that the model under-predicts the peak count rates. The opposite would be expected if coincidence dead-time was important since this should reduce the count-rates [134].

To assess the ability of the ACCES-calibrated digitizer model to replicate the experimental data, a new set of simulations is conducted using the optimised values. The results of these simulations were plotted against the experimental data in Figure 2.11. To quantify the accuracy, a mean absolute percent difference is calculated for each head separation and each type of count rate and presented in Table 2.8. Additionally, the results for the manually calibrated digitizer model are presented in Table 2.9.

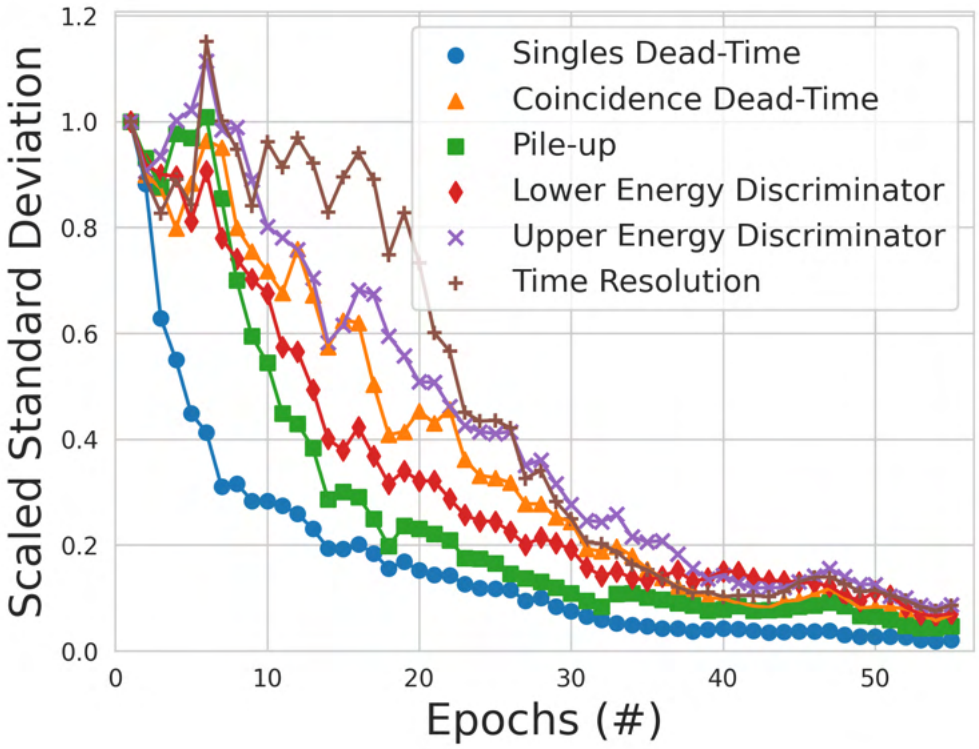

Figure 2.10: The standard deviation of the parameter values tried by ACCES. A lower standard deviation represents lower uncertainty in discovering the optimal parameter values [45].

Table 2.7: Calibrated digitizer parameter values [45].

Parameter	**Calibrated Value**	**Uncertainty**
Singles Dead-Time (ns)	1070	±16.7
Coincidence Dead-Time (ns)	10	±54.7
Pile-Up (ns)	498	± 9.31
Lower Energy Discriminator (keV)	284	±10.1
Upper Energy Discriminator (keV)	1020	±30.2
Time Resolution (ns)	17	±0.347

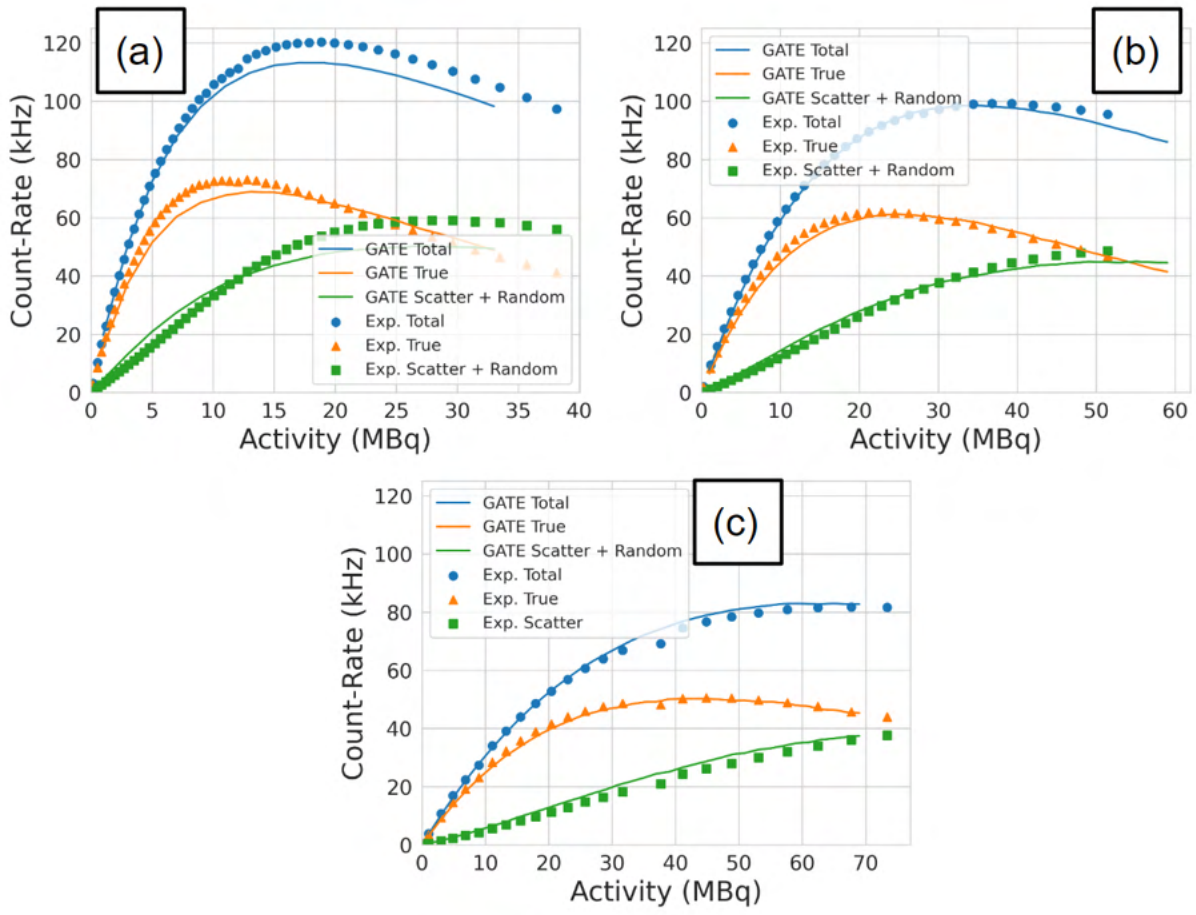

Figure 2.11: The optimised GATE model count-rates are plotted against the experimental data for the (a) 250 mm experiment, (b) 525 mm experiment, (c) and the 800 mm experiment [45].

The average mean absolute percent differences for the 250 mm, 525 mm, and 800 mm were 7.55%, 4.30%, and 5.48%, respectively. The separation which was closest to the phantom experienced the highest error between the simulation and experiment. This could be caused by the closer separation amplified differences between the source position in the simulation versus the experiment. In addition to this, the ACCES-calibrated model improves the match between simulation and experiment compared to a manually-calibrated digitizer model which produced a mean absolute percent difference in the count rate response of 17.78%, 15.42%, and 21.75%. This represents ACCES producing a calibration

Table 2.8: Mean absolute percent differences in the count-rates of the ACCES-calibrated digitizer model [45].

ACCES-Calibrated Results

Head Separation	**250 mm**	**525 mm**	**800 mm**
Total Count-Rate	4.76%	1.71%	2.43%
True Count-Rate	4.55%	2.85%	2.25%
Scatter + Random	13.33%	8.33%	11.77%
Average Error	7.55%	4.30%	5.48%

Table 2.9: Mean absolute percent differences in the count-rate of the manually calibrated digitizer model [45].

Manually Calibrated Results

Head Separation	**250 mm**	**525 mm**	**800 mm**
Total Count-Rate	11.50%	4.48%	6.25%
True Count-Rate	17.86%	15.68%	10.85%
Scatter + Random	23.98%	26.10%	48.15%
Average Error	17.78%	15.42%	21.75%

which achieves a nearly three times better agreement with the experiments. This is a significant improvement and one accomplished without guiding the optimiser to these solutions. Overall, this calibration represents an agreement with the experiment that would be sufficient for the GATE model to be used as a predictive tool to generate data representative of real experiments and improves upon the performance of the manually-calibrated digitizer.

2.3.2 Large Modular Array (LaMA)

The LaMA was developed at the PIC in the early 2000s, growing out of the need to have a system with flexible geometry to image large industrial systems. Initially, three CTI/Siemens ECAT ring scanners were dismantled to retrieve block detectors, with each block consisting of an 8 x 8 BGO crystal array and crystal dimensions of 6.25 mm x 6.75 mm x 30 mm. The BGO crystals are optically coupled to a 5 mm thick glass light guide to four photo-multiplier tubes (PMTs) 100 mm in length [97]. Using the extracted 192 block detectors, a system of 48 boxes comprised of 4 block detectors each was designed. Each box was approximately 360 mm in width, 95 mm in height (including spacers), and 460 mm in thickness with each of the four block detectors spaced 90 mm apart from centre-to-centre. A single box is shown on the left side of Figure 2.12 and a stack of four boxes are shown on the right. These form the fundamental building blocks of LaMA and up to 32 boxes can be connected to a single coincidence processor unit to form a new system. These boxes can be placed in any configuration, allowing for flexible geometries to be designed [35]. Since the LaMA is reconfigured for each experiment, the performance characteristics of the camera are difficult to predict, making it challenging to estimate the spatiotemporal resolution of the trajectories, design optimised geometries, and select ideal tracer activities.

To address these issues, a Monte Carlo model of LaMA is created and validated in a configuration with two stacks of four boxes separated by 500 mm [48]. This is one of the simplest geometries that can be created. Thus by characterising its performance and validating the model on a small scale, the behaviour of the system when it is scaled up will remain accurate. The model is created using GATE v9.1 [56, 115, 90]. The performance of the camera is characterised following the industry-standard NEMA protocol and the pulse-processing digitizer model is calibrated through evolutionary simulation and validated against the experimental characterisation [90, 42, 92].

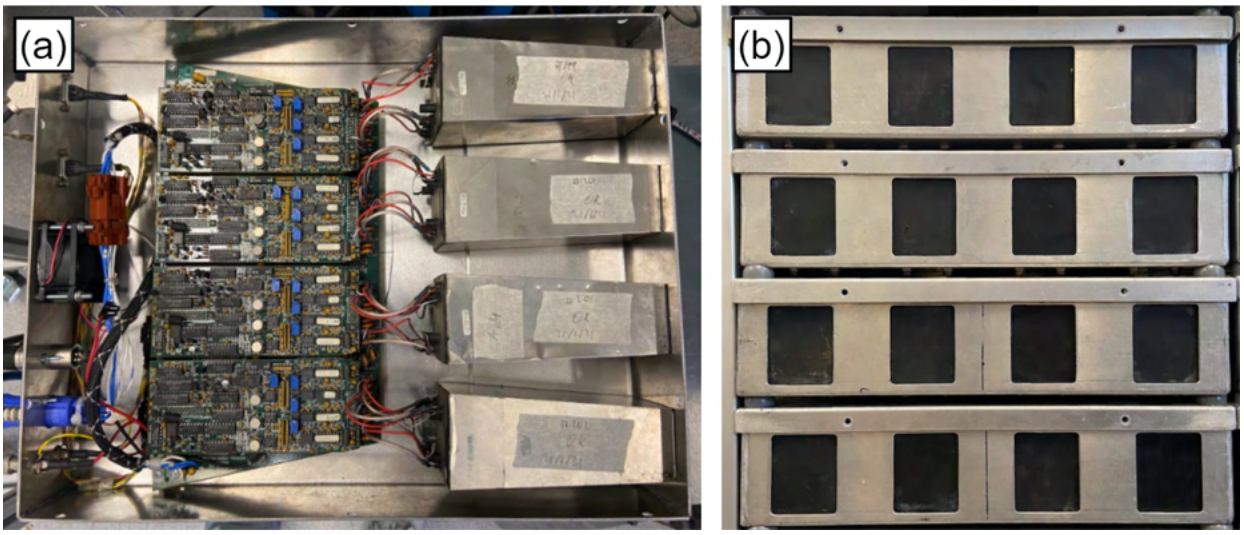

Figure 2.12: (a) A view inside a single box where four block detectors are mounted. Boxes are the building blocks of the LaMA. (b) A stack of four boxes. The geometry used in this work are made of two stacks of four boxes separated by 500 mm.

366 **Methods**

367 In the same way that the performance characteristics of the ADAC Forte were measured,
368 this is done for LaMA as well [44]. These characteristics are compared to those observed in
369 the GATE model of LaMA and used as the basis for calibrating the digitizer model. For
370 the spatial resolution tests, the source used is a 1 mm sphere of anionic exchange resin,
371 volumetrically activated with fluorine-18 in a solution of water produced by the University
372 of Birmingham MC40 cyclotron [95]. For imaging, the source was placed in a small plastic
373 sample holder and fixed to a block of polystyrene foam at six locations ranging from the
374 centre of the FOV and locations at 1/4th of the FOV. The polystyrene foam was chosen
375 because of its low density, and thus low linear attenuation coefficient, which will not induce
376 a large amount of scattering [129]. These locations and the source activities at the time of
377 the experiment are listed in Table 2.10. The source and the LaMA geometry are shown in
378 Figure 2.13.

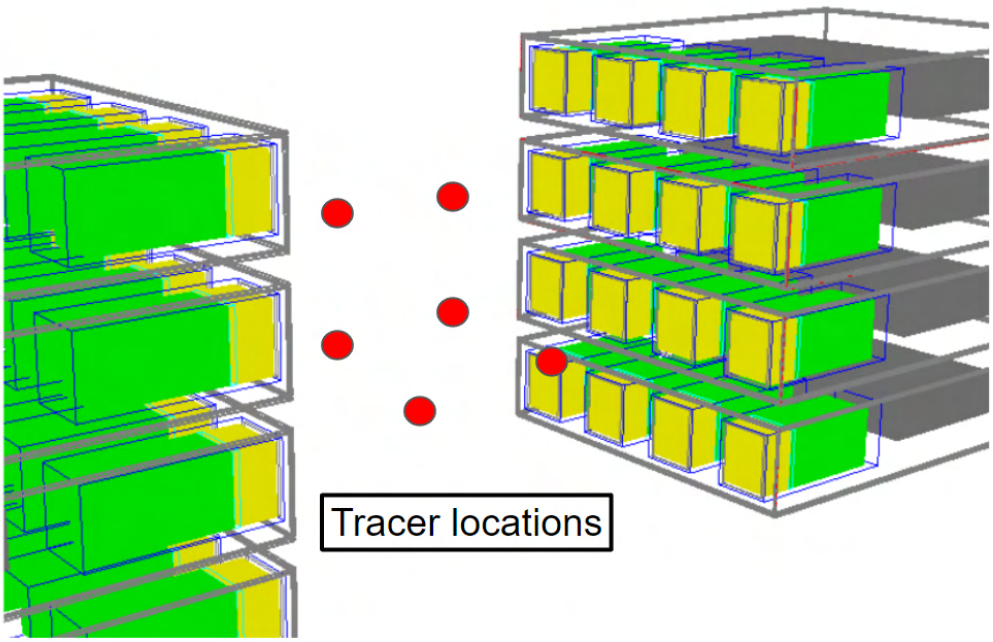

Figure 2.13: The spatial resolution tests are conducted using 1 mm diameter resin beads [48].

Table 2.10: Spatial resolution test parameters.

X-Position (mm)	Y-Position (mm)	Z-Position (mm)	Initial Activity (MBq)	End Activity (MBq)
0	0	0	2.26	2.07
0	77.5	0	2.06	1.91
86.25	0	0	1.88	1.69
0	0	125	1.39	1.20
0	77.5	125	1.19	1.05
86.25	0	125	1.00	0.86

The count-rate experiment measures the response of the detector to a source over several half-lives. The total, true, and corrupted (scattered plus random) LoRs count rates were extracted using the NEMA protocol and recorded at regular intervals as a function of source activity [48]. The phantom is a hollow, HDPE cylinder measuring 120 mm in height and 50 mm in diameter. The inner cavity, which measures 100 mm in height by 12 mm

in diameter, is filled with a solution of fluorine-18 and water. Initially, the activity of the phantom was approximately 80 MBq and this activity was chosen such that the expected count rates exceeded the maximum rate at which LoRs can be recorded by the camera. The phantom was imaged over several half-lives until the activity reached near that of the background. The phantom and detector geometry is shown in Figure 2.14.

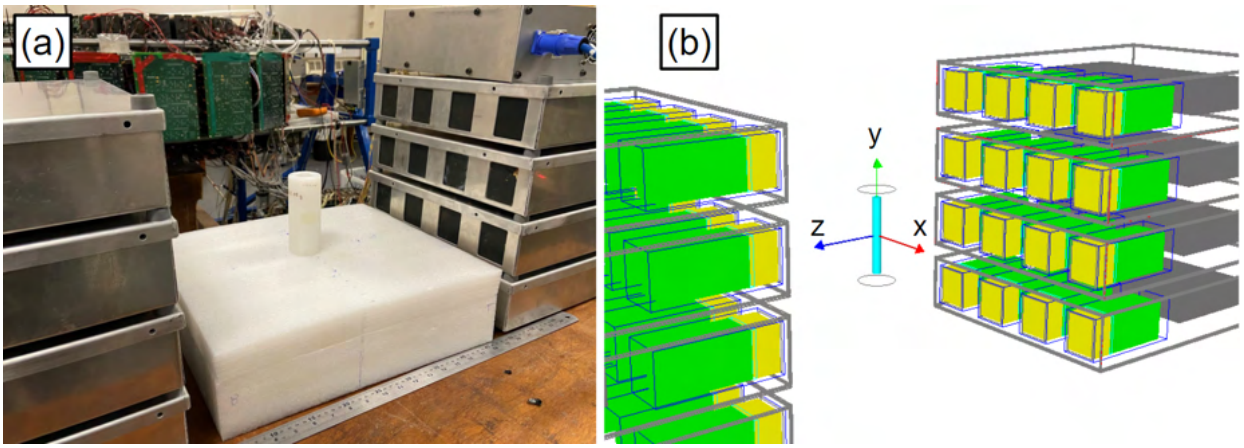

Figure 2.14: (a) The count-rate experiment is conducted with the high-density polyethylene phantom placed in the centre of FOV and imaged over several half-lives. (b) a GATE model of the same experiment is conducted [48].

The design of a GATE model for the LaMA presents several challenges since it must be easily customised to rapidly prototype new geometries, only allow specific coincidences to be formed between connected boxes, and be able to emulate the noise, data buffer, and spatial blurring inherent in the system. Achieving these goals is accomplished using the tools available in the GATE software in addition to data post-processing.

Since any LaMA geometry is built using boxes, to build a model of the LaMA, only a single box needs to be described which can later be copied, translated, and rotated to any position and orientation using GATE's generic repeater function. Importantly, the four ECAT951 block detectors are included in each box and the 8 x 8 BGO crystal array in each block is defined as the 'Sensitive Detector' through which GATE records the interactions of

particles. The model of a single box is shown in Figure 2.15 from various viewpoints. The order in which the repeated boxes are listed determines their volume number, which will become important later when defining which pairs of boxes are valid for recording coincidences.

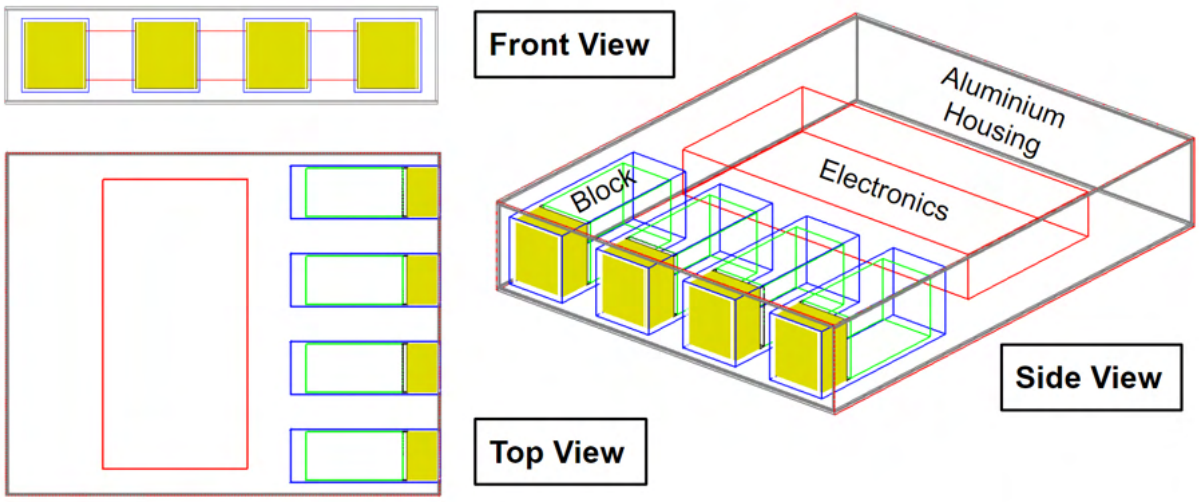

Figure 2.15: A model of a box for the LaMA consisting of four ECAT951 block detectors each with an 8x8 array of BGO crystals. The box is shown from various viewpoints and has major components labelled [48].

The digitizer for the LaMA is shown in Figure 2.16. It is a typical digitizer model for a PEPT detector system, but an additional post-processing stage is implemented to capture aspects of the detector possible using GATE [74]. This forces the detector to only record coincidences between pairs of boxes which are associated with each other in the real coincidence processor, implements a random spatial blurring to match the experimental spatial resolution, and implements a bandwidth which limits the rate at which LoRs are written to file. Random spatial blurring is needed due to the end of the LoR being centred in the crystal. By randomly distributing the ends of an LoR within the crystal, a continuous image on which it is easier to measure the spatial resolution is produced. The post-processing steps for the LaMA GATE model are the last three steps shown in Figure 2.16 before the

final simulated detector response is produced.

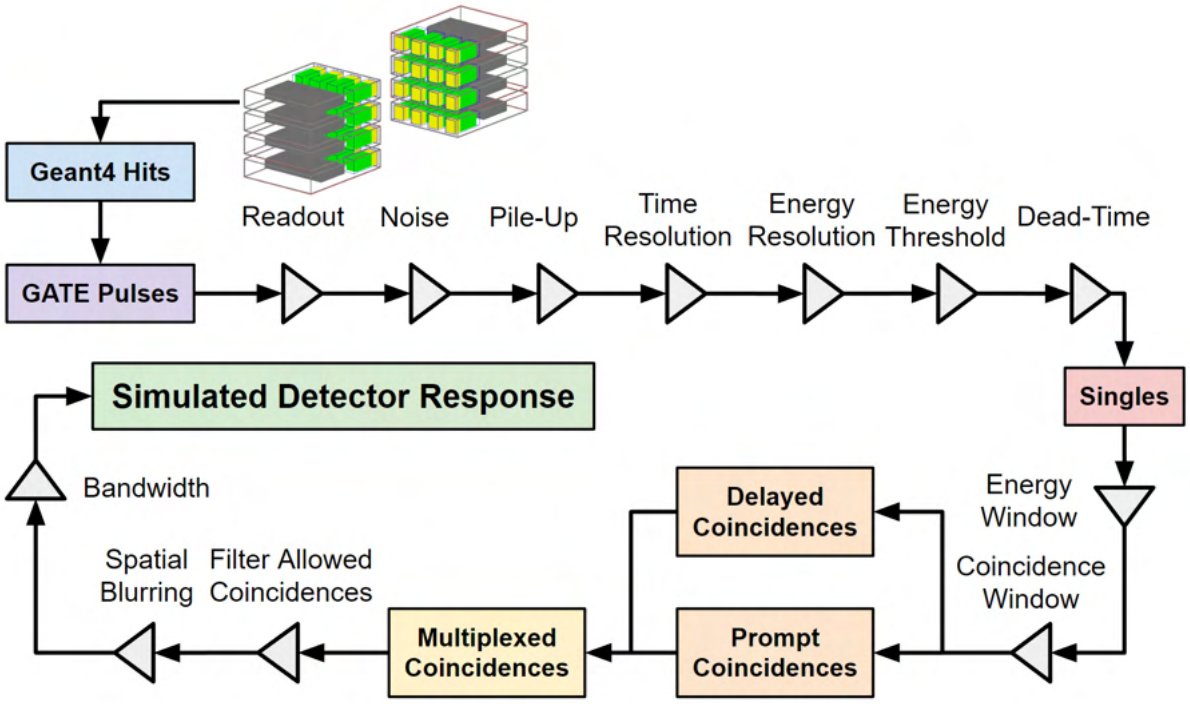

Figure 2.16: The pulse-processing model of the LaMA [48].

Even if every detail of the LaMA were known, there are differences between simulation and experiment which require aspects of the digitizer to be calibrated [71, 102]. In this case, six stages of the digitizer are chosen to be calibrated because of the assumed effect they have on replicating the count-rate response. These are the noise frequency, pile-up time, time resolution, lower-level energy discriminator, upper energy discriminator, and the non-paralysable singles dead-time. The noise frequency is the rate at which random events are generated simply by having the detectors running and is a combination of the background activity and electronic noise. The singles dead-time is a non-paralysable dead-time model, described by Equation 2.3 which limits the rate of recording single events [134].

$$\frac{\lambda_{out}}{\lambda_{in}} = \frac{1}{(1 + \lambda_{in}\tau)} \tag{2.3}$$

The objective cost function used for assessing the match to the real performance characteristics, shown in Equation 2.4, is the product of the percent differences between the experiment and simulation's total, true, and corrupted count-rate response over a range of source activities, calculated using Equations 2.5-2.7 [48].

$$\varepsilon = \varepsilon_{Tot}\varepsilon_{True}\varepsilon_{Corrupt} \tag{2.4}$$

$$\varepsilon_{Tot} = 100 \sum |(R_{Tot_{exp}} - R_{Tot_{sim}})/R_{Tot_{exp}}| \tag{2.5}$$

$$\varepsilon_{True} = 100 \sum |(R_{True_{exp}} - R_{True_{sim}})/R_{True_{exp}}| \tag{2.6}$$

$$\varepsilon_{Corrupt} = 100 \sum |(R_{Corrupt_{exp}} - R_{Corrupt_{sim}})/R_{Corrupt_{exp}}| \tag{2.7}$$

For this optimisation, 100 epochs with 100 parameter value combinations per epoch were used as the terminating criteria because this provides greater than 10 times the factorial of the number of free parameters in the optimisation, sufficiently constraining the problem [51]. Each combination is simulated over 10 different activities, ranging from 2 MBq to 80 MBq, until 5 million events are generated at each activity. The bounds of the parameters and their initial values are provided in Table 2.11. The optimisation is conducted on the University of Birmingham's HPC system, BlueBEAR, on Icelake cores with 16 GB of memory each [133]. After the optimisation is finished, the calibrated parameters were extracted and a new set of simulations was conducted with 20 activities over the same activity range until 30 million events are generated in order to reduce statistical error. These simulations are presented and compared in Section 2.3.2.

Table 2.11: Digitizer parameter bounds and initial parameter values for calibration [48].

Parameter	**Lower Bound**	**Upper Bound**	**Initial Values**
Singles Dead-Time (ns)	0	10000	5000
Noise Frequency (ns)	1000	10000	5500
Pileup Time (ns)	0	1000	500
Lower Level Discriminator (keV)	0	400	200
Upper Level Discriminator (keV)	700	2200	1450
Time Resolution (ns)	0	30	15

Results

For spatial resolution, six tracer positions were imaged over several minutes then the FWHM of the one-dimensional projection is extracted. The FWHM at these positions is used to compute the transverse and axial spatial resolutions in the centre of the FOV and at 1/4th of the FOV. Next, the experiments are reproduced in simulation and the crystal blurring is adjusted until the best match between the experiment and simulation was achieved. The crystal blurring which best agrees with the experiment is approximately twice the crystal dimensions (6.25 mm by 6.75 mm), as evidenced in Fig. 2.17. The experimental and simulated results are presented in Table 2.12 and compared through their percent differences [48].

Following the spatial resolution characterisation and crystal blurring calibration, the count rate response experiments were analysed to be used as a comparison for the ACCES optimisation. The optimisation took approximately three days to complete. At the end of the optimisation, the final mean parameter values were extracted. The value for these parameters and their uncertainty are provided in Table 2.13. Additionally, the history of these parameters during the optimisation (uncertainties and mean values) is presented in

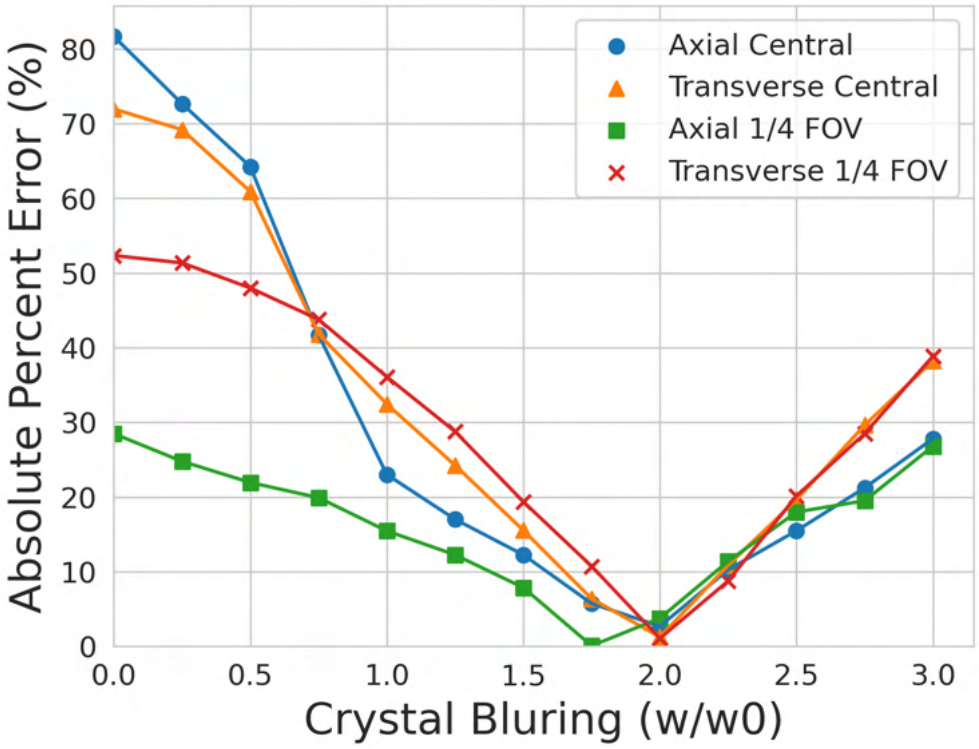

Figure 2.17: The spatial blurring is calibrated by finding the crystal blurring that minimises the absolute percent error [48].

Table 2.12: Results and comparisons of the spatial resolution tests for the experiment and simulation [48].

Spatial Resolution	**Experiment (mm)**	**Simulation (mm)**	**Percent Error (%)**
Central Transverse	8.26	8.38	1.44
1/4 FOV Transverse	8.95	9.37	4.72
Central Axial	21.83	22.34	2.32
1/4 FOV Axial	24.09	23.80	-1.20

Figure 2.18, demonstrating that before the end of the optimisation, each parameter reaches a stable value, meaning that the parameters have been calibrated [48].

All parameters produced reasonable calibrations within the upper and lower bounds given to the optimiser. Interestingly, the optimised value for the time resolution falls within the 12 ns ±2 ns measured in a previous characterisation of the LaMA [71]. This provides further evidence that the calibrated parameters correspond to physical reality and are global solutions, not simply local solutions. The parameter with the highest uncertainty is the lower-level energy discriminator. This is likely due to the relatively small impact of this parameter on the calibration. For example, the singles dead-time has a very strong effect on the simulation. This is the case because it has a strong pressure to be calibrated since it is applied to nearly all events which are detected. Meaning small changes in the calibrated values will cause large differences in the simulated count-rate response. Similarly, the lower energy discriminator acts on a smaller number of events which have undergone scattering and as a result has a somewhat higher uncertainty. However, the upper energy discriminator can only be applied to events that have piled up. Since the upper energy discriminator is set to 1990 keV, this means that at least four 511 keV events must be grouped together and this happens only a limited number of times in a simulation. As a result, there is not a strong pressure to calibrate this value resulting in a higher uncertainty but still an adequately calibrated value [48].

Table 2.13: Digitizer calibration results and uncertainty [48].

Parameter	**Calibrated Values**	**Uncertainty**
Singles Dead-Time (ns)	6630	±57.9
Noise Frequency (ns)	1970	±9.39
Pileup Time (ns)	637	±6.91
Lower Level Disc. (keV)	324	±7.99
Upper Level Disc. (keV)	1990	±85.5
Time Resolution (ns)	13.7	±0.084

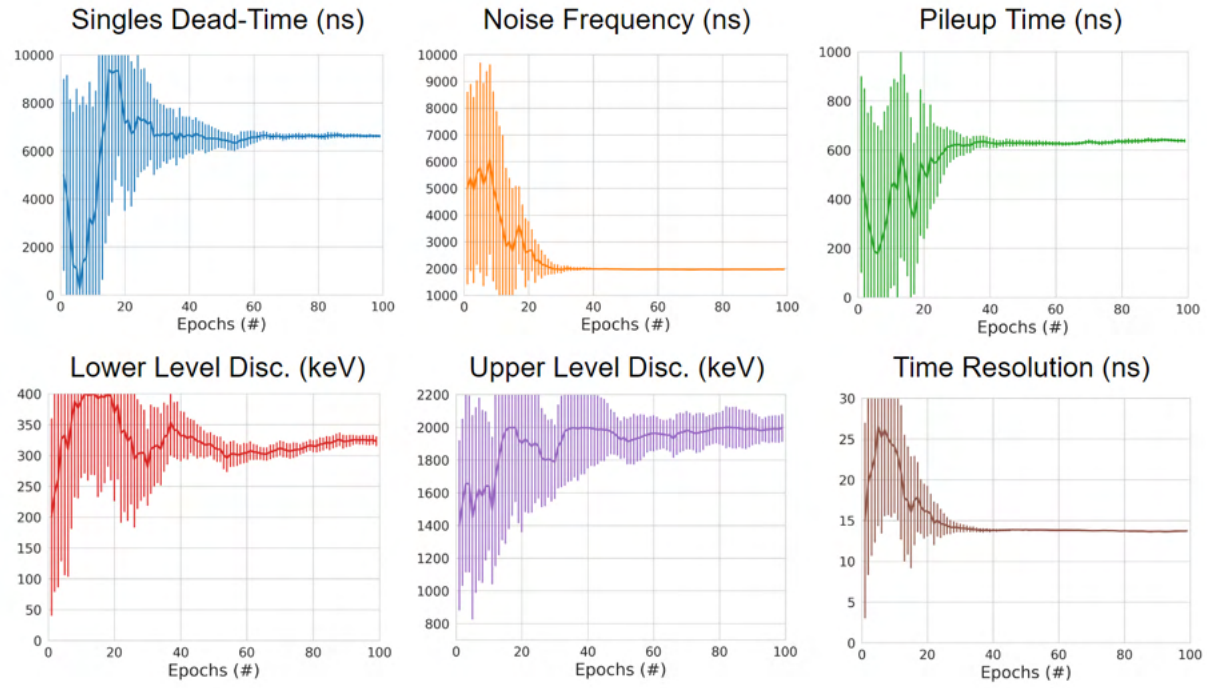

Figure 2.18: Subplots showing the mean values of calibrated parameters with the error bars as the standard deviation of solutions [48].

When the new set of simulations is conducted with the calibrated digitizer, the results match the experiment to a mean absolute difference of 3.41% over all three count-rates with the total, true, and corrupted count-rates being 2.31%, 2.18%, and 5.72%, respectively. The experimental and simulated count-rate response is shown in Fig. 2.19. To quantify the calibration further, it is also important to observe how the fraction of true and corrupted counts behave as a function of source activity. These results are presented in Fig. 2.20, showing that their behaviour is approximately the same over all activities, with the true and corrupted count fractions reconstructed to 1.91% and 3.72% error, respectively [48].

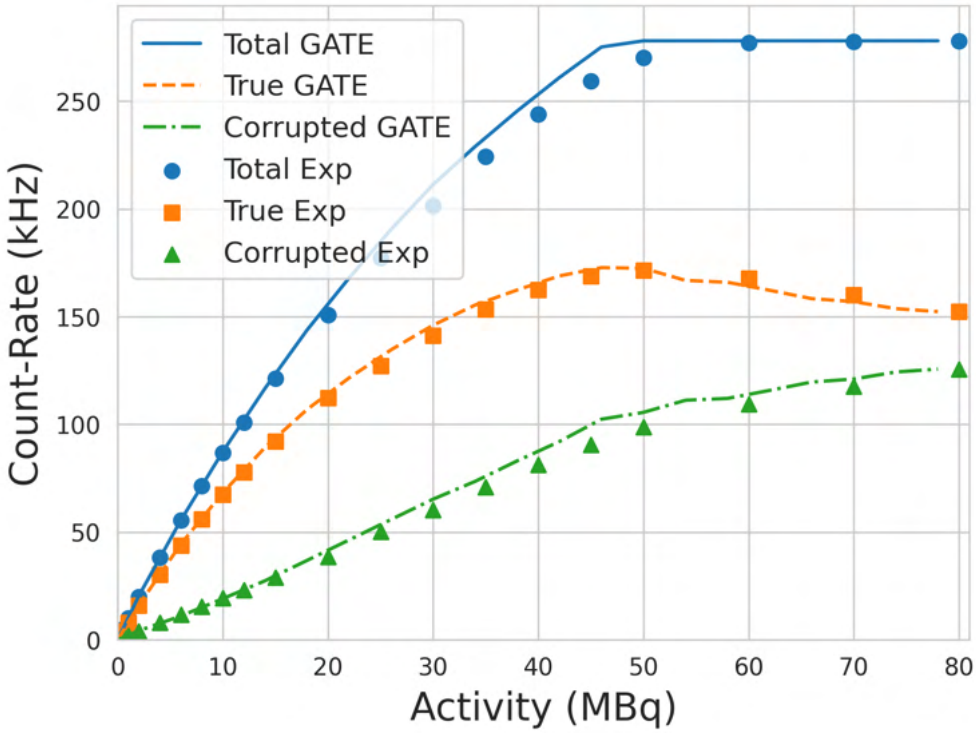

Figure 2.19: Results of the count-rate experiment and comparison of the ACCES calibrated GATE model of the LaMA [48].

2.4 Discussion

In this Chapter, the development of two new GATE models for PEPT detector systems and their calibration to experimentally observed performance characteristics using a novel, autonomous calibration method was presented. These GATE models are particularly useful for the PIC since these two systems are the most used and by having a validated model, a better understanding of their performance in different configurations can be gained [97]. Moreover, the ability to simulate the response of the system with a model of a PEPT experiment and PEPT tracer presents an opportunity to optimise the data acquisition through the use of tailored detector geometries and tracer activities.

In PEPT experiments, a balance between the quantity and quality of LoRs is desired.

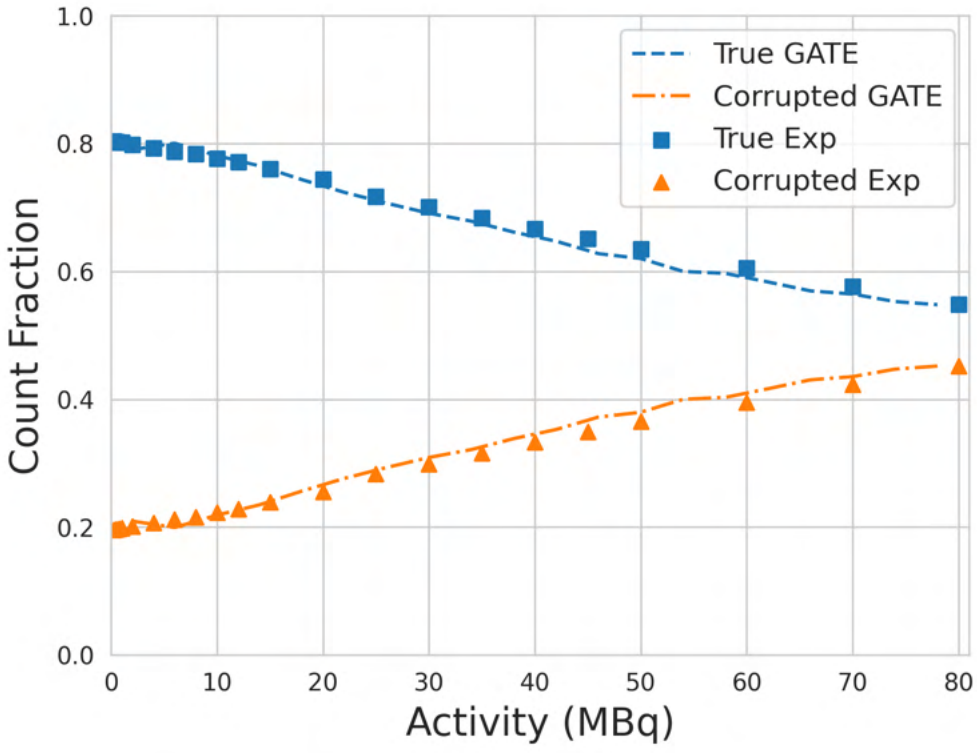

Figure 2.20: Results of the count-rate experiment in terms of the true and corrupt count fractions and comparison with the ACCES calibrated GATE model of the LaMA.

High count-rates are needed to create a sufficient temporal resolution of tracers along their trajectory, but increasing the source activity can result in poor count-rates due to excessive dead-time. Additionally, the fraction of corrupted events also increases at high count-rates, degrading the spatial resolution of reconstructed trajectories [125, 148]. As evidenced by the experiments, the detector response is shown to be a complex relationship between the source activity and detector head separation, which will be explored in more detail in Chapter 4. Using this model, a prediction of the optimal source activities and head separation can be made.

Moreover, the use of evolutionary algorithms to calibrate a digitizer has been proven to be a viable and useful strategy [45]. Calibrating digitizers has often been a time-consuming process and is not guaranteed to produce optimal calibration. However, by comparing the

real performance characteristics and simulated performance characteristics generated from a set of candidate solutions and applying an evolutionary algorithm to search the optimal set autonomously, both of these problems are solved. No active attention needs to be given to the calibration once the ACCES optimisation is created. Users need only to check in once the optimisation is finished after a few days to validate the final set of solutions. If the initial bounds of the search contain the optimal solution, the number of simulations per epoch is sufficient to learn the response of the detector to each free parameter and produce an optimal set of candidate solutions.

To assist users in developing their own optimisations using ACCES, an example of this method with test data is included within the GitHub repository found here. This example uses a simulated count-rate response of the ADAC Forte GATE model with prescribed parameter values in the digitizer as the ground truth response, then uses ACCES to calibrate two parameters, the singles dead-time and time resolution, to match the ground truth response. Two parameters were chosen because this is a more complex optimisation than a single parameter, yet easier to visualise than an optimisation with three or more parameters. The prescribed values for the singles dead-time and the time resolution are 1000 ns and 15 ns, respectively. The methodology in this simple example follows the same as that described in Section 2.3.1. The results from this optimisation in Figure 2.21 show the optimal parameter was determined to be 995.016 ns for the singles dead-time and 15.022 ns for the time resolution, which matches the prescribed parameters.

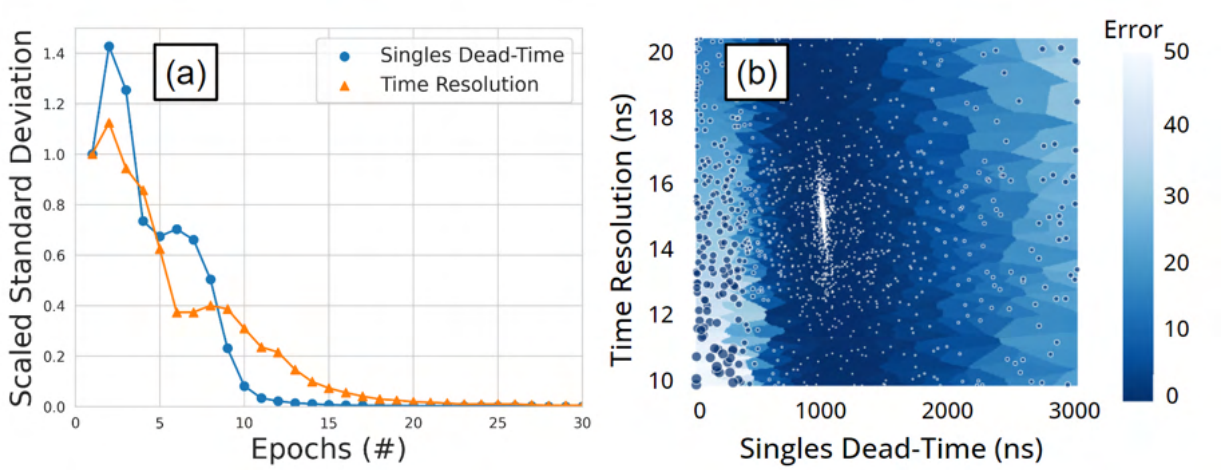

Figure 2.21: An example of an optimised GATE model with two free parameters. (a) The scaled standard deviation for each parameter over several epochs of simulations (b) a Voronoi diagram of the parameter combinations shows the solution converges to the optimal parameters. Each point is a candidate solution and the larger the point or lighter blue the Voronoi plot the higher the error in the cost-function evaluation [45].

Chapter Three

Benchmarking PEPT Algorithms

3.1 Introduction

One of the advantages of using PEPT over conventional, optical-based methods of imaging is that 511 keV gamma-rays can penetrate through opaque media and thick-walled systems [100]. This allows for tracking in a variety of systems which cannot be studied using other methods [147]. However, when using real data to track a PEPT tracer, this can present a challenge to verify whether or not the trajectory generated by a PEPT algorithm is an accurate reflection of the real tracer's behaviour. Additionally, since there exist multiple PEPT algorithms and these have been tested through a variety of disparate experiments using different types of detectors, comparisons between algorithms can not easily be made.

A variety of methods have been developed to test the accuracy of PEPT algorithms. One such method involves using a static, tracer placed in the centre of a PEPT detector system. A PEPT algorithm is then used to find the tracer's location several times and the three-dimensional standard deviation from the mean position of the tracer is given as the accuracy. [100]. This is equivalent to Equation 1.16 described in Section 1.4.3. However, as shown previously in Figure 1.5, the accuracy of a PEPT algorithm improves proportionally

to the inverse square root of the number of LoRs used per location [100, 97]. Though, with this, a trade-off in the temporal resolution is required, since fewer positions can be generated with a larger sample size of LoRs [100, 97].

A similar experiment can be conducted using a moving source with a known trajectory. For example, with a tracer placed on the edge of a horizontally rotating disk moving at a constant known angular velocity, the motion of the tracer can be described using sine waves [100]. By fitting the trajectory to this form and taking the average vertical component as a mean position, the instantaneous differences can be calculated and used to compute a three-dimensional standard deviation from the expected trajectory. However, while these provide a measure of accuracy, both methods are subject to systematic bias in the tracer position and uncertainty. Since the true position of the tracer is not precisely known and a wobble in the rotating disk can distort the expected trajectory, the bias cannot be determined so these methods are limited in their usefulness [47].

One way to attempt to resolve this is to conduct a simultaneous complementary measurement. In previous work, a PEPT measurement of jet flow through a baffle in a transparent system with simultaneous high-speed video of the tracer has been compared [69]. These measurements showed that the tracer trajectories were comparable and the reconstructed velocity fields differed only slightly. However, since there still exists uncertainty in both measurements this makes it hard to quantify the exact accuracy of a PEPT algorithm. Therefore, to solve this problem, an approach where the trajectory is known absolutely is needed.

The goal of this Chapter is to use GATE models to create a data set where the exact position of the tracer is known at all times, thereby creating a framework for comparing the performance of PEPT algorithms. Since GATE simulations can produce LoRs which emulate those produced by PEPT detector systems and the position of the source is prescribed, this

makes these models ideal for comparing PEPT algorithms. Producing this common data set of benchmark tests and collaborating with the creators of different PEPT algorithms, provides, for the first time, a quantitative and comparative assessment.

The benchmark tests have been carefully designed to allow a comprehensive, quantitative and, fair comparison of PEPT algorithms [149]. As different algorithms will inherently possess strengths and weaknesses, the benchmarking framework is comprised of a number of individual tests, each designed to evaluate a specific aspect of an algorithm's performance. Since not all PEPT algorithms are designed for multiple tracer tracking, these tests are divided into two parts: one assessing an algorithm's ability to locate a single PEPT tracer, and the other, the ability to locate and distinguish multiple tracers.

The analysis judges an algorithm's performance based on a number of quantities, including both spatial and temporal resolution, the number of tracers successfully tracked, and resilience to noise. Each individual test is performed using two different PEPT detector system geometries, which are the ADAC Forte and Siemens ECAT [44, 57]. This is done to ensure that algorithms developed for a ring-shaped detector are not unfairly disadvantaged by being tested exclusively on a planar detector, and vice-versa.

In order to make the comparisons of algorithms fair, no users of their respective algorithms were permitted to know the ground-truth locations of the tracers. To aid other developers of PEPT algorithms in using these benchmarks, the GATE data set and the Python functions used to compare the various algorithms against the ground truth positions of the simulated tracers can be found here.

3.2 Description of Algorithms

3.2.1 Birmingham Method

The Birmingham Method, developed at the University of Birmingham is the oldest, yet most widely-used PEPT algorithm [100]. The fundamental principle of the method is, for a given set of LoRs, some of the LoRs will be true LoRs, which lead back to the PEPT tracer, and some will be corrupted through scattering or are random LoRs. The unique feature of the Birmingham Method is that it can discard a specified fraction of the LoRs, ideally leaving only the true LoRs to locate the tracer [100].

The algorithm considers a sample containing a user-defined number of consecutive LoRs, N_{LoRs}, from a PEPT data set. Each LoR in the sample is defined by two points which are the positions the coincident gamma rays interact with the detector. The smallest perpendicular distance between the LoRs defined by these points is then determined and called the minimum distance point (MDP), (x, y, z). If all of the LoRs in the sample lead back to a single point, then the $\delta(x, y, z)$ is exactly 0. However, because of the limitations of spatial resolution, movement of the tracer, and incorporation of corrupted LoRs in the sample, $\delta(x, y, z)$ will have some positive value.

When all of the LoRs in a sample are used, the MDP will likely be distorted from the actual location of the PEPT tracer. Since the MDP will still likely be close to the actual location of the tracer, the Birmingham Method considers the LoR which has the largest $\delta(x, y, z)$ to be the most likely corrupted LoR. This LoRs is removed and the MDP is recalculated. This process is repeated until only a user-specified fraction of LoRs from the original sample remains. The fraction that remains in the sample is called f and must be provided by the user. The best value for f is one which removes all of the corrupted LoRs and leaves only the true LoRs. Determining the optimal value for f, often called f_{opt},

is difficult since it is PEPT detector specific, dependent on the position and activity of the tracer, and is affected by the amount of scattering in the experimental system. There are ways to determine f_{opt} using static or moving tracers which will be discussed in Chapter 4, but most commonly a conservative estimate is used [18]. The role of f in the Birmingham Method is demonstrated in Figure 3.1.

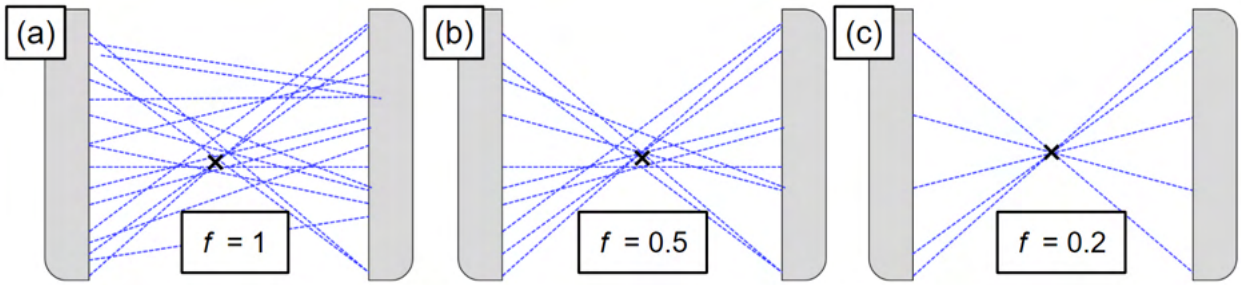

Figure 3.1: The Birmingham Method removing corrupt LoRs illustrated. In (a) all the LoRs are considered, in (b) half of the LoRs are considered, and in (c) a fifth of the LoRs are used [149].

3.2.2 G-Means Method

The G-Means method was developed at the University of Tennessee, Knoxville to enable multiple tracer tracking [139]. Instead of using a fixed number of LoRs per sample, this method begins by tracing LoRs from a given time step onto a three-dimensional grid of voxels within the FOV of the PEPT detector system. Voxel sizes and the time step are prescribed by the user and typically range from 0.5 mm - 2 mm and 1 ms - 100 ms. As LoRs are summed into the voxels grid, the voxels which are near the tracers will have a higher fraction of the LoRs passing through them. Initially, a check is made to ensure that the maximum voxel value is above a user-specified threshold value. If this check is passed, the grid of voxels is then high-pass filtered based on a fraction, usually 0.25-0.5, of the peak line-crossing value. This filters out the voxels which do not have a large number of LoRs

passing through them by setting their value to 0. The remaining voxels are then clustered via G-Means clustering to determine the positions of tracers within that time step [38].

G-Means is a modified k-means algorithm that uses statistical checks to overcome the limitation of traditional k-means which requires the number of clusters in the data set to be specified [2]. PEPT tracers are identified as follows: Firstly, each voxel is treated as a number of points positioned at the centre of each voxel. The number of points is equal to the number of LoRs which pass through the voxel. Initially, it is assumed that only one cluster is present and a three-dimensional Gaussian function is fit to all of the points. At this stage, an Anderson-Darling (A-D) test [4] is used to determine the 'goodness-of-fit' to the Gaussian form. If the cluster fails the test, then subsequently two Gaussian functions are fit to the points and the A-D test is performed again. By repeating this process, the natural number of clusters becomes apparent when all clusters pass the A-D test. The centroids of each cluster are taken to be the positions of the PEPT tracers.

After all of the tracers are identified within each time step from a given measurement, they are linked into continuous trajectories using a nearest-neighbours algorithm [139]. A linking method is needed in multiple tracer tracking methods since the ID of the tracer cannot be obtained through PEPT alone. This linking method has been seen to offer sufficient results when the between-frame displacements of tracers are shorter than the average tracer spacing. If between-frame displacements are larger, such as in the case of fast-moving flows and/or low data acquisition rates, a more robust linking algorithm may be used that employs the prediction of tracer position based on the fitting of tracer positions from multiple consecutive frames [104, 140].

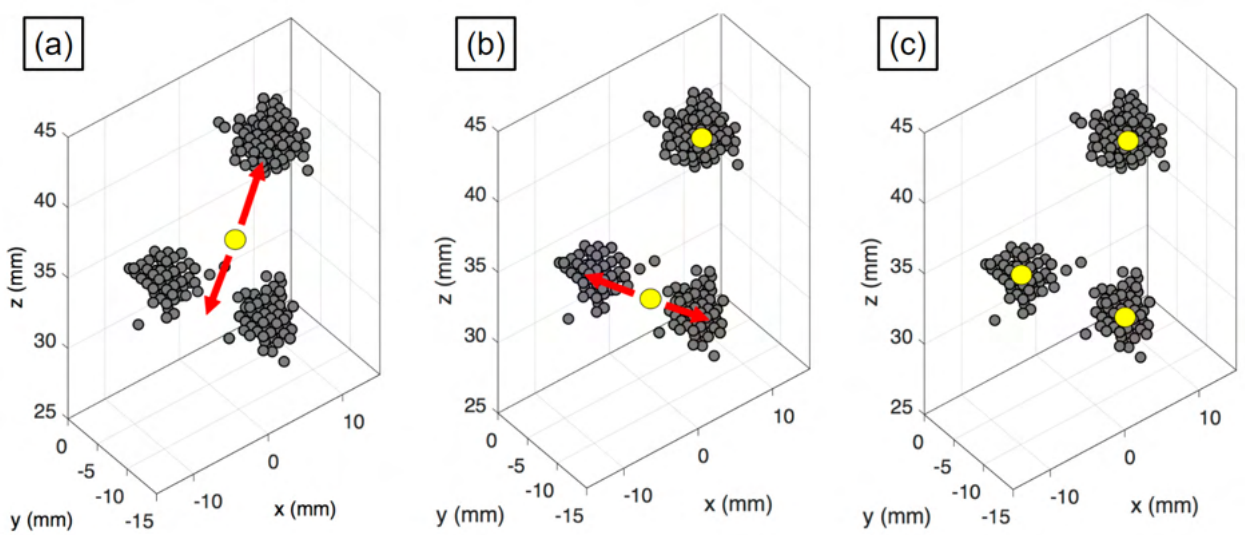

Figure 3.2: Visual depiction of clustering via G-Means. Clustering begins by treating the data as one cluster (a). This is then split into two clusters (b) and then three (c). Further splitting is rejected at this point and assumed that the algorithm has found the total number of tracers [149].

3.2.3 Spatiotemporal B-Spline Reconstruction

Spatiotemporal B-spline reconstruction [73] (SBSR), is a technique developed at Stanford University. The goal of this algorithm is to find the trajectory that is most consistent with a sparse set of LoRs. Unlike most prior PEPT algorithms which were developed with industrial applications in mind, the SBSR method's intended purpose is focused on biomedical applications, having recently been demonstrated through tracking the migration of single cells in live subjects [60]. To this end, the algorithm is designed to track tracers using a minimal number of LoRs. In biomedical applications, this is needed due to the small size of the tracers, and hence typically low activity, used for single-cell tracking and blood flow studies [73].

In this method, the trajectory of a PEPT tracer is modelled as a piece-wise three-

dimensional spline of the continuous temporal variable t [66]. Each of the three spatial coordinates is modelled as a weighted sum of spline basis functions, $B_i(t)$. Therefore, each trajectory, $\vec{r}(t)$, is represented by three sets of basis coefficients, $\vec{a}_i = (a_i^x, a_i^y, a_i^z)$, shown in Equation 3.1.

$$\vec{r}(t) = \sum_{i=1}^{N} \vec{a}_i B_i(t). \tag{3.1}$$

The number, N, of basis functions can be adjusted by the user for each data set. Once the spline basis functions are set up, a convex optimisation algorithm is used to find the spline weights that best match the recorded LoRs. The objective function used for this reconstruction is the sum of the squared distances between each recorded LoR and the time-matched location along the estimated spline trajectory. For instance, if an event was detected at time t_0 by the scanner, the algorithm will compute the estimated position of the source along the spline trajectory at the exact same time t_0, then compute the distance between this point and a single LoR. This distance is then squared and added to the objective function. A decrease in the objective function means that the reconstructed spline trajectory lies closer to the LoRs. As the objective function is convex, the optimisation is guaranteed to converge to a unique trajectory.

The SBSR algorithm also accounts for scattered and random coincidences. These events can have a large impact on the estimated position of the source because the distance penalty in the objective function is squared. Therefore, the distance penalty is capped beyond a maximum threshold, which is set according to the spatial resolution of the PET scanner. For the benchmark tests reported later in this article, the reconstructed data used a distance penalty threshold of 8 mm. An example demonstrating the basic method of the SBSR algorithm is shown in Figure 3.3

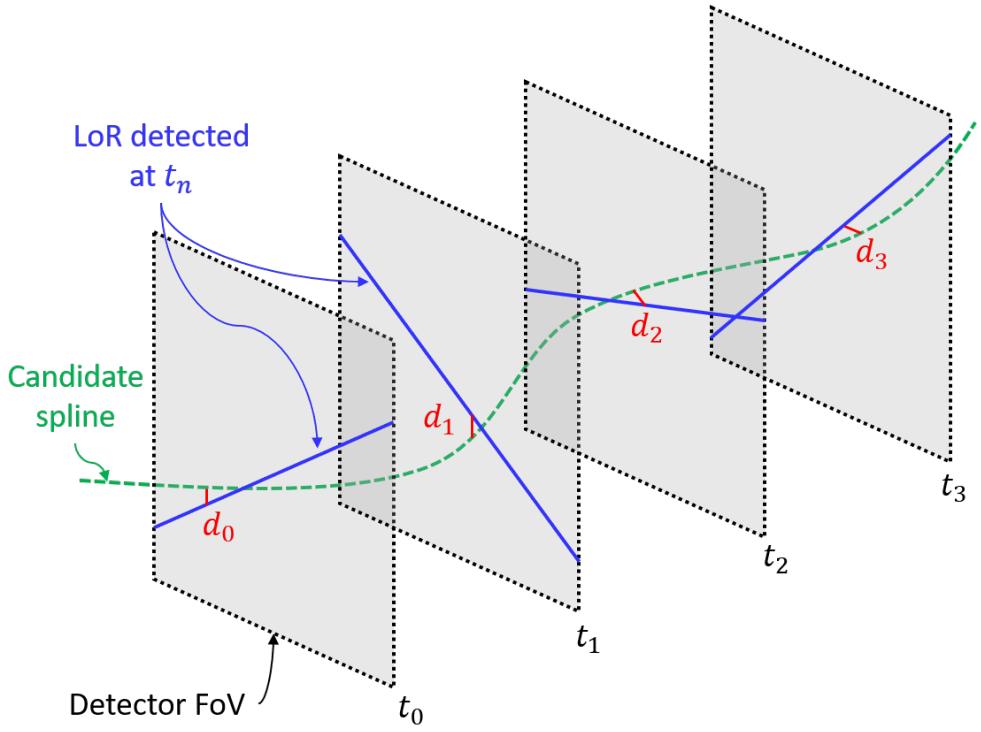

Figure 3.3: Simple illustration of the SBSR method. The algorithm uses a least-squares approach to determine the best-fitting three-dimensional spline for a given set of points [149].

3.2.4 Feature Point Identification Method

The feature point identification (FPI) method was developed at the University of Tennessee, Knoxville as a successor to the G-Means method [138]. The FPI method begins similarly to the G-Means method by using a voxelised grid in which the crossings of LoRs in a sample are summed. Since a voxelised grid is the three-dimensional version of two-dimensional pixels, techniques which were originally developed for image processing can similarly be adapted for voxels to identify and track tracers [117, 23].

In treating the voxel grid as a three-dimensional image, where $N^t(x, y, z)$ is the number of line-crossings at voxel position (x, y, z) during a time step t. The image can be first

smoothed via a convolution-based filter. Typically, this is performed with a box-car kernel over a cubic region of width $2w+1$ where w is the smoothing size, using Equation 3.2:

$$N'^{t}(x,y,z) = \frac{1}{(2w+1)^3} \sum_{i=x-w}^{x+w} \sum_{j=y-w}^{y+w} \sum_{k=x-w}^{z+w} N^{t}(i,j,k) \tag{3.2}$$

In addition to this, convolution with a Gaussian kernel for image smoothing can also be applied. This is useful in cases in which significant background noise is expected due to a large number of scattered events [140]. Examples of the raw line-crossing grid and the smoothed image are seen in Figure 3.4.

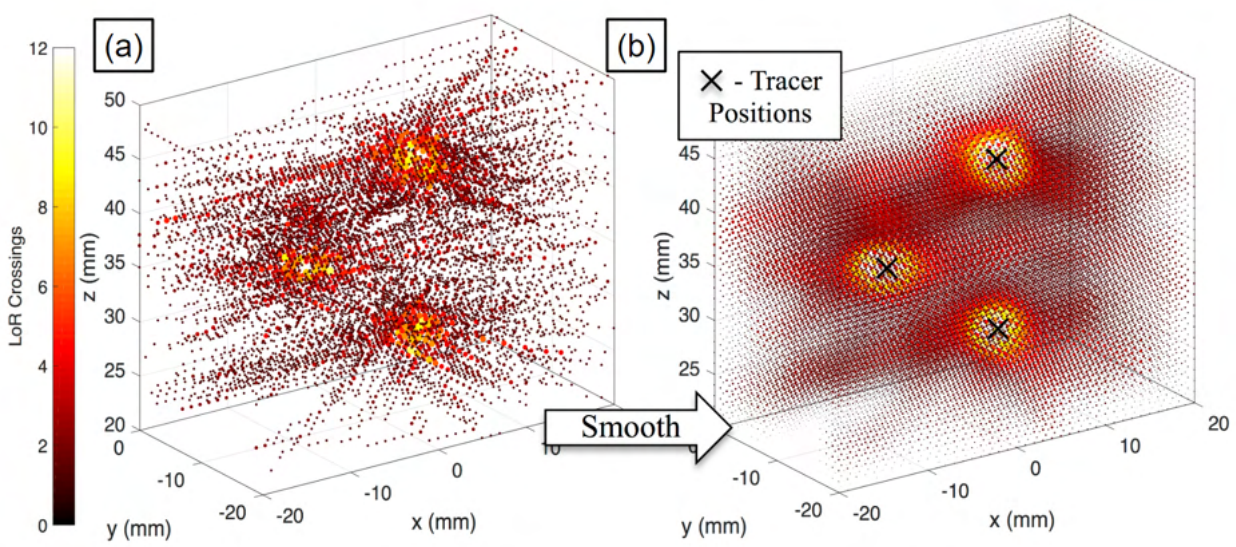

Figure 3.4: Visual depiction of the FPI method, showing (a) unfiltered line-crossing grid and (b) a filtered grid [149].

Next, tracer positions are estimated as local intensity maxima in the smoothed image N'^{t}. These local maxima are taken to be voxels with intensities in the upper r^{th} percentile of a given time frame and a line-crossing value greater than any other voxel within a neighbourhood of $2w+1$. The value r is adjusted based on the amount of noise in an image and the subsequent effective 'brightness' of each tracer relative to the background.

The parameter w serves as an apparent tracer image size, where the image of each tracer is a cube of high LoR density centred at the maxima. However, w also limits the between-tracer separation that can be resolved in a given time step. Below a separation of w voxels, the tracer that appears brightest will be detected, and the others will be ignored but will bias the detected position of the brightest tracer toward the centroid of the tracers involved. Thus, if two tracers approach closely, the one having the higher activity is likely to be followed while the measured history of the other tracer will be split into two. If possible, w should be chosen such that it is smaller than expected tracer-tracer separations.

Final tracer positions are calculated using one-dimensional Gaussian fits in each spatial direction of each identified cluster. This removes a pixel-locking defect that biases calculated positions toward centres or edges of voxels that were observed in the G-Means algorithm. Further, in order to link tracer positions from multiple frames into a coherent trajectory, a nearest-neighbour or tracer location prediction linking algorithm can be used [104, 140].

3.2.5 Spherical Density Method

The spherical density method (SDM) was developed at the University of Cape Town and represents a somewhat different approach than other PEPT algorithms [94]. The SDM requires the user to provide information regarding the size of the tracer and also the positron range [94]. The additional parameters remove sources of uncertainty involved in other methods, allowing the more targeted removal of erroneous LoRs, and thus may potentially yield higher accuracy.

The method's operation is predicated on two main assumptions: (1) a positron's mean free path, λ, between emission and annihilation is smaller than the tracer radius and (2) the

highest and most uniform concentration of annihilation centres lies within a radial distance r^* of the emitted positron.

The first assumption relies on the positron range, which is typically on the order of a millimetre in water. PEPT tracers are often approximately this size and significantly denser than water, meaning the majority of positron annihilation will occur within the tracer [88]. The second assumption seeks to relate the method of tracer activation to the spatial distribution of positron annihilation. With directly activated tracers, the tracer is uniformly activated, meaning the positron-emitting isotope is evenly distributed throughout the volume of the tracer. However, with indirect activation, depending on the exact method used, the positron-emitting isotope may be mostly concentrated on the surface. Whether the tracer is volumetrically activated or surfaced activated can affect the spatial distribution of positron annihilation, which surface-activated tracers having a larger radial distance in which positrons can annihilate.

Using these two assumptions, the SDM algorithm searches for a concentration of LoRs which pass within a distance of each other. The distance between LoRs is determined by the radius of the tracer, mean free path of positrons in the material being imaged, and the method of tracer activation (volumetric or surface activation). This process can then be repeated for all LoRs until a number of concentrations, called sets, s_n, are identified. The set possessing the largest number of LoRs is taken to represent the tracer location. In cases where multiple sets possess the same number of remaining LoRs, the algorithm simply chooses the first such set. As can be seen in the present case, it is typically unimportant which of these sets is chosen, as all true LoRs belonging to the chosen set are members also of the other sets with the same maximum number of remaining LoRs. The SDM is demonstrated in Figure 3.5.

In order to perform multiple tracer tracking, the algorithm is run as described above,

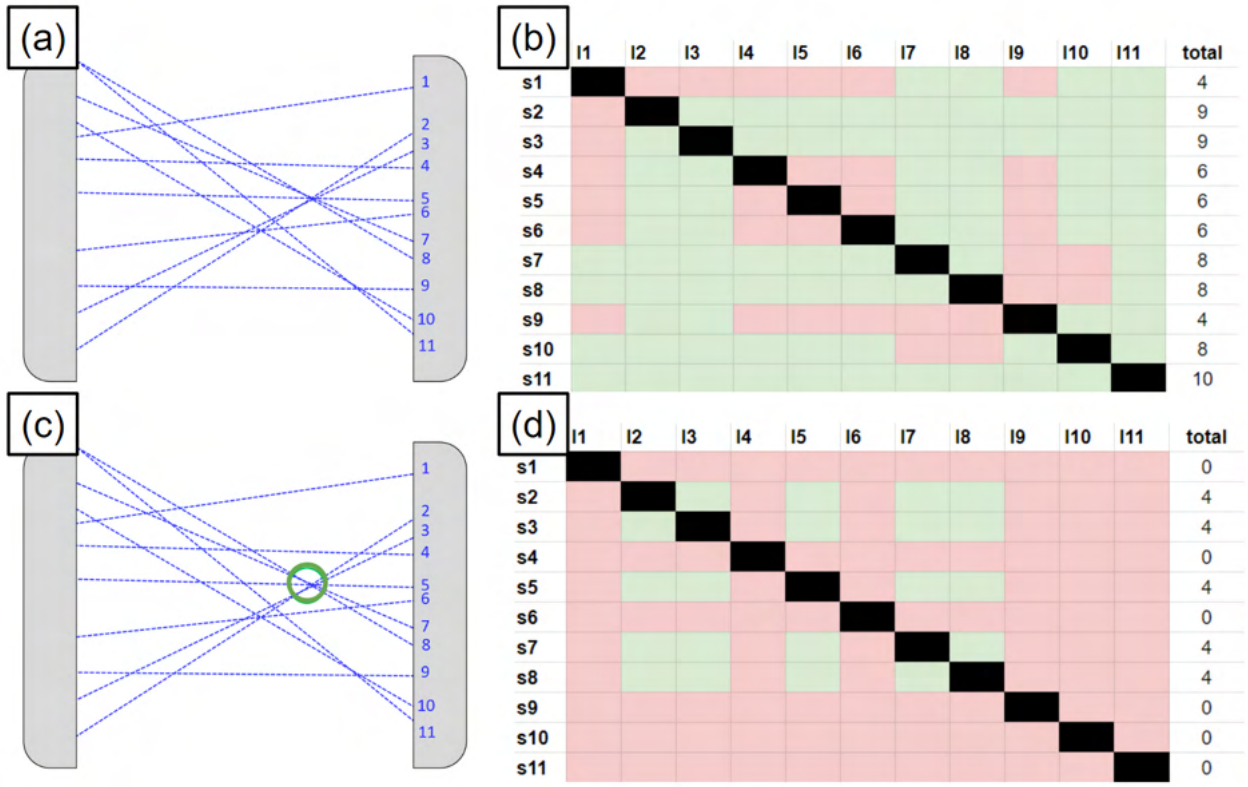

Figure 3.5: Visual depiction of the SDM: (a) first a sample of LoRs in the FOV is shown, (b) then the LoRs which interact with each other are identified, (c) next a region of interest where the most LoRs interact is located, (d) then the LoRs in this region are used to locate the tracer [149].

and the centroid of the set containing the largest number of true LoRs is taken as the position of the first tracer. The LoRs corresponding to this set are then removed from the data set, and the algorithm is re-applied to the remaining LoRs. This process is then repeated until all N_t tracers within the system have been detected. In order to associate computed centroids from a given sample with the subsequent sample of centroids, the algorithm computes the absolute separation distances between each new centroid and those from the previous sample. If the time step between successive frames is suitably small, and the tracer velocities suitably low, it can be assumed that for two successive computed centroids, the tracer would not have moved significantly from its initial computed position. On this basis, the algorithm can

associate each newly computed centroid with the nearest neighbour from the previous set.

3.2.6 PEPT using Machine Learning

The PEPT using machine learning algorithm (PEPT-ML) was developed at the University of Birmingham in 2020 and builds upon modern machine learning techniques [91]. It was built specifically to extract tracer trajectories even in the presence of significant amounts of corrupted LoRs and is capable of locating sources even with low numbers of LoRs per sample. These features help PEPT-ML to locate both small and fast-moving tracers. An important focus in its development was handling non-ideal practical aspects of PEPT experiments, such as unknown numbers of tracers that may collide, leave the field of view, or decay. Accordingly, PEPT-ML introduced a number of general-purpose techniques for identifying tracer trajectories even after intersections or collisions. It was implemented in the context of a broader Python library that also includes auxiliary tools, a user-focused interface, and a straightforward generation of interactive three-dimensional plots.

The PEPT-ML algorithm starts by subdividing a data set of LoRs into samples of equal length or optionally variable length. For each sample, all pairs of LoRs which pass within a user-defined value, d_{max}, from one another are transformed into cutpoints. A cutpoint is the halfway point of the segment connecting a pair of LoRs at their point of closest approach. This is a very different approach than voxel-based algorithms, generating a cloud of point data instead. To cluster these point clouds, PEPT-ML uses an open-source, high-performance implementation of Hierarchical Density-Based Spatial Clustering of Applications with Noise (HDBSCAN) [11, 84].

Another important aspect of this algorithm is the use of 'second-pass' clustering. This exploits the fact that tracer locations found by PEPT algorithms are, more or less, scattered

around the real tracer position. However, it is possible to improve the trajectory by averaging the positions. This is achieved by reapplying the HDBSCAN algorithm to a sample of the cluster centres. This treats the PEPT locations as a point cloud, similar to the cutpoints, and finds the centroid, yielding a better estimate of the tracer location. Additionally, typically only approximately 30 locations are needed in this step. This is effectively a smoothing filter, but when the temporal resolution is high with respect to the tracer velocity this allows second-pass clustering to retain fidelity in tracer trajectory. Further, second-pass clustering can be beneficial by excluding erroneous tracer positions which are further away from other centroids. While this method is more computationally expensive than many other algorithms, due to the high-performance, parallel implementation, PEPT-ML is still very competitive speed-wise [91].

In the case of multiple tracers the reconstructed positions must be split into individual trajectories. The current trajectory separation approach uses a spatiotemporal-based segregation step and signature-based linking. The segregation stage involves a single-linkage clustering algorithm which connects all points and then removes paths longer than a maximum distance. However, points that are in the same spatial position at different times must be differentiated. To do this, a custom spatiotemporal metric is used, which limits which locations can be connected. A second stage, making use of the cluster sizes, can then be employed if it is expected that tracers may collide or intersect with one another. This signature is the number of cutpoints around each tracer location. When tracers are of differing activities, it is expected that more of the LoRs in a sample will belong to the higher activity tracers. Since the number of cutpoints produced by a sample of LoRs scales quadratically with the number of LoRs meeting the max distance cutoff, this means that even small differences in the activity can be registered and used to correctly piece together tracer trajectories. The basic workflow of the PEPT-ML algorithm is shown in Figure 3.6 [91].

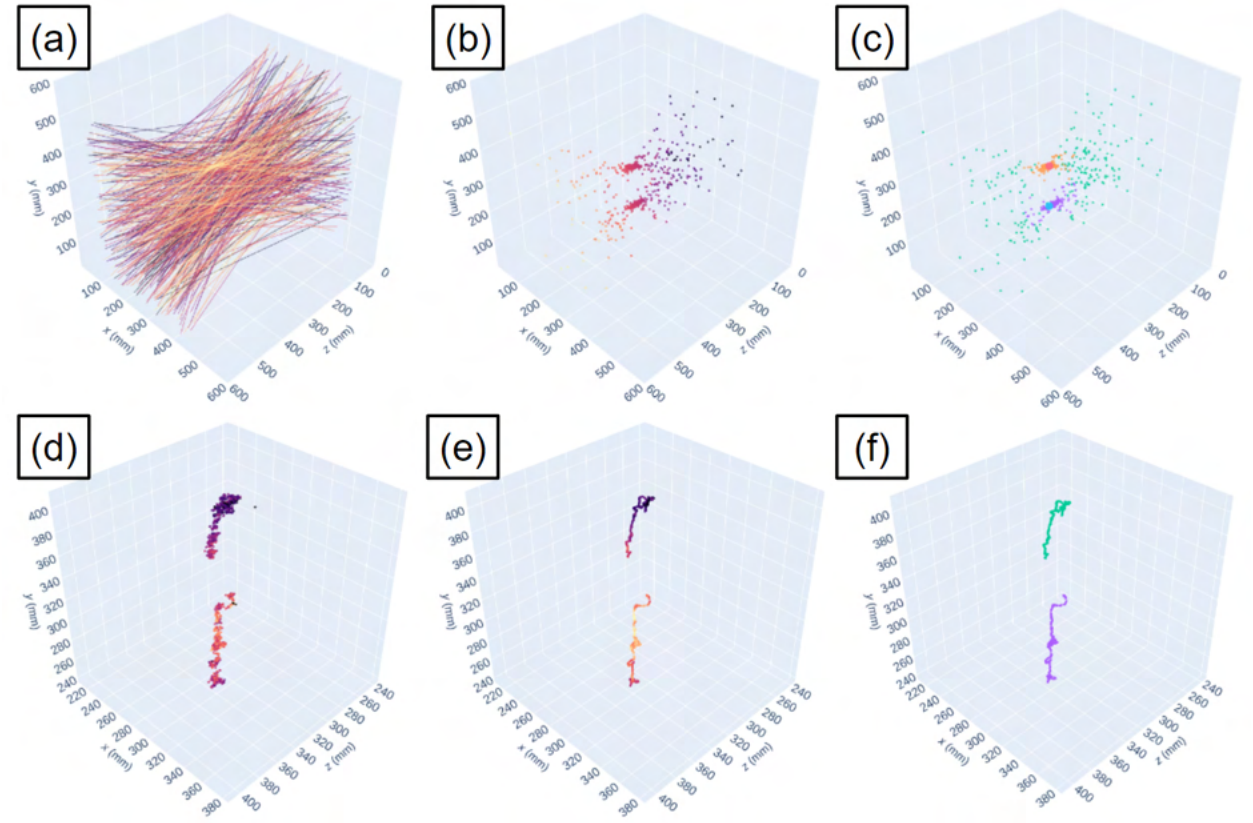

Figure 3.6: Illustration of steps used by PEPT-ML in tracking two 1 mm diameter tracers in a bubbling fluidised bed, imaged at the University of Birmingham. Panel (a) depicts 400 LoRs from the data set. Panel (b) shows their corresponding cutpoints, computed for pairs of LoRs which are closer than $d_{max} = 0.5$ mm. Panel (c) colour-codes the cutpoints by the cluster IDs assigned by HDBSCAN: green represents noise, purple is cluster one and orange is cluster two; centroids are shown as large circles, blue and red, respectively. Panel (d) depicts the tracer locations found after the first pass of clustering of 80,000 LoRs; panel (e) shows the corresponding tracer locations after the second pass of clustering. Both panels (d) and (e) colour-code points by their cluster size. Panel (f) illustrates the separated tracer trajectories, colour-coded by the trajectory ID [149].

3.2.7 PEPT using Expectation Maximisation

The PEPT expectation-maximisation algorithm (PEPT-EM) was developed at the University of Edinburgh in collaboration with the University of Birmingham, expanding on the original concept of the Birmingham Method [82]. This method clusters LoRs into 'inlier' and 'outlier' sets similar to how the parameter f is used to separate true and corrupted LoRs. Using a maximum-likelihood approach based on Bayesian statistics, the LoRs in a set of data can be allocated to one of many clusters which are iteratively recalculated until convergence [82]. Scattered or random LoRs are assigned to a larger outlier cluster.

The PEPT-EM method is novel in that it is based on a physical model of positron emission and LoR generation. For a set of K tracers, each tracer has a centre defined as $\vec{x}_k$ where $k = 1, 2, \cdots, K$ and emits positrons that annihilate a short distance away from the centre at y_k, where the probability of annihilation at y_k is defined by a Gaussian distribution around the centre $\vec{x}_k$ with a variance σ_k^2. In addition, the relative weight of each cluster is defined by ρ_k to characterise their relative activity.

This sets up the problem as a Gaussian mixture, which can be solved using expectation maximisation. Present in all samples of LoRs is an outlier cluster where $k = 0$, centred at the middle of the FOV and with a variance of the order of the size of the FOV. The inverse of this variance is termed α, such that $\alpha^{-3/2}$ is of the order of the volume of the FOV. In practice, this value is a constant that requires fine-tuning. The weight of this cluster, ρ_0 is also defined by normalising ρ such that $\sum_{k=0}^{K} \rho_k = 1$. Using this information, the log-likelihood of finding each cluster with parameters $\vec{x}_k, \sigma_k^2, \rho_k$ in a set of lines $\mathcal{L}$ as Equation 3.3. With an initial guess of the clusters, the solution is optimised iteratively until the likelihood is maximised.

$$L\left(\{\vec{x}_k, \sigma_k^2, \rho_k\}_{1\leq k\leq K} \middle| \mathcal{L}\right) = \sum_{\ell\in\mathcal{L}} \log\left(\rho_0\alpha + \sum_{k=1}^{K} \rho_k \sigma_k^{-2} \mathrm{e}^{-D^2(\vec{x}_k,\ell)/2\sigma_k^2}\right). \tag{3.3}$$

In order to use expectation maximisation, unobserved, or latent, variables are introduced. These variables are used to simplify the model. The iterative nature of the algorithm requires alternating between the calculation of the latent weights and maximising the likelihoods until convergence is achieved. The convergence of the inlier and outlier cluster is illustrated in figure 3.7.

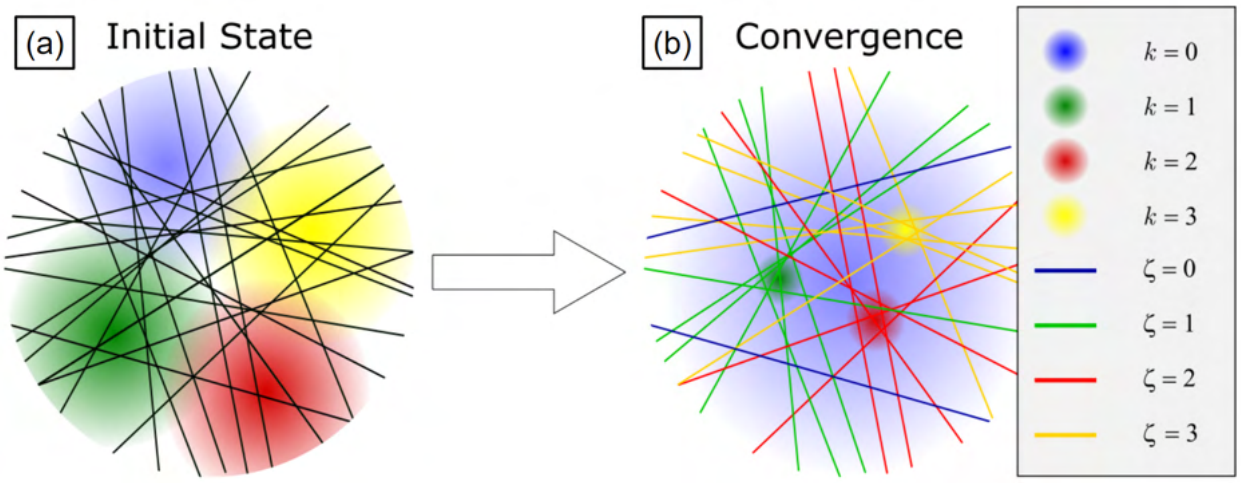

Figure 3.7: (a) A schematic of the initial clusters generated by the PEPT-EM algorithm, showing clusters $k = 0, 1, 2, 3$. (b) After convergence, the large outlier cluster, $k = 0$, accounts for detected LoRs that have undergone scatter or are due to random coincidences, while smaller clusters converge using their most-likely LoRs [149].

The number of iterations to reach convergence can vary depending on the number of clusters and the initial guess. Typically for the first frame, a number of the order of 1000 iterations is used. However, when tracking trajectories, convergence can be achieved using fewer iterations provided that the time interval between batches of LoRs is sufficiently small that the tracer centroids have only moved a small distance. Consequently, by updating the position of individual clusters trajectory separation is implicitly achieved.

3.3 Description of Benchmarks and Data Analysis

The single tracer tests are designed to test the minimum number of LoRs needed for position reconstruction, the maximum tracer velocity, robustness to noise, and the ability to track tracers near the edge of the FOV. All algorithms are tested on these metrics. However, multiple tracer tests are also conducted using the G-Means, FPI, SDM, PEPT-ML, and PEPT-EM algorithms. The multiple tracer benchmarks are designed to test the minimum separation distance between tracers needed to resolve both tracers, the ability to distinguish false positives from real tracers, test the trajectory linking of moving tracers, and compare the ability to track large numbers of tracers.

For all simulated tests, 10 MBq of fluorine-18 is prescribed as the activity of a 2 mm diameter, volumetrically activated glass tracer in air, unless otherwise specified. The particle emitted by the source are positron and the GATE physics list used is 'emstandard'. Within a short distance of emission, the positrons will annihilate with an electron and produce two 511 keV coincident gamma rays. In instances of multiple tracers, a total activity of 10 MBq is divided equally between tracers. The tracer size, the isotope used, and total activity was chosen to reflect common practices in PEPT experiments [148].

3.3.1 Single Tracer Tests

Minimum LoRs

The first test, and the most fundamental, is designed to evaluate the accuracy of PEPT algorithms as a function of the number of LoRs used, as well as their ability to locate tracers using small numbers of LoRs. The test simply requires the reconstruction of a single, static, point-like tracer of a prescribed 'activity', i.e. a number of LoRs ranging between $N = 10$ and

$N = 10,000$. For each sample size, N, 10 repeat tests are performed, each with a different random set of LoRs. The tracer location is fixed at 0, 0, 10 mm near the centre of the FOV. Placing the tracer exactly in the centre of FOV is avoided since this could result in randomly distributed noise giving the appearance of higher-than-expected accuracy.

The algorithms must use all LoRs in a sample to locate the tracer once. Thus, for each N, 10 repeats are used to assess the average error from the prescribed tracer location as well as the standard deviation of the errors. The performance of each PEPT algorithm is judged on the accuracy of the mean reconstructed tracer position using Equation 1.14 and if a location is able to be returned at all for low numbers of LoRs.

Maximum Velocity

This test establishes the ability of a PEPT algorithm to track fast-moving tracers. It involves reconstructing the trajectory of a constant velocity moving tracer following a one-dimensional 'sawtooth' path. The test is repeated for 10 different velocities ranging from 0.04 m/s to 20 m/s, a span designed to represent and possibly exceed the range of velocities that might be encountered in a real PEPT experiment. The range of motion is between ± 10 mm and the tests are conducted separately for axial and transverse motion with respect to the PEPT detector systems. An example of this trajectory is shown in Figure 3.8.

The sawtooth path is specifically chosen to introduce significant accelerations and discontinuities and thus further challenge the algorithms. Since the trajectory of the tracer is reconstructed and there are several positions comprising a trajectory, there is a trade-off between spatial and temporal resolution. Thus, to compare the results of the different algorithms, which may use different amounts of LoRs per location, they are compared through the combined spatiotemporal resolution of the trajectory which is the mean position error divided by the inverse square root of the number of locations found. This penalises low

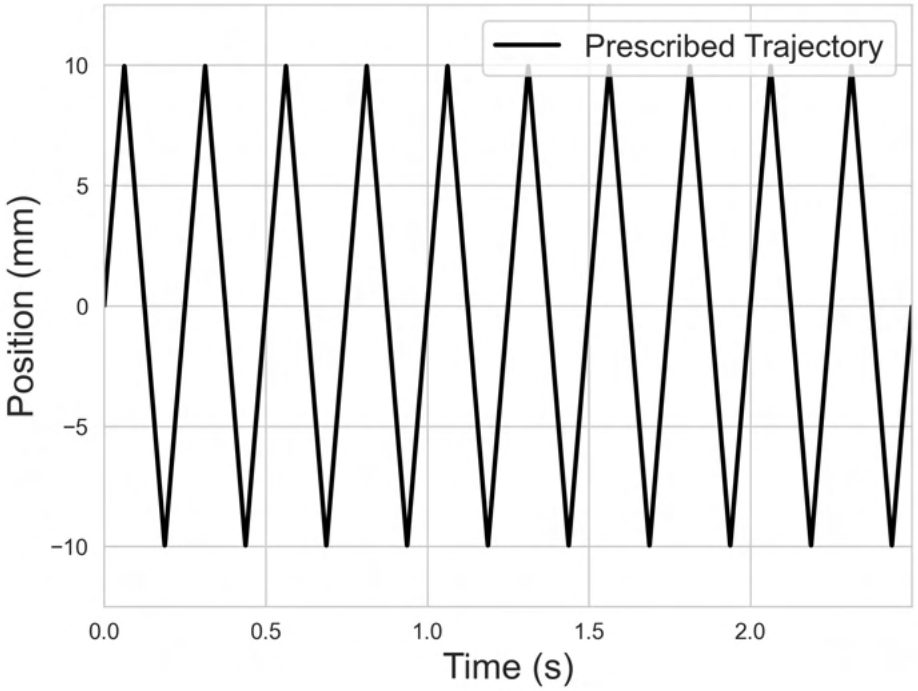

Figure 3.8: An example of the sawtooth trajectory used in the velocity test.

accuracy and low numbers of returned positions. The location error is again calculated as per Equation 1.14 using the prescribed position for the tracer and used with Equation 1.22 to determine the spatiotemporal resolution. Each algorithm is judged on the spatiotemporal resolution as a function of the tracer velocity.

Robustness-to-Noise

The goal of this test is to establish the ability to handle noisy data. In real PEPT experiments, noise is caused by random and scattered LoRs [148]. To produce realistic noise, scattered LoRs are generated by placing the GATE-simulated static source at the centre of a 50 mm diameter sphere of scattering material near the centre of the FOV at 0, 10, 0 mm. The physical properties of the sphere are varied to elicit different levels of scatter, and thus different levels of noise. The scatter media used are air, polyethylene, water, polyvinyl

chloride (PVC), aluminium, sodium iodide (NaI), cadmium zinc telluride (CZT), stainless steel, copper, and bismuth germanate (BGO). The details of these materials, including their densities and attenuation coefficients, are listed in table 3.1. Figure 3.9 illustrated the set-up of the test.

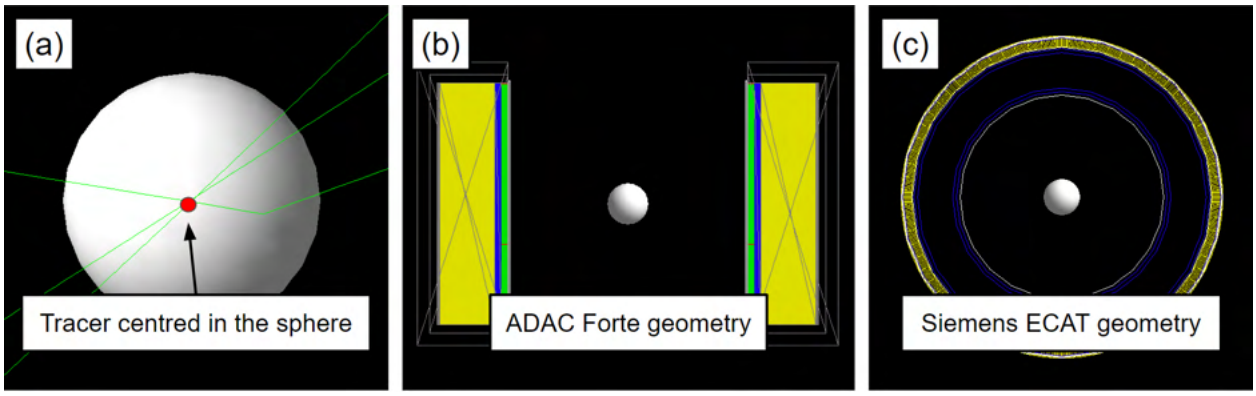

Figure 3.9: Figure illustrating the (a) scatter sphere, (b) scatter sphere in the simulated ADAC Forte PEPT detector system, and (c) the scatter sphere in the simulated Siemens ECAT PEPT detector system [149].

From the simulation, 10,000 LoRs are generated and provided to the users of the PEPT algorithms. Users have the choice to provide a single position for the tracer or an average position. The algorithms are judged on their ability to accurately locate the tracer through Equation 1.14. The results are presented as the spatial accuracy as a function of the attenuation coefficient of the sphere.

Field-of-View

In this test, a simple, linear tracer trajectory is simulated, with the tracer starting outside the PEPT detector system, passing through the centre of the FOV, and then out the other end. The tracer moves at a constant 100 mm/s velocity starting at -250 mm from the centre of the FOV and finishing at 250 mm. The goal of this test is to establish the ability of the

Table 3.1: List of materials and properties used in the robustness to noise benchmark tests [129, 149].

Material	**Density** (g/cm^3)	**Cross-section** (cm^2)	**Attenuation Coeff.** (cm^{-1})
Air	0.00129	0.0806	0.000111
Polyethylene	0.960	0.0986	0.0946
Water	1.00	0.0960	0.0960
PVC	1.65	0.0890	0.147
Aluminium	2.70	0.0837	0.226
NaI	3.67	0.0933	0.342
CZT	5.68	0.0896	0.509
SS304	7.92	0.0832	0.659
Copper	8.96	0.0827	0.741
BGO	7.13	0.1501	1.070

algorithms to track tracers not only in the centre of a detector's FOV, where the spatial and temporal resolution are the greatest but also near the edges where sensitivity and the fraction of true LoRs may be reduced. Figure 3.10 illustrates the trajectory of the simulated tracer through the FOV of the PEPT detector systems.

The PEPT algorithms are judged on their ability to accurately locate the tracer with the spatial error calculated from Equation 1.14 using the prescribed positions of the tracer. In addition to the mean error and the standard deviation of the errors, the spatiotemporal resolution is calculated using Equation 1.22 and the instantaneous spatial errors are presented as a function of tracer position.

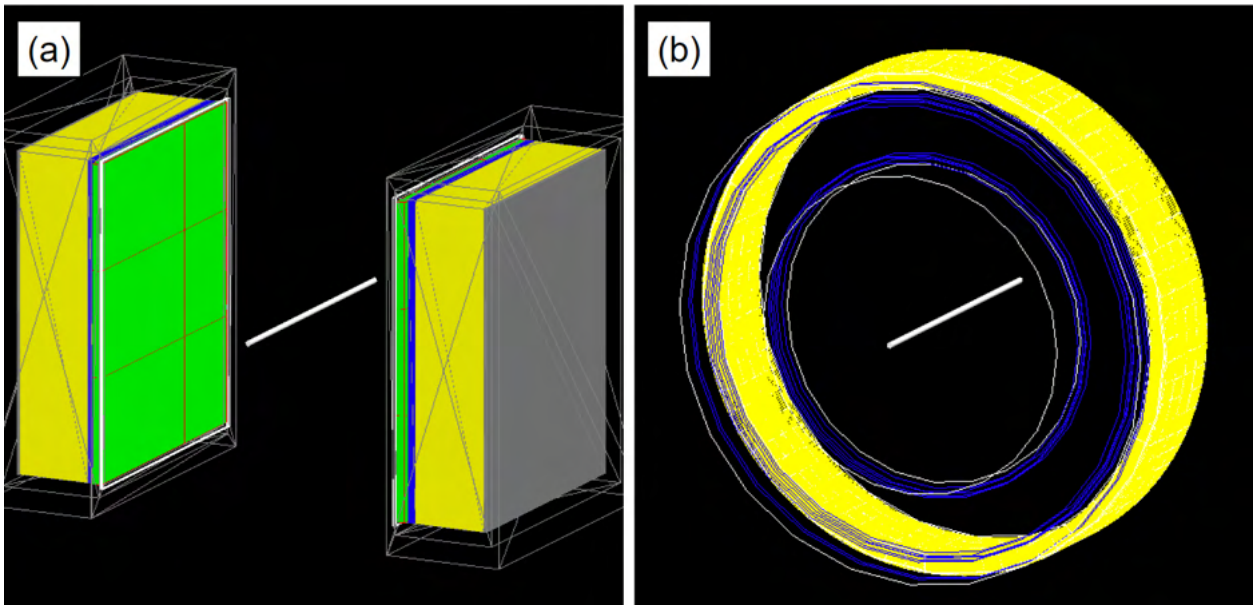

Figure 3.10: Figure illustrating the trajectory of the tracer moving through the centre of the FOV for the (a) ADAC Forte and (b) Siemens ECAT [149].

3.3.2 Multiple Tracer Tests

Minimum Tracer Separation

In this test, a pair of 2 mm diameter tracers, each 5 MBq, are separated by a variable centre-to-centre distance. The value of this distance is varied from 30 mm down to 2 mm, which is the distance the two simulated tracers are in contact, representing collision. Two sets of simulations are conducted with one set having the tracer separated axially with respect to the PEPT detector system and another transverse. The axial direction in the Forte is defined as perpendicular to the detector face, while in the ECAT it is parallel and vice-versa for the transverse direction. Spatial resolution is typically worse in the axial direction due to oblique LoRs, presenting additional challenges.

The purpose of this test is to assess the ability of a given algorithm to resolve two tracers in close proximity, as opposed to falsely merging them into a single location which

has been seen in other PEPT experiments, [49]. This is an important aspect for PEPT algorithms when multiple tracers are used since during an experiment tracers may pass near to each other or collide.

In each simulated case, 10,000 LoRs are provided to the users and either the two trajectories, or two average positions are returned. To assess their ability to distinguish the tracer it is first noted if the algorithm returns a merged position or two separate positions, then the mean spatial error is calculated using equation 1.14.

False Positives

In this test, the algorithms must locate an array of 4, 8, and 16 tracers arranged, respectively, to form a square, a cube, or a 'tesseract'. These geometries are deliberately chosen to include high degrees of symmetry to create areas of overlap between LoRs from different tracers which may be falsely interpreted as tracer positions. In the case of the square arrangement, for example, high concentrations of intersecting LoRs at the centre of the square could be falsely interpreted as a 5th tracer location. In the case of a cube or tesseract this issue is exacerbated, as well as introducing other such intersection points. This test is therefore designed to assess an algorithm's resilience to false tracer locations. Figure 3.11 shows the relative positions of the tracers for each of the three shapes.

All arrangements share a vertex separation of 10 mm, meaning that no two tracers lie closer than this distance. This specific minimum separation distance is chosen to be greater than the detector resolution and corresponds to the separation above which the tracers are not likely to be merged. This ensures that the test does not unfairly re-examine the abilities tested previously. The algorithms are compared based on the presence of any false positive in the returned positions and also the mean error for the tracers, calculated using Equation 1.14.

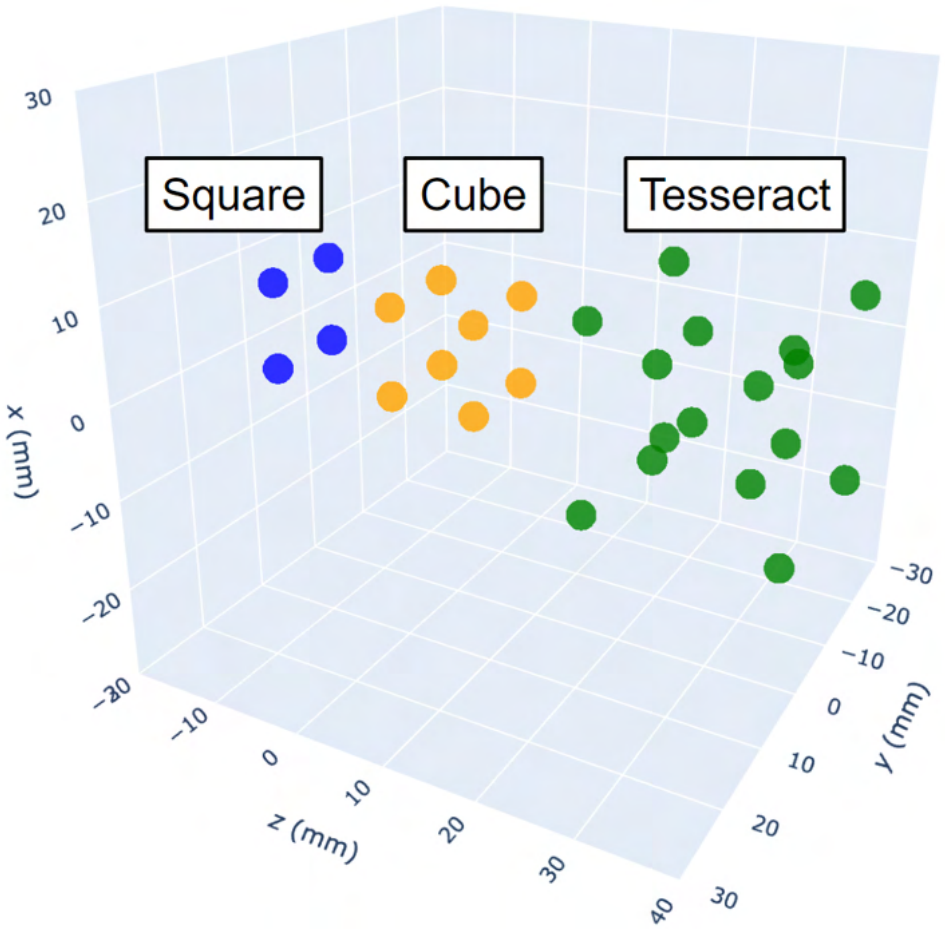

Figure 3.11: Figure illustrating the positions of tracers for the false positive test. The three tests correspond to groups of 4, 8 and 16 tracer arrangements, respectively, in a square, a cube and a tesseract. Note that in reality each group of tracers are simulated separately and centred at 0, 0, 0 mm.

Trajectory Linking

This test aims to assess not only an algorithm's ability to locate a tracer but also to follow the tracer across multiple time steps. To test this, three tracers are placed on the edge of an imaginary sphere and rotated about two axes. Each tracer is separated from the others by a fixed distance of 34.5 mm, moving continuously across the surface of a sphere of diameter 40 mm at a constant velocity of 0.8 m/s. This specific motion was chosen because while

the tracers are always separated spatially, their trajectories will intersect if plotted without respect to time. As such, this is a test not only of the raw location accuracy of a method but also of the associated trajectory separation. Figure 3.12 shows the initial position of the three tracers and their trajectory history.

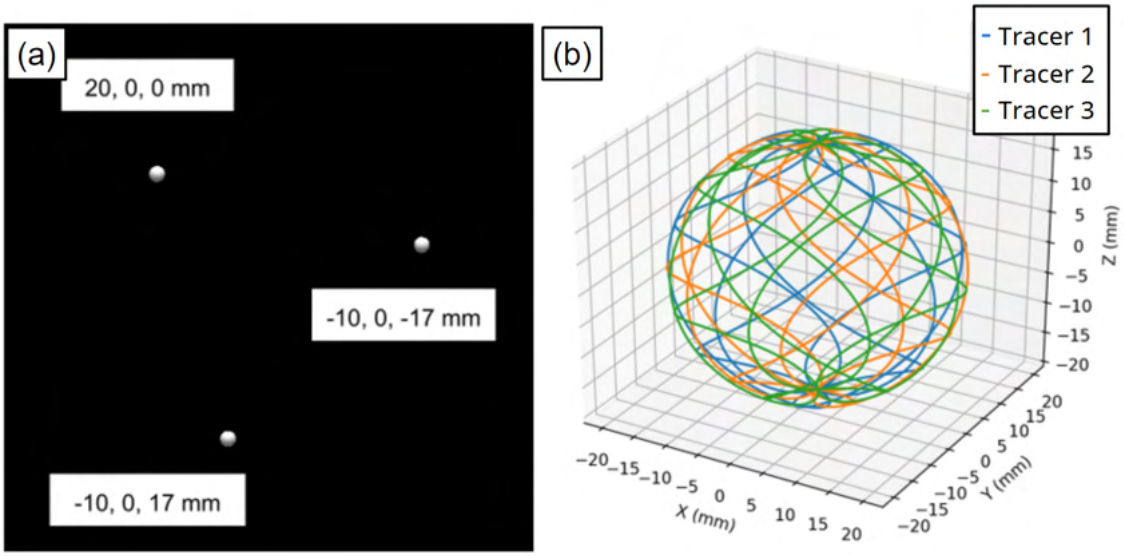

Figure 3.12: Figure illustrating the (a) initial position of the three tracers and (b) their trajectory history for the trajectory linking benchmark [149].

To assess the performance of the PEPT algorithms in this test, the reconstructed trajectories are compared individually against the relevant prescribed positions. In particular, it is observed whether the positions in a trajectory all correspond to the same simulated tracer. Additionally, the mean spatial error and standard deviation of the errors are calculated using Equation 1.14. Further, the spatiotemporal resolution is calculated and compared among the PEPT algorithms.

Large Numbers of Tracers

The eighth test is intended to assess limitations on the number of tracers, N_t, that can be tracked simultaneously by a given algorithm, evaluating how the algorithm's accuracy is

influenced by increasing N_t. The test consists of nine distinct data sets, each containing a different number of tracers placed in random positions throughout the FOV of the PEPT detector system. The tests are deliberately conducted with pseudo-random numbers of tracers such that those attempting the benchmarking tests were not able to make an educated guess as to the expected number of locations in a given data set, and thus deliberately tune their algorithms to more precisely detect this number. For each set of tracers simulated, 10,000 LoRs were provided to users of the PEPT algorithms and a single position or average position was returned for each tracer.

Two aspects of the codes' performance are directly tested: first their ability to locate large numbers of tracers with no prior information regarding the specific number to be located, and second the accuracy of tracer locations. The results of PEPT algorithms' performance are presented as the percentage of tracers correctly found and also the mean spatial accuracy calculated using Equation 1.14.

3.4 Results

The results of each benchmark test are presented in this Section in the same order as they were described in Section 3.3. The goal of these tests is not to show whether one algorithm is superior to all others, but rather to show their relative strengths and weaknesses. Beginning with the minimum LoRs test, the results of each algorithm are shown in Figure 3.13.

From Figure 3.13, it is shown that the spatial accuracy increases when more LoRs are used to locate the tracers. This holds true for all of the algorithms tested. Interestingly, when the results are plotted on a log-log plot, the precision of all algorithms show error decreases proportionally to the inverse square root of the number of LoRs used, which matches previous studies with the Birmingham Method [100, 97]. This is because for stationary tracers, larger

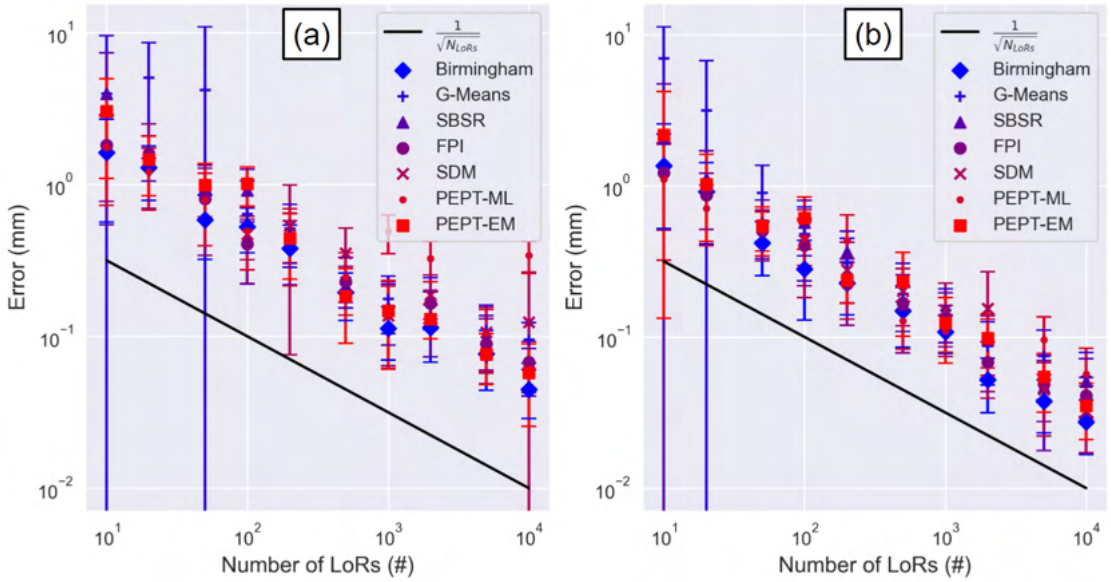

Figure 3.13: (a) The spatial accuracy of the PEPT algorithms as a function of the number of LoRs used in the ADAC Forte geometry and (b) the results in the Siemens ECAT geometry. Additionally, the proportionality to the inverse square root of the number of LoRs used is demonstrated [149].

sample sizes increase the spatial resolution since there is more information about where the tracer is located, with error decreasing for static tracers proportionally to the inverse square root, i.e. the standard error. However, the proportionality constant of different algorithms appears to be different.

Despite being the oldest algorithm, the Birmingham method achieves the highest overall accuracy in both PEPT detector system geometries and performs consistently well across all tests. In many cases, it is observed to outperform even the most recent algorithms, though PEPT-ML and FPI are observed to yield better results in some cases, but within the margin of error of the Birmingham Method. PEPT-EM behaviour is somewhat more unpredictable. It is believed that in the case of having only 10 LoRs, no outlier cluster could be identified, and PEPT-EM solely relied on the alpha parameter being well-tuned to reject

LoRs that are far away from the fitted cluster, resulting in relatively poor performance. 5 of the 7 algorithms managed to produce a location using only 10 LoRs, though in all cases the error was > 1 mm. For $N \leq 50$, SDM did not return a result and while the G-Means algorithm returned values which in many cases were erroneous, leading to significant error bars, reflecting the large standard deviation from the mean location. Ultimately, all algorithms were found to yield sub-millimetre accuracy with as few as 100 LoRs and a clear relationship between the number of LoRs used and the spatial accuracy was identified.

In the maximum tracer velocity test, the PEPT algorithms are observed to exhibit a somewhat richer phenomenology than the previous benchmark. The results for this test are shown in Figure 3.14. The Birmingham Method PEPT algorithm, though typically producing the best results in the previous, stationary-source tests, now exhibits the largest location errors in a significant majority of cases. Conversely, the PEPT-EM algorithm, which produced some of the weakest results for the stationary source in the ECAT geometry, now produces the smallest errors for all but two cases. In the ADAC geometry, while PEPT-EM performs strongly for slow-moving tracers its accuracy falls off rapidly as velocity increases, to the point where at high velocities it gives the largest error of all algorithms tested. This is likely due to the fact that at the highest velocities the algorithm cannot meaningfully identify both an 'outlier' and 'inlier' cluster.

PEPT-ML consistently performs well across both PEPT detector system geometries and all velocities. The SBSR algorithm also performs extremely well, in particular in the ADAC geometry, and in the case of fast-moving tracers. This may be expected since the SBSR algorithm was designed to accurately locate moving tracers using minimal numbers of LoRs. All algorithms degrade in spatiotemporal resolution when the tracer velocity increases. The effects of this are demonstrated through Figure 3.15, showing that at high velocities the trajectories are sparser and more inaccurate. This also shows the effect PEPT detector systems have on the quality of data. The ECAT geometry has a higher sensitivity than the

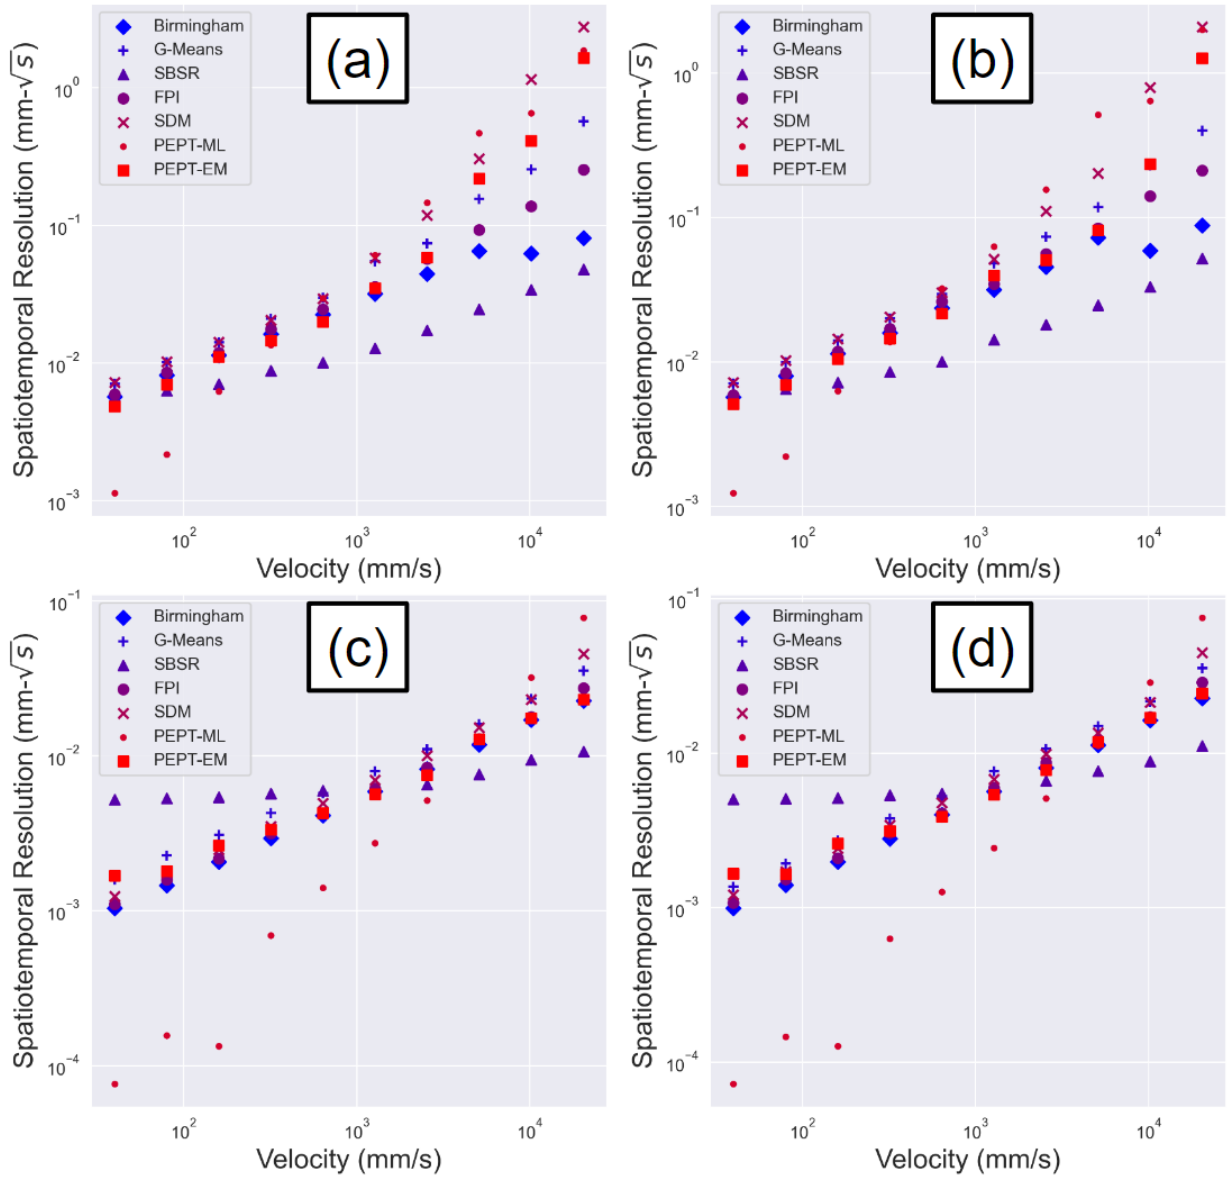

Figure 3.14: The error in PEPT detected positions increases when tracers move at a higher velocity. (a) The Forte detector geometry and movement in the axial direction and (b) shows Forte detector geometry and movement in the transverse direction. (c) The ECAT detector geometry and movement in the axial direction, while (d) shows the ECAT detector geometry and movement in transverse direction [149].

513 Forte, thus producing more LoRs and better trajectories.

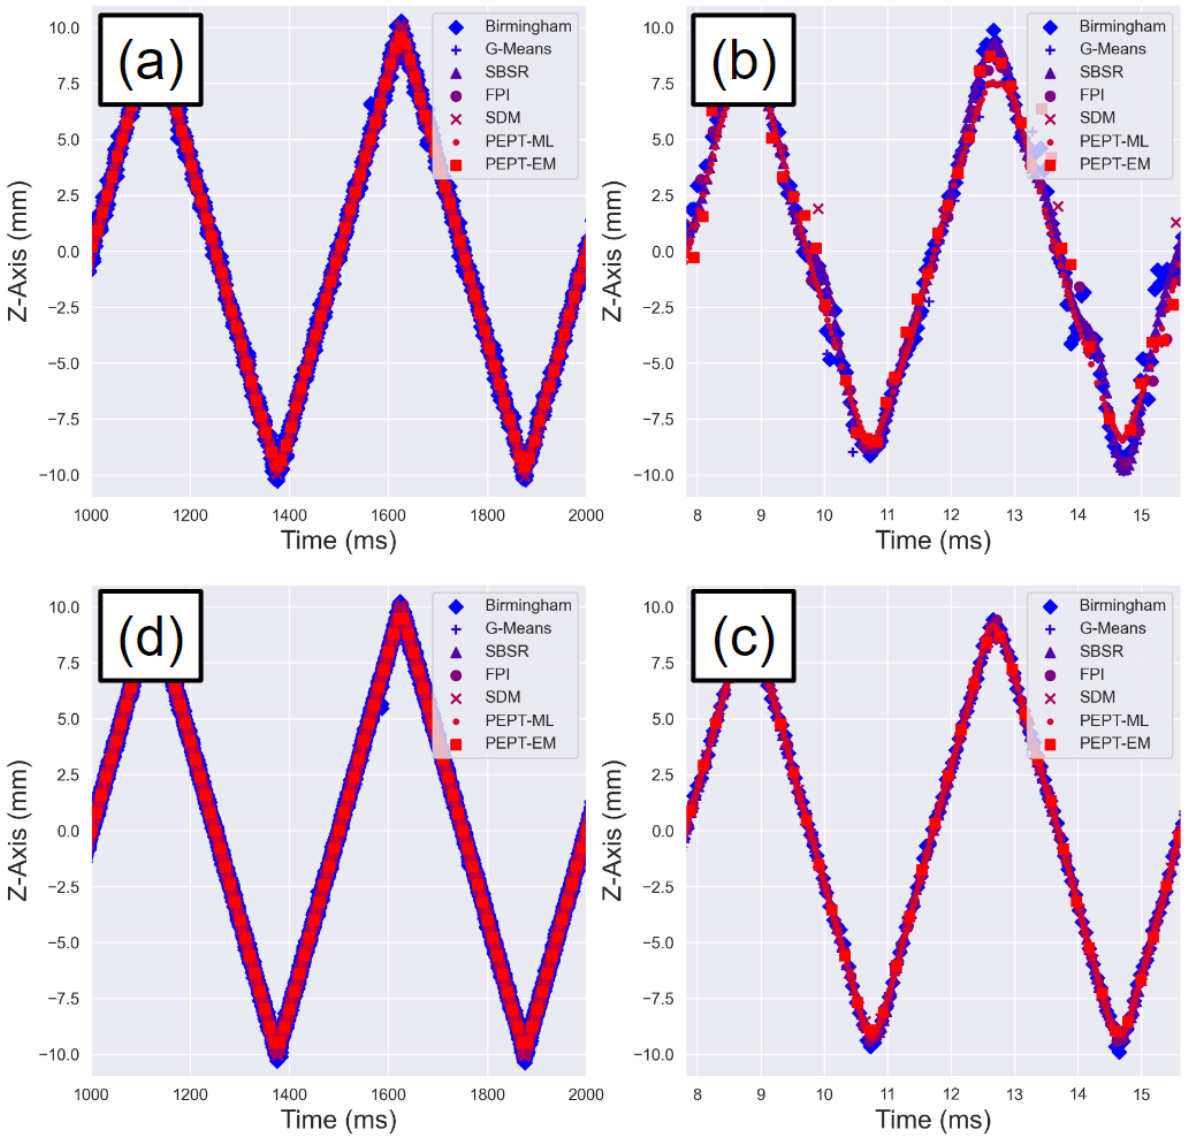

Figure 3.15: Example trajectory for the velocity tests over two periods of motion. (a) Forte geometry and for $v_{tracer} = 80$ mms^{-1}. (b) Forte geometry for $v_{tracer} = 10240$ mms^{-1}. (c) ECAT geometry and for $v_{tracer} = 80$ mms^{-1}. (d) ECAT geometry for $v_{tracer} = 10240$ mms^{-1} [149].

Ultimately this benchmark demonstrates that tracer velocity plays a role in PEPT algorithm performance, with higher velocity tracers not able to be tracked as well as static tracers [18]. This is to be expected, especially at higher velocities since the tracer will have

moved a non-negligible distance relative to its size in the time it took to collect the LoRs in single sample [97].

In the robustness-to-noise test, unlike the two preceding tests which produced relatively clear trends and generally straightforward results, the results are more complex. While the overall trend is a general increase in error when more noise is present, the individual behaviours of several algorithms are observed to be somewhat erratic. This is perhaps most pronounced in the PEPT-ML and PEPT-EM algorithms, which for different scatter spheres can yield both the highest and lowest accuracy of all tested algorithms. The full results are presented in Figure 3.16.

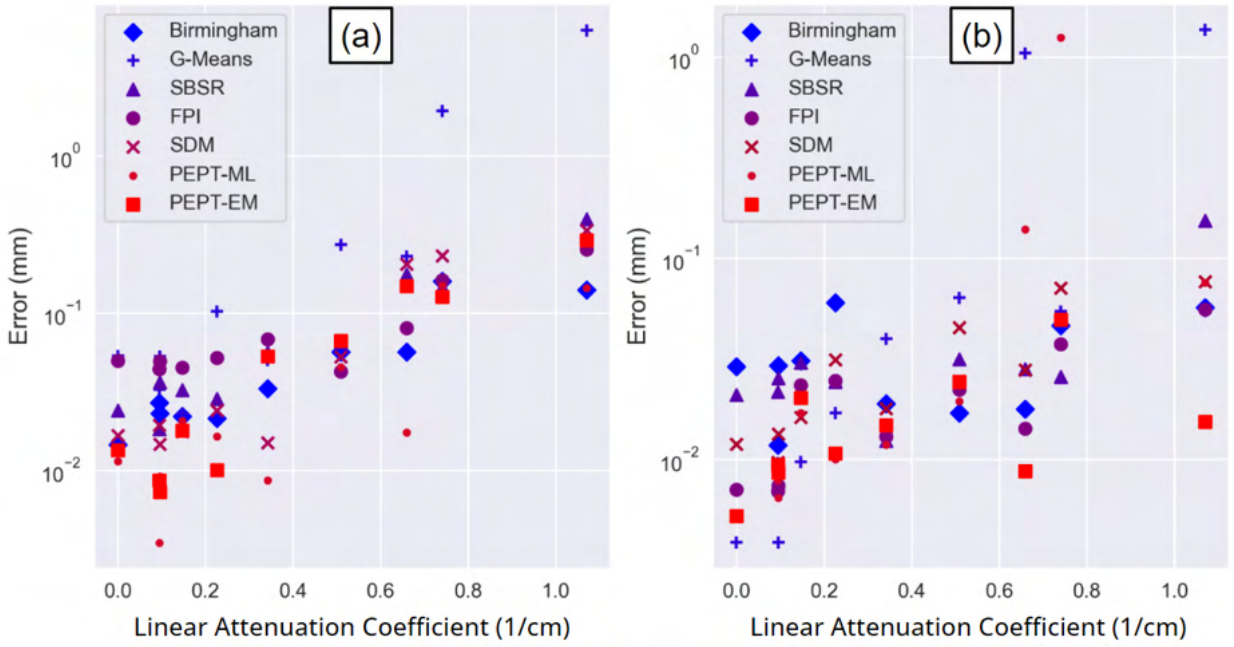

Figure 3.16: The error in PEPT detected positions increases when more LoRs are corrupted by scattering. (a) Forte detector geometry and (b) ECAT detector geometry [149].

In terms of yielding a continuous trend, only the G-Means algorithm offers any such predictability, suggesting it is the least sensitive to noise. However, on average, it yields the lowest accuracy. The FPI, SDM, and SBSR methods all yield relatively consistent and strong

results across both detector geometries. The Birmingham method shows both relatively good consistency and accuracy, though it is somewhat less reliable in the ECAT, perhaps due to incorrect tuning of the f parameter.

The origins of the variation in location error are not entirely clear. The most likely explanation is that under high-noise conditions, random fluctuations in the LoR density field may be comparable in magnitude to that registered at the true tracer location and that certain algorithms are more susceptible to this. These findings raise the question as to whether ability or reliability is the most important factor when choosing an algorithm. If the former, then PEPT-EM would be the most sensible choice but would run the risk of producing erroneous tracer locations. If the latter, then the FPI, SDM, or SBSR method would provide a safer choice. A deeper investigation into the origins of the observed inconsistencies in location would make a valuable future study and is addressed for the Birmingham Method in Chapter 6.

In the FOV benchmark test, in terms of overall accuracy, all codes are observed to achieve sub-millimetre precision for both PEPT detector system geometries with respect to their mean spatial error results, shown in Figure 3.17 along with the average temporal resolution. These metrics are combined into the spatiotemporal resolution of the trajectory in Figure 3.18.

In the ADAC geometry, the Birmingham Method and PEPT-ML algorithms are observed to provide the highest accuracy, and in the ECAT geometry, the FPI and PEPT-EM algorithms perform similarly well. Whereas the Birmingham Method, FPI, and PEPT-EM algorithms all follow relatively consistent trends, PEPT-ML's accuracy varies more erratically. Nonetheless, the PEPT-ML algorithm displays significantly higher temporal resolution than the other codes whilst maintaining a comparable spatial resolution, as is illustrated by the spatiotemporal resolution, shown in Figure 3.18. Comparable performance from many

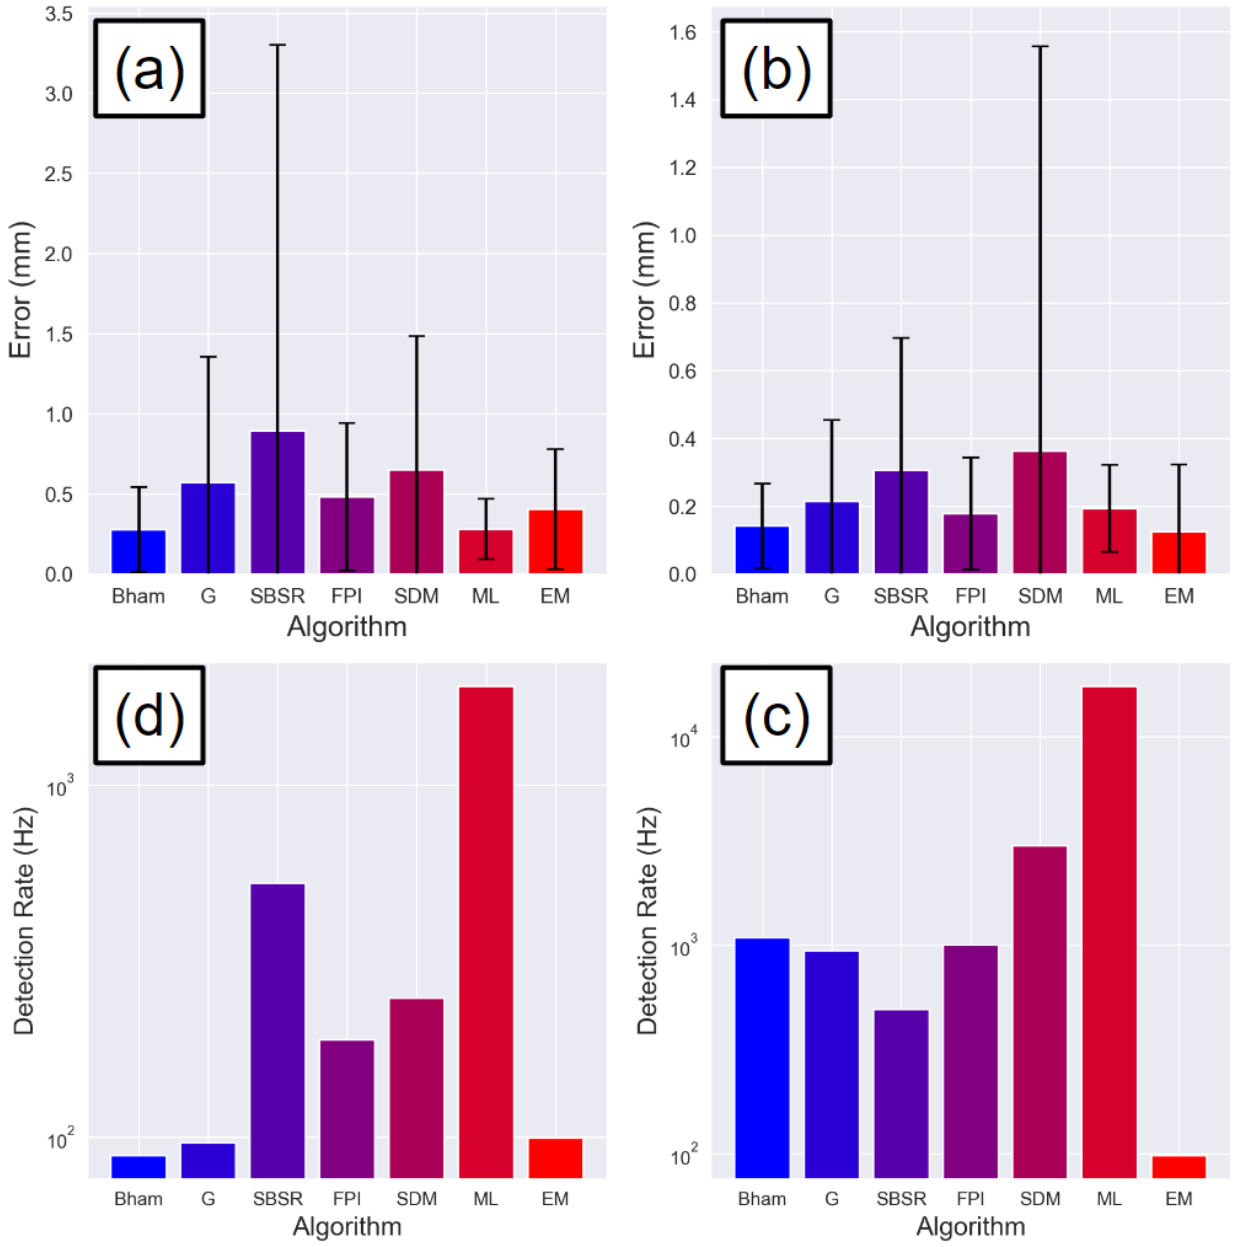

Figure 3.17: The mean error and standard deviation of the spatial errors for the PEPT algorithms for the FOV benchmark in the (a) ADAC Forte and (b) Siemens ECAT geometries, as well as the temporal resolution for the PEPT algorithms in the (a) ADAC Forte and (b) Siemens ECAT geometries [149].

554 of the methods is seen, with especially good results from PEPT-ML. This may again be due
555 to the use of second-pass clustering to significantly improve spatial and temporal resolution

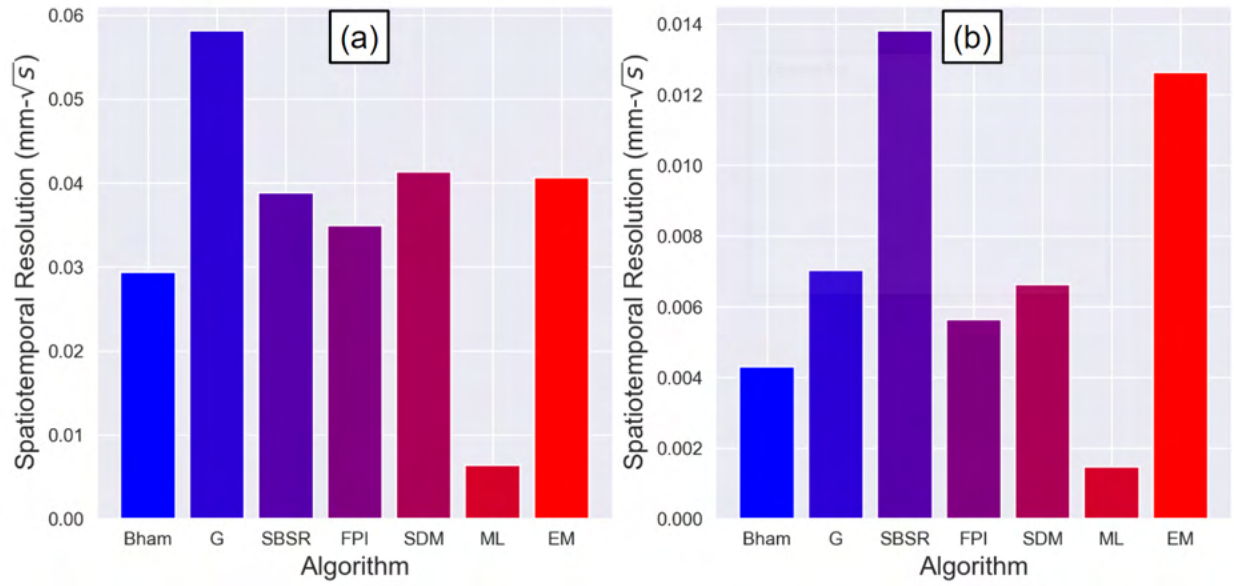

Figure 3.18: The spatiotemporal resolution of the PEPT algorithms in the (a) ADAC Forte and (b) Siemens ECAT [149].

for low-acceleration tracers.

In addition to these metrics, how the spatial accuracy changes as a function of the source positions within the FOV is also tested. The results of this test are shown in Figure 3.19.

In general, spatial error values of the majority of algorithms follow an approximate U-shape, with a minimum at the centre of the detector and a maximum near its edge. This is likely due to the higher proportion of random events near the edges of the detectors, where it is more likely that one-half of an annihilation-pair of gamma rays will hit the detector and the other miss. The corresponding increase in corrupted LoRs will tend to 'pull' the measured location toward the centre of the field of view. Further, for the algorithms which use time slice instead of a fixed number of LoRs, such as FPI and G-Means, the lower rates of LoRs near the edge of the PEPT detector system will cause an increase in the error. However, the PEPT-ML algorithm does not exhibit this behaviour, perhaps due to the second-pass

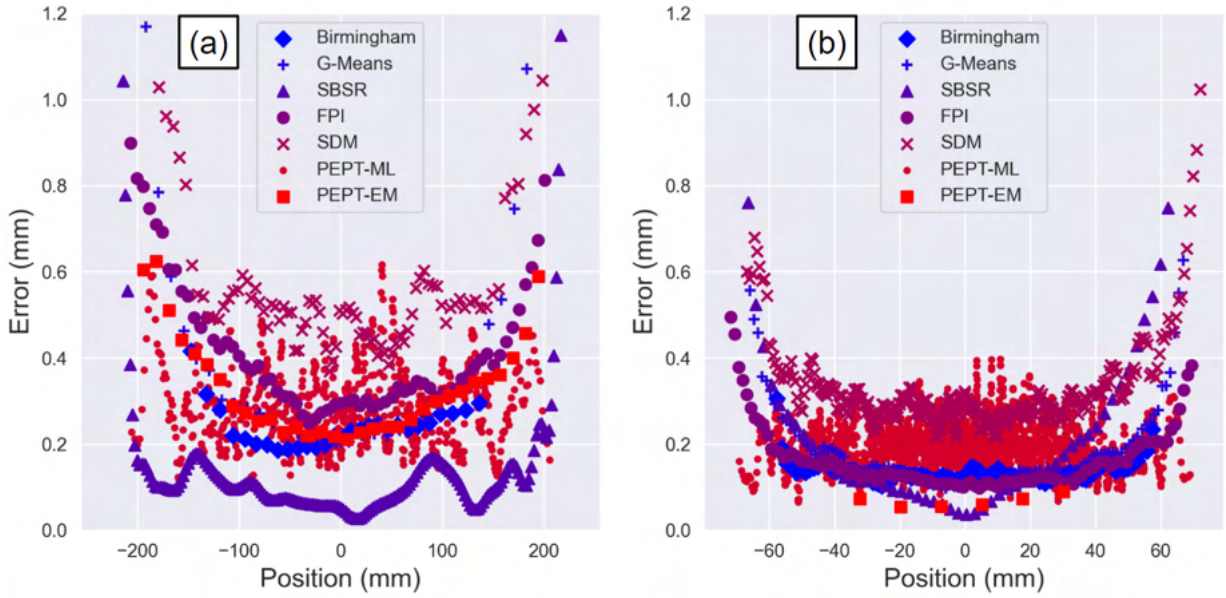

Figure 3.19: The spatial accuracy of the PEPT algorithms varies as a function of the source position in the (a) ADAC Forte and (b) Siemens ECAT [149].

clustering which combines multiple locations into a smoothed trajectory.

Additionally, there exists a slight asymmetry to the 'U'-shape. This is because in the Forte geometry, the detectors are fixed to a large aluminium gantry ring about which they are able to rotate. Similarly, in the ECAT geometry, tungsten shielding on one end of the ring acts to partially shield the detectors from rays emanating from outside of the FOV. The material in both cases increases the proportion of detected LoRs corrupted by scattering, increasing the error in the PEPT reconstructed position.

Overall, this test shows not only the importance of the detection rate of LoRs in the accuracy of the algorithms but also the effect the PEPT detector system geometry can have on the accuracy of the returned tracer locations.

The next four tests cover the results of the PEPT algorithms which are capable of tracking multiple tracers simultaneously. These algorithms are G-Means, FPI, SDM, PEPT-

ML, and PEPT-EM. The first multiple tracer test is the minimum separation distance in which two tracers can be resolved. The most interesting result from this test is that the PEPT-EM, PEPT-ML, FPI, and SDM algorithms are all capable of successfully resolving the two separate tracers up to and including the point of contact. This is a potentially highly consequential finding, as it had previously been assumed that the PEPT technique was incapable of imaging tracer collisions. This observation, therefore, potentially opens up a number of valuable new applications for PEPT.

All five codes tested are observed to achieve sub-millimetre accuracy across the majority of the tested tracer separation values, with the highest overall accuracy being achieved by PEPT-EM and the lowest by the SDM. The full results from this test are shown in Figure 3.20.

The accuracy levels achieved by all codes are observed to remain relatively independent of the tracer separation for comparatively large distances (> 5 mm), and sharply increase after this point in all cases except for PEPT-EM, which maintains a remarkable degree of precision, approximately < 100 μm, even to the point of contact. This impressive consistency is perhaps due to the fact that, unlike other PEPT algorithms, the Gaussian mixture model underlying PEPT-EM is based on a more comprehensive mathematical model of the physics underlying PEPT. That is, PEPT-EM seeks to maximise the probability of there being exactly 2 tracers clusters, rather than trying to identify N clusters and their locations. This makes the task simpler as only the positions and sizes of the clusters need to be calculated.

It is also of note that in many of these results, there appears to be a gap of roughly 0.1 mm between the accuracy achieved by PEPT-EM and PEPT-ML and that of the G-Means and FPI methods. This may be due to a small discretisation error caused by the use of voxels in the latter methods. For this exercise, both G-Means and FPI used 1 mm grids for

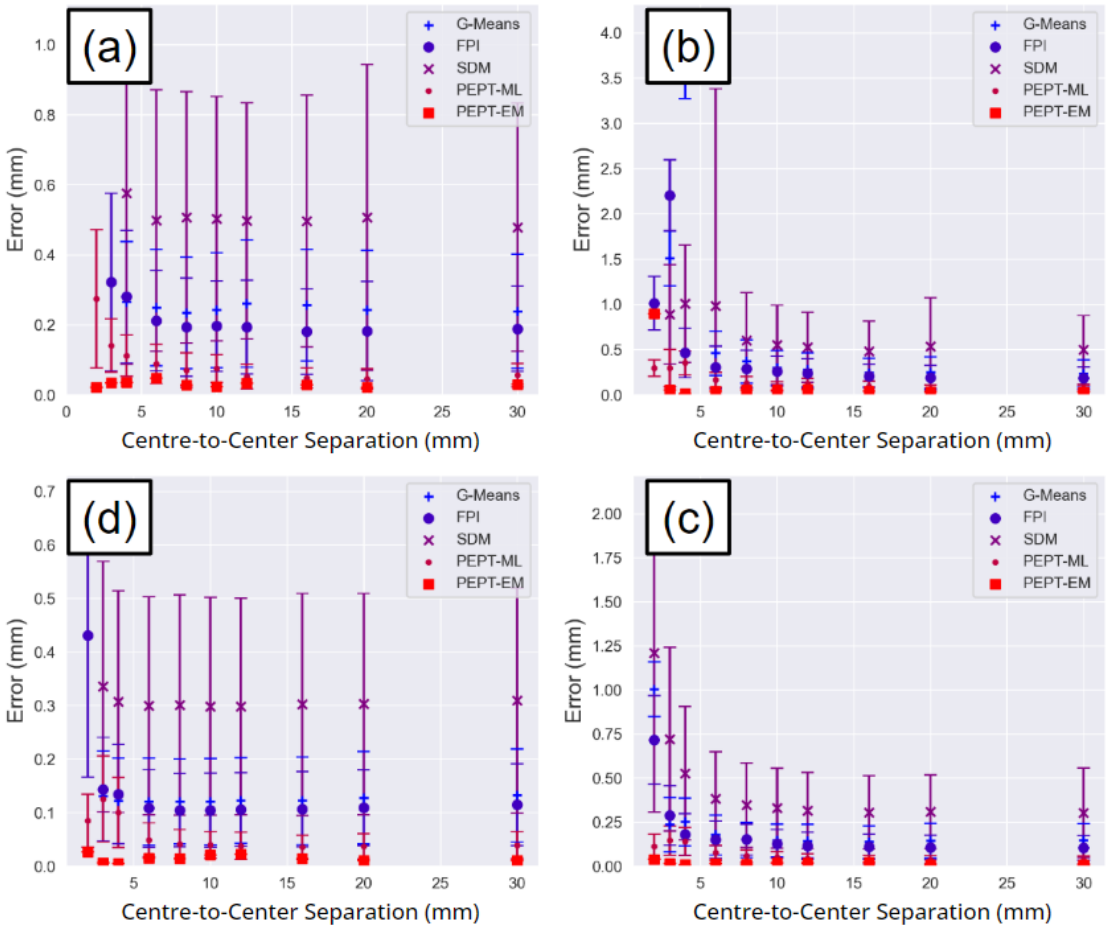

Figure 3.20: The measured error in PEPT location for the case of two static tracers separated by various centre-to-centre distances ranging from 2 mm (tracers in contact) to 30 mm. (a) Forte detector geometry and separation in the transverse direction. (b) Forte detector geometry and separation in the axial direction. (c) ECAT detector geometry and separation in the transverse direction. (d) ECAT detector geometry and separation in the axial direction [149].

line density tallying.

In the false positives benchmark, PEPT algorithms are thought to be susceptible to returning false locations caused by overlap between LoRs generated from different tracers. In this test, multiple tracers in highly symmetrical positions were placed in the FOV and the

objective for the algorithms is to accurately locate the positions of the tracers while ignoring areas where a false tracer position is generated. However, all codes, except for SDM, located all tracers in all cases. The generally-observed increase in location error for the tesseract, however, suggests an influence from the 'cross-talk' between tracers, though this seemingly manifests itself as a shift in detected positions as opposed to an outright false location. The SDM algorithm, on the other hand, produced a large fraction of locations which were between tracers in the areas where LoRs overlap.

For all algorithms tested, we observe a general decrease in accuracy with increasing tracer number, as shown in Figure 3.21. At higher N_t, all algorithms experience a more significant decrease in accuracy and, in the case of the ECAT geometry at least, a relative decrease in the differentiation between different algorithms.

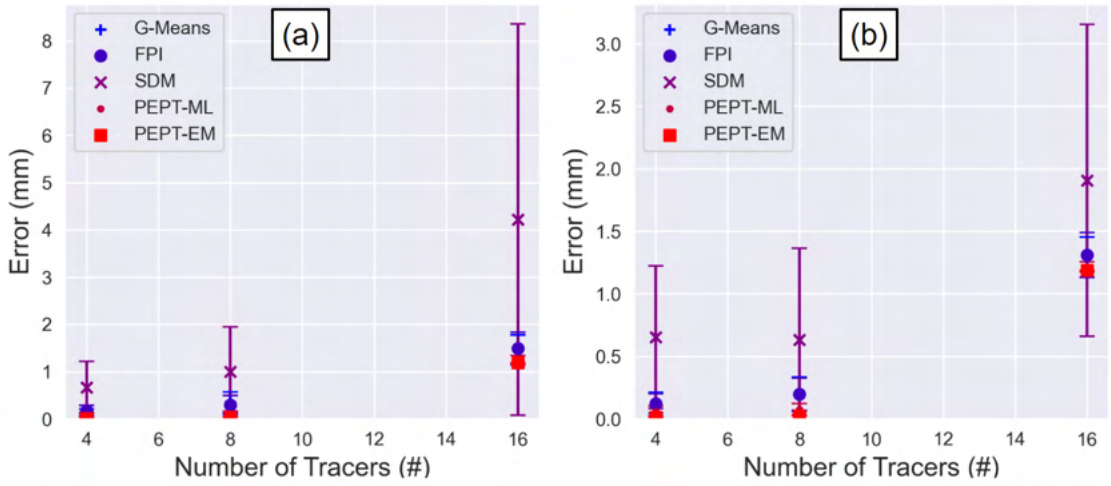

Figure 3.21: The measured error in PEPT location for the case of the false locations benchmark [149].

While this trend of increasing error could simply be due to an increase in tracer number, it is notable that a similar trend is not observed in the results of testing large numbers of tracers, where tracer positions are entirely random. This suggests that it is indeed the symmetry of the systems modelled here producing an additional confounding

effect. Specifically, it is observed that the measured tracer locations produced by SDM are biased toward the centre of the arrangement of tracers, as evidenced in Figure 3.22. There also exists a different form of error observed between the two PEPT detector system geometries for the algorithms where a greater spread of errors corresponds to the direction of lower spatial resolution (in the x-axis (axial) for the ADAC Forte and the z-axis (axial) for the ECAT).

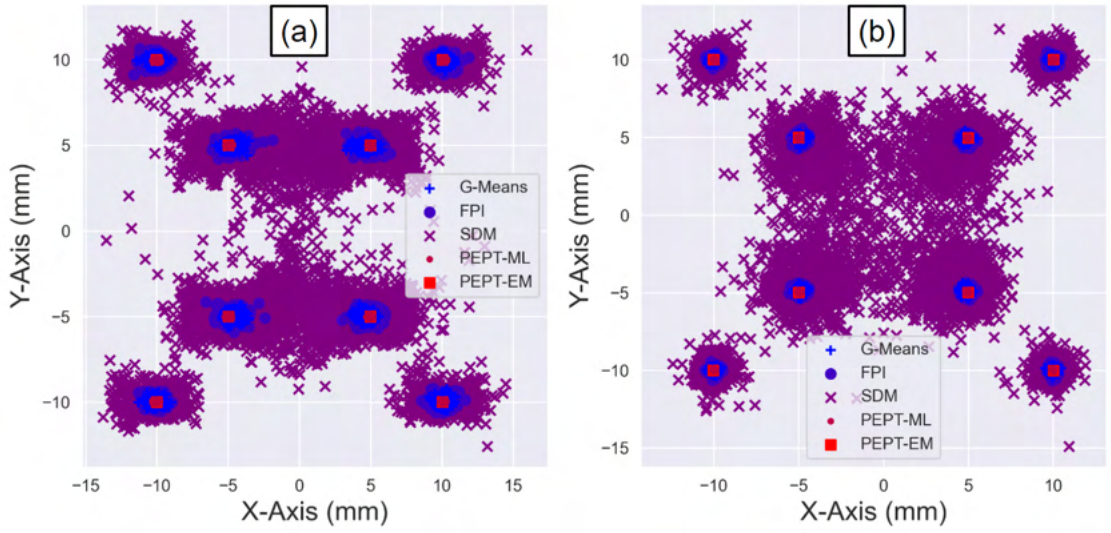

Figure 3.22: The individual positions of the tracers given by each algorithm for the arrangement of 16 tracers in the (a) ADAC geometry and (b) ECAT geometry. Some positions given by the SDM algorithm in particular are shown to be biased toward the centre of the arrangement of tracers [149].

Ultimately, this test shows that lower numbers of tracers and lower degrees of symmetry produced more accurate results. All algorithms are observed to achieve sub-millimetre accuracy, with PEPT-ML and PEPT-EM both providing resolution on the scale of 100 microns in the 8 tracer case, and the latter achieving accuracy of below 10 microns in the 4 tracer case. It is also interesting to note that there seemingly exist 3 distinct groups in the observed results, corresponding to 3 broad groups of algorithms: mesh-free clustering methods, meshed clustering methods, and iterative methods. The small difference in measured

error between the meshed methods (G-Means and FPI) and mesh-free clustering methods (PEPT-ML and PEPT-EM) potentially arises due to the effect of pixelation. The relatively increased error in SDM may be caused by the iterative nature of the algorithm. As data are being discarded for each additional tracer detected, it stands to reason that the precision of locations will decrease with a higher number of tracers due to the discarding of true LoRs.

The next benchmark seeks to test the trajectory-linking aspect of PEPT algorithms. The results show that all of the tested algorithms were able to distinguish and link the separate paths followed by the three moving tracers. In other words, all PEPT algorithms yield trajectories with consistent IDs, avoiding the common pitfalls of tracers switching IDs at points of intersection or being confused with other tracers at some point in their trajectory. Additionally, all algorithms produced sub-millimetre accuracy in this case of three moving tracers.

For both detector geometries, PEPT-ML and PEPT-EM produced both the highest and most consistent accuracy. The full results of this benchmark in terms of overall spatial accuracy and spatiotemporal resolution are shown in Figure 3.23.

Despite relatively large errors and the largest variability of all codes tested in the ADAC geometry, in the ECAT geometry, the G-Means algorithm shows a capability similar to that of PEPT-ML and PEPT-EM and even outperforms its successor, FPI. As discussed previously, however, the simple measure of overall error does not account for the frequency of detection. If instead the spatiotemporal resolution is considered, an improved relative performance from FPI is observed. The FPI algorithm also outperforms the PEPT-EM in the ECAT geometry. PEPT-EM offers very high spatial accuracy but produces a lower temporal resolution, which degrades the spatiotemporal resolution.

Ultimately, the results of this test show all algorithms were capable of reliably reconstructing and delineating the trajectories of multiple moving tracers, even when the

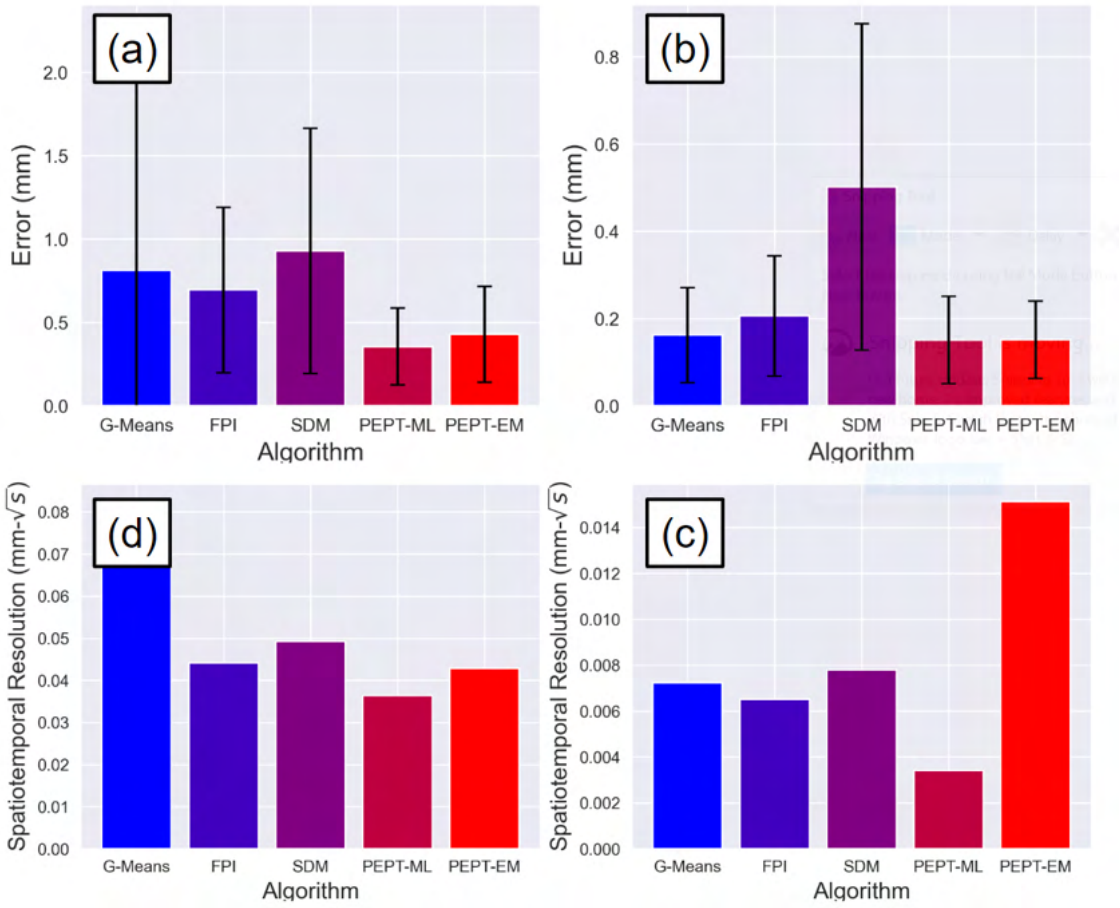

Figure 3.23: The results of trajectory linking test for the PEPT algorithms capable of tracking multiple tracers in the ADAC Forte and Siemens ECAT geometries. (a) spatial accuracy in the Forte, (b) spatial accuracy in the ECAT, (c) spatiotemporal resolution in the Forte, and (d) spatiotemporal resolution in the ECAT [149].

trajectories are observed to cross one another, though the PEPT-ML algorithm offers the highest and most consistent accuracy.

The final benchmark test examines the ability to locate large numbers of tracers within the FOV. This pushes the limits of how many tracers are able to be tracked simultaneously and also adds noise to the sample since many true LoRs from a particular source will likely pass near other tracers. Of all the algorithms tested, the results show that the majority are

able to locate all or nearly all of the tracers successfully. The full results are shown in Figure 3.24 in terms of the percentage of tracers successfully found and the mean accuracy of their returned positions.

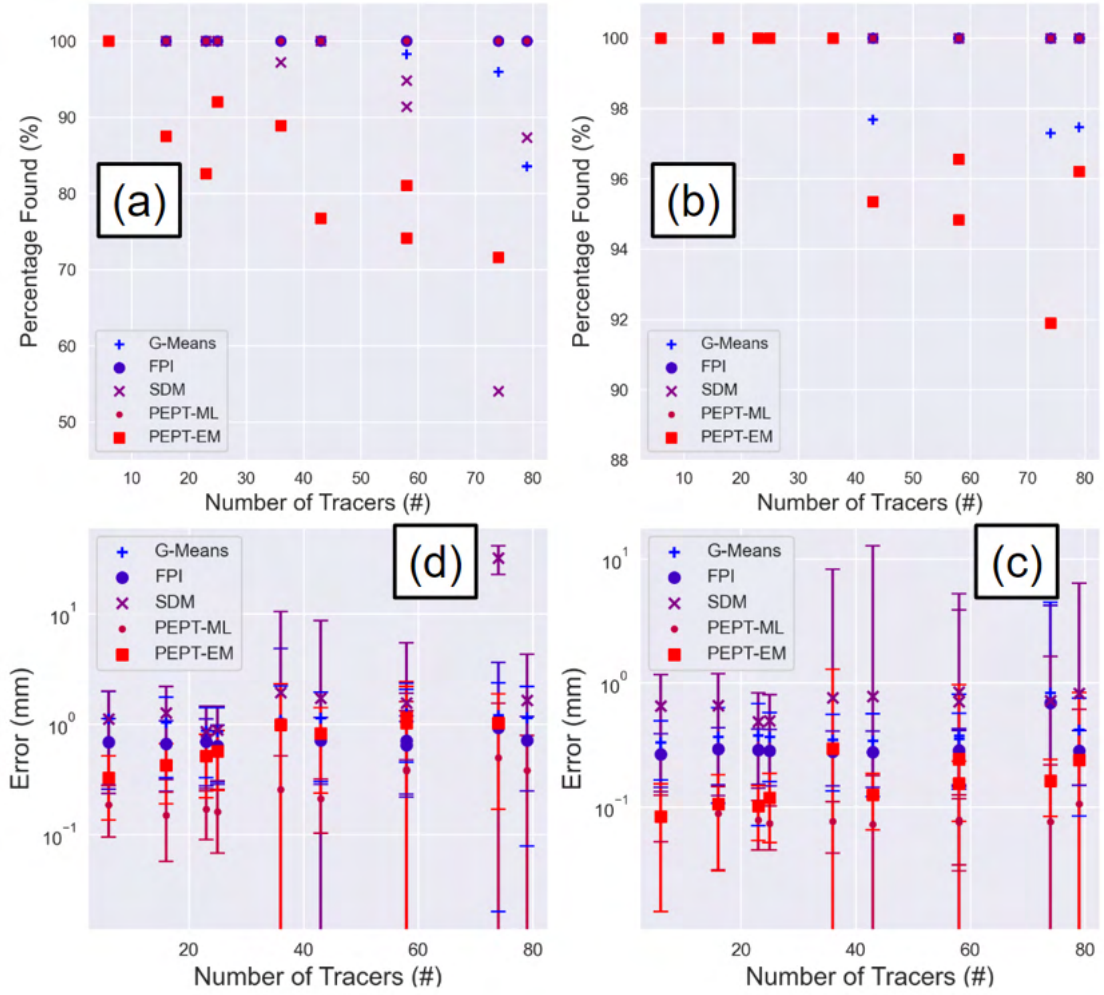

Figure 3.24: The results of the maximum number of tracers for the PEPT algorithms in the ADAC Forte and Siemens ECAT geometries. (a) percentage of tracer found in the Forte, (b) percentage of tracer found in the ECAT, (c) spatial accuracy in the Forte, (d) spatial accuracy in the ECAT.

For the ECAT geometry, the SDM, FPI and PEPT-ML algorithms successfully locate all tracers in all tested cases, with the latter two also receiving a perfect score for the ADAC geometry. The G-Means algorithm also performs relatively well, typically losing only

a relatively small fraction of tracers in cases below 40 tracers. The PEPT-EM algorithm performs well for comparatively low tracer numbers in the ECAT geometry, but loses tracers in almost all ADAC test cases, the fraction of lost tracers generally increasing with N_t.

In terms of the accuracy with which the successfully-located tracers are found, the algorithms perform relatively similarly, with PEPT-ML and PEPT-EM generally producing the highest overall accuracy, though PEPT-EM in many cases is detecting a smaller fraction of the total tracer number present than the other algorithms, thus potentially skewing this metric in its favour by not including poorly located tracers.

Perhaps the most striking feature of this plot, however, is the consistency in average spatial accuracy across the full range of tracer numbers explored. The FPI method, for example, can seemingly detect 79 tracers just as precisely as it can detect 6! Indeed, all algorithms tested are found to consistently produce millimetre- or submillimetre accuracy across practically all data sets. This observation bodes well for future research using large numbers of PEPT tracers, a possibility that to date has yet to be properly explored and exploited.

3.5 Discussion

The results from the PEPT algorithm benchmark tests offer the first comparative assessment of PEPT algorithms. Perhaps the most notable finding is that of the diverse suite of PEPT algorithms, there does not exist a single 'best' algorithm. Rather, all exhibit particular strengths and weaknesses. As such, to get the most out of the PEPT technique, it may be advisable to use different algorithms for different goals. For example, when tracking slow-moving and relatively active single tracers, the highly computationally efficient Birmingham Method is likely the ideal algorithm. Conversely, for applications requiring the tracking of

large numbers of tracers, or where the imaging of tracer-tracer collisions is desirable, instead, one of the more modern, but more computationally intensive, algorithms such as PEPT-ML, PEPT-EM or FPI should be employed. If tracer activity is limited, or the detected count-rate is otherwise limited or inconsistent, the SBSR method may be the optimal choice.

This international collaboration not only demonstrates the strengths and weaknesses of different approaches but also moves the community of PEPT researchers forward in terms of collaboration and future development of algorithms. To this end, the FPI method has been incorporated into the 'pept' Python package, along with the Birmingham Method and PEPT-ML algorithm. This means that researchers who want to use these algorithms only need to install one package and the tools developed within can be used interchangeably between them. For example, in the Python package implementation of the FPI method, both a fixed number of LoRs can be used as well as the time slice approach it was originally developed to use. Similarly, overlapping windows of LoRs, second-pass clustering using HDBSCAN, and trajectory separation methods are now shared between the algorithms. This is a major development in standardising the PEPT technique, moving away from individuals working on similar problems in isolation and towards communities developing best practices and shared projects.

Additionally, without the use of the simulation methods presented in this thesis for modelling PEPT tracers and detectors, this work would not have been possible. It is only through the development of GATE models for the ADAC Forte and adaptation of the pre-existing model of the Siemens ECAT that the ability to model PEPT experiments in detectors which are the same as those used at real PEPT facilities could be achieved. Moreover, the simulated experiments provide a direct comparison of the PEPT reconstructed positions of tracers to their prescribed location in the GATE simulation, creating, for the first time, a quantifiable accuracy of PEPT algorithms that is not biased or contains uncertainties on the same magnitude as the measurement.

To make the most use of these benchmarks and the PEPT algorithm comparison functions, they have been made public and freely downloadable from a data repository. This ensures that future researchers who develop PEPT algorithms can create a comparison of their algorithms to those tested in this work.

Chapter Four

Optimising Tracer Activity and Detector System Geometry

4.1 Introduction

As has been shown in previous Chapters, GATE models of PEPT detector systems, tracers, and experiments are capable of producing realistic synthetic data. The data can be treated the same as experimentally acquired LoRs and used in conjunction with PEPT algorithms. Since simulated experiments offer precise control of tracer properties, source locations, and experimental geometries, with the simulations themselves able to be rapidly changed through parameterised scripts, this allows for testing variations of PEPT experiments [44]. Moreover, through the techniques developed in this thesis, the effect of variations in PEPT experiments can be quantified in terms of the spatial, temporal, and spatiotemporal resolution of reconstructed tracer trajectories [149].

This Chapter seeks to develop a technique to optimise two practical aspects of PEPT experiments: tracer activity and detector geometry. These are important for PEPT because they influence the rate at which LoRs are recorded and the quality of the LoRs in terms

of the fraction that can be used to locate the tracer [44]. Tracer activity determines the production rate of coincident gamma rays from positron annihilation, whereas the detector geometry affects the fraction of these rays which interact with the detector [148]. While the ideal rate at which coincident gamma rays emitted from a point-like tracer will interact with a PEPT detector, R, can be determined by finding the solid angle covered by the area of the scintillation crystals, or geometric efficiency, ε_g, the probability of interaction within the crystal volume, or intrinsic efficiency, ε_i, and the tracer activity, A, using Equation 4.1, it is non-trivial to determine how the complex geometry of an experimental system affects this [44, 47]. Moreover, the highly non-linear response of the detector caused by the electronic pulse-processing chain further complicates this [44].

$$R = \varepsilon_g \varepsilon_i^2 A \tag{4.1}$$

The current protocols for running PEPT experiments do not often allow for optimisation of these factors [148]. Since the University of Birmingham's Positron Imaging Centre (PIC) is a user-facility, usually a single week or only a few days are allotted to groups running experiments [98]. Setting up the equipment and beginning data acquisition as soon as possible is more important than taking several days to establish the optimal experimental protocol. In many instances, this can be acceptable, owing to the fact that there is a fairly wide band of conditions which will yield tracer trajectories of sufficient resolution that will provide the information researchers are interested in [100, 47]. However, some experiments have been unsuccessful because PEPT tracers were not active enough (or too active) for the detector geometry used. Additionally, in the experiments that were successful, it is possible that more about the system could be learned if better resolution trajectories were extracted.

Therefore, starting an experiment using the optimal conditions is desired. By running GATE simulations prior to real experiments, not only can more combinations of tracer

activities and detector head separations be tested than is possible through real experimentation, but also the results of the simulations can be analysed and optimal settings established before experimental equipment is moved to the PEPT facility. For this reason, it would be beneficial to use GATE simulations prior to PEPT experiments in order to explore a range of possible ways to run the experiment. These simulations would save the considerable amount of time and resources needed for physical experimentation and still provide helpful information that will optimise data acquisition.

In the following Sections, one method of finding these optimal settings will be demonstrated using a GATE model of the ADAC Forte and a hypothetical experiment [44]. In the simulation, an experimental geometry representative of a continuous blender used in the pharmaceutical industry is created and an analytical trajectory of a tracer moving through the blender is simulated [153]. The blender is placed in the centre of the field-of-view (FOV) and the tracer makes one pass through the system. In this example, only the tracer activity and head separation are changed, with the LoRs recorded by the GATE model processed using the PEPT-ML algorithm to form a trajectory [91]. Using the prescribed position of the tracer, the reconstructed trajectory is analysed for spatial resolution, temporal resolution, and the combined spatiotemporal resolution [149]. Finally, the results of a parameter sweep of tracer activity and detector head separation are compared and discussed, with a method of selecting the optimal settings demonstrated.

4.2 Methods

4.2.1 GATE Simulation

The simulated PEPT experiment is modelled on a continuous blender used in the pharmaceutical industry to blend powders of a drug formulation [156]. This blender is proprietary and specifics of the model, such as the impeller design are omitted from this work. However, in practice, these could be added if the model were being used to inform real experiments. In this case, the geometry is generalised and consists of three main parts: an aluminium outer shell, a powder-filled interior, and an axially rotated impeller shaft. The aluminium outer shell is 80 mm in radius and 5 mm wall thickness, the inside of the mixer is filled with microcrystalline cellulose (MCC) power with a bulk density of 1 g/cm^3, and the middle of the blender is a hollow aluminium impeller shaft 25 mm in radius with a 5 mm wall thickness [128]. The blade, in this case, is not modelled due to the proprietary design but is only expected to slightly contribute to the overall amount of scattering in the system. Some scattering is expected, but due to the low volume of the blade compared to other components, like the shaft and walls, it is negligible. The mixer is tilted upward at 15 degrees to match how it is typically used in industry [83]. An image of the GATE-modelled blender is shown in Figure 4.1.

The tracer is modelled as a positron-emitting fluorine-18 source. The geometry of the tracer is a solid MCC sphere, 0.5 mm in radius, with a density of 1.58 g/cm^3. In reality, the motion of the tracer is a spiralling trajectory as it is blended with the other powders in the formulation in order to make a homogeneous mix [156]. This motion is mimicked in the GATE simulations by prescribing a helical trajectory through GATE's Generic Move function. A series of times and positions for the tracer are read from a file. To create the trajectory file, first, the motion for the tracer is modelled as a helix where the tracer moves in

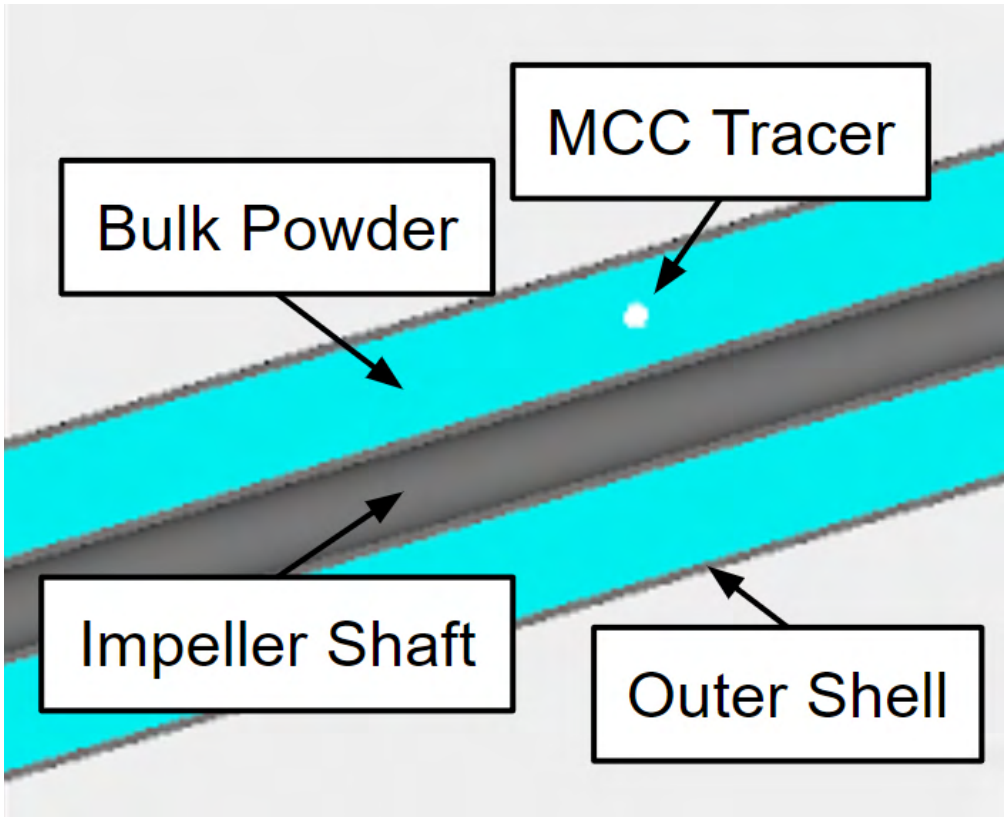

Figure 4.1: A cross-section of the continuous blender GATE model.

a circle with respect to the XY-plane (plane perpendicular to the ground and detector faces) and with a constant velocity in the Z-direction (direction parallel to the ground and detector faces). The radius of the circle is 50 mm to coincide with the centre of the blender. Once the helical trajectory is created, the trajectory is rotated 15 degrees in the X-axis by using a three-dimensional rotation matrix. The rotation matrices are shown in Equations 4.2-4.4. In this way, the tracer moves through the centre of the GATE-modeled blender, throughout its length, and moves at a constant velocity of 200 mm/s. In the future, a trajectory captured from a DEM simulation of the same system can be used for realistic tracer motion [46]. The original helical trajectory and the trajectory which has been rotated 15 degrees are shown in Figure 4.2.

$$R_x(\theta) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\theta & -\sin\theta \\ 0 & \sin\theta & \cos\theta \end{bmatrix} \tag{4.2}$$

$$R_y(\theta) = \begin{bmatrix} \cos\theta & 0 & \sin\theta \\ 0 & 1 & 0 \\ -\sin\theta & 0 & \cos\theta \end{bmatrix} \tag{4.3}$$

$$R_z(\theta) = \begin{bmatrix} \cos\theta & -\sin\theta & 0 \\ \sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{4.4}$$

The geometry of the continuous blender and the MCC tracer are imaged using the GATE model of the ADAC Forte. This PEPT detector system has two detector heads which can be separated up to 800 mm to accommodate a variety of experimental equipment [103]. Often it is unknown what tracer activity and detector head separation is optimal for a given experiment such as this. By using this GATE model it is possible to test a range of different experimental conditions and directly observe their effect on the quality of the trajectories. In this work, a matrix of different experimental conditions is simulated which explores the parameter space more quickly and thoroughly than is possible from real experiments. The blender and PEPT tracer are shown inside the FOV of the ADAC Forte in Figure 4.3

The detector separation is varied from 200 to 650 mm in 50 mm increments and the tracer activity is varied from 1 MBq to 30 MBq in increments of 1 MBq. This is equivalent to 300 different combinations of the head separation and tracer activity, representing, at a minimum, 10 different real experiments if initially a 30 MBq source was used and allowed to decay to 1 MBq. Each experiment would take approximately 9 hours to complete, assuming

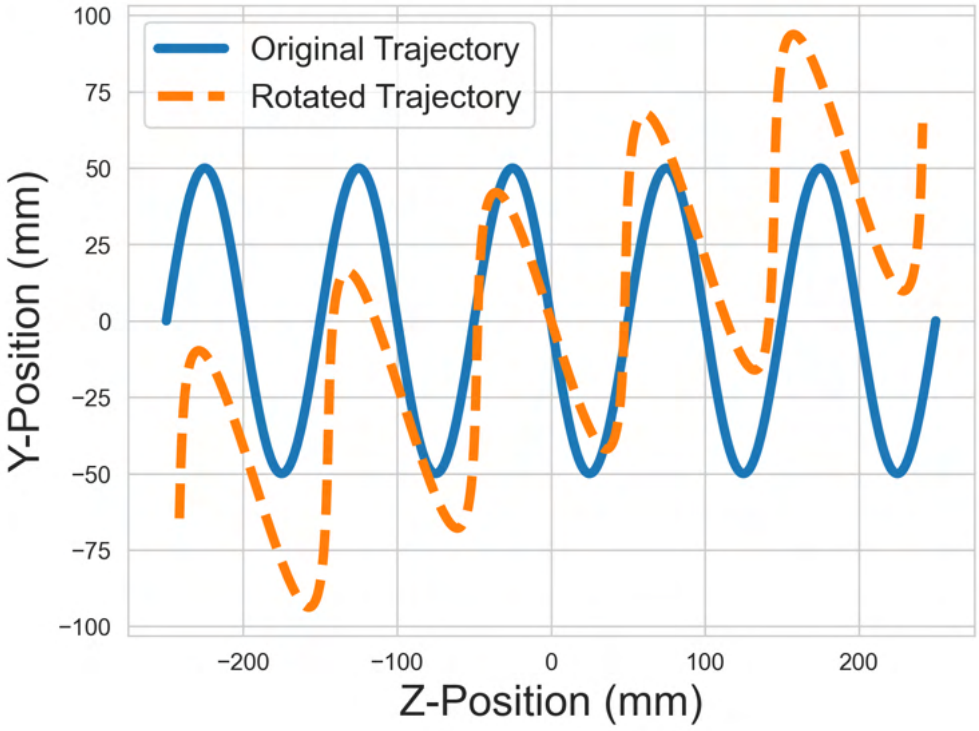

Figure 4.2: The original helical trajectory and the trajectory which has been rotated 15 degrees to move through the continuous blender.

fluorine-18 was the positron-emitting isotope used.

For the tracer to move along 500 mm of the blender this takes approximately 8.25 seconds. This is set as the length of the data acquisition. In the middle of the simulated time, the tracer passes through the centre of the FOV. A time slice of 0.0001 seconds is prescribed so that the tracer only moves 0.02 mm between slices, a distance considerably smaller than the tracer radius. This ensures that the trajectory of the tracer appears continuous. The data output is set to coincidences so that LoRs can be generated and the same Forte geometry and digitizer settings are the same as those prescribed in Chapter 2.

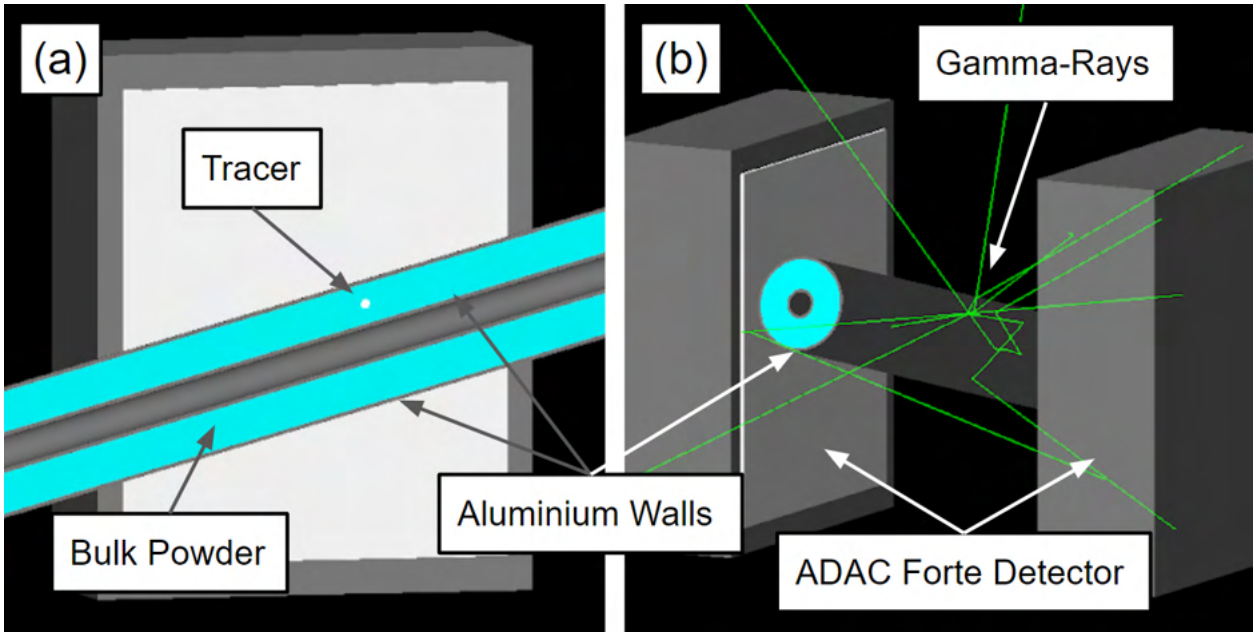

Figure 4.3: (a) A cutaway of the continuous blender in the ADAC Forte. (b) The full blender in the Forte FOV with the MCC tracer emitting positrons which form back-to-back gamma rays.

4.2.2 Data Post-Processing

Once all simulations have been completed, the detected LoRs are passed through a PEPT algorithm to extract the tracer trajectories. In this work, the PEPT-ML algorithm is used to transform samples of LoRs into trajectories [91]. For a full understanding of the PEPT-ML algorithm, please refer to Section 3.2.6. In processing the results of these simulations, a common set of algorithm parameters were used to generate trajectories. This ensures all simulated results are treated equally but could result in the non-optimal parameters being used. A sample size of 200 LoRs is used to compute the cutpoints with only cutpoints falling within a range of 0.15 mm considered for clustering. Additionally, only the 15% most densely clustered cutpoints are used to calculate the position of the tracer. Second-pass clustering is similarly applied, but with 70% of the densest detected positions considered for clustering.

To compare the detected and prescribed trajectories, both the position error and the frequency of detection must be taken into account. Low values in position error denote good spatial resolution, whereas high values of detection frequency denote good time resolution. The position error in this work is given as the mean three-dimensional distance between the detected and prescribed positions, calculated using Equation 1.14. The frequency of detection is calculated as the number of detections in a trajectory, divided by the time elapsed from the first detection to the last, calculated using 1.17. Additionally, to get an overall measure of the quality of the trajectories, the spatiotemporal resolution is calculated using Equation 1.22, which combines the spatial and temporal resolutions [149]. From previous work, it is known that spatial resolution decreases proportionately with the inverse square root of the temporal resolution [100, 103, 97].

4.3 Results

The first thing to note from these simulations is that the count-rate versus activity curves produced by each head separation have a complex relationship. Due to the differing geometric efficiencies of the detector geometry and the count-losses caused by dead-time effects, the count-rates do not follow a linear relationship with geometric efficiency or tracer activity [44]. Instead, from Figure 4.4, it can be seen that the lower head separation and higher tracer activities do not always produce high LoR count-rates.

Since a common sample size of 200 LoRs is used to find a tracer location and each simulation runs for the same amount of time, this means that the count-rates are essentially a surrogate for the temporal resolution. To better visualise how the count-rates affect the temporal resolution, a two-dimensional grid can be used for each tested combination of head-separation and tracer activity and then coloured by the temporal resolution. The temporal

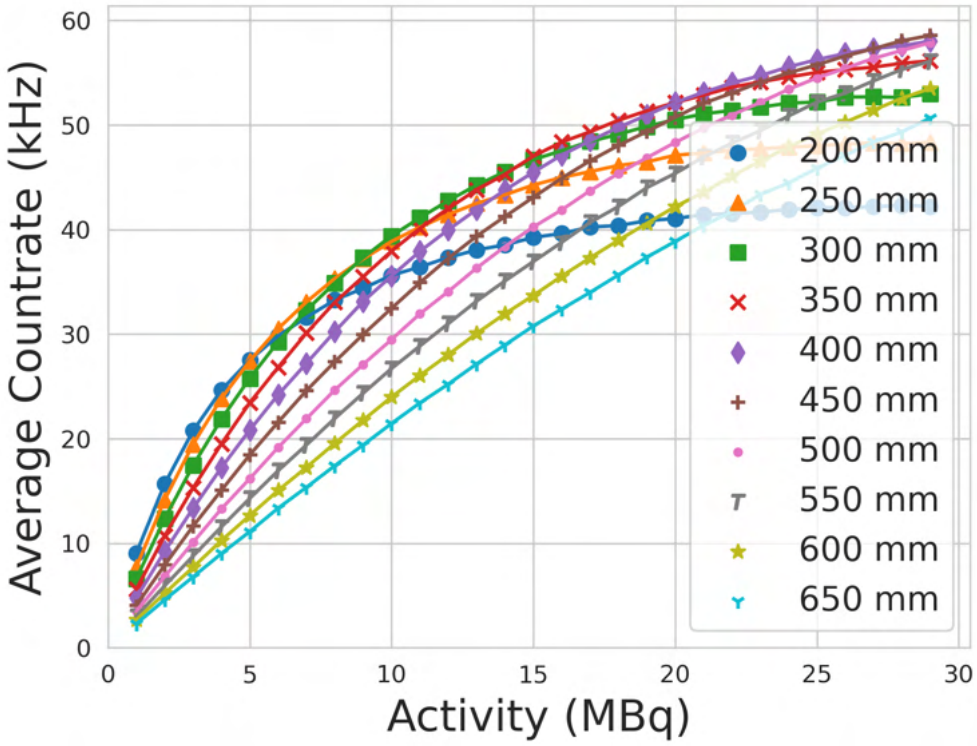

Figure 4.4: The total count-rate response of the ADAC Forte to variations in the head separation and tracer activity.

resolution is shown in Figure 4.5.

In terms of the temporal resolution, there is a clear range of conditions which produce the optimal number of tracer locations. These conditions are the combinations of the head separation and tracer activity which produce the highest LoR count-rates. Additionally, it appears that the best combinations are a balance between the head separation and tracer activity, with the optimal head separation decreasing as the activity on the tracer decays.

When the spatial resolution results are taken into account, the relationship becomes more complex. Figure 4.6 shows the results of comparing the PEPT reconstructed trajectories to the GATE-prescribed positions. Again a similar relationship between the head separation and tracer activity is demonstrated, but with somewhat more erratic results,

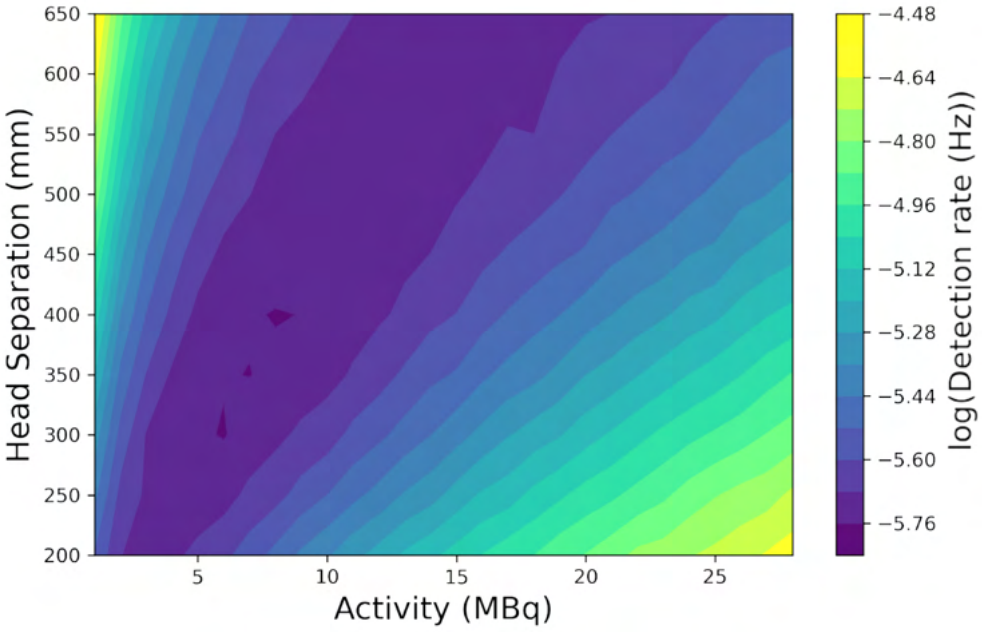

Figure 4.5: The temporal resolution of each combination of detector head separation and tracer activity.

perhaps owing to only using a single pass of the tracer through the blender, and with the low separation and high tracer activity combinations showing the lowest spatial resolutions, likely caused by the large fraction of corrupted LoRs being recorded.

To better understand the relationship between the head separation and tracer activity as it relates to the temporal and spatial resolutions it is necessary to plot the PEPT reconstructed trajectories against the GATE-prescribed trajectory. A plot of the highest separation and lowest activity, lowest separation and highest activity, and a medial separation and activity is shown in Figure 4.7. In this plot, a few things about the effect head separation and tracer activity have on the quality of trajectories become evident. First, in the high head separation and low activity case (1 MBq and 650 mm) the trajectory is sparsely populated and slightly pulled toward the centre of the FOV due to the low detection rate of LoRs. Secondly, the low separation and high tracer activity case (30 MBq and 200 mm) shows a relative improvement in the number of locations in the trajectory but is

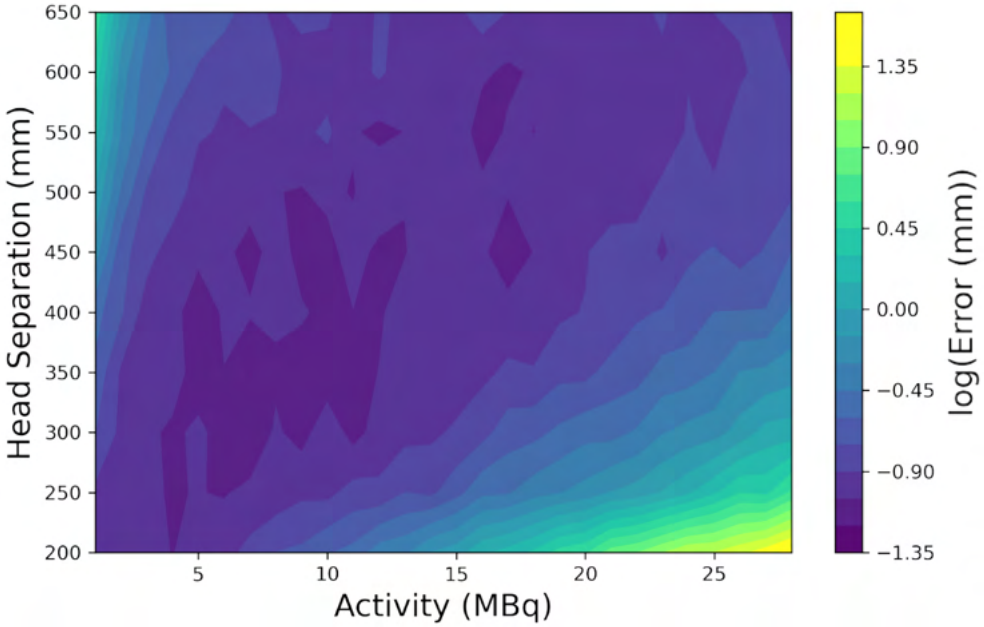

Figure 4.6: The spatial resolution of each combination of detector head separation and tracer activity.

substantially distorted on the ends due to the large fraction of corrupted LoRs generated from random coincidences, leading to degradation in the spatial resolution. Finally, it is only when a balance is struck between the head separation and tracer activity (7 MBq and 350 mm) that the trajectory is densely populated and accurate to the prescribed positions of the tracer.

Moreover, when the spatiotemporal resolution from each simulation is calculated, a clear optimal set of conditions can be identified. The spatiotemporal resolution for each simulation is shown in Figure 4.8. From this plot, the optimal head separation and tracer activity are approximately 300 mm - 350 mm and 6 MBq - 8 MBq. It is at these settings that both the LoR count-rate and the accuracy of the PEPT algorithm are maximised, leading to the best possible reconstruction of the tracer trajectory.

However, during an experiment, the activity of the tracer is always decaying. While

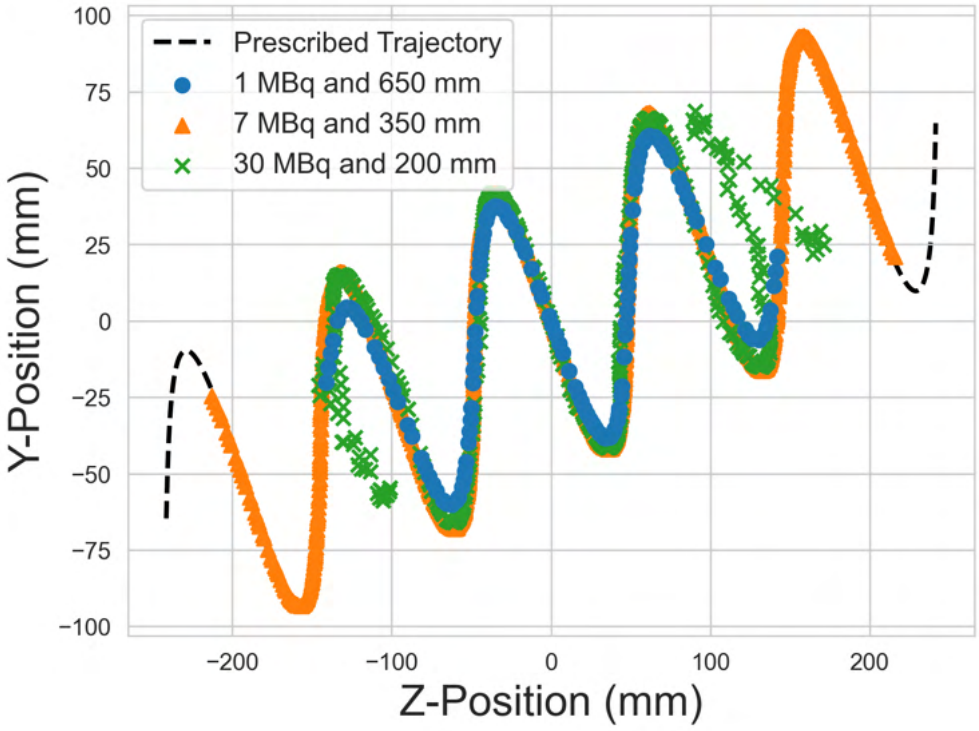

Figure 4.7: The PEPT reconstructed trajectories compared to the GATE prescribed trajectory for three combinations of a head separation and tracer activity.

fluorine-18 has a half-life of approximately 109 minutes, PEPT experiments can typically last anywhere between 30 minutes and 3 hours. Because of this, if the experiment was started with the previously identified optimal head separation and tracer activity, by the end of the experiment the tracer will have likely decayed below the optimal activity range. This means that a larger head separation and higher tracer activity are needed at the beginning of the experiment.

The tracer activity will decay regardless of what happens, but the head separation is controllable and can be adjusted during an experiment. Therefore, by finding an equation which describes the optimal head separation as a function of the tracer activity, once the optimal initial activity is found, the heads of the detectors can be slowly moved in as the activity decays. Using Figure 4.8, the optimal head separations can be found for each activity

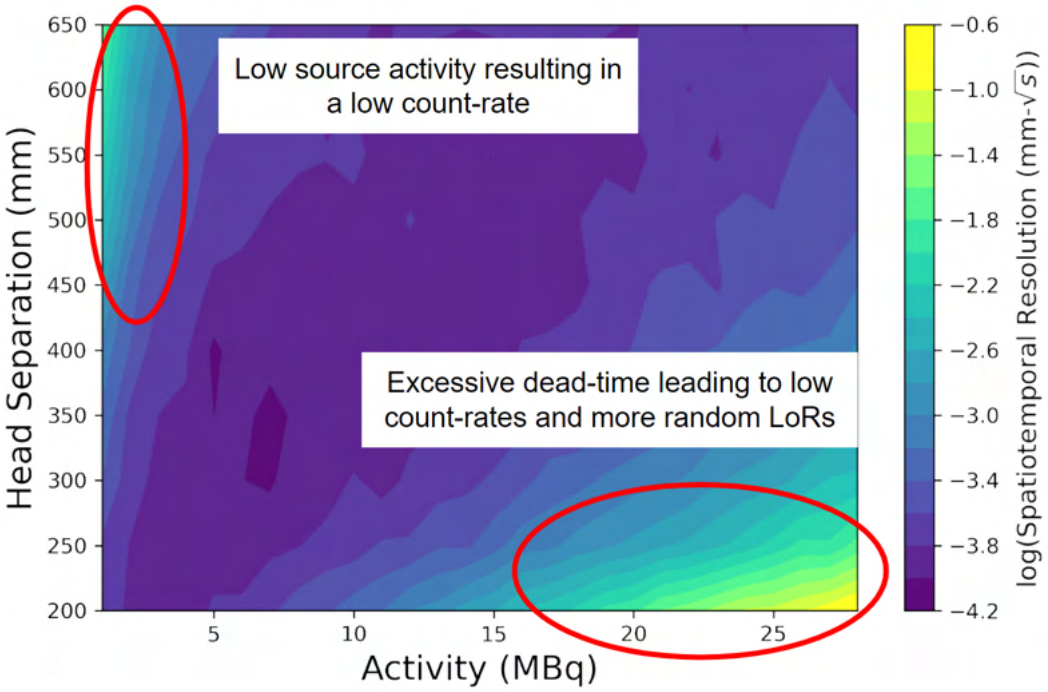

Figure 4.8: The spatiotemporal resolution of the PEPT reconstructed trajectories.

and an equation is fitted to determine this relationship. The fit appears to be more complex than a linear relationship and roughly takes the form of a power law. By assuming this form, generalised by Equation 4.5, the extracted data points can be used to find the best fitting parameters, where S is the head separation, A is the tracer activity, and a, b, c are parameters. Figure 4.9 shows the extracted optimal combinations of head separation and tracer activity overlaid on the spatiotemporal resolution with the best-fit line and optimal power law parameters.

$$S = a(A - b)^c \tag{4.5}$$

From this, optimal head separation can be defined as a function of activity. Next, the optimal starting activity can be found by setting the desired length of time to run the experiment and then integrating the spatiotemporal resolution associated with each optimal

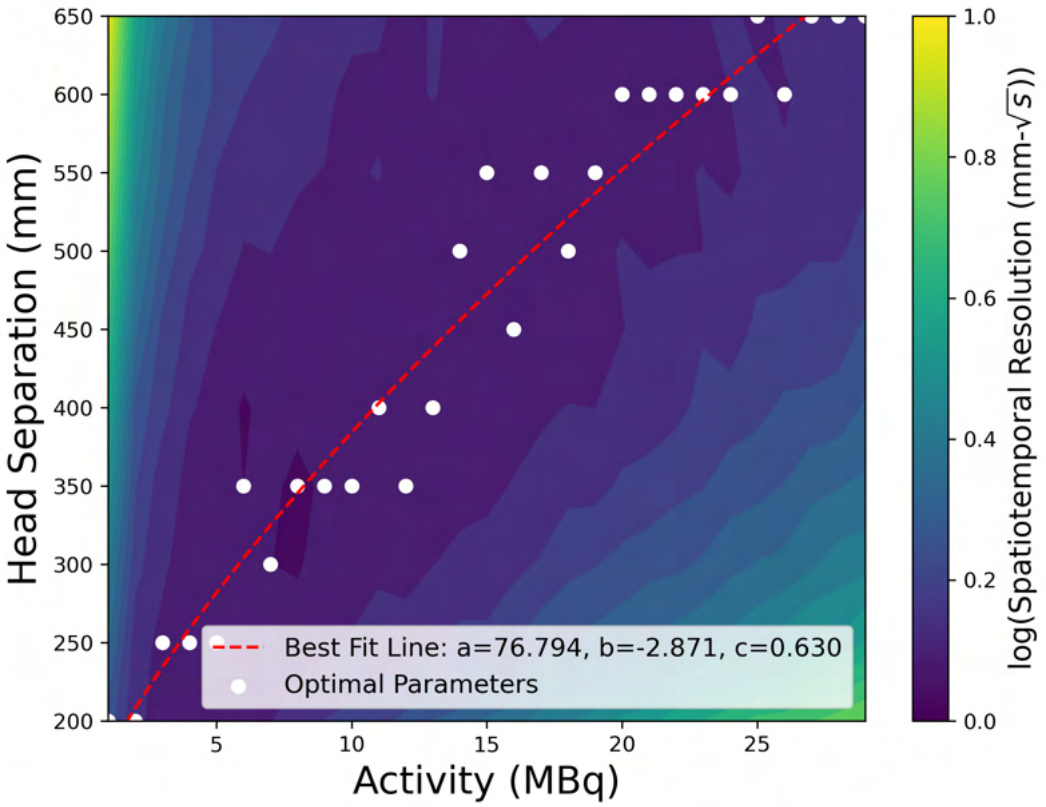

Figure 4.9: The spatiotemporal resolution of the PEPT reconstructed trajectories with the overlaid fitted equation describing the activities and head separation which maximise the spatiotemporal resolution.

activity over an initial activity and the ending activity at the end of the experiment. This yields an average spatiotemporal resolution and can be calculated using Equation 4.6. The optimal spatiotemporal resolution is extracted from Figure 4.8 and plotted as a function of activity in Figure 4.10.

$$\bar{\zeta} = \frac{1}{A_i - A_f} \int_{A_i}^{A_f} \zeta(A) dA \tag{4.6}$$

The average spatiotemporal resolution as a function of the initial activity for different lengths of time running a PEPT experiment is shown in Figure 4.11. With this, it can

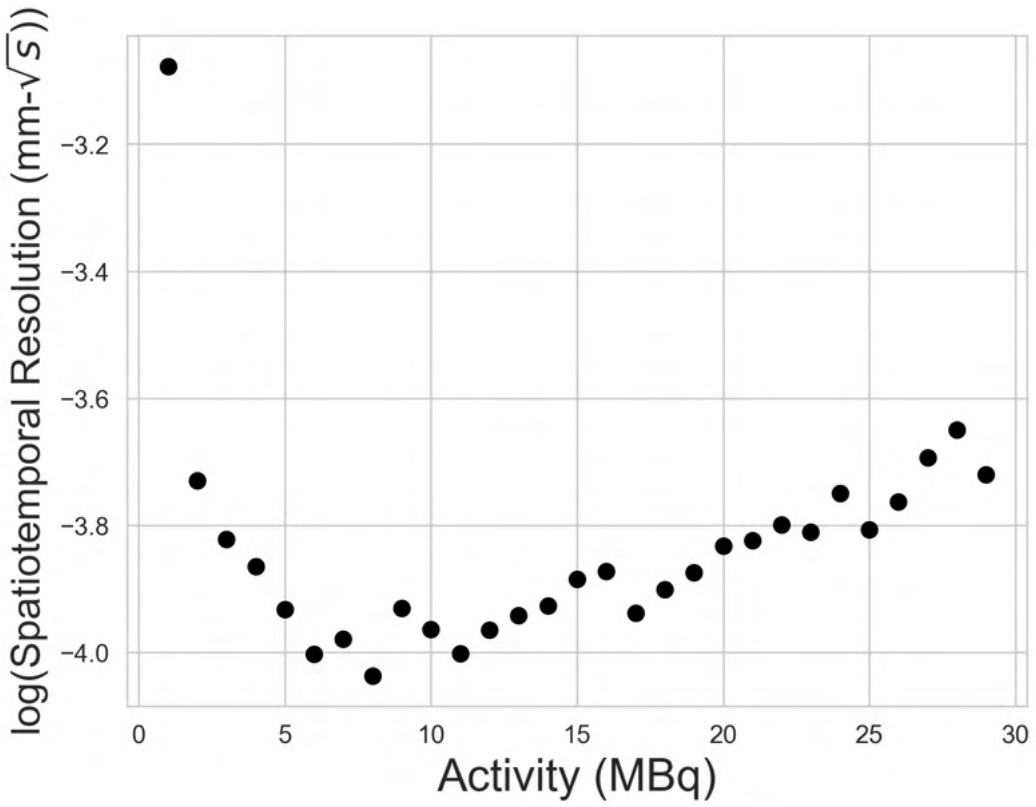

Figure 4.10: The optimal spatiotemporal resolution of the PEPT reconstructed trajectories as a function of activity.

now be determined what the initial tracer activity and detector head separation of a PEPT experiment should be. Further, using Equation 4.5 and the fitted parameters from Figure 4.9, as the source decays the head separation can be adjusted.

4.4 Discussion

This now creates a powerful new tool for PEPT users to be able to get the maximum spatiotemporal resolution trajectories from an experiment. This type of optimisation has not previously been possible due to the considerable amount of time needed for physical experimentation. However, by using GATE simulations to produce realistic estimates of

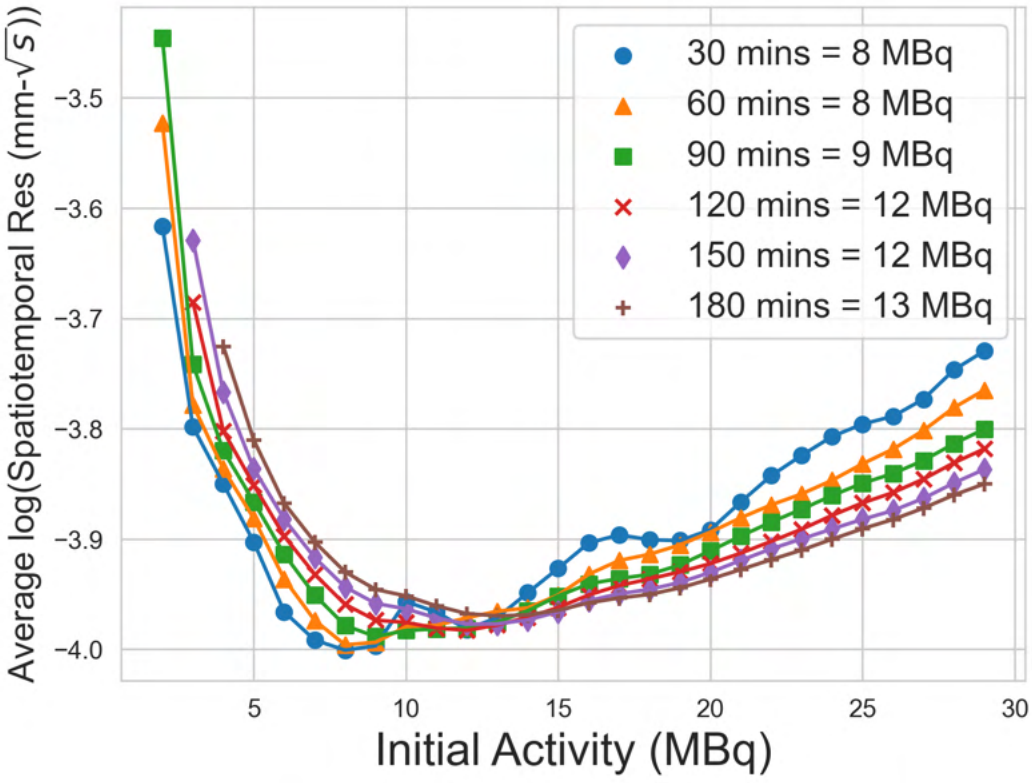

Figure 4.11: The average spatiotemporal resolution of the PEPT reconstructed trajectories as a function of initial tracer activity for different lengths of a PEPT experiment with the optimal initial activities listed in the legend.

data produced using an experiment, variations in the tracer activity and detector geometry can be rapidly tested. Additionally, through the methods demonstrated in this Chapter, a way to quantify the increase in spatiotemporal resolution has been demonstrated. This provides a justifiable reason why particular tracer activities and detector head separations were used during a PEPT experiment, making the technique more rigorous.

While this technique was demonstrated in one specific PEPT experiment, the methods demonstrated in this Section can similarly be applied to any PEPT experiment, given a model of the detector and experimental geometry. The ADAC Forte was used in this work because it is the most widely used PEPT detector system at the PIC and a validated GATE model

exists [103, 44, 149, 45]. However, validated GATE models of other PEPT detector systems, such as the Siemens ECAT and the Large Modular Array (LaMA) could similarly be used [57, 72]. While the geometry of the Siemens ECAT cannot be changed, the modular design of the LaMA provides more customisation in geometry than is possible with the ADAC Forte [71, 97]. Therefore, this work provides a method for establishing an optimal LaMA geometry through a quantifiable comparison of proposed geometries.

It should also be noted that several types of optimisation can be conducted for different aspects of the same experiment. For example, while it was shown in this Chapter that tracer activity and the detector head separation can be optimised, in the next two Chapters, it will be shown how the GATE simulations can also be used to optimise the length of an experiment, methods of correcting PEPT measurements, and also the optimisation of parameters in PEPT algorithms [46, 47]. These other factors can be established independently and combined, thus creating an optimised workflow for PEPT which includes aspects of the experimental design, use of algorithms, and data post-processing. As the use of GATE simulations for the optimisation of PEPT experiments is further developed and matures, it is possible that similar work can become standard practice and be implemented before all PEPT experiments.

Chapter Five

Optimising Experiment Length and Measurement Corrections

5.1 Introduction

This Chapter addresses establishing the length of time needed to collect adequate statistics of tracer behaviour and, additionally, correcting reconstructed tracer behaviour from known distortions [46]. When running PEPT experiments, it is often unknown if the tracers will have generated enough data for the behaviour of the system to be adequately characterised [144]. This is complicated by the fact that tracer motion is stochastic, meaning the specific path a tracer may take is unpredictable as it undergoes collisions with other discrete particles and walls in a granular system [147]. By chance or because of the tracer properties, a PEPT tracer may only explore certain regions of the experimental system, potentially leading only to that particular tracer behaviour being characterised and not being extendable to the bulk behaviour of all particles [150].

This would violate a core assumption of many PEPT experiments, ergodicity, meaning that if the experiment is run long enough with identical particle species, then the average

behaviour of one particle reflects that of all others [37]. On a short time scale, all particles will reflect local behaviour, but through averaging over a longer time, the behaviour of an individual tracer should approximate the bulk behaviour [143]. The time needed to reach ergodicity is difficult to know prior to the experiment [144]. Additionally, tracer size, density, and shape can lead to segregation effects, resulting in different behaviour for particles of different species [150].

Furthermore, when analysing the reconstructed trajectory of the tracer, especially in systems that are opaque and without complementary imaging techniques, it can be difficult to know if the reconstructed fields of the systems correspond to the real behaviour of the tracer or are influenced by a distortion in the measurement [49]. Examples of these fields are the velocity field or more sophisticated measurements like the granular temperature fields [148]. Since in real experiments, these measurements have no basis of comparison to the real tracer behaviour, what PEPT provides is all that can be known. This is different in GATE-simulated experiments because the tracer trajectories are prescribed, meaning the fields from PEPT trajectories can be directly compared to the fields generated by the prescribed trajectories, providing insight into distortions from PEPT measurements [49, 46, 47]. It is then possible that if the real behaviour of the tracer is known it could then be used to correct the PEPT measurement, bringing it back in line with reality [46]. However, if this method is to be used to correct experimental measurements, then the GATE-simulated trajectories must closely approximate reality.

The goals of this Chapter are to demonstrate how realistic tracer trajectories prescribed in a GATE simulation can be used to optimise the length of time needed to run an experiment such that the ergodic assumption is valid and, additionally, to develop a framework for correcting PEPT reconstructed fields which are informed by the prescribed behaviour of the tracers. To do this, a new method of using highly-detailed Lagrangian trajectories of particles from a discrete element method (DEM) simulation as the basis for

GATE-simulated PEPT experiments is created [64, 46]. For the first time, this will allow a meaningful and direct comparison between PEPT and DEM data since the prescribed tracer position provides a means of direct comparison between a known trajectory and its PEPT-reconstructed measurement.

First, a DEM simulation of a system representative of a typical PEPT experiment is created and individual particle trajectories are extracted. In this case, the experimental system is a small rotating drum, known commercially as the GranuDrum, which is used for powder flow characterisation [30]. Then this trajectory is prescribed as the motion of a positron-emitting tracer in a simulated PEPT experiment using calibrated GATE model of the ADAC Forte dual-headed PEPT detector system [44, 45]. The GATE-simulated PEPT experiment reproduces the behaviour of the PEPT detector system, produces output quantitatively similar to that of real detectors, and can be processed using a PEPT algorithm to locate the tracer. To assess the difference between the DEM trajectory and the PEPT reconstructed trajectory, the trajectories are compared point-by-point to compute an average two-dimensional spatial error. Then, both trajectories are post-processed using the same treatment to reconstruct Eulerian, or time-averaged, fields which describe the system behaviour [148]. Since values generated from DEM are the prescribed data, PEPT reconstructed values can be compared to assess to what extent the behaviour of the modelled system is captured and, crucially, to what extent the PEPT-reconstructed trajectories differ from the model DEM, potentially allowing for the measurement to be corrected.

5.2 Methods

5.2.1 DEM Simulation of a Rotating Drum

To generate realistic tracer trajectories, first, a DEM simulation must be created. DEM is a simulation technique for the numerical modelling of granular systems by solving Newton's equations of motion in discrete time-steps [25]. This method resolves the Lagrangian motion of the particles, detects contacts, and calculates contact forces between particles. DEM is therefore a powerful tool to predict granular behaviour. In this work, the system simulated is the GranuTools Granudrum [30]. The GranuDrum is a type of rheometer, or flow characterisation equipment, for measuring properties of granular media [78]. It consists of a small, thin, rotating drum inside a box with a camera to measure the dynamic free surface which allows for the calculation of the angle-of-repose [120]. An image of the GranuDrum is shown in Figure 5.1. The reason why this system was chosen is that the experimental volume is relatively small, which reduces the computational resources needed, and the device has good mixing properties, which allows for the tracer to explore all areas of the system sufficiently within a relatively short amount of time, thereby reducing the length of time needed for the DEM simulation.

The DEM software LAMMPS Improved for General Granular and Granular Heat Transfer Simulations (LIGGGHTS) is used to simulate the GranuDrum behaviour [64]. The DEM model is created with a hollow rotating cylinder of the same dimensions as the GranuDrum and filled with simulated mono-disperse spheres of microcrystalline cellulose (MCC). MCC is a commonly used material in the pharmaceutical industries which has been used in PEPT experiments in previous work [155]. A geometric description of the GranuDrum can be found in Table 5.1 and a side-by-side comparison of the real and simulated GranuDrum rotating drum is shown in Figure 5.2.

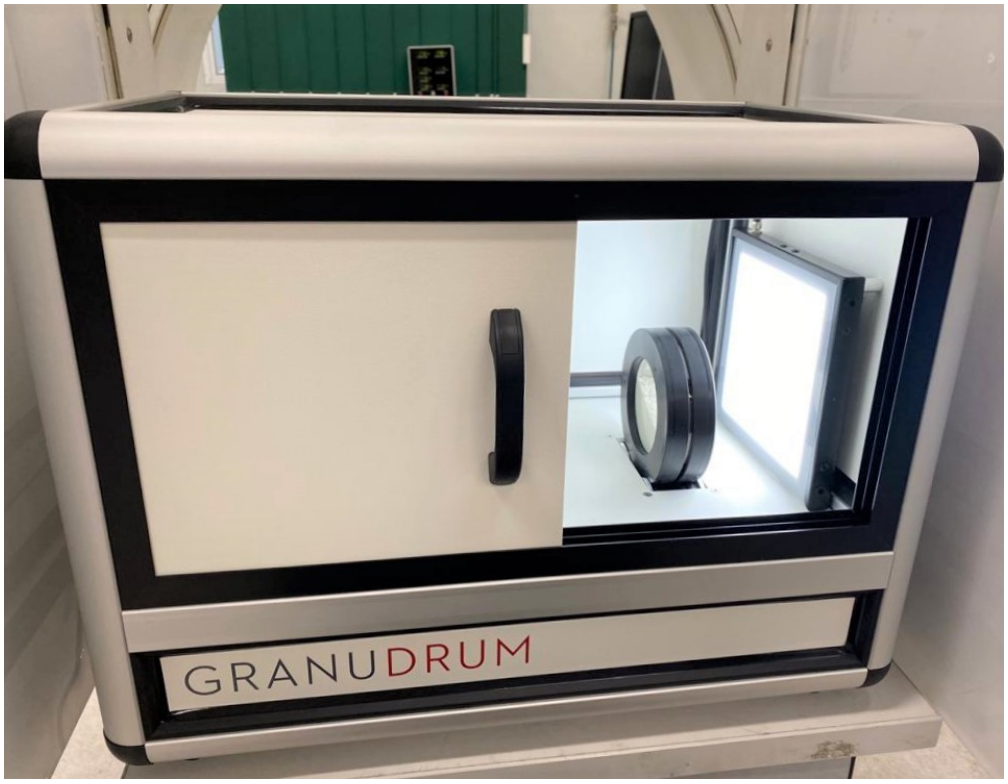

Figure 5.1: An image of the GranuTools GranuDrum [46].

The drum is filled with 10,000 mono-disperse MCC spheres of diameter 1.8 mm and rotated at 45 rotations per minute (RPM) for 30 minutes. At this RPM, the MCC particles are within the cascading flow regime, where mixing is expected to be strong [87]. Higher regimes such as cataracting and centrifuging can result in the crushing of particles or particles centrifuging around the edge of the drum, whereas lower speed regimes such as slipping and slumping result in the particles sliding over each other rather than properly mixing [101, 152]. The positions of all particles are recorded and used to extract trajectories for comparison with PEPT-detected trajectories of a single tracer. At the end of the simulation, the trajectories of all 10,000 particles have been recorded at a time-interval of 1 ms, as this keeps the distance travelled by the fastest particles (moving at approximately 1 ms^{-1}) between successive locations smaller than the radius of the MCC sphere itself. This frequency is also representative of the achievable temporal resolution of the ADAC Forte PEPT detector

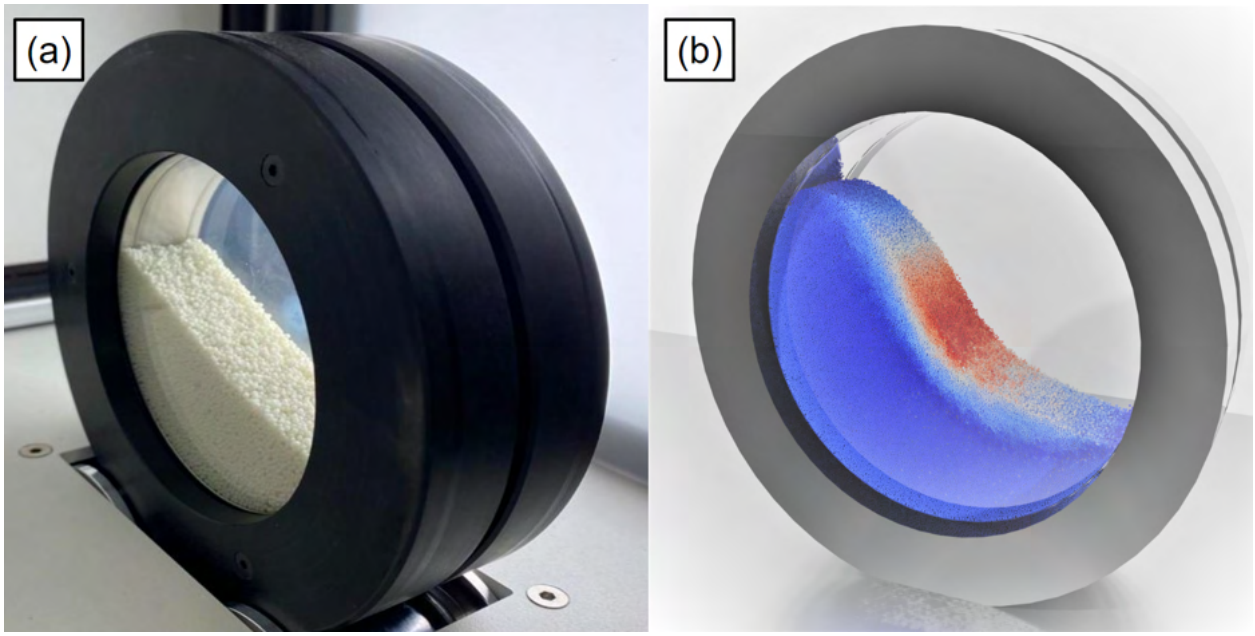

Figure 5.2: (a) An image of the GranuTools GranuDrum. (b) Simulated GranuDrum visualisation coloured by particle velocity [46].

system used to image the flow.

The material properties of the DEM particles used in the simulation are shown in Table 5.2. The first four parameter values in the table were determined using the Autonomous Characterisation and Calibration using Evolutionary Simulation Software, ACCES [92] The remaining parameters were obtained from published values [128].

Table 5.1: GranuDrum Description [46].

Characteristic	**Values**
Drum Diameter	84 mm
Drum Thickness	20 mm
Glass Wall Thickness	3 mm

Table 5.2: Measured material properties of MCC particles prescribed to the LIGGGHTS simulation [46].

Simulated Property	**Values**
Particle-Wall Friction	0.32
Particle-Particle Friction	0.32
Coefficient of Restitution	0.3
Rolling Friction	0.0025
Young's Modulus	5 x 10^6 Nm^{-2} [53]
Poisson's Ratio	0.3 [77]
Density of MCC	1580 kgm^{-3} [128]

5.2.2 GATE Simulation of a Rotating Drum

Following the simulation of the DEM model, the raw DEM data were converted into a format that can be used in a GATE simulation. Two things are needed from this data: (1) individual trajectories of particles, and (2) a three-dimensional array representing the packing density of the flowing MCC powder which can be used to approximate the scattering of gamma-rays in the drum. To perform both of these tasks, a Python library, DEM2GATE, was created [46]. GATE has capabilities to read from files containing trajectory information and can import complex geometries from the three-dimensional image, called voxelised geometries. The novel work described here is the process of extracting this information from a LIGGGHTS simulation and converting it to a GATE readable format. The workflow for this task is as follows:

1. Tracer Trajectories

    (a) Extract particle position from LIGGGHTS;

(b) Convert trajectories to GATE Placements Format;

(c) Prescribe trajectory to GATE source.

2. Voxel Array

(a) Calculate a three-dimensional array of packing density;

(b) Export array as a MetaImage;

(c) Define GATE material definitions;

(d) Create a GATE range translator;

(e) Place voxel array in GATE simulation.

The DEM2GATE software is free and open-source and can be used by interested researchers who wish to apply these techniques to their own work. The software can be found in the Positron Imaging Centre's DEM2GATE GitHub Repository.

The DEM simulation of the GranuDrum tracks the position of all 10,000 simulated particles simultaneously. For every time step, the positions are updated based on the current forces acting on the particles. At regular intervals of 1 ms, the positions of all tracers are written to the VTK file. To extract a DEM trajectory, the VTK files are read into Python using a VTK reader, allowing a user to view and select specific data. Next, the time and x, y, and z positions for a specified particle are extracted and written into the GATE placements file format. The Placements file is a text file consisting of a header describing the format of the data accompanied by the data columns. GATE uses this file to identify where a source should be placed inside the simulation during a particular time step by selecting the time and corresponding positions that are closest to the current GATE time step. Finally, the particle trajectory is prescribed in the GATE simulation by declaring that the source position should be read from this placement file via the Generic Move function.

The second set of data to extract from the simulated rotating drum is a 3-dimensional array of the packing density, which is used to produce a voxel array for the GATE simulation. In this representation of the GranuDrum, each voxel has a value corresponding to the number of particles it contains; when all voxels are the same size, this provides a measure of packing density. To generate a voxel array of the flowing particles in the GranuDrum, an empty array is first created which breaks up the drum volume into a 50 x 50 x 50 grid over the simulation space of the LIGGGHTS simulation, -0.048 to 0.048 m in the x and z, then -0.002 to 0.0221 m in the y-direction. This array is created in Python using NumPy arrays binning the VTK particle data. Then, for each time step, the number of particles in each voxel is summed, creating a time-averaged packing density. This voxel array is normalised and stored in a GATE-readable MetaImage format.

To map the normalised voxel value to the real corresponding bulk density, the maximum bulk density of the MCC particles must be calculated by multiplying the particle density, $\rho = 1580\ kgm^{-3}$, with the packing density. In this case, a loose random packing fraction of 0.6 was used, yielding a bulk density of $\rho = 948\ kgm^{-3}$ [93]. This is then equated to the maximum voxel value. All other voxels are mapped to 50 discrete combinations in between the bulk density and the density of air, e.g 20% of the bulk particle density and 80% air density would lead to $\rho_{20/80} = 189.9\ kgm^{-3}$. This is accomplished using a GATE range translator which uses the voxel value to choose a material from the list of predefined material mixtures.

Once the voxel array has been created, shown in Figure 5.3, this must be converted to a GATE-readable format. This is accomplished by saving the voxel image as the medical imaging format, MetaImage, followed by creating a range translator to map the voxel values to predefined materials in the GATE material database. The three-dimensional NumPy array of packing density is converted to a MetaImage by using the SimpleITK Python package and specifically the GetImageFromArray function. To map the voxel values to the predefined

materials, a GATE range translator is applied which uses the voxel value to choose a material ranging from bulk MCC to air. To ensure that a smooth transition from bulk MCC to air occurs, 50 material combinations between bulk MCC and air are generated. Each material has density calculated using the respective fractions of bulk MCC and air.

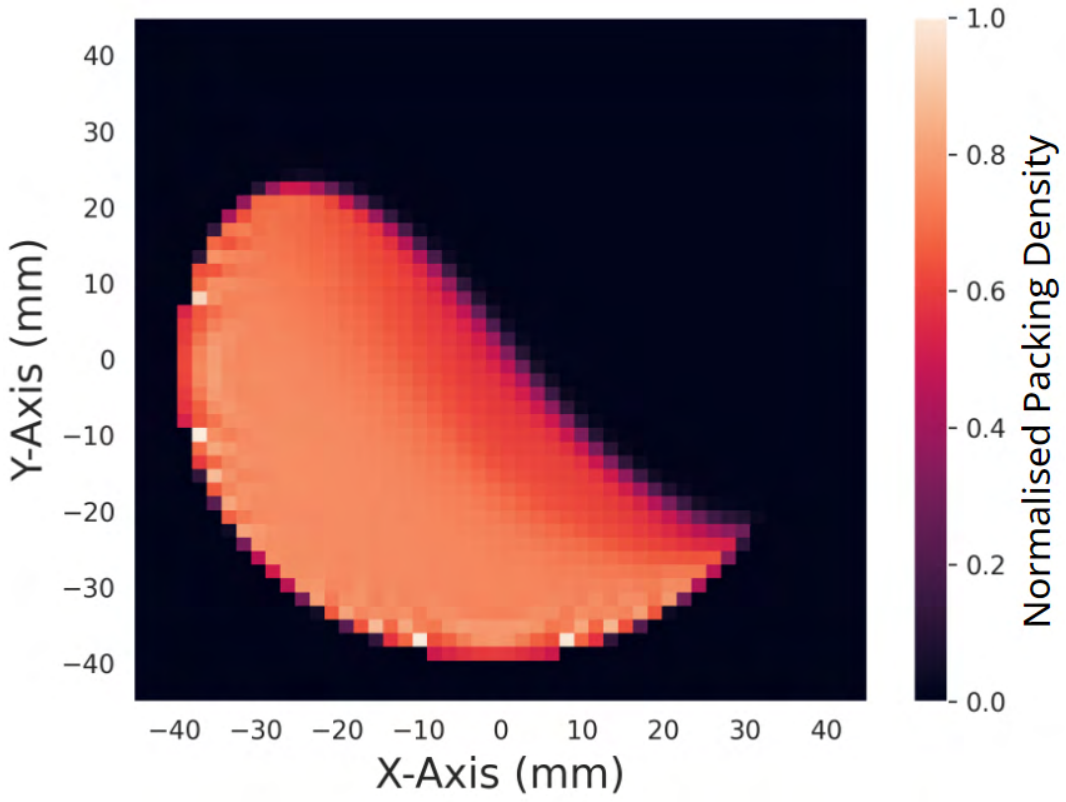

Figure 5.3: The packing density of the GranuDrum from a slice of the three-dimensional voxel array [46].

Next, the tracer trajectories and voxel array, along with a geometric representation of the GranuDrum structure, are placed in the centre of the field-of-view (FOV) of the validated model of the ADAC Forte detector [44]. The geometric model is shown in Figure 5.4. Two scenarios are simulated: a single DEM trajectory in air without the GranuDrum and a single DEM trajectory in the GranuDrum geometry. These represent a best-case (minimal scattering) and a realistic case (considerable scattering) for a PEPT experiment,

176 respectively. It is worth noting that the velocity of the tracer in both cases remains the
177 same, and the tracer follows the same trajectory.

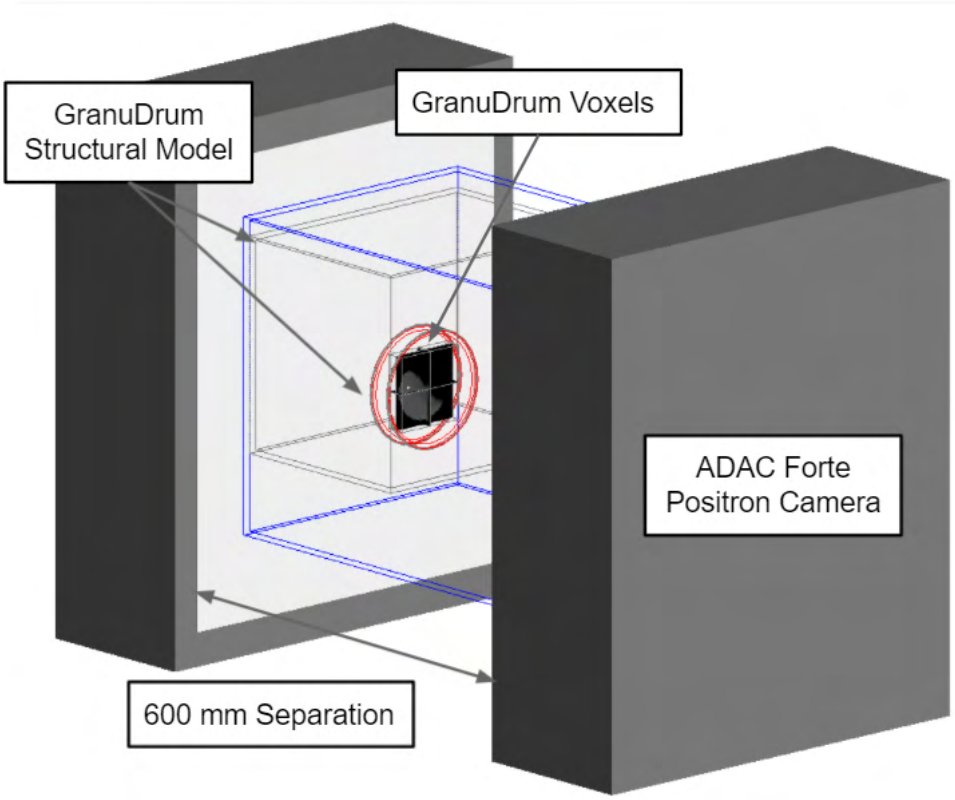

Figure 5.4: An annotated visualisation of the GATE-simulated GranuDrum [46].

178 The geometric model of the GranuDrum follows the description given in Section 5.2.1.
179 To be consistent with the real GranuDrum, the rotating drum inside the GranuDrum struc-
180 ture is translated 145 mm closer to one side of the ADAC Forte detector. This has the effect
181 of reducing the count-rate detected from the tracer due to the lower geometric efficiency.
182 The trajectory is co-registered with the voxel array such that the DEM trajectory is fully
183 within the voxels at all times. This case will have considerably more scattered LoRs due
184 to the added material between the tracer and detector. In turn, it will result in both a
185 lower detection rate for the tracer and less accurate detection, caused by the incorporation
186 of scattered LoRs in PEPT algorithms, thus providing a full, true-to-life representation of a

real PEPT experiment.

Once the detector model is applied and set to a detector separation of 600 mm, the next step is to define the source. The DEM particle is a solid 1.8 mm diameter sphere of MCC with density $\rho = 1580\ kgm^{-3}$. This material definition is added to GATE and modelled as a volumetrically activated, positron-emitting source of 10 MBq activity. A 10 MBq tracer has been shown to result in a count-rate near the peak true LoR count-rate for the ADAC Forte at 600 mm separation [44]. The positrons are prescribed an energy spectrum corresponding to that of fluorine-18. Then, using the Generic Move function, the DEM trajectory which has been converted to a GATE placements file is prescribed as successive positions of the source. This allows the GATE-modelled tracer to take the same path as the DEM tracer, with the addition of emitting positrons.

To run GATE simulations efficiently, the inherent parallelisation of GATE is leveraged by splitting the simulation into smaller jobs, rather than simulating the whole 30-minute trajectory in one simulation. The recorded LoRs are then combined later. The simulations are run on BlueBEAR, the University of Birmingham's High-Performance Computer [133]. The 30-minute simulation is broken into 200 smaller jobs, each running for 9 seconds of the simulation. The output of the simulation contains LoRs which can be processed with a PEPT algorithm to extract a trajectory of the tracer.

5.2.3 Data Analysis

Several PEPT algorithms exist, as has been mentioned in Chapter 3, and their performance has been assessed in other work [149]. The PEPT-ML algorithm is used in this work to cluster samples of the recorded LoRs and locate the tracer position [91]. The sample size of LoRs is prescribed as $N_{LoRs} = 200$. A moving window of LoRs is used to provide more

samples of LoRs, with the latter 50% of LoRs in a given sample used in the next sample. All cutpoints falling within 0.15 mm of the LoR pair are saved for clustering. For each sample of cutpoints, the densest 15% of points are considered core points of the cluster and their centroid is calculated. After this, the reconstructed tracer location second-pass clusters with a sample size of 30 points, with 70% of the densest points being the core points and their centroid calculated. To produce nearly as many second-pass locations as the first-pass method, a moving window of one less than the sample size is applied, in this case, a sample size of 29. An example of a sample of LoRs and a reconstructed tracer trajectory from the GATE simulation are shown in Figure 5.5.

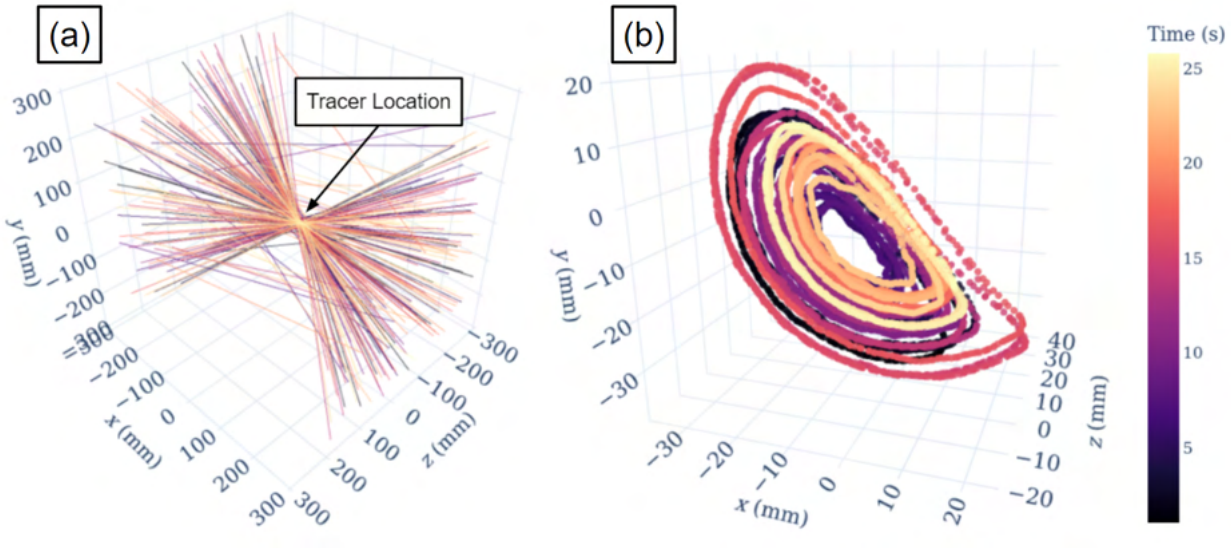

Figure 5.5: (a) A sample of 200 LoRs and (b) a PEPT reconstructed trajectory of approximately 25 seconds of the tracer flowing in the GranuDrum [46].

Once a PEPT trajectory has been produced, there are two ways to assess the accuracy of the PEPT-detected trajectory. The first method is to compare the position of the tracer compared to the prescribed position, averaging the mean three-dimensional error using Equation 1.14. However, while looking at the point data of a trajectory is useful for characterising the uncertainty in the position of an individual tracer, an individual tracer does

not necessarily explore the whole phase-space of particles within the system, particularly if it has not been exploring the system for a long enough time [144]. A more robust method of comparison would be to compare a reconstructed field of the system to a field produced directly from the DEM simulated particles. For example, an Eulerian flow field generated from a simulated PEPT measurement, such as the velocity field, can be compared to the field reconstructed using the model DEM data. Comparison of macroscopic fields shows not only to what extent PEPT trajectories follow an individual tracer but also how well the bulk behaviour of the system is reconstructed from a single tracer. This type of comparison is more significant for PEPT users because it is a measure of how the uncertainty in PEPT trajectories actually affects the measurement of a bulk system behaviour.

In this work, both the overall spatial error of the PEPT trajectory and the effect of this on the measurement of bulk system behaviour is examined. The generation of PEPT reconstructed fields also offers an opportunity to take advantage of the recent developments in the PEPT-ML library which uses time-step interpolation and polynomial fitting of positions in a Savitzky-Golay filter to derive tracer velocity [116]. The reconstructed Eulerian fields for the PEPT trajectory in air and PEPT trajectory inside the model of the GranuDrum are compared against fields generated by averaging all DEM trajectories.

All Eulerian fields are generated using the same procedure, whereby for each time step of a trajectory, an instantaneous measurement of a system field is assigned to bins spanning the drum volume. For simplicity and because the GranuDrum has a small drum thickness in comparison to its diameter, only two-dimensional fields are computed. The number of bins in each dimension as well as the value range over which the function is applied is provided by the user, and then the mean, standard deviation, or counting is performed on the values falling within each bin. The fields calculated in this work are velocity components, velocity magnitude, acceleration components, acceleration magnitude, granular temperature, and occupancy. All fields are computed on the same 30 x 30 grid spanning ±45 mm in the x

and y directions centred on the centre of the GranuDrum, producing a field with 3 x 3 mm resolution. Velocity is calculated using a Savitzky-Golay filter with a window size of y points as a step within the PEPT-ML algorithm and is written into the trajectory as an appended column of data, thus the velocity components and magnitude are pre-computed and only need to be binned. The velocity fields are calculated by finding the mean values within each bin. Acceleration is calculated by numerically differentiating the velocity by time according to Equations 5.1 and 5.2. The acceleration fields are calculated by finding the mean values within each bin.

$$a_x = \Delta v_x / \Delta t \tag{5.1}$$

$$a = \sqrt{a_x^2 + a_y^2} \tag{5.2}$$

Granular temperature is a measure of the squared velocity fluctuations [143]. This is shown in Equations 5.3 and 5.4, through which the 1-dimensional and two-dimensional granular temperatures are calculated. To compute the granular temperature field, first, the x and y velocity fields are calculated and the bin number for each instantaneous velocity is recorded. Using this information, the mean velocities are subtracted from the instantaneous velocity to produce the fluctuating velocity according to the tracer location within the system. Then the granular temperature can be calculated using the squared velocity fluctuations and multiplying by the mass of the tracer [143].

$$T_x = m(v_{xi} - \bar{v}_x)^2 \tag{5.3}$$

$$T = \frac{1}{2}(T_x + T_y) \tag{5.4}$$

The occupancy measures how long a particle spends within a region of the system. This is useful since occupancy is an indirect measurement of particle density in different areas of the systems, which can be used to infer particle dynamics in industrial systems [100, 87, 119]. In this work, the DEM and PEPT trajectory positions are recorded in constant intervals of time and continuously tracked, thus the occupancy percentage for a bin can be calculated by counting how many times a particle has been found in that bin, multiplying by the time-step, then dividing by the difference in time between the last and first particle position in the trajectory. The occupancy is calculated using Equation 5.5.

$$O = 100\frac{N_{points}t_{step}}{t_f - t_0} \tag{5.5}$$

All Eulerian fields for DEM and PEPT are binned in 3 x 3 mm bins, computed over the same range, and using the same number of bins. To compare the fields, the mean bin percent difference is calculated. The percent difference is calculated by finding the absolute difference between the PEPT reconstructed fields and baseline DEM model fields, then dividing by the DEM model value and multiplying by 100. The individual bin percent differences are then averaged as shown in Equation 5.6.

$$\varepsilon_{Mean} = \frac{100}{N_{bins}} \sum_{N=1} N_{bins} \frac{|X_{Experiment} - X_{Model}|}{X_{Model}} \tag{5.6}$$

However, to avoid outlier bin errors distorting the mean percent difference due to a low number of tracer passes, a threshold is set using the occupancy so that only bins which have an occupancy higher than 0.00225% are considered. This value was chosen by looking at the cumulative probability distribution of the bin occupancy, as shown in Figure 5.6. An occupancy of 0.00225% immediately precedes a sharp rise in the plot, thereby excluding only the bins which are not adequately explored and do not contribute significantly to the

286 overall system behaviour. The results of this threshold applied to the occupancy are shown
287 in Figure 5.7.

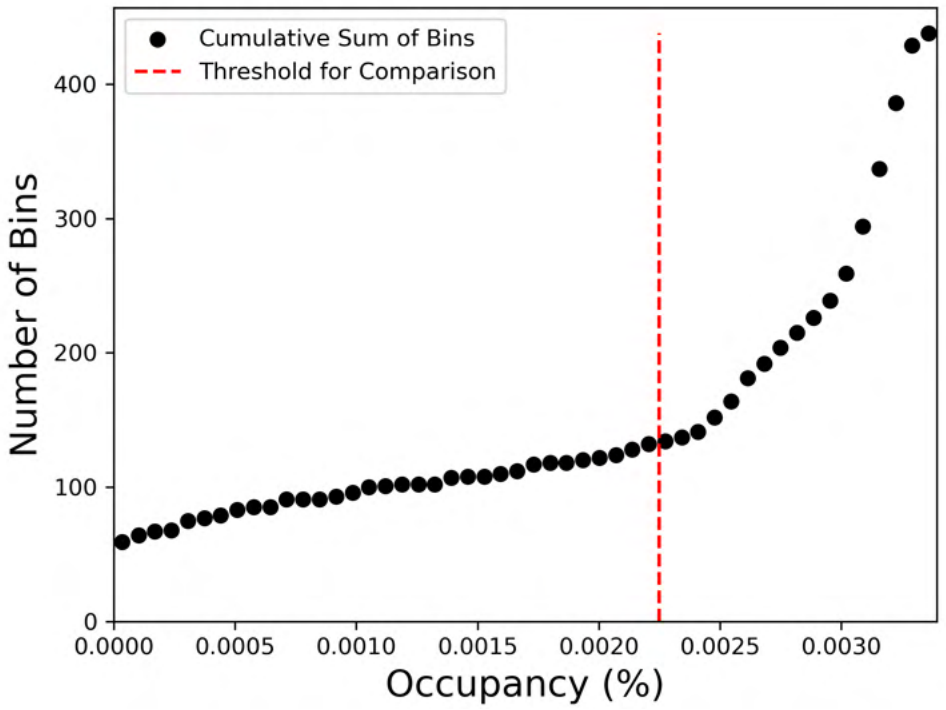

Figure 5.6: The cumulative probability of occupancy in the GranuDrum used to determine which bins to consider for comparison [46].

288 5.3 Results

289 The results can be summarised according to the three main components of this work: (1)
290 analysis of the DEM simulation, (2) comparison of the PEPT reconstruction with the DEM
291 model, and (3) demonstration of a new method to correct PEPT reconstructed values based
292 on the model DEM.

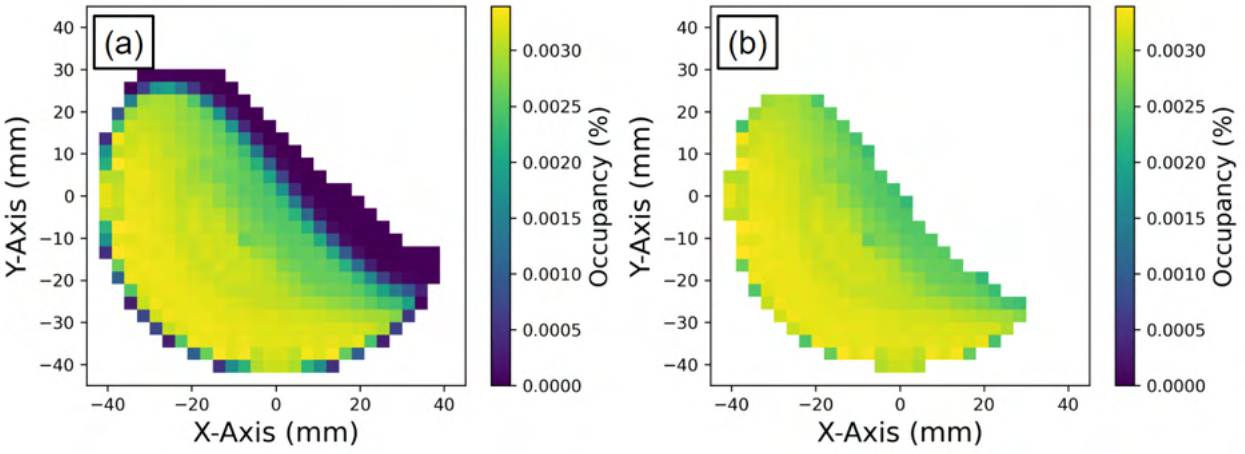

Figure 5.7: (a) A threshold is applied to remove bins that the tracer explores only a few times. (b) The bins left after thresholding are the only bins used for comparison in later analysis [46].

5.3.1 Analysis of the DEM Simulation

Using the full 30-minute trajectories extracted from the DEM simulation of the monodisperse GranuDrum, the Eulerian fields were generated. In Figure 5.8, the velocity components and magnitude field are shown. This reveals an active, fast-moving region on the free surface and a passive and slow-moving region in the bulk. These are both well-known features of rotating drums operating at the Froude number of 0.1, demonstrating that the simulation provides sensible flow behaviour [28].

Numerically differentiating the tracer velocity with respect to time yields acceleration. The acceleration magnitude field is produced and shown in Figure 5.9. The shoulder and toe regions experience the highest accelerations. The shoulder of a rotating drum is the region at the top of the cascade where particles depart from the wall [87]. The toe of a rotating drum is the region at the bottom of the particle cascade where particles impact the wall [87]. Since particles in the shoulder falling down the cascade are accelerated by gravity and particles in the toe are abruptly decelerated by crashing into the drum wall this is expected.

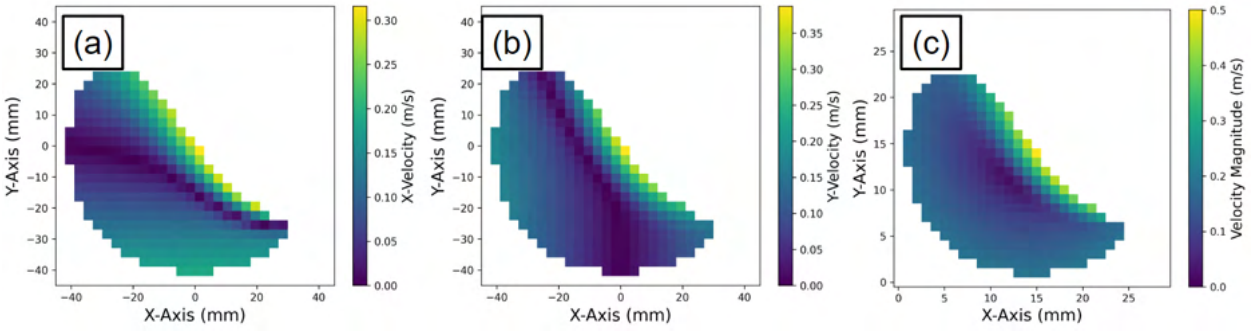

Figure 5.8: Velocity fields of particles in the mono-disperse DEM simulation of the GranuDrum: (a) mean particle velocity in the x-direction, (b) mean particle velocity in the y-direction, and (c) mean particle velocity magnitude [46].

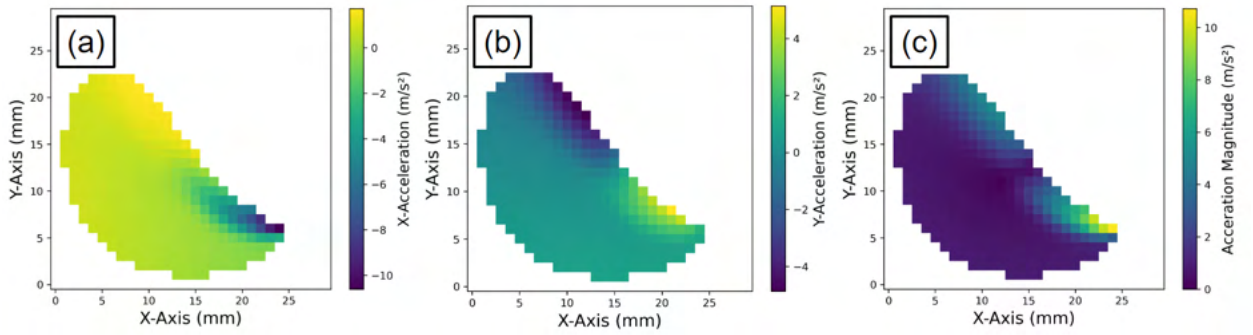

Figure 5.9: Acceleration fields of particles in the mono-disperse DEM simulation of the GranuDrum: (a) mean particle acceleration in the x-direction, (b) mean particle acceleration in the y-direction, and (c) mean particle acceleration magnitude [46].

The granular temperature is shown in Figure 5.10. The granular temperature is highest in the active region of the drum and especially in the toe. This is because collisions in this area are frequent and energetic, leading to large fluctuations in velocity [151].

These fields and the previously calculated occupancy field were then generated for every minute of the simulation time for individual DEM particles. These are compared to the previously shown fields produced after 30 minutes using all of the particles which represent the steady-state behaviour of the drum. This is done to investigate how the measured system

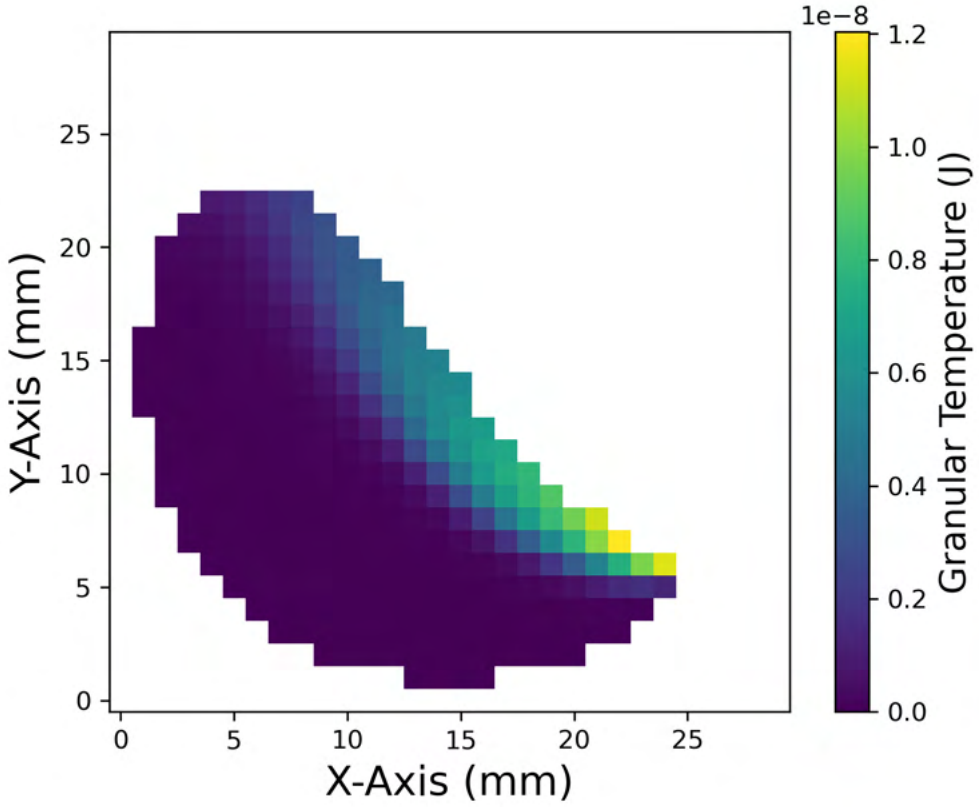

Figure 5.10: Granular Temperature in the mono-disperse GranuDrum DEM simulation[46].

behaviour reconstructed from a single particle compares to the time-average bulk behaviours of the system. This shows how long is needed for one particle to approximate the behaviour of all particles, finding at what point the ergodicity is achieved to a specified level of agreement. The results of this comparison are shown in Figure 5.11.

Across all measured fields, the differences from steady-state fields are found to decrease logarithmically with the measurement time. Derived measurements of lower-order than acceleration, such as velocity, are found to reach a high level of accuracy relatively quickly; 1% mean difference after 10 minutes of measurement time. Higher-order derived measurements like acceleration, which rely on the velocity measurement, take longer to reach the same level of accuracy. For all fields, the mean percent differences are under 10% by 30 minutes. These results show for this particular system not only can system behaviour be reconstructed from

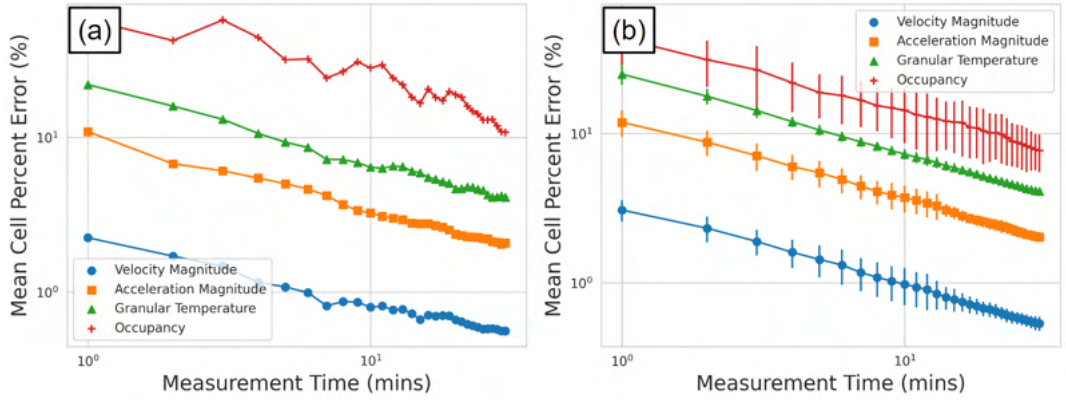

Figure 5.11: Estimating the time needed to ensure adequate statistics are collected: (a) a single tracer compared to the bulk behaviour, (b) average of multiple tracers with error bars representing standard deviation of the percent differences [46].

single particles, but also the length of measurement time needed to reach a desired level of accuracy can be estimated.

5.3.2 PEPT Reconstruction

After running the GATE simulations of a particle in air and of a particle in the simulated mono-disperse GranuDrum, the reconstructed PEPT trajectories are compared to the prescribed DEM trajectory. For a 2-second long segment of the simulation, the positions for the DEM tracer and the two PEPT reconstructed trajectories are plotted and then compared point-by-point to compute a mean spatial error. Further, the instantaneous spatial errors are plotted according to the tracer positions in Figure 5.12. The tracer in air and the tracer in the GranuDrum voxels over the 30-minute trajectory are tracked to a mean spatial accuracy of 0.175 mm and 0.412 mm, respectively. Additionally, the detection rates of the two tracers are 889 Hz and 355 Hz, respectively. This demonstrates that DEM2GATE can be used to reconstruct particle behaviour and then quantify the spatial and temporal resolution of the reconstruction.

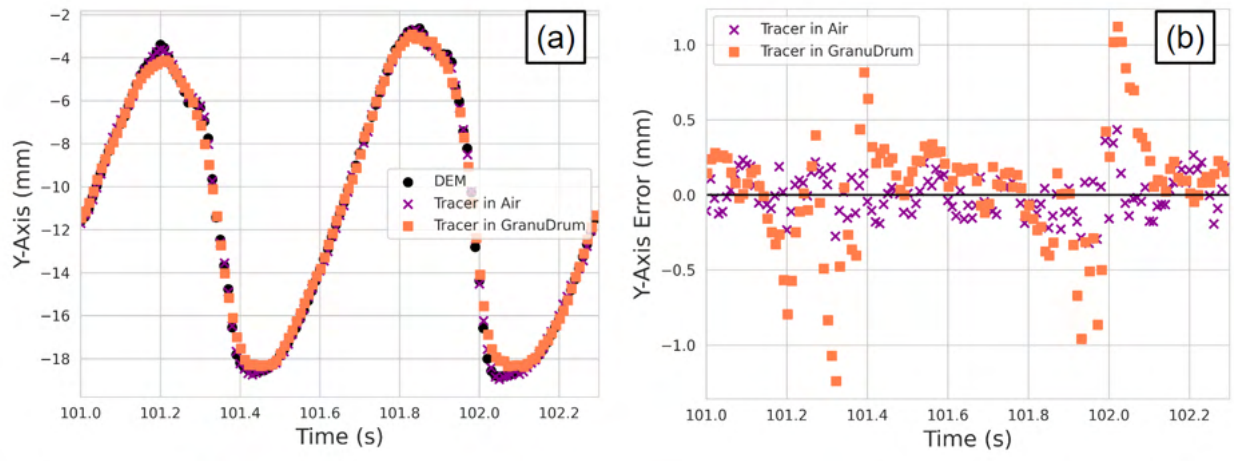

Figure 5.12: The prescribed DEM trajectory, simulated tracer in air, and simulated tracer in the GranuDrum: (a) the positions of the tracer, (b) the error in the PEPT reconstruction [46].

The spatial errors can be broken down further into the mean spatial error in each bin, shown in Figure 5.13. Figure 5.13 provides information about where spatial errors occur and indicates why this happens. Higher errors occur when the tracer is falling from the shoulder of the drum or crashing into the toe region. By referring back to Figure 5.9, these are the two areas with the highest accelerations.

When a bin-by-bin comparison between the spatial error in PEPT measurement and the model DEM acceleration is plotted in Figure 5.14, the relationship between the two is shown to be a linearly increasing spatial error with increasing tracer acceleration. When a straight line is fitted to the data, the y-intercept is the mean error that is to be expected for a tracer at rest or at constant velocity. The y-intercept value is higher for the tracer in the GranuDrum voxels, most likely due to the increase in scattered LoRs as this has been shown to increase PEPT error in other work and thus demonstrates that this is an important relation with real physical significance [149].

Additional analysis of the PEPT trajectories is performed by reconstructing the Eu-

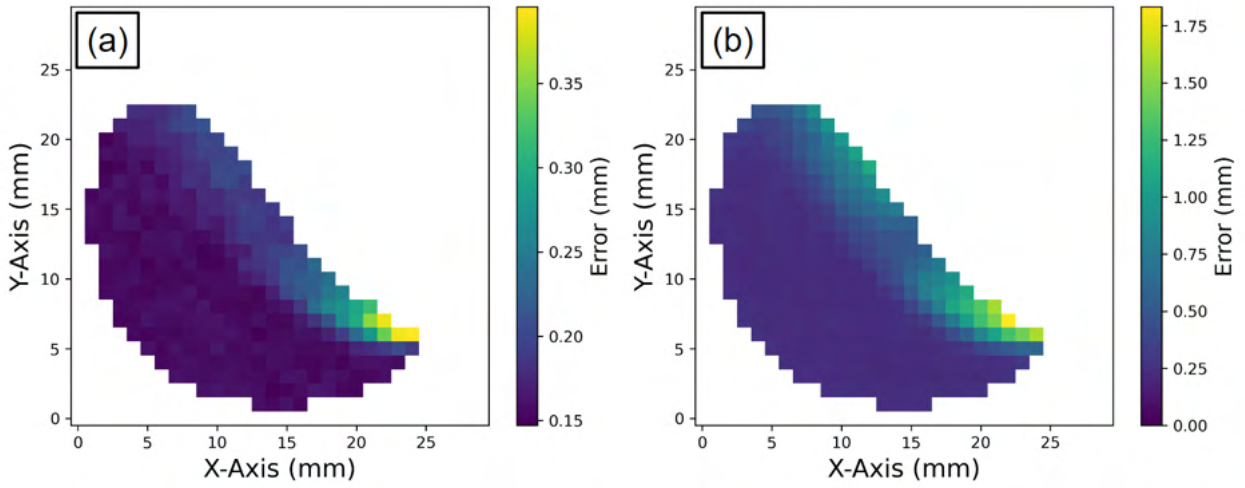

Figure 5.13: The binned mean errors of the PEPT reconstructed trajectory: (a) the mean error of the tracer simulated in air, (b) the mean error of the tracer simulated in the GranuDrum [46].

lerian fields. These are calculated in the same manner as the DEM fields and are calculated every minute and then compared to the DEM values calculated using all the model DEM trajectories after 30 minutes of the simulation. The results of comparing the Eulerian fields are shown in Figure 5.15. Figure 5.15 shows that for the velocity, acceleration, and occupancy fields the PEPT reconstruction aligns with the model DEM fields, but for granular temperature, the fields do not align and do not get more accurate with longer measurement times.

5.3.3 Correcting PEPT Measurements

As shown in Figure 5.15, the PEPT reconstructed granular temperature fields do not agree with the model DEM fields. These fields rely on the fluctuating velocity components and, as such, they are particularly prone to errors caused by both low temporal resolution and the inherent uncertainty in PEPT measurements. The low temporal resolution of a PEPT

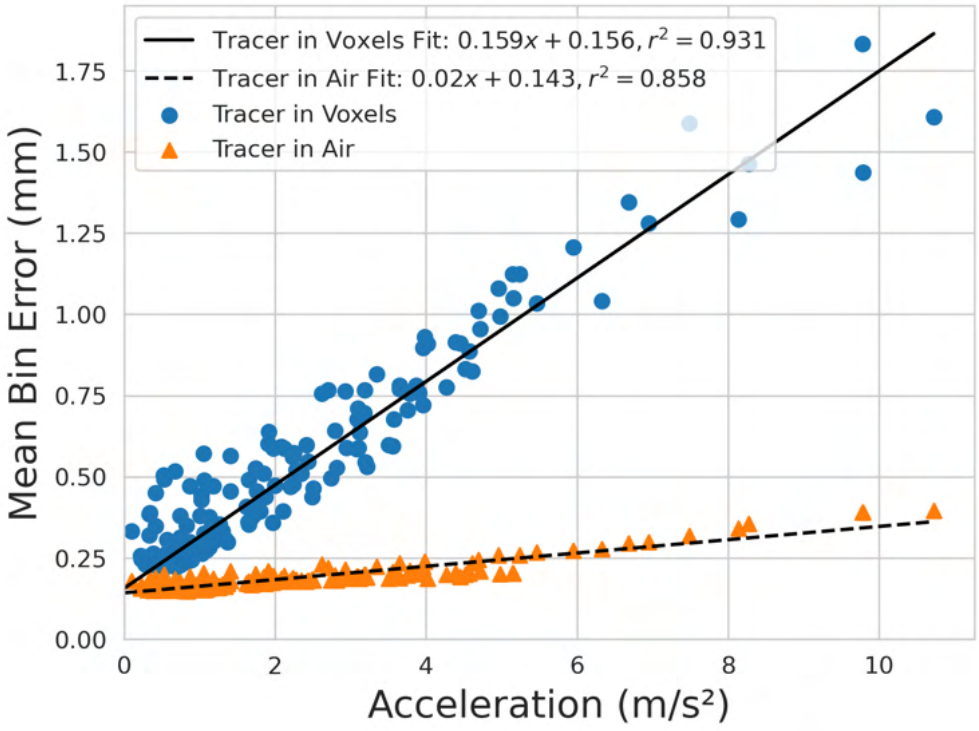

Figure 5.14: The PEPT spatial error increases linearly with acceleration, with the y-intercept being the mean error for a constant velocity tracer; the y-intercept is higher for the PEPT tracer in voxels because of the increased scattering of LoRs [46].

measurement in comparison with the behaviour of the DEM tracer smooths out fluctuations in velocity while uncertainty in the PEPT measurement induces fluctuations in velocity. These two factors come together to produce errors in the PEPT reconstructed fields.

However, when the PEPT and model DEM fields are compared bin-by-bin, there exists a correlation between them, shown in Figure 5.16. When the square root of the granular temperature (square root since this is calculated using a squared velocity term) for the PEPT reconstruction and model DEM are compared, a straight line can be fitted to the data. The fitted equations show that bins with low granular temperature in the model DEM fields are over-predicted by PEPT, as evidenced by a positive y-intercept. Further, at high granular temperatures in the model DEM, PEPT measurements under-predict the values,

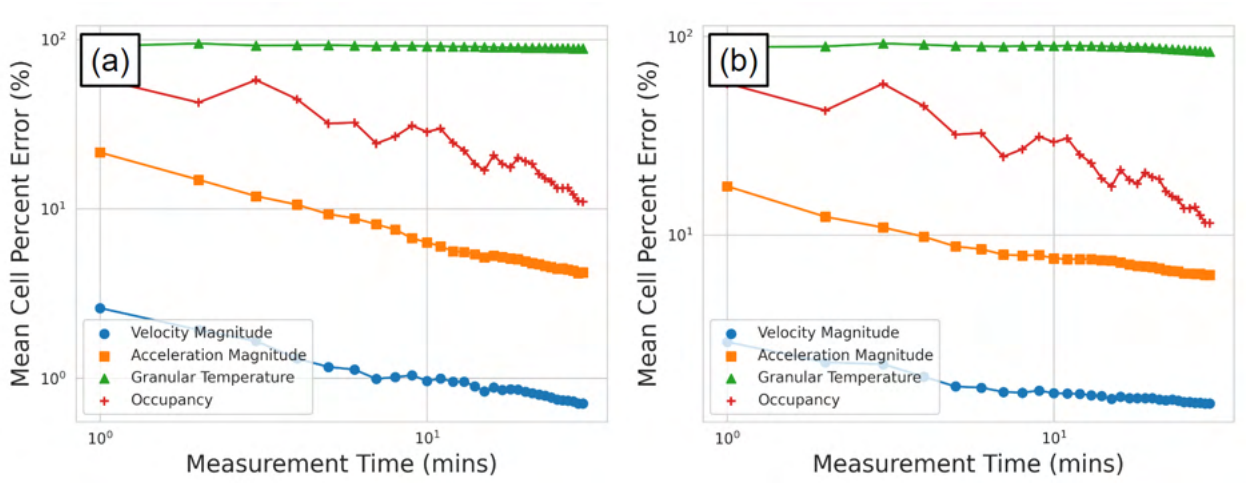

Figure 5.15: Estimating the time needed to ensure adequate statistics using results from a PEPT tracer: (a) in air and (b) in the GranuDrum. Notice the Granular Temperature is not improving for either case with additional measurement time [46].

as evidenced by a slope gradient of less the 1.

Not only does this result provide insight into a possible source of error when comparing PEPT measurements with DEM models, but also shows that the PEPT measurements have predictable errors which can be corrected. To correct the PEPT measurement, the terms of the linear model can be rearranged and applied to the PEPT reconstructed fields. When this is done, the mean bin percent difference over the measurement time aligns more closely to what is expected from the model DEM simulation, as shown in Figure 5.17. This correction method is robust because it is not dependent on the overall behaviour of the particle within the system, only its velocity fluctuations. Thus, to produce a correlation between the PEPT reconstructed values and model DEM values, a fully calibrated DEM model is not required, only a model which produces velocity fluctuations across the range which might be expected in the experiment.

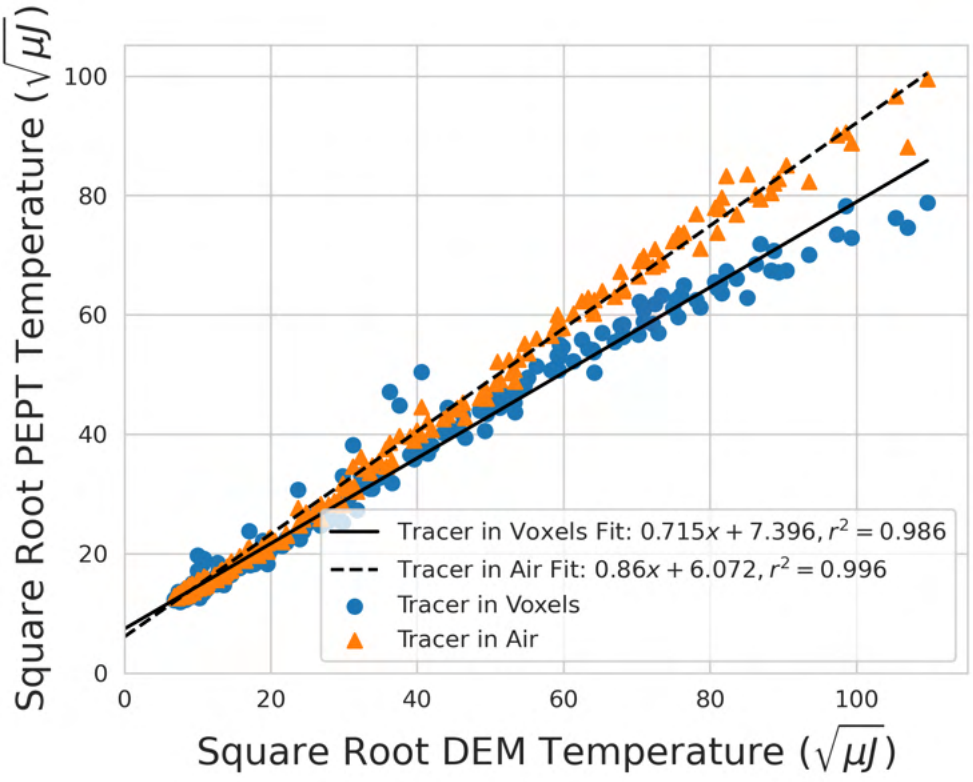

Figure 5.16: A correlation between the PEPT reconstructed and model DEM fields can be used to correct the PEPT-measured granular temperature [46].

5.4 Discussion

In this Section, the use of novel tools in the open-source DEM2GATE library to combine DEM simulation with GATE-simulated PEPT experiments was demonstrated. A LIGGGHTS DEM simulation of mono-disperse MCC particles in a model of the GranuTools GranuDrum was simulated and the DEM simulation used as the basis of a GATE-simulated PEPT experiment [64, 30]. From this DEM simulation, Lagrangian trajectories of individual particles were extracted and, additionally, a volume representing the time-averaged particle density throughout the rotating drum was produced.

A GATE simulation was created using the extracted DEM trajectories to move a

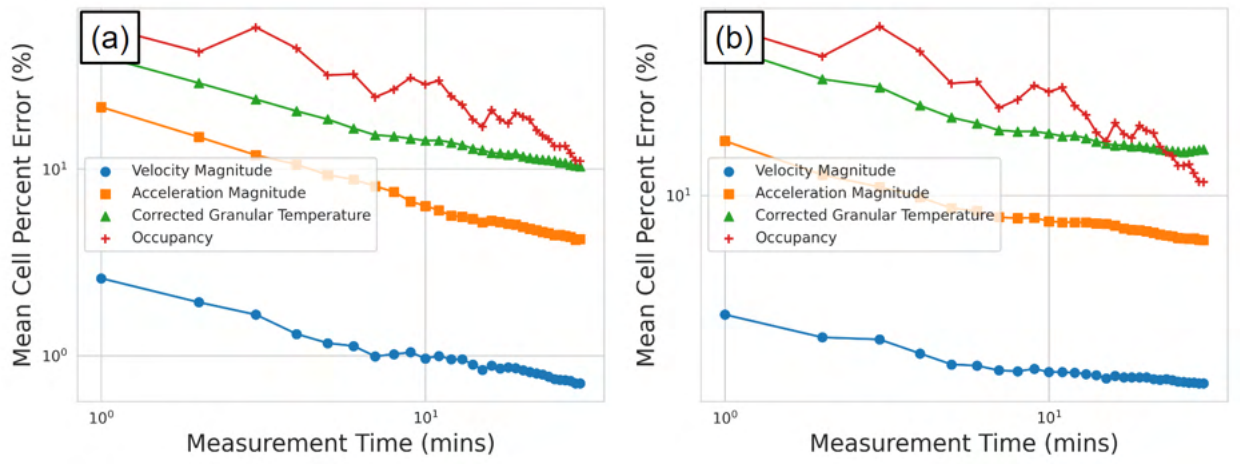

Figure 5.17: The mean bin percent differences for the PEPT reconstructed tracers: (a) in air, and (b) in the GranuDrum. When the correction for granular temperature is applied, the mean bin percent differences for these fields are improved, reaching approximately 10% after 30 minutes [46].

positron-emitting source modelled on a PEPT tracer. In a separate GATE simulation, the volume of particle density and a structural model of the GranuDrum were added to produce realistic gamma-ray scattering that would be expected in a real experiment. The synthetic LoRs collected over the 30-minute-long trajectory were processed using the PEPT-ML algorithm, yielding a PEPT-tracked version of the trajectory in both GATE-simulated experiments [91]. The reconstructed trajectories were compared with the prescribed model DEM trajectories through point-by-point comparison, as well as through two-dimensional Eulerian fields. This analysis showed that there were two main drivers of spatial error, particle acceleration and gamma-ray scattering, with the effect of these being quantified.

Another important finding was that the accuracy of reconstructed fields is shown to depend on the length of measurement time. Specifically, the DEM data show that accuracy improves logarithmically for the reconstructed fields and that higher-order derived measurements, such as acceleration fields, take longer to reach a desired level of accuracy compared

to lower-order derived measurements, like velocity fields. The reason for this was assumed to be because of error propagation in the higher-order measurements. After 30 minutes of a single-particle GATE-simulated PEPT measurement velocity, acceleration, and occupancy fields were reconstructed to within 10%. However, the granular temperature PEPT measurements produced mean bin percent errors of approximately 50% and above. A correction for the PEPT data was developed by comparing the PEPT and DEM values for these two measurements on a bin-by-bin basis, finding they were linearly correlated. Low model DEM values in these fields were over-predicted in the PEPT measurement, and then under-predicted for higher values. When the PEPT measurement correction was applied, the granular temperature for both GATE-simulated PEPT experiments after 30 minutes of measurement time reached approximately a 10% mean bin error.

This work shows, to the extent that DEM simulations replicate the real behaviour of particles in granular media, that GATE-simulated PEPT experiments faithfully reproduce the PEPT-measured Lagrangian particle trajectories with high spatiotemporal resolution [46]. The Eulerian fields reconstructed from the single-particle GATE-simulated PEPT experiments are capturing the dynamics produced from full-field DEM data, thus suggesting that the PEPT experiments capture real system behaviour. Regarding the length of a PEPT experiment, there has been work to try and relate the necessary duration to key system parameters, finding that duration scales with the system size and inversely with mean velocity [144]. However, these findings were made using only the DEM trajectories of particles and do not take into account the spatial and temporal resolution limitations of the PEPT technique. Additionally, there has never previously been an attempt to correlate PEPT-measured quantities to known DEM quantities, with this work being the first of its kind in demonstrating that it is possible to correct a PEPT measurement. The methods introduced in this Chapter can be similarly applied to any DEM-simulated granular system, provided that the DEM output has been verified to represent the realistic flow behaviour of the system. A trajectory

435 extracted from the DEM simulation can be used to assess whether the expected trajectories
436 from a real PEPT experiment are of good enough resolution to track the real granular mo-
437 tion, and thus would be sufficient to calibrate a DEM model. Ultimately, this work provides
438 a way to estimate how long an experiment should be run to reach a desired level of accu-
439 racy and also provides a framework for correcting PEPT measurements of behaviour like the
440 granular temperature.

Chapter Six

Optimising PEPT Algorithm Parameters

6.1 Introduction

Many PEPT algorithms discussed and evaluated in Chapter 3 have user-prescribed parameters which are used to improve their accuracy [100, 138, 94, 91]. For example, the PEPT-ML algorithm requires users to prescribe the maximum distance between LoRs which could form a cutpoint for clustering [91]. Additionally, the feature point identification (FPI) PEPT algorithm requires users to prescribe a voxel threshold to filter out voxels which contain few crossing of LoRs so that tracers can be identified from the background noise [138]. While these algorithms often come with default settings, to get the best results the parameters should be tuned [149]. However, determining the optimal values is subjective to the user because there are not methods of comparing the PEPT-reconstruction to the known tracer behaviour in real systems [149].

This Chapter seeks to provide a method of establishing the optimal values for free parameters in PEPT algorithms by creating GATE simulations of PEPT experiments with prescribed tracer motions. Since the trajectory of the tracer is prescribed, this provides a quantitative method of comparing the accuracy of PEPT reconstruction when certain

parameter values are used [47]. While similar methods of parameter optimisation could be demonstrated with any PEPT algorithm with free parameters, this Section will cover the optimisation of the Birmingham Method's N_{LoRs} and f parameters. The Birmingham Method is the most used PEPT algorithm and requires only these two parameters to function [100]. The N_{LoRs} parameter specifies how many LoRs should be included in a sample while the f parameter specifies the fraction of LoRs in the sample which should be used to calculate the tracer position.

The Birmingham Method (BM) PEPT algorithm was designed such that only a fraction of the LoRs in a sample should be used to find the tracer position and the most likely corrupted LoRs removed [100]. Corrupted LoRs may form when scattering in one or both of the detected rays and the formation of random LoRs is caused by two unrelated rays being detected within the coincidence window [125, 148]. By initially using all of the N_{LoRs} to estimate the tracer position, the LoR furthest from this point can then be removed and the remaining LoRs are recycled to update the position. This iterates until only a user-specified fraction of the LoRs remain, f. The parameter values used are typically conservative so they can be applied to many different systems and produce reasonable trajectories [100].

It is known that thicker and denser systems will cause more scattered LoRs and that more active tracers will cause more random LoRs [44]. However, little has been done to develop ways to understand how this affects PEPT measurements and, further, predict the values for f and N_{LoRs} which will maximise the spatial and temporal resolution of the Birmingham Method. Existing methods rely on using a static point-like tracer and measurements of the standard deviation of the tracer position, rotating disk studies where the tracer positions can be determined based on the rotation rate and radius of the disk, or by estimating the location error by indirect means [100, 18, 16]. However, these methods are oversimplified and not able to directly measure the spatial error of the tracer.

This Section seeks to build upon these methods by using both PEPT measurements and GATE simulations to investigate how the spatial accuracy of the Birmingham Method is affected by the parameters f and N_{LoRs} under a range of experimental conditions, including both stationary and moving tracers. The GATE model of the ADAC Forte is used to generate synthetic LoRs from recreations of real experiments [44, 45]. These serve as further validation for the GATE model and are later expanded to test a range of conditions not explicitly considered during the experiments. The values for f and N_{LoRs} which maximise the spatial accuracy for each PEPT trajectory are extracted as a function of the tracer position. To utilise this information, a new version of the Birmingham Method, called the Dynamic-BM (DBM), is developed which dynamically changes the fraction of LoRs remaining in the sample and the sample size of LoRs based on an estimate of the tracer position. This new method is expected to provide both higher spatial and temporal resolution than the original Birmingham Method algorithm since it can use the optimal parameter combination on a sample-by-sample basis.

6.2 Methods

6.2.1 Static Tracer Experiments

The tracer chosen for this experiment is a 1 mm diameter sphere of anion exchange resin labelled with fluorine-18. Fluorine-18 is an ideal positron-emitting isotope for PEPT because it has one of the lowest energy spectra for positrons and thus a low range [57]. The anion exchange resin adsorbs fluorine-18 ions from a solution of radioactive water produced on-site at the Positron Imaging Centre [96]. On the day of the experiment, the tracer was activated with an initial activity of 2.8 MBq and placed inside a 0.5 ml plastic vial for handling. According to a recent characterisation of the ADAC Forte, a tracer of 2.8 MBq will produce

less than 5% random LoRs, meaning nearly all LoRs not intersecting the tracer location will most likely have undergone scattering before being detected [44].

To attenuate the 511 keV gamma-rays, the vial is placed in the centre of an 800 ml cylindrical glass beaker filled with 500 ml of bulk density attenuating material. The inner diameter of the beaker is 100 mm and filled to the height of 65 mm, confirming that a volume of material of approximately 500 ml is used. The wall thickness of the beaker is 5 mm. Once filled, the beaker mass was measured. The initial mass of the beaker was subtracted to find the total mass of the attenuating material. The density of the materials, ρ, is calculated by dividing the mass by the volume. The materials and their properties are listed in Table 6.1. Materials like air and high-density polyethene (HDPE) have low linear attenuation coefficients, μ, while steel and copper have larger coefficients, meaning they will attenuate a larger fraction of the gamma rays [129]. A small sample of some of these materials and the filled beaker is shown in Figure 6.1. The fact that materials are bulk and porous is not expected to have an effect in this experiment since the diameters of the materials are much smaller than the diameter of the beaker. The attenuation of gamma rays in this experiment will approximate those of a continuous attenuation medium.

During the experiment, the beaker is placed in the centre of the field-of-view (FOV) of the ADAC Forte [103, 44]. The two heads of the Forte are set to their maximum separation of 800 mm to achieve the most uniform illumination. Each material is imaged until more than 1,000,000 LoRs have been collected, which is enough to locate the tracer several times using the Birmingham Method across a range of N_{LoRs} in the sample size. Since the position of the tracer is not known exactly, the standard deviation of the PEPT detected position, the PEPT uncertainty, σ, is used to quantify the performance of different f and N_{LoRs} parameter combinations. This is calculated using Equation 1.16, which is the three-dimensional standard deviation of the detected positions.

Table 6.1: List of the bulk materials and their attenuating properties [129, 47].

Material	ρ_{bulk} (g/cm^3)	μ/ρ_{bulk} (cm^2/g)	μ (cm^{-1})	**Attenuation (%)**
Air	0.00129	0.0806	0.000111	19.4
HDPE	0.890	0.0931	0.0828	64.8
MCC	1.421	0.0915	0.130	78.0
Glass	1.661	0.0858	0.143	80.7
Steel	4.425	0.0832	0.368	98.0
Copper	5.025	0.0827	0.415	98.7

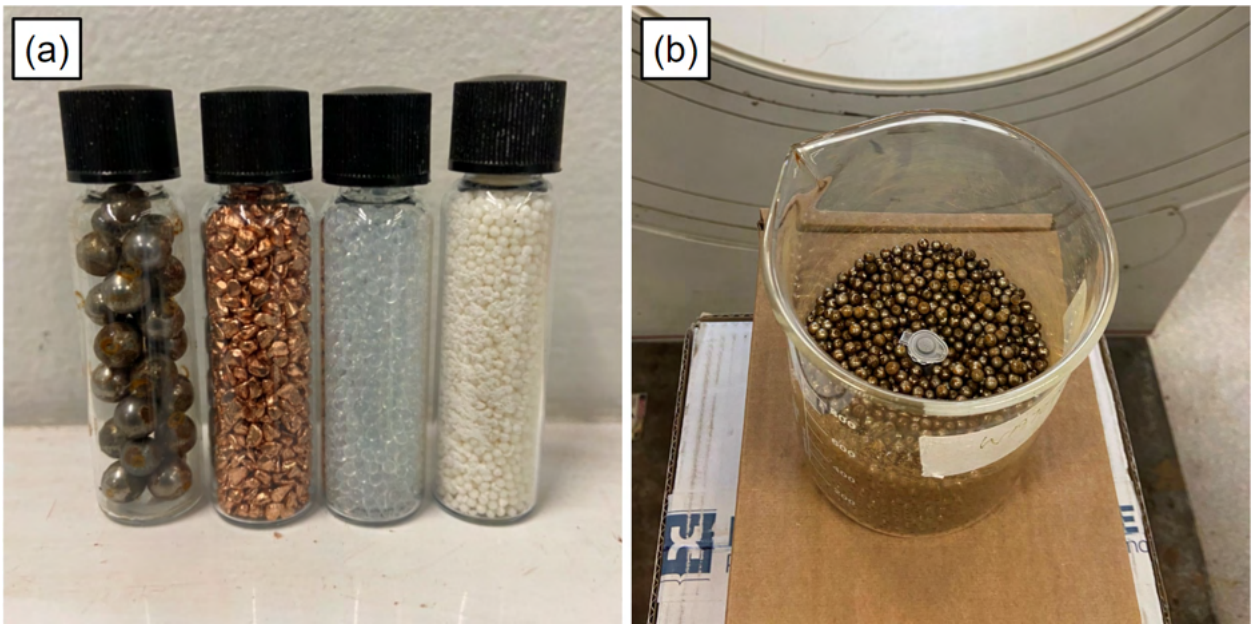

Figure 6.1: (a) A small sample of four of the materials used to induce attenuation. From left to right: steel, copper, glass, and MCC. (b) The beaker is filled with steel balls and the source is placed in the centre of the field-of-view (FOV) of the Forte [47].

6.2.2 Moving Tracer Experiment

In previous work, the spatial accuracy for static tracers has been shown to improve when more LoRs are used per sample [149]. However, in real PEPT experiments, the tracer typically moves, limiting the N_{LoRs} per sample which can be used. This is because if the tracer moves more than a few millimetres between the detection of the first and last LoRs in the sample the PEPT-detected position will incur higher spatial errors [100, 46]. Thus there should always exist an optimum N_{LoRs} per sample for moving tracers.

To investigate this in a system representative of a PEPT experiment, an attritor mill is placed near the centre of FOV and a tracer is fixed to the end of the impeller as it rotates in the mill at 100 RPM. Mills of this type are used across a wide range of industrial and pharmaceutical applications and have been studied in the past using PEPT [153, 26]. Due to the thick steel walls and predictable circular rotation of a particle fixed to the impeller, this system is ideal to investigate the effect of N_{LoRs} on the spatial accuracy of the Birmingham Method. In other work, similar rotating systems have been used to better understand how changing the parameters f and N_{LoRs} affects PEPT measurements [100, 18]. The attritor mill and a schematic of the mill dimensions are shown in Figure 6.2.

The tracer used for this experiment is a 1.2 mm diameter microcrystalline cellulose (MCC) bead activated with a solution of fluorine-18 and water to an initial activity of 22 MBq. A tracer of this activity is ideal for PEPT experiments in the Forte since this is approximately the activity which will produce the highest true LoR count-rate before dead-time and random LoRs degrade the measurement [44]. This tracer is taped to the end of the upper impeller in the attritor mill at a radius of 63 mm. At 100 RPM the tracer will rotate at a constant velocity of 660 mm/s. The mill is then placed near the centre of the FOV of the ADAC Forte at a head separation of 510 mm. The mill is imaged over approximately 1 minute (i.e. 100 rotations) to develop good statistics of the tracer locations as it rotates.

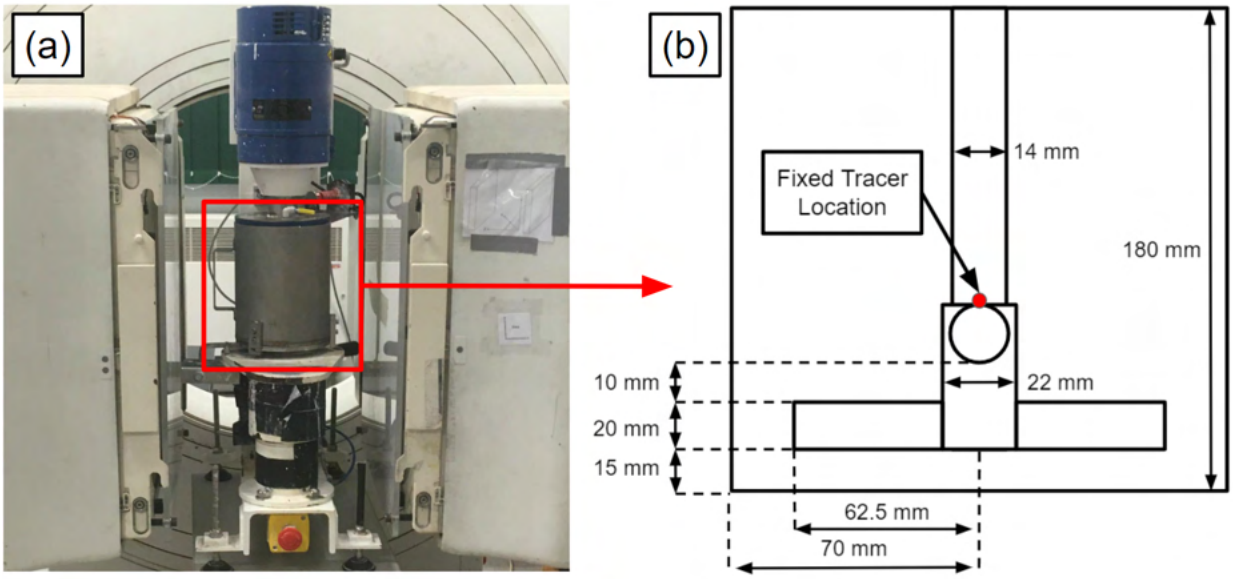

Figure 6.2: (a) The attritor mill near the centre of the FOV of the ADAC Forte. (b) A schematic of the dimensions of the mill and the tracer fixed to the impeller [47].

To assess the performance of different f and N_{LoRs} combinations using the Birmingham Method, first, each directional component of the PEPT trajectory is fitted to a sinusoidal equation as a function of time, t, as shown in Equation 6.1. The amplitude, A, in the x and z components should be approximately 63 mm and in the y component, 0 if the mill is perfectly oriented with respect to the detector axes. Since the mill is likely not perfectly level a sinusoidal equation is fitted to the y-component as well. The rotation rate, ω, is approximately 100 RPM, and ϕ is the phase shift which depends on the initial position of the tracer. The PEPT deviation from this motion is calculated by comparing each PEPT detected position to the predicted position using the fitted equations. In this way, the mean spatial error of the trajectory is calculated using Equation 1.14.

$$x_{Fit}(t) = A\sin(\omega t + \phi) + c \tag{6.1}$$

6.2.3 GATE Simulations

The GATE model of the ADAC Forte used in this work has been validated using a characterisation of the detector when it was installed at the PIC and calibration experiments following an industry-standard protocol, showing agreement between simulation and experiment to within 10% across all tested metrics [44]. This GATE model of the ADAC Forte is first used to repeat the experiments described in Sections 6.2.1 and 6.2.2 for comparison and further validation. After this, the GATE model and recreated attritor mill are used to expand the tested conditions to observe how the optimal combination of f and N_{LoRs} changes throughout the system. These optimal parameters are used in Section 6.2.4 to develop the DBM algorithm.

Following the experiments, the tracer, geometry, and detector for the static tracer experiments are replicated in GATE. This starts by creating a radioactive tracer defined as a 1 mm diameter resin sphere, emitting positrons with a fluorine-18 energy spectrum [88]. The tracer is placed inside a geometric model of the glass beaker which has been filled with 500 ml of bulk-density material. A cross-section of the beaker and source geometry is shown in Figure 6.3. To replicate the attenuation of each material, new material definitions are added to the GATE material database which is described in Table 6.1. The tracer is prescribed an activity equal to that of the tracer activity at the beginning of each data acquisition. Initially, the tracer was approximately 2.8 MBq, but by the end of all experiments decayed to approximately 1.3 MBq. For each simulation, the activity is adjusted accordingly to compensate for decay. Finally, the tracer and beaker are placed in the centre of the FOV of the GATE model of the ADAC Forte. The detector separation is set to 800 mm. While the exact positions of the tracers simulated in GATE are known, Equation 1.16 is still used to calculate the PEPT uncertainty to compare the real experiment and simulations directly.

Similar to before, the moving tracer experiment is recreated using the GATE model

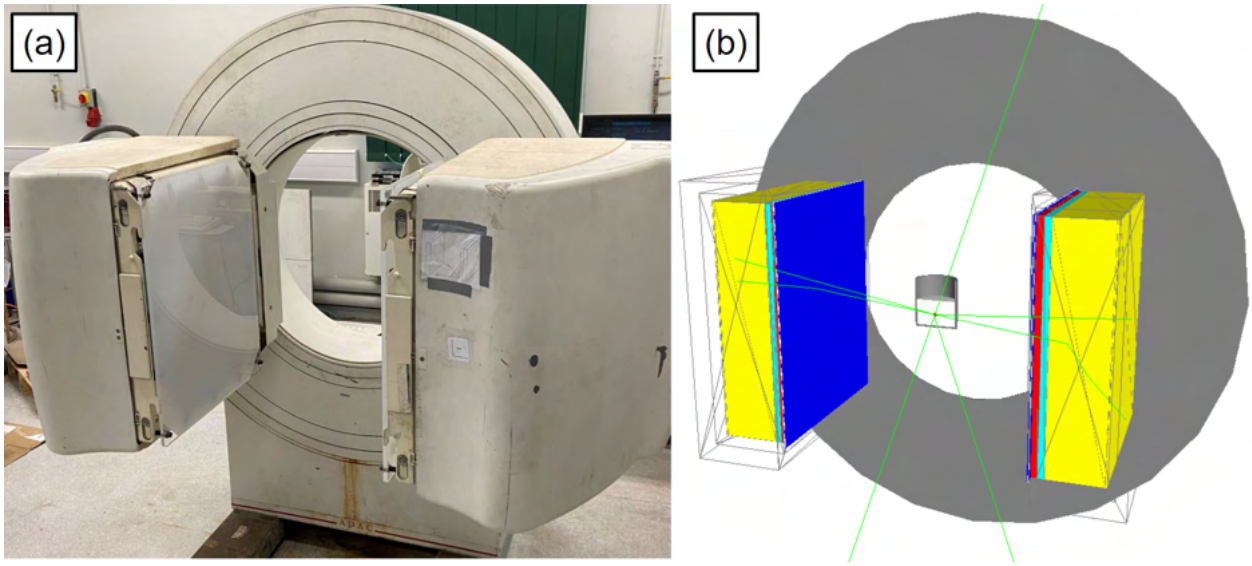

Figure 6.3: (a) The ADAC Forte dual-headed positron camera at the Positron Imaging Centre (left). GATE model of the Forte with a cutaway of the experimental geometry and source in the centre of the FOV (right) [47].

of the ADAC Forte. The head separation is set to 510 mm, which is as close to the system as the detectors could be set. A lower head separation results in higher geometric efficiency, meaning LoRs can be collected at a higher rate. At the time of the experiment, the tracer was measured to approximately 20 MBq. This is modelled in GATE as a 1.2 mm diameter sphere emitting positrons with a fluorine-18 energy spectrum.

The mill dimensions can be found on the right-hand side of Figure 6.2 and are recreated in the GATE simulation as shown in Figure 6.4. It consists of a 70 mm diameter vessel lined with polyethylene and a stainless-steel impeller with 4 pins. The polyethylene liner is 10 mm thick and the outer stainless-steel wall is 5 mm thick. The inside of the vessel is filled with air to match the experimental conditions.

The whole mill is rotated about the y-axis at 100 RPM to induce tracer motion. This is achieved through discrete rotations of every simulation time-step of 0.0005 s. For the tracer fixed at a 63 mm radius, this results in a change in position of approximately 0.33 mm

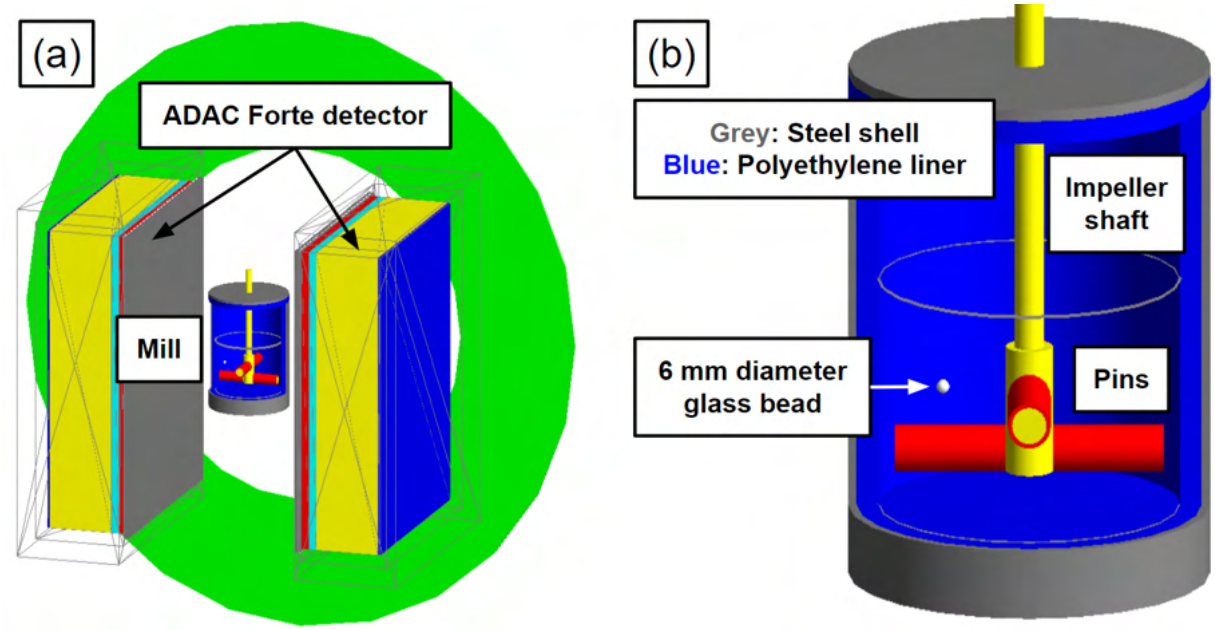

Figure 6.4: (a) The GATE model of the ADAC Forte PEPT detector system. (b) A cut-away of the attritor mill. In the cut-away, the major components are labelled [47].

per time-step. This is smaller than the tracer diameter and can thus be safely used to mimic continuous tracer motion [46]. The GATE simulation produces LoRs which are processed in the same way as real PEPT data. A range of different f and N_{LoRs} values are used to locate the tracer. The PEPT trajectory is then compared back to the GATE-prescribed tracer positions to calculate an error based on Equation 1.14. The GATE simulated tracer position is known exactly and there is no variability in its rotation rate or system vibration present. As such, the spatial accuracy of the GATE simulated tracer is expected to be somewhat higher than in the real experiment.

6.2.4 Dynamic Birmingham Method

In PEPT experiments, the scattering environment and detector sensitivity change as a function of the tracer position. The amount of corrupted LoRs in a sample affects the optimal

f while the sensitivity affects the optimal N_{LoRs}. However, the Birmingham Method uses constant parameter values. This means users must choose a conservative parameter combination that will work over the whole data set. Inevitably, this will return trajectories with lower spatial accuracy than is theoretically possible. To solve this problem, a variant of the Birmingham Method is developed called Dynamic-BM (DBM), 'Dynamic' because it can dynamically change f and N_{LoRs} to the optimal values determined by a GATE simulation as the tracer moves through the system.

Initially, the DBM uses constant values of f and N_{LoRs} to estimate the position of the tracer. Then, for each initially-estimated tracer position, the optimal parameter values are looked up from a table produced by Monte Carlo simulation. To use this table, the predicted 3-dimensional location of the tracer and other optional information is input, then the closest simulated position is found and the optimal parameters are output. This approach relies on simulating a sufficiently high number of tracer positions within the system so that the change in parameter values between adjacent simulated positions is smooth.

To test the DBM, the attritor mill is simulated using GATE with the tracer placed in a range of initial positions. To match the experimental conditions, the mill is rotated at 100 RPM and the tracer activity is set to 22 MBq. Each tracer position is a new GATE simulation. The initial tracer positions are seeded in one quadrant of the mill to take advantage of the symmetry of the system. The positions are created in regular intervals in the x and z direction from 0-60 mm in 20 mm steps and the y-direction from $\pm$50 mm in 12.5 mm steps. The positions falling outside of the system or intersecting the impeller are removed, leaving a total of 84 positions. Since the impeller is rotating, the position of the tracer is important for the optimal f and N_{LoRs} values since it will cause different amounts of scattering depending on its rotated angle. In the simulation, the rotation of the impeller is prescribed so it is known, but during a PEPT experiment, this can be more difficult to ascertain. However, the ADAC Forte can record readings from an optical switch directly into

the data file, allowing the impeller rotation angle to be recorded throughout an experiment [99]. The 360° is divided into 30° increments and the optimal f and N_{LoRs} are calculated over this range.

To use the look-up tables, first, the amount the mill is rotated must be determined. This is done by multiplying the time by 360° and dividing by the rotation period, 0.6 s. Since all rotations beyond 360° are duplicates of previous rotations, a modulo operation is performed. Then all the simulated positions within the nearest rotation are compiled and the values of f or N_{LoRs} for the position closest to the PEPT estimated position are extracted. Using this method, the optimal f and N_{LoRs} of any point within the system can be determined. Due to the change in N_{LoRs} per sample, the number of detected positions in a trajectory processed using the Birmingham Method and the DBM may vary. To ensure the changing parameters are used at the appropriate position within the data set, the f and N_{LoRs} are linearly interpolated as a function of time and the average time of the LoRs in each sample is used to compute the exact parameter values.

To assess the performance of the DBM, it is compared to both the default Birmingham Method parameter values of 0.05 f and 250 N_{LoRs} as well as the best constant parameter values. This represents the default and the best possible Birmingham Method performance, showing how dynamically changing the parameter values can improve PEPT algorithms.

6.3 Results

6.3.1 Model Validation

Each of the static tracer experiments described in Sections 6.2.1 and 6.2.2 produces a unique fraction of attenuation. Some of the coincident 511 keV gamma-rays attenuated via Compton

scattering are inadvertently passed to the PEPT algorithm as corrupted LoRs. When the attenuating medium is air, the amount of corrupted LoRs is relatively low, but as the medium becomes more attenuating, a larger fraction of LoRs are corrupted. This is shown in Figure 6.5 where the LoRs from the air, glass, and copper experiments are plotted.

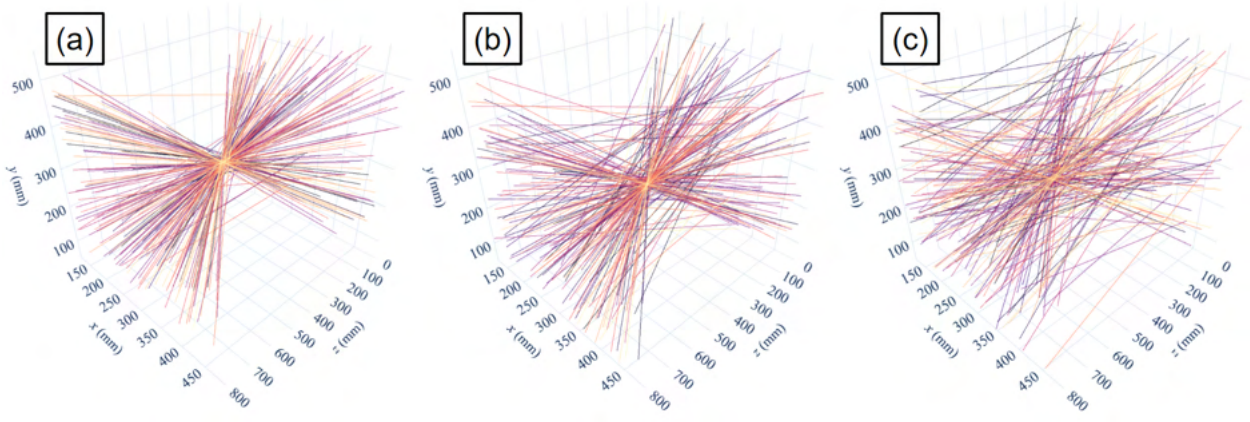

Figure 6.5: Three samples of 200 experimentally acquired LoRs from the static tracer experiment for (a) air, (b) glass, and (c) copper show that the amount of corrupted LoRs in the sample increases with more attenuation [47].

Since these materials produce different fractions of corrupted LoRs, the optimal f for each experiment should be a unique value. The LoRs from the real experiment and simulations were both processed using the Birmingham Method under a range of f and N_{LoRs} and the variation of the standard deviation in position for different combinations of the two parameters is plotted as colour variation in Figure 6.6. These plots show that for static tracers the optimal value of f is decreased when more corrupted LoRs are present and the standard deviation in the position decreases as more N_{LoRs} per sample are used. The experiment and GATE simulations closely agree across all the parameter combinations, both in the optimal values for f and in the standard deviation in the tracer position. For a static tracer, the optimal value for f remains constant and the standard deviation in position will always decrease with greater N_{LoRs}. The optimal values for f across all the materials tested

for both the experiment and the simulation are shown in Figure 6.7. Additionally, in this figure, the values for f which minimise the uncertainty are plotted against the fraction of true LoRs demonstrating that f must decrease when more attenuation occurs. Moreover, this shows that the values for f which minimise position uncertainty are approximately equivalent to the fraction of true LoRs in the sample. This, until now, has been an assumption of the Birmingham Method, but this work now provides the first direct evidence that this assumption is true. The experiment and simulation provide similar values in both the overall PEPT uncertainty across all the parameter combinations tested and also the values of f which minimise the uncertainty for a given N_{LoRs}.

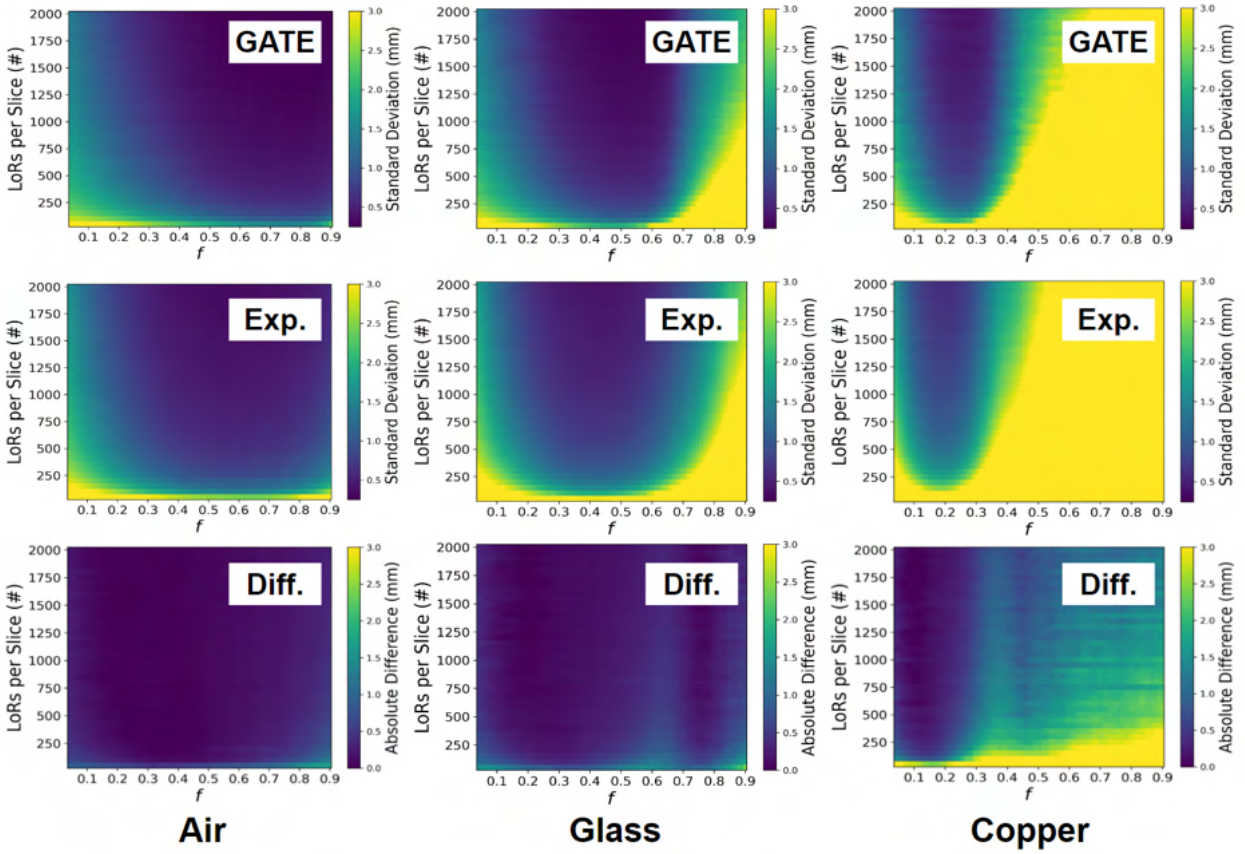

Figure 6.6: Position uncertainty in the static tracer experiments and simulations for air (left), glass (middle), and copper (right), and the absolute difference between the simulation and experiment [47].

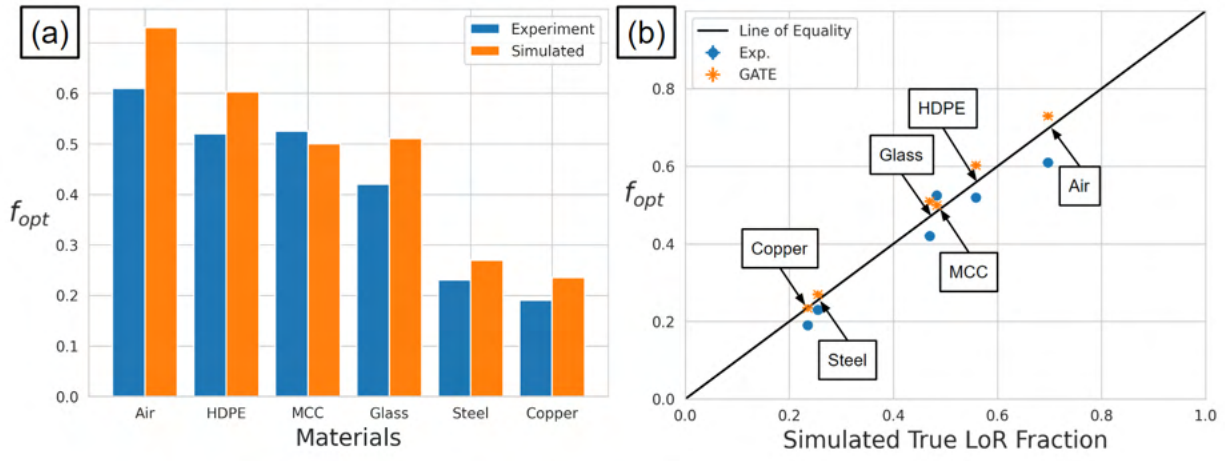

Figure 6.7: (a) A comparison of the optimal values for f across all the materials tested. (b) A comparison of the optimal values for f across all the materials as a function of the fraction of true LoRs in the sample [47].

A similar analysis was conducted for the moving tracer experiment. A 22 MBq tracer was fixed to the impeller of an attritor mill and rotated at 100 RPM. In this case, a position error was calculated using Equation 6.1 for the experiment and simulation. A range of constant f and N_{LoRs} values are used with the Birmingham Method to assess the spatial errors produced under different parameter combinations. Trajectories of the experiment and simulation are presented in Figure 6.8 showing that the rates of detection are approximately the same. In Figure 6.9, the parameter values for f and N_{LoRs} are varied to assess their effect on the spatial error of the reconstructed PEPT trajectory and find the optimal combination. The experimental plot has a generally higher error because the error was calculated using fitted functions, assuming perfect circular motion, whereas the GATE simulations benefit from having analytical functions describing the tracer motion. From Figure 6.9, the optimal values for f and N_{LoRs} for the real experiment are determined to be 0.275 and 1400, respectively. Similarly, the optimal parameters for the simulation are determined to be 0.25 and 1300, respectively. The optimal parameters and their percent errors are shown in Table 6.2.

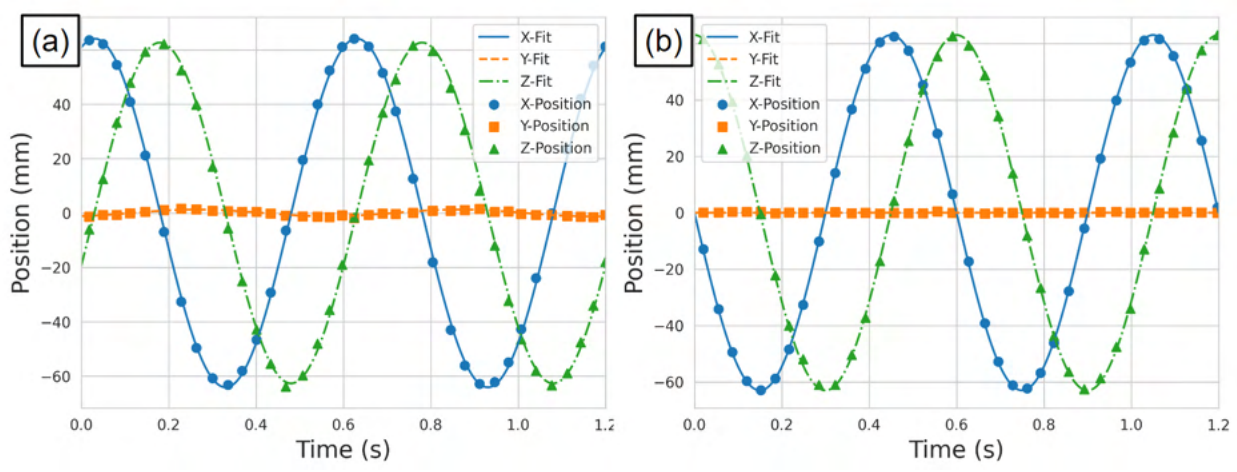

Figure 6.8: Trajectories for the (a) experimental tracer in the attritor mill under a constant f of 0.25 and N_{LoRs} of 1500 compared to the (b) simulated tracer trajectory under the same parameters [47].

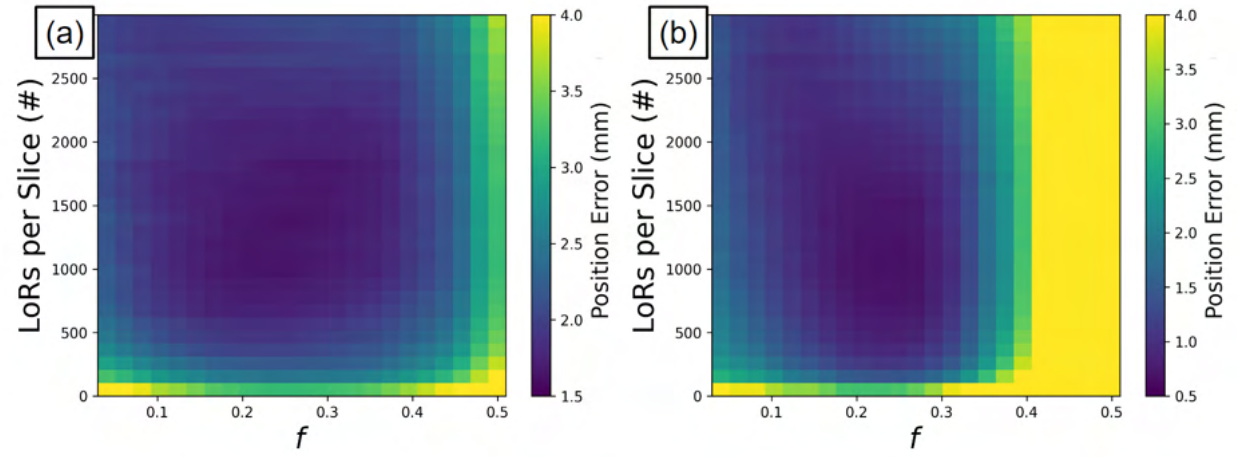

Figure 6.9: The parameter values for f and N_{LoRs} are varied for the (a) real experiment and (b) GATE simulation of the moving source, affecting the spatial error of the reconstructed PEPT trajectory [47].

6.3.2 Dynamic Parameter Optimisation

For each of the tested positions in the expanded GATE simulation of the attritor mill, the optimal values for f and N_{LoRs} were recorded as the tracer moves through the system. These

Table 6.2: Comparisons of the experiment and simulation in the optimal parameters for the moving tracer experiment [47].

Method	**Optimal** f	**Optimal** N_{LoRs}
Experiment	0.275	1400
Simulation	0.25	1300
Percent Error (%)	-9.09	-7.14

optimal parameters were calculated for every 30° rotation to update their values continuously. Figure 6.10 shows the change of optimal parameters at two different degrees of rotation, one where the least amount of attenuation occurs and another where the tracer is behind the impeller shaft, with respect to the detectors, where the most attenuation occurs. As the mill rotates, the optimal parameters fluctuate, becoming more or less conservative when more corrupted events are recorded or the count-rate decreases, such as when the tracer passes in front or behind the impeller shaft.

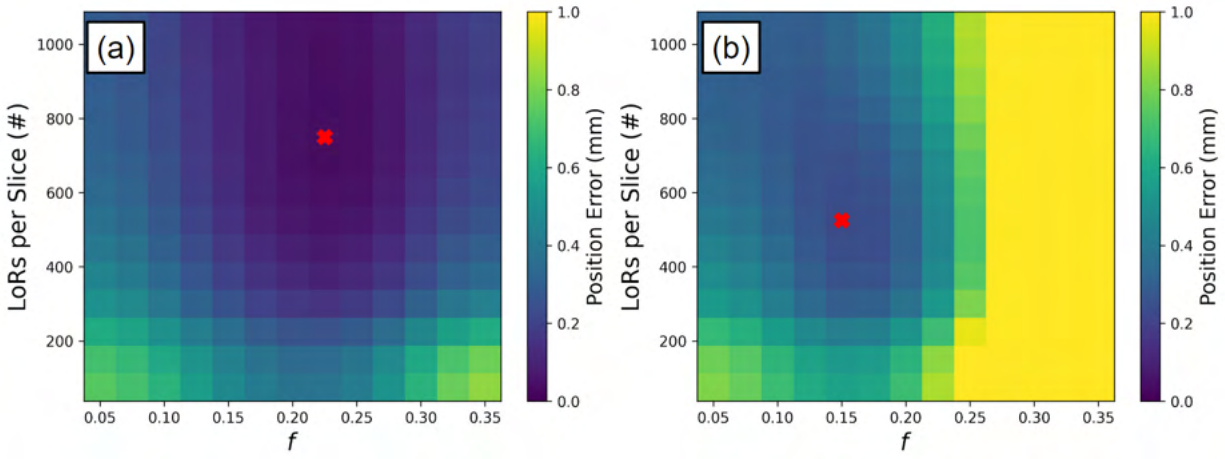

Figure 6.10: (a) The results of a parameter sweep when the tracer is beside the impeller shaft where the least amount of LoRs have been attenuated. (b) The effects of the tracer passing behind the impeller shaft where a larger fraction of LoRs are attenuated [47].

After the optimal parameters have been extracted, these are used to inform the DBM

algorithm. Each simulated trajectory is reprocessed with an estimate of the optimal parameters and the mean spatial error is calculated over the trajectory. These are compared to the trajectories extracted using the default and the best constant parameters. The results of this comparison are in Table 6.3. The mean spatial error of the default Birmingham Method parameters, best constant parameters, and the DBM are 2.20 mm, 0.544 mm, and 0.517 mm, respectively. While the default parameters can reconstruct the trajectory, when tailored parameters extracted from Monte Carlo simulation are used, in this case at least, the errors decrease to nearly a quarter of their original values. When dynamic parameters are used, the errors decrease by 76.5% over the default parameters and by 4.03% over the best constant parameters. A histogram of the percent changes in spatial error between the trajectories produced with the best constant parameter and the DBM is shown in Figure 6.11. This plot shows the DBM increases the accuracy of nearly all trajectories and also some individual trajectories are improved by over 10%. The trajectories which are the most improved are from areas of the system around the impeller blades where the local optimal parameters deviate the most from the best constant parameters.

Table 6.3: Mean trajectory comparisons using different algorithm parameter methods [47].

Parameters	**Spatial Error (mm)**	**STD (mm)**	**Locations (N)**
Default	2.20	1.48	1337
Best Constant	0.544	0.372	223
Dynamic	0.517	0.3481	236

6.4 Discussion

This Chapter demonstrates that realistic synthetic data produced by GATE simulations can be used to optimise PEPT algorithm parameters, representing a departure from previous

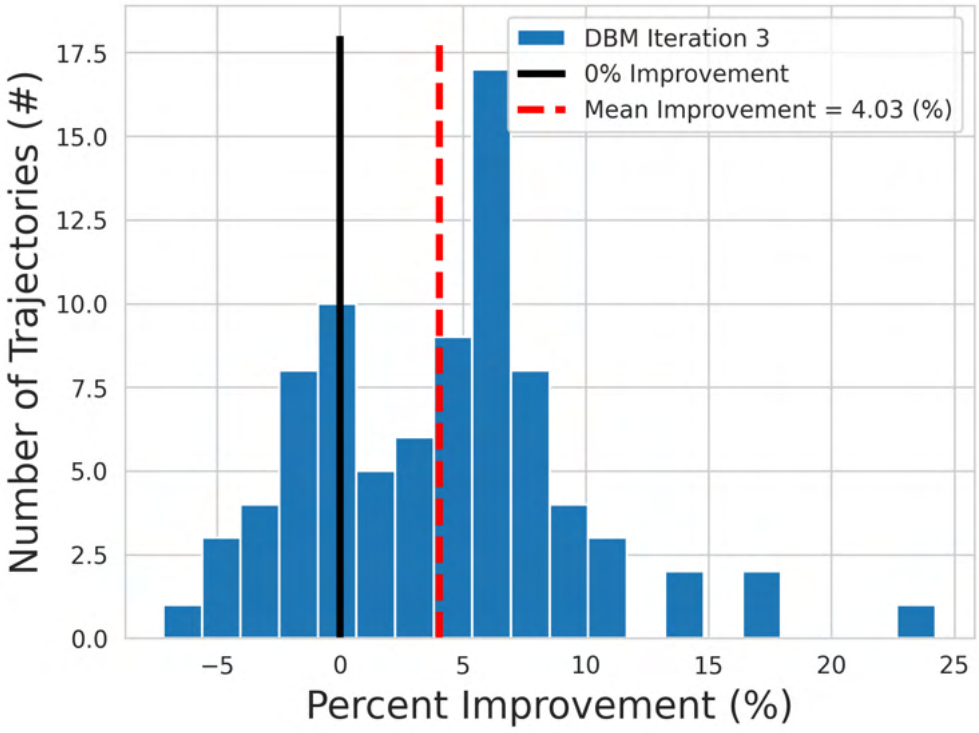

Figure 6.11: The change in spatial accuracy for each trajectory. The mean improvement is 4.03% over the best constant parameters [47].

approaches to algorithm optimisation. Optimisation of PEPT algorithms has previously been limited in use and often impractical in terms of the time and resources needed to characterise the PEPT detector system response to a tracer in an experimental system and quantify how algorithm parameters affect the quality of reconstructed trajectories. In particular, using simulated data has been shown to produce the same optimal parameter values as carefully controlled real experiments and to be useful in identifying the optimal PEPT algorithm parameters which lead to the best reconstruction of simulated tracer trajectories in conditions which could not be directly tested.

The optimisation of the PEPT algorithm parameters was explored by using a validated GATE model of the ADAC Forte PEPT detector system and a model of an experimental system which would be studied using PEPT. Previously, the use of static tracers or predictably

moving tracers to find the optimal parameters values for PEPT algorithms has been used in other work [18, 16, 100]. However, these methods lack the ability to compare the actual location of the tracer to its PEPT reconstruction, instead relying on indirect measurements such as the standard deviation of the positions. As such, they are susceptible to bias. By using GATE simulations, the prescribed position of the tracer can be compared to its PEPT reconstructed location.

Additionally, since GATE simulations can be rapidly changed, a much wider variety of conditions can be tested than would be possible in physical experimentation. In real experiments, typically only one location can be tested at a single activity [100]. However, since the LoR count-rate and the fraction of corrupted LoRs in a sample are sensitive to the tracer position, this approach can only provide a rough estimate of the optimal parameters [44]. By using GATE simulations in which the location and activity can be changed rapidly, the complex three-dimensional nature of the optimal parameters can be discovered and also the state of the experimental system taken into account to dynamically update the PEPT algorithm parameters on a sample-by-sample basis [47]. Doing this achieves the optimal spatial and temporal resolution of a PEPT trajectory. Moreover, using the methods presented here, GATE simulations conducted in conjunction with PEPT experiments can be used to remove the guesswork of choosing f and N_{LoRs}. This means users of PEPT algorithms do not need to be experts to produce good trajectories, making the Birmingham Method PEPT algorithm more rigorous and more accessible.

This work is the first to develop these techniques and apply them to real and hypothetical PEPT experiments. The results show that these methods of optimisation are worthwhile and would provide valuable improvement in PEPT experiments if applied. Due to the time and resources needed for physical experimentation, optimisations of experiments have only seen limited use. However, the use of realistic GATE simulations is now a viable option when a validated GATE model of the PEPT detector system and recreation of the

experimental geometry is used [47]. When combined with a DEM tracer trajectory, these simulations become even more valuable as a stand-in for real data [46]. While this work was conducted using the Birmingham Method PEPT algorithm and the ADAC Forte detector system, a similar workflow could be applied to other PEPT algorithms which have free parameters, such as the Feature Point Identification PEPT algorithm, and other detector systems for which a validated GATE model exists, such as the Siemens Inveon [138, 74].

Chapter Seven

Measuring Tracer Size with PEPT

7.1 Introduction

In the previous Chapters, the development of GATE models of PEPT detector systems, applying these models to benchmark PEPT algorithm performance, and using simulations of PEPT experiments to optimise data acquisition and trajectory reconstruction has been discussed. However, simulating the response of a PEPT detector system is not only useful for improving the existing ways of using the PEPT technique, but also for exploring new uses. As will be shown in this Chapter and next, the ability to create a model of an experiment where all aspects are controlled and precisely known creates opportunities to improve PEPT such that different types of information can be extracted and, additionally, to apply PEPT to new types of experiments. This Chapter will cover a novel PEPT methodology to infer the size of a volumetrically activated tracer as it decreases in size and loses some of its activity to the system in which it is contained.

The advantage of using simulations rather than real experiments to develop new uses of PEPT is two-fold: (1) the conditions in the simulation are known absolutely, thus providing a concrete way to compare the results of the PEPT reconstruction to the prescribed values,

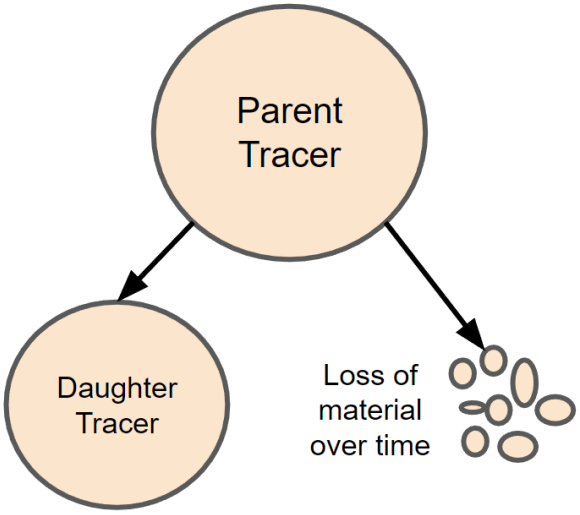

Figure 7.1: A volumetrically activated PEPT tracer which loses material during an experiment through attrition, dissolution, or other mechanisms produces a smaller daughter tracer and the lost material disperses into the system. Both the daughter tracer and the lost material may contain positron-emitting material.

and (2) the limited time and resources of PEPT facilities like the University of Birmingham's Positron Imaging Centre (PIC) or hospitals do not need to be expended on experiments for which the value of the results are uncertain or difficult to measure.

In PEPT experiments, often the experimental system is opaque, dense, or fast-moving which means that tracer size cannot typically be obtained through optical measurements [148]. Moreover, during a real experiment, it is not necessarily impossible to stop the experiment, physically extract the tracer, measure its size, reintroduce it to the system, and then continue the experiment, but carries with it more complexity and the risk of disturbing the measurement. Handling radioactive material also carries with it a radiation safety risk to personnel. Additionally, it is beneficial to provide a proof-of-concept before applying a new methodology to experimental data, thus verifying that it works as intended. Therefore, using simulated results showing that a methodology produces the expected results demonstrates

that it can be a valuable new tool. In the next Sections, why these new methodologies are useful extensions of the PEPT technique, the development of the methodologies using GATE simulations, and the results of proof-of-concept studies will be presented. For the purposes of this Chapter, the loss of material from a particle will be referred to as particle attrition, although there are several distinct mechanisms at play in different processes. Attrition specifically refers to when edges of particles suffer abrasion or when a weak plane of a particle is stressed to fracture [12]. However, the dissolution of soluble particles could similarly be studied using the technique presented in this Chapter.

7.2 Background

In many areas of engineering and physical sciences, the size of particles and the particle size distribution (PSD) play a role in many aspects of a process, such as understanding the physical and chemical properties of a material, reaction rates of solids, and flow behaviour [130]. Changes in particle size over time can be caused by attrition, fracture, dissolution or loss of material through other mechanisms such as a phase change [130]. One area where particle sizes are especially important is in the mineral and chemical processing industries [22]. Over-processing particles during mixing in solid-liquid agitated vessels can lead to particle damage, diminished product quality, and expending more energy than needed in a process, which increases costs [110]. Monitoring the changes in PSD is typically a manual process involving sample collection for offline analysis [6]. In addition to offline analysis real-time process monitoring of PSD using optical or acoustic methods is increasingly used, which provide the ability to make rapid changes to process control [80, 52].

However, these methods do not provide insight into where or how the loss of particle material occurs. To fill in these gaps, discrete element method (DEM) simulation can be

used to provide more insight into the behaviour of particles over time and why material is lost [33, 34]. This insight is gained because a single particle and its loss of material over time can be observed in simulation. While the attrition rates calculated from DEM simulations and real experiments can be compared as a method of validating the results, there does not currently exist a way to directly validate the attrition experienced by individual particles, limiting its usefulness [40]. To fully understand the causes of attrition, a method of direct observation is needed.

One method of tracking the trajectory of a single particle through an opaque system is through PEPT [148]. If this method could be extended to also infer the size of tracer particles over time, then this information could be used to better understand attrition rates and, crucially, determine where attrition occurs within a system. This Chapter seeks to extend the Birmingham Method PEPT algorithm such that a tracer size can be measured while it is being tracked within a system [100]. For a detailed description of the Birmingham Method, refer to Section 3.2.1. A parameter of the Birmingham Method, f, which is used to discard LoRs which do not emanate from the tracer, is studied as a surrogate for measuring the amount of activity coming from the tracer compared to the amount of activity which has been lost to the background. By knowing the initial conditions of the tracer, observing the activity loss, and assuming the tracer shape, its size can be estimated.

In the next Sections, this method of measuring tracer size through the Birmingham Method f parameter is investigated through GATE simulations. The GATE simulations are designed such that a volumetrically activated tracer of some initial size and activity is placed in the centre of a cylinder of water, then, over time, the tracer is decreased in size, its activity decreased proportional to the volume of material lost, and the lost activity uniformly transferred to the water cylinder. This essentially creates background noise which emulates the loss of tracer material to its surroundings through attrition.

7.3 Methods

7.3.1 Extending the Birmingham Method to Measure Tracer Size

When positrons emitted from a PEPT tracer annihilate, they produce two coincident gamma rays near the tracer which can be recorded with a PEPT detector and used to triangulate the tracer position [100]. However, only gamma rays which reach the detector without undergoing scattering can be used, since the line-of-response (LoR) formed by scattered gamma rays no longer intersect the tracer location. When PEPT tracers that are volumetrically activated lose material through attrition, some of the activity of the tracer is lost and dispersed into the system [148]. In addition to scattered LoRs, LoRs which come from fragments of the attrited tracer cannot be used to locate the tracer since they also do not intersect the tracer location.

A useful feature of the Birmingham Method PEPT algorithm is the inclusion of a user-defined parameter that specifies the fraction of LoRs in a sample to discard such that they are not used in locating a tracer [100]. This is included because, in a sample of LoRs, some LoRs will have scattered before reaching the PEPT detectors, be formed from two unrelated events, or originate in the background. Ideally, this parameter should be set to a value corresponding to the fraction of LoRs that only come from the tracer. For static tracers, the value of f which minimises the position uncertainty of the tracer is typically termed f_{opt} and is approximately the fraction of LoRs in a sample which emanates from the tracer [47].

If a tracer loses some of its activity to the background, this would present as a decrease in f_{opt} since the fraction of corrupted events in a sample would increase proportionally to the amount of activity lost. By observing changes in the measured f_{opt} of the tracer, it is possible this could be used to infer changes in the tracer size due to the loss of material

through attrition or another mechanism such as dissolution. To do this, when the tracer is first introduced to the system and before it starts to lose material, the initial f_{opt}, termed f_{opt_0}, should be measured. This value corresponds to the fraction of LoRs which lead back to the tracer when all of the activity is contained within the tracer volume. As the tracer loses material, the measured f_{opt} will decrease as a result of more LoRs originating in the background. The ratio of the newly calculated f_{opt} compared to its initial f_{opt_0} then reveals the fraction of activity left on the tracer. To translate the fractional volume into a tracer size measurement, an assumption about the tracer shape must be made. For spherical tracers, the volume, V, of a tracer is proportional to the cube of the tracer radius, r, shown in Equation 7.1. Therefore, if a tracer were to decrease in radius by half, then the volume would decrease eight-fold. Therefore, using this relationship, the measured fractional volume of the tracer can be converted into a measurement of the tracer radius, which is calculated using Equation 7.2 and the initial tracer radius, r_0.

$$V = \frac{4}{3}\pi r^3 \tag{7.1}$$

$$r = r_0 \left(\frac{f_{opt}}{f_{opt_0}} \right)^{\frac{1}{3}} \tag{7.2}$$

While the uncertainty in position is one metric which can be used to find the f_{opt}, another that seemingly produces similar results, if not better, is to measure the average residual of the LoRs to the PEPT calculated position. The Birmingham Method locates tracers by minimising the distance between the estimated location of the tracer, called the minimum distance point (MDP), and the closest approach of each LoR to this point. For the purposes of this work, the sum total of distances between each LoR used to locate the tracer and the MDP is called the residual, R. This metric is calculated by dividing the sum of the distance from each LoR used to calculate the position of the tracer, d, by the amount

of LoRs used to locate the tracer, which is the sample size, N, multiplied by f. The equation used for calculating the average residual per LoRs is Equation 7.3.

$$R = \frac{1}{Nf} \sum_{n=1}^{Nf} d_n \tag{7.3}$$

When plotted against the position uncertainty in Figure 7.2, this provides similar and smoother results. This is especially clear when the derivative of these plots with respect to the change in f is shown in Figure 7.3. For the derivative of mean residual per LoRs, there is a maximum in the plot representing the value of f which produces the largest change in the residual. This can be interpreted as the value of f which results in using LoRs which are most likely corrupted because they are far away from the PEPT reconstructed position. Therefore, finding the maximum rate of change of the residual per LoR reveals the f_{opt}.

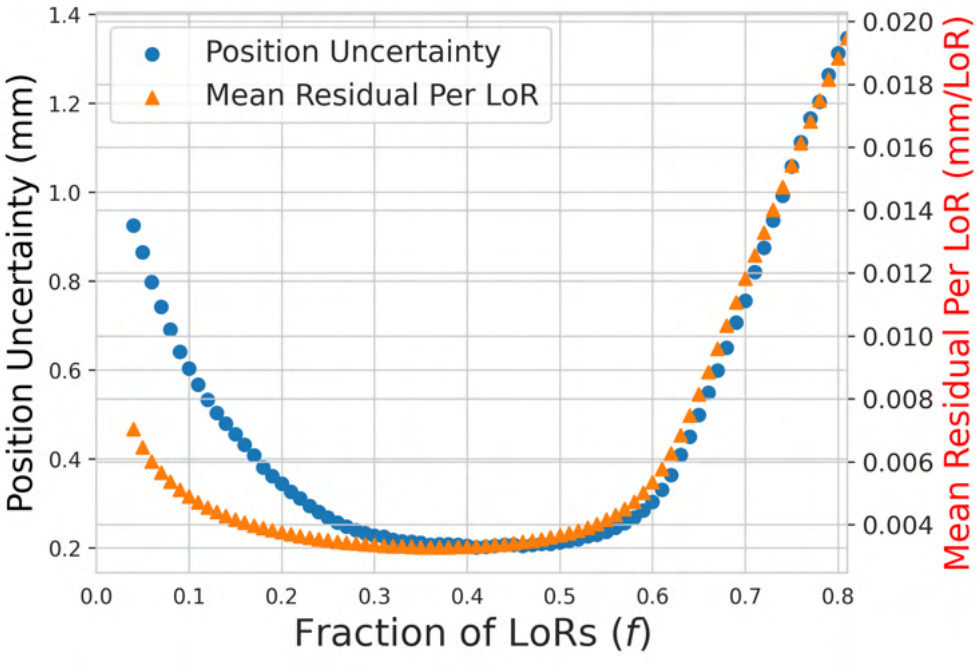

Figure 7.2: The position uncertainty and the mean residual of LoRs as a function of f using the Birmingham Method.

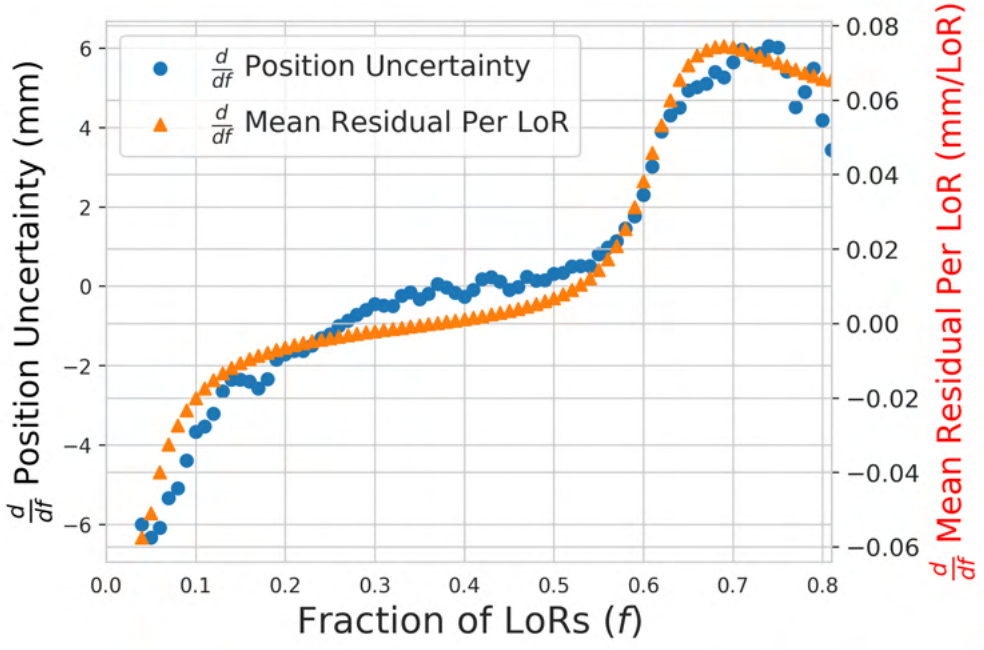

Figure 7.3: The rate of change in the position uncertainty and rate of change of in the mean residual of LoRs as a function of f using the Birmingham Method.

7.3.2 GATE Simulation

To test the method of determining tracer size by measuring changes in f_{opt}, a test data set is needed. This data could be collected from real experiments or produced from GATE simulations of the PEPT detector system and tracer. GATE simulations are used in this case since they provide careful control over the data through user-prescribed parameters. The GATE software v9.1 is used in this work for the Monte Carlo simulations [56]. To produce the data set, a GATE model of the ADAC Forte PEPT detector system is used [44, 45]. For full details of this PEPT detector system, please refer to Section 2.3.1. The ADAC Forte is comprised of two detector heads, each with a 16 mm thick wide-area sodium iodide scintillation crystal with an active area for detecting gamma rays measuring approximately 380 mm x 510 mm [103]. For the simulated experiment, the detector heads are separated by 500 mm and a cylinder of water, measuring 50 mm in radius and 200 mm in height, is

144 placed in the centre of FOV. A diagram of the simulated experiment is shown in Figure 7.4.

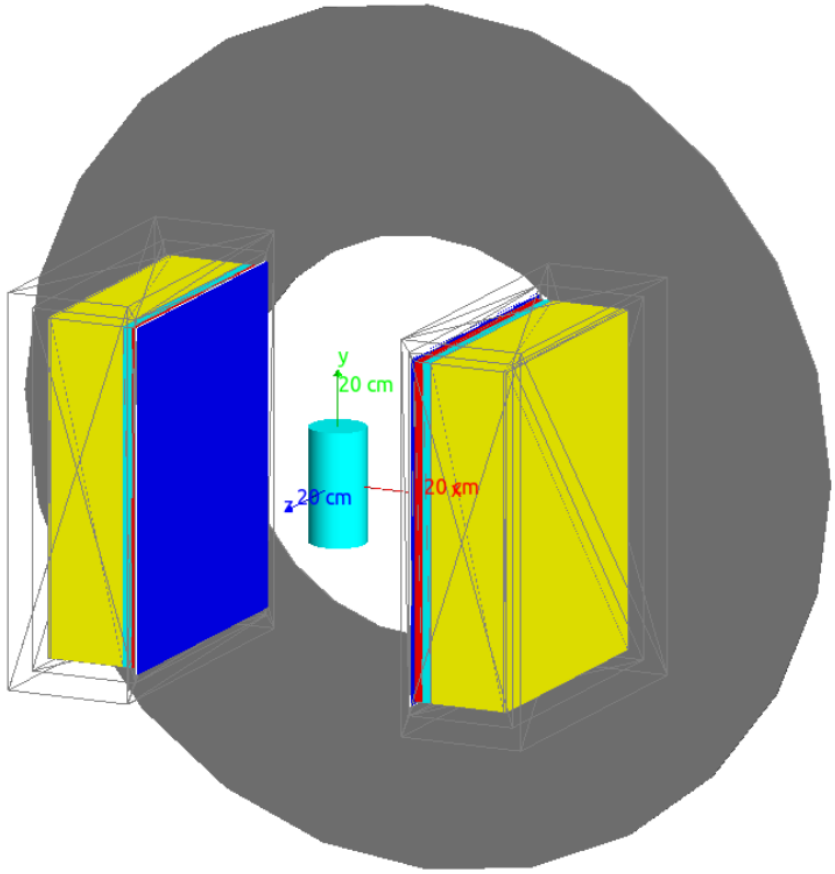

Figure 7.4: A diagram of the GATE simulation for testing the reconstruction of tracer size.

145 In the centre of the cylinder is a simulated PEPT tracer consisting of a sphere of
146 graphite and initially measuring 0.5 mm in radius. Graphite is used for this example because
147 it is a mineral which can be volumetrically activated in a cyclotron, yet is soft enough that it
148 can be easily sheared during the course of a PEPT experiment, providing a realistic method
149 of testing tracer size measurement in the future [149, 124]. The tracer is modelled as being
150 volumetrically activated with 10 MBq activity. For the purposes of the simulation, two
151 back-to-back gamma rays are produced by the tracer rather than positrons. This will not
152 accurately capture the expected positron range from a real tracer. This simplification was
153 used to reduce the amount of computation needed and additionally, because the positron

range in water is approximately 1 mm and is not expected to significantly affect the results [88].

For the first GATE simulation, all of the activity was contained within the static tracer in the centre of the water cylinder. Then, in subsequent GATE simulations, to emulate the loss of tracer activity caused by attrition, some of the activity of the tracer is transferred to the surrounding water cylinder and the volume and size of the tracer are decreased relative to the amount of activity transferred. The volume of the tracer is simulated over an order of magnitude, ranging from the initial volume to 0.1 of the initial volume in 10 increments. A decrease in the initial tracer volume by an order of magnitude will result in a decrease in the tracer radius by a factor of 0.464 to the initial radius. The total amount of activity present in the tracer and water cylinder is kept constant at 10 MBq to keep the proportion of the contribution from random LoRs relatively constant. The decay of the tracer could be simulated, as would be expected in a real experiment, but a correction factor for the contribution from random LoRs would be needed.

For all simulations, approximately two million LoRs are collected and a large sample of 2000 LoRs per PEPT-reconstructed location is used. When applying the Birmingham Method to this data, approximately 1000 locations are generated. The mean residual per LoR is calculated from the 1000 locations for each value of f from 0.04 to 0.8 in increments of 0.01. The derivative of the mean residuals per LoRs is then computed over the range of f for each of the different amounts of tracer attrition. The value of f corresponding to the maximum derivative of the mean residual per LoR is assumed to be the value of f_{opt} for the tracer. The f_{opt} values are then compared to the f_{opt} when no activity has been lost from the tracer, providing a measure of the volume lost. The attrited tracer size is then calculated using Equation 7.2.

7.4 Results

Using the LoRs generated from the GATE simulations, the mean residual per LoR as a function of the fraction of LoRs used to locate the tracer as well as the rate of change in the mean residual per LoR are plotted in Figure 7.5. Part (a) of this plot shows as the tracer decreases in size by losing activity to the water cylinder the rise in the mean residual per LoR occurs at lower values of f. This is expanded on in part (b) where the rate of the change in the mean residual per LoR is plotted. For each tracer size simulated, there is always a peak in the rate of change of the mean residual per LoR. Moreover, these peaks occur at intervals seeming to correspond proportionally with the volume lost from the initial tracer volume.

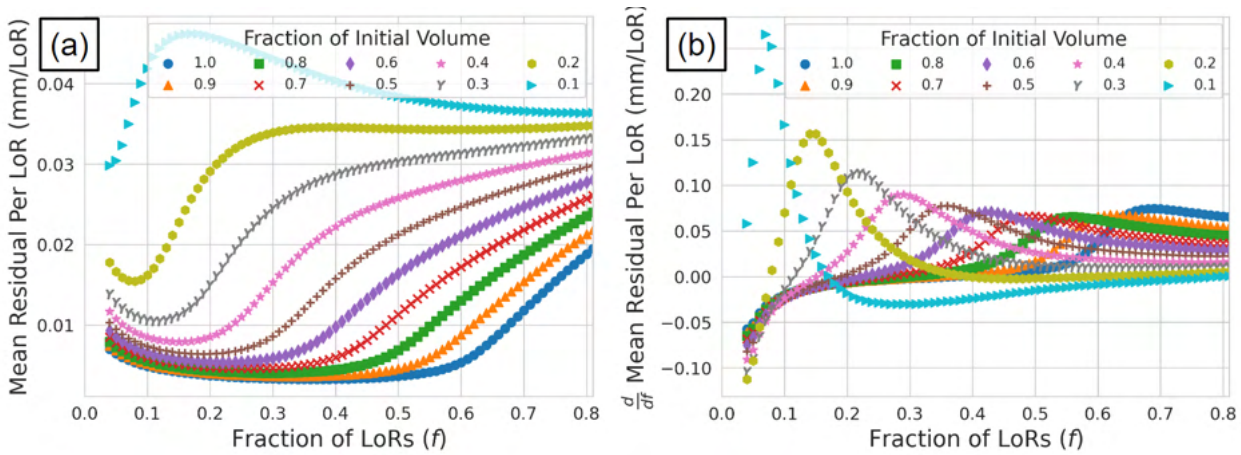

Figure 7.5: (a) The mean residual per LoR as a function of the fraction of used LoRs in a sample for different levels of tracer attrition (b) the gradient of the mean residual per LoR as a function of the fraction of used LoRs in a sample for different levels of tracer attrition. The gradient plot shows a global peak which corresponds to the f_{opt} associated with each amount of tracer attrition.

When the tracer is fully intact, meaning all of the activity is contained within the tracer, the f_{opt} is measured to be approximately 0.69. After half of the tracer volume has been lost, the f_{opt} falls to approximately 0.36, which is nearly half of the previous value.

By using the extracted values of f_{opt}, the fraction of f_{opt} compared to the f_{opt_0} for the fully intact tracer yields the reconstructed tracer volume, shown in Figure 7.6 with the error bars representing the $\pm$ 0.01 f in the extraction of the peak from Figure 7.5. To compute the estimated tracer radii, Equation 7.2 is used, allowing for comparison of the estimated radii to the GATE-simulated tracer radii, shown in Figure 7.7 with the uncertainty in the f carried through.

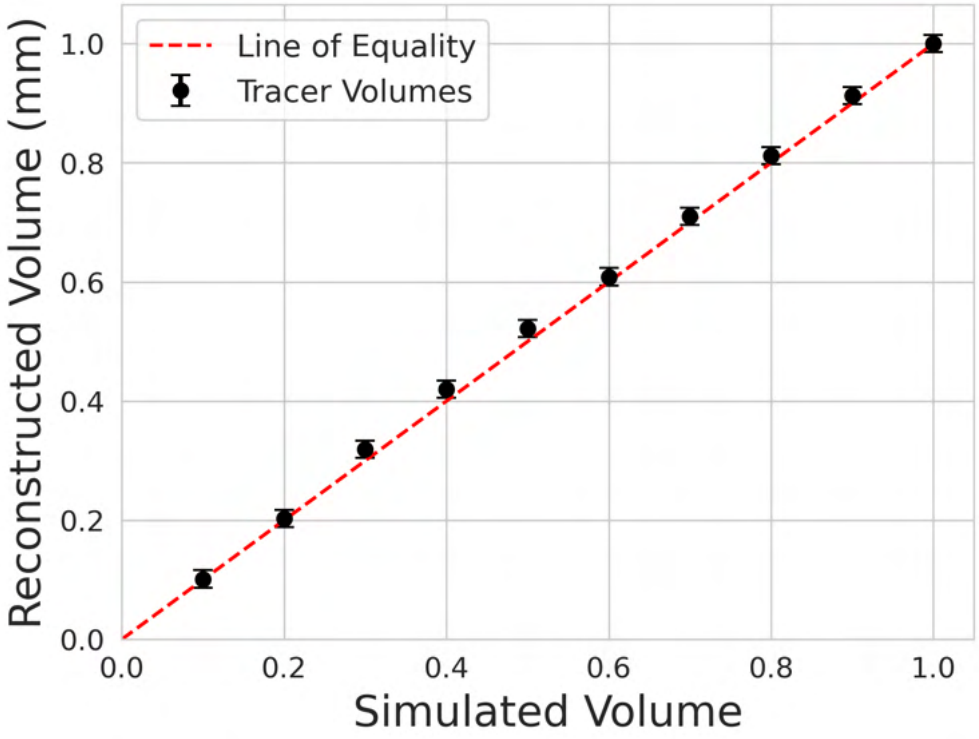

Figure 7.6: The PEPT reconstructed tracer volumes compared to the GATE simulated volumes.

The results of Figures 7.6 and 7.7 show that both the reconstructed tracer volume and radius agree within the uncertainty of the measurement. Figure 7.7 also shows that the uncertainty in the reconstructed tracer radius increases as the tracer becomes smaller. This is caused by the increment in f becoming a larger fraction of the extracted f_{opt} as the f_{opt} for the attrited tracer decreases.

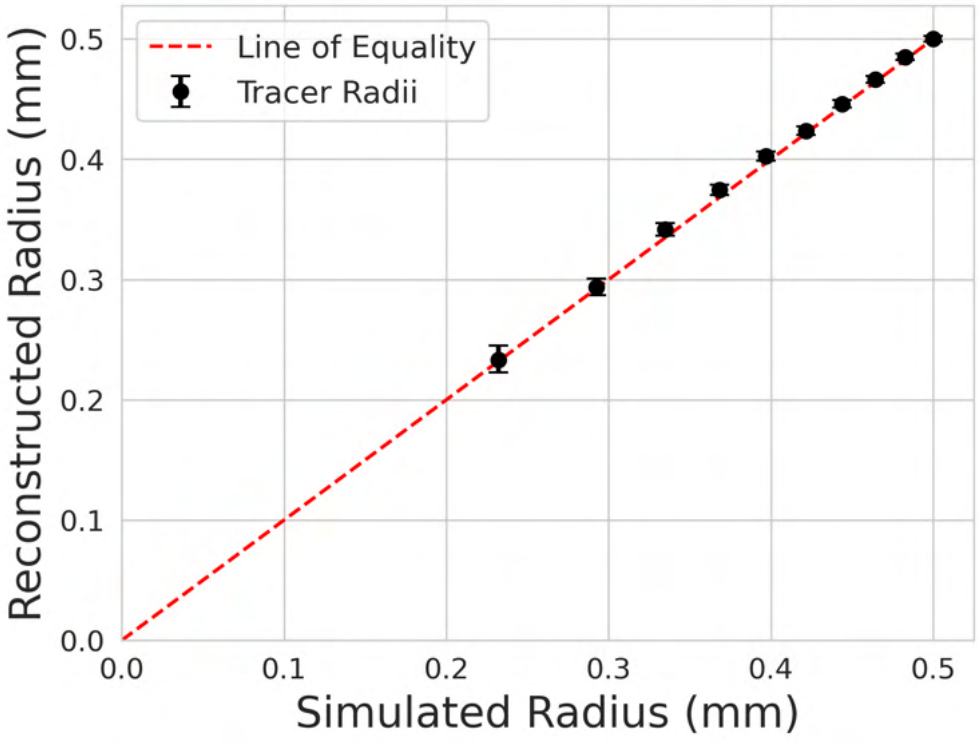

Figure 7.7: The PEPT reconstructed tracer radii compared to the GATE simulated radii.

7.5 Discussion

This Chapter has shown through simulated results that it is possible to infer tracer size through PEPT measurements. This type of information has never before been able to be generated from PEPT measurements. Additionally, this technique can be used with the same PEPT algorithm, the Birmingham Method, which has been used in decades of PEPT research [100, 149]. Importantly, this novel measurement of tracer size can be generated alongside the time and position of a tracer. If this could be used in real experiments where the loss of particle material occurs, then not only could the loss of material over time be analysed, but also the areas where it occurs could be identified.

The results of this work are an important first step towards realising this capability. It should be noted that this proof-of-concept was simulated under ideal conditions. For

example, the tracer was uniformly activated, is a spherical static tracer, and the activity of the tracer disperses to the surrounding system instantaneously and uniformly. Nevertheless, now that the measurement of tracer size has been shown to be within the realm of possibility, future development of the technique can be carried out which will extend the technique to more realistic conditions. Additional GATE simulations could be created in the future that test this technique using tracers of different shapes, moving tracers, and non-uniform activity distributions within the tracer or the experimental system.

Ultimately, this technique should be tested using real data in carefully crafted experiments which allow the size of the tracer and its loss of activity to be validated through independent techniques. One such type of experiment where this could be used is in the exfoliation of thin layers of graphene from graphite in a high-shear aqueous system [124, 105]. Graphene is a useful material for a variety of applications, but industrial-scale production remains elusive [154, 126]. However, graphene occurs naturally in graphite and can be produced by introducing graphite particles to a high-shear aqueous system and extracting thin exfoliated layers of graphene sheets [105]. The natural carbon-12 in a graphite particle can be irradiated in a cyclotron beam to produce carbon-11 which decays through positron-emission [13]. By using a graphite tracer in a PEPT experiment it could be possible to measure tracer size and relate the exfoliation rates of graphene sheets to the hydrodynamic conditions the tracer experiences in the system.

Chapter Eight

General Discussion

8.1 Summary of Work

Overall, the work contained within this thesis presents a new way of using simulation to inform and optimise experiments in the field of positron emission particle tracking (PEPT). Monte Carlo simulations of radioactive sources emulating PEPT tracers, the transport of radiation through geometries representing experimental systems, and recording interactions with models of PEPT detector systems have been shown to capture all of the aspects of real experiments [44, 46, 45, 47, 149, 148]. Geant4 Application for Tomographic Emission (GATE) v9.1 was used to perform these simulations [56, 55]. The GATE software provides the tools to realistically model PEPT experiments, including complex geometries and materials, radioactive decay, detector modelling tools, and control over the pulse-processing behaviour of detectors [115]. Comparing the GATE-simulated data to experimentally acquired data shows minimal quantitative differences and no qualitative differences [44, 47]. Simulations of PEPT experiments using GATE thus provide a powerful way of generating synthetic data to better understand the performance and limitations of PEPT algorithms, identify optimal experimental parameters or algorithm settings, and developing new PEPT

methodologies [149, 47].

The first step toward realising the utility of GATE simulations in PEPT research, discussed in Chapter 2, was to create GATE models of the most commonly used PEPT detector systems. These PEPT detector systems were the ADAC Forte and the Large Modular Array (LaMA) [97]. To create of GATE models of these systems, not only did their geometry need to be recreated in GATE, but also their performance characteristics [90]. This was achieved by measuring their spatial resolution and count-rate response [44, 45]. To calibrate the GATE model of the Forte, both a manual calibration and a novel autonomous method of calibration using an evolutionary algorithm were used [44, 45]. The manual calibration method was found to produce a reasonable agreement with the performance characterisation experiments but was subjective to the combinations of digitizer parameter values a user chose, such as the dead-time and threshold values, and took a considerable amount of time to calibrate [44]. By using an evolutionary algorithm to perform the digitizer calibration, it was hypothesised that the GATE model's digitizer parameters could be optimised to best match a set of target output values, in this case, the count-rate response. In this way, using an evolutionary algorithm to calibrate the digitizer removes user bias in parameter selection, speeds up the time to calibrate the digitizer, and improves the resulting calibration beyond what a manual calibration could produce [45].

In order to interface the GATE simulation with an existing evolutionary algorithm, CMA-ES, a Python library was developed, called the Autonomous Calibration and Characterisation via Evolutionary Software (ACCES) v0.2.2 [92]. ACCES uses a Python script defining the GATE simulation and populates the user-defined free parameters of the digitizer with candidate solutions generated by CMA-ES, autonomously launches the simulation, analyses the results with a cost function, and then uses CMA-ES to generate improved candidate solutions for the next generation of simulations [92, 42]. When ACCES was applied to simultaneously calibrate 6 free parameters of the digitizer, the optimisation took approx-

imately 4 days to complete, which was significantly less time than the approximately one month of time the manual calibration took [45]. In addition to the improvement in the time needed to calibrate the model, the average absolute percent differences for the 250 mm, 525 mm, and 800 mm head separations were found to be approximately 7.55%, 4.30%, and 5.48%, respectively [45]. Compared to the manually calibrated version of the models, which produced mean absolute percent difference in the count rate response of approximately 17.78%, 15.42%, and 21.75% for the 250 mm, 525 mm, and 800 mm head separations, the ACCES-calibrated model improves the agreement between simulation and experiment by approximately a factor of three [45]. The results show that evolutionary simulation being used to calibrate a GATE model's digitizer is a viable and useful new tool which can provide better quantifiable reasons why digitizer parameter values were selected for the model and decreases the effort required by the user to calibrate new models.

Additionally, a GATE model for the LaMA was created and the same performance characterisation experiments were conducted and ACCES was used to calibrate its digitizer model. For these experiments, a simple configuration of the LaMA was used such that by characterising its performance and validating the model on a small scale, the scaled-up behaviour of the system will remain accurate. Creating a GATE model for the LaMA which can recreate the experimental data presented several new challenges since it must be easily customised to rapidly prototype new geometries and only allow specific coincidences to be formed between connected boxes of block detector, features that were not needed in the GATE model of ADAC Forte [39]. To solve the issues of creating customised geometries, the generic repeater functionality was used to create a single fully detailed module which could be copied into user-specified positions in the simulation to create any arbitrary configuration. Additionally, since the method of forming coincidences could not be specified in the simulation, a limitation of the GATE software, a post-processing script was written to remove coincidences from the output file which could not be produced in reality between

boxes not connected.

Following a manual calibration of the spatial resolution through tuning the crystal blurring, the count rate response experiments were analysed to be used as a comparison for the ACCES optimisation. The optimisation took approximately three days to converge. When the new set of simulations was conducted with the calibrated digitizer, the results match the experiment to a mean absolute difference of 3.4% over all three count-rates with the total, true, and corrupted count-rates being 2.3%, 2.2%, and 5.7%, respectively. One of the most interesting findings of the ACCES-calibrated digitizer was that the optimised value for the time resolution falls within the 12 ns $\pm$ 2 ns measured in a previous characterisation of the LaMA, providing evidence that the calibrated parameters correspond to physical reality and are global solutions, not simply local solutions [71].

The ACCES-aided calibrations of both the ADAC Forte and LaMA produced GATE models which closely match the experimental performance characteristics [45]. From this work, the results were two new validated GATE models of PEPT detector systems and the creation of a novel method of calibration which could be applied in the future to other GATE models. With these GATE models, it is now possible to produce realistic synthetic data which could be used to better understand the quality of data produced by proposed experiments and, indeed, to optimise experiments.

In Chapter 3, the use of validated GATE models to compare the performance of PEPT algorithms using realistic synthetic data was discussed. Using GATE models to produce test data sets for PEPT algorithms is ideal because in simulation the position and movement of sources can be precisely prescribed [115]. With this information, the PEPT reconstructed locations of the sources can be quantitatively compared, which is not possible through real experiments. PEPT algorithms have been produced by different authors, at different institutions, and their performance was tested through a wide range of different experiments [100,

139, 73, 138, 94, 91, 82]. However, never before has there been an attempt to understand their relative strengths and weaknesses using a common suite of test data in which the position of the source is known precisely. The work in Chapter 3 presents the first study of its kind in which the performance of 7 different PEPT algorithms were compared in terms of their spatial, temporal, and spatiotemporal resolution on a simulated data set [149]. The algorithms tested include the Birmingham Method, G-Means, Spatiotemporal B-Slice Reconstruction (SBSR), Feature Point Identification (FPI), Spherical Density Method (SDM), PEPT using Machine Learning (PEPT-ML), and PEPT using Expectation Maximisation (PEPT-EM) [100, 139, 73, 138, 94, 91, 82].

The simulated tests in this work were carefully designed to allow a comprehensive, quantitative and, fair comparison of PEPT algorithms [149]. In total, 8 individual tests were designed, each to evaluate a specific aspect of a PEPT algorithm's performance. Additionally, the tests were split into two groups, one to test single tracer tracking and another to test multiple tracer tracking since not all PEPT algorithms are designed for multiple tracer tracking. The single tracer tests analysed the accuracy of position reconstruction using different numbers of LoRs, the relationship between spatiotemporal resolution and tracer velocity, the robustness to noise in the data set, and the ability to track tracers near the edge of the FOV. In addition to these four tests, the multiple tracer tracking tests analysed the minimum separation distance allowable to resolve two tracers, the ability to distinguish false locations from real tracer locations, the robustness of distinguishing distinct trajectories of moving tracers, and the ability to track large numbers of tracers within the FOV. Each test was performed using two different PEPT detector system geometries, which are the ADAC Forte and Siemens ECAT [44, 57]. This was done to ensure that algorithms developed for a ring-shaped detector are not unfairly disadvantaged by being tested exclusively on a planar detector, and vice-versa. The full results of these tests can be found in Chapter 3.

The results of the PEPT algorithm tests show that there is no single best algorithm

and that each algorithm has its own strengths and weaknesses [149]. Perhaps the best approach is to use different algorithms for different goals. The Birmingham Method is computationally efficient and good for tracking slow-moving, relatively active single tracers, while PEPT-ML, PEPT-EM or FPI are better for tracking large numbers of tracers or resolving collisions between tracers. On the other hand, the results suggest that the SBSR method is the optimal choice for limited tracer activity or low detected count-rate. Moreover, the ability to compare the performance of these algorithms would not have been possible without the use of methods for modelling PEPT tracers and detectors developed in this thesis. The development and validation of the GATE model for the ADAC Forte and the adaptation of the pre-existing model of the Siemens ECAT allowed for the modelling of PEPT experiments in detectors that are similar to those used in real PEPT facilities [97]. The simulated experiments provide a direct comparison between the PEPT reconstructed positions of tracers and their prescribed location in the GATE simulation, which makes it possible to quantitatively measure the accuracy of PEPT algorithms, which is achieved for the first time, without introducing biases or uncertainties that are as significant as the measurements.

The work of Chapters 4, 5, and 6 sought to demonstrate the optimisation of experimental parameters, length of time for data acquisition or treatment of the trajectory data, and use of PEPT algorithm parameters to extract as much useful information from a PEPT experiment. When conducting real PEPT experiments, it is not always possible to explore a range of conditions to find the optimal conditions for the experiment due to the limited time and resources of PEPT facilities [148]. This is further complicated because the exact behaviour of the tracers in an experiment is not known. Thus, elucidating the best treatment of the data with the PEPT algorithms themselves or the post-processing also cannot typically be determined because there is no way to compare the data to the ground truth. It is in these situations where the simulation of PEPT experiments using validated GATE

models which produce realistic synthetic LoRs are demonstrated to be useful.

The first practical aspect of PEPT experiments that is shown to be capable of being optimised through simulation, covered in Chapter 4, is the tracer activity and geometry of the detector. The tracer activity and head separation are two factors which can be readily customised before a PEPT experiment and the head separation can also be changed during the experiment [103]. These aspects of an experiment play a role in the rate at which LoRs can be recorded and the quality of the LoRs. Finding the optimal settings is difficult to predict due to the non-ideal nature of PEPT detectors, tracers, and experimental geometries, but through simulation, a wide variety of combinations for these parameters can be tested and analysed. This was demonstrated through a GATE simulation using the ADAC Forte PEPT detector system [45]. In this simulation, a trajectory of a tracer moving through a continuous blender used in the pharmaceutical industry was prescribed and a range of tracer activities and head separations were tested to observe their effect on the quality of the PEPT reconstructed trajectory in terms of its spatial, temporal, and spatiotemporal resolution [83].

The results from these simulations reveal a complex relationship in count-rate versus activity curves for each head separation [44]. This leads to lower head separations and higher tracer activities not always producing a high LoR count-rate. The count-rates were found to be proportional to the temporal resolution of the PEPT reconstructed trajectories. However, when the spatial resolution was also taken into account in the spatiotemporal resolution, a clear optimal set of conditions was identified, approximately in the range of 300 mm - 350 mm and 6 MBq - 8 MBq [149]. However, in real experiments, the tracer activity cannot be kept constant due to radioactive decay. This means that if the experiment was started with these optimal parameters, after only a short amount of time the decay in tracer activity would result in non-optimal conditions [96]. Therefore, a method of determining the optimal starting parameters was also demonstrated which takes advantage of the expected source decay over the length of the experiment by finding parameters which would produce the best

average spatiotemporal resolution. In this simulated case, if the experiment were expected to take 30 minutes to complete, the optimal starting activity should be set to approximately 8 MBq. However, if the experiment were expected to take 180 minutes, the optimal starting activity would be approximately 13 MBq.

Ultimately, the research in Chapter 4 demonstrates an effective new tool for PEPT users to obtain the most precise trajectories from an experiment by optimising the experimental parameters which are determined through simulated results. Although the technique was shown in one particular PEPT experiment, the methods are presented generally such that they could be adapted to any PEPT experiment in which there exists a proper model of the PEPT detector system and experimental setup. This method is therefore beneficial as it eliminates the need for physical experimentation to determine optimal parameters, saving time and resources for PEPT facilities.

In Chapter 5, a separate type of PEPT experiment optimisation was demonstrated. This Chapter used realistic tracer motion from the discrete element method (DEM) simulation to prescribe the movement of a source and used the occupancy of the flowing media to create a voxelised volume to induce realistic gamma-ray scattering [46, 64]. A DEM simulation of a system representative of a typical PEPT experiment was created and a single DEM particle trajectory was extracted. In this case, the experimental system chosen was a small rotating drum, commercially known as the GranuDrum, used for power flow characterisation [30]. The GranuDrum system was recreated in GATE simulation and placed in the centre of the FOV of the ADAC Forte [46]. Then the extracted DEM particle trajectory was prescribed as the movement of a source approximating a PEPT tracer. Using the LoRs recorded by the simulated PEPT detector system, the tracer trajectory was reconstructed using the PEPT-ML algorithm and the two-dimensional velocity, acceleration, occupancy, and granular temperature fields were reconstructed along with an analysis of the spatial error between the GATE-prescribed and PEPT-reconstructed trajectories [91, 46].

This work produced two important results in addition to demonstrating that DEM can be combined with GATE simulations for realistic modelling of PEPT experiments and that the main driver of spatial error in PEPT measurements is tracer acceleration, rather than velocity, as was previously thought [149, 46]. These were the accuracy of PEPT reconstructed two-dimensional fields increase logarithmically with measurement time and distortions in PEPT measurements derived from fluctuating velocity terms can be predicted in simulation and corrected using correlations [143, 46].

When considering that longer measurement time results in diminishing returns in accuracy, this can be exploited by a user to determine how long is needed to achieve a desired level of consistency with the bulk system's steady-state behaviour without spending unnecessary time on the measurement. Previous work has attempted to connect the necessary experiment duration to key system parameters, finding that system size and mean velocity were important considerations [144]. However, these findings were based solely on DEM particle trajectories and do not take into account the spatial and temporal limitations of the PEPT technique [144]. In particular, it was found that PEPT reconstructed fields take longer to reach the same level of agreement with bulk particle behaviour than single DEM trajectories. This is due to lower temporal resolution in PEPT trajectories, which smooths velocity fluctuations and adds random fluctuations caused by measurement uncertainty [46]. This did not cause major differences in reconstructed fields except for when the granular temperature was compared. Because the calculation of granular temperature relies on fluctuating velocity terms, it was found to be susceptible to distortions caused by smoothing when the tracer velocity was high and random fluctuations induced by PEPT reconstruction when the tracer velocity was low. Importantly, when comparing the granular temperature calculated directly from the DEM to when it was calculated by from PEPT-reconstructed trajectories, a correlation was discovered which could be used to correct the PEPT measurement. What makes this particularly useful, is that by combining a DEM simulation with a

GATE simulated model of the experiments, these correlations may be used to correct real PEPT measurements [46].

Further, in Chapter 6, an investigation into how GATE simulations could be used to improve the selection of PEPT algorithm parameters was conducted. While it is known that thicker and denser experimental systems generate more scattered LoRs and more active tracers will cause more random LoRs, little has been done to develop methods to predict PEPT algorithm parameters which will maximise the spatial and temporal resolution trajectories. However, since GATE simulations can produce realistic data similar to that generated from PEPT experiments, it was hypothesised this could be used to identify optimal parameter values. To test this, real experiments using the ADAC Forte and the GATE model of the Forte were used to generate LoRs from static tracer experiments and a recreated PEPT experiment involving an attritor mill [103, 44]. Then, the Birmingham Method was applied with different values of the algorithm's f and N_{LoRs} parameters, and the optimal parameters selected by comparing the PEPT reconstructed trajectories to the known positions or GATE-prescribed trajectories [100]. In both the real experiments and GATE simulations, the amount of attenuation was found to be directly correlated with the values of f that minimised the position uncertainty [47]. Moreover, the experimental and simulated results produced approximately the same values for the optimal value of f. When the fraction of true LoRs in the data was measured in the GATE simulation, it was found to correspond to the value of f which minimised the uncertainty, providing the first experimental evidence that f corresponds to the fraction of true LoRs in a sample [100, 97].

In the moving tracer experiment, a tracer was fixed to the impeller of a real attritor mill and rotated, with the same conditions recreated in the GATE simulation [47]. In this experiment, the spatial error was calculated by directly comparing the expected trajectory from the real experiment and the prescribed trajectory from the experiment. The optimal values for f and N_{LoRs} for the real experiment were determined to be 0.275 and 1400,

respectively. Similarly, the optimal parameters for the GATE simulation were determined to be 0.25 and 1300, respectively. These results demonstrated the extracted optimal values for the experiment and simulation are in reasonable agreement, providing further validation of the GATE model. When the GATE simulations were expanded to place the tracer in different positions within the attritor mill it was possible to create a fully three-dimensional map of the optimal parameters.

Ultimately this work resulted in the demonstration of a variant of the Birmingham Method which can dynamically change its parameter values for f and N_{LoRs} to use the optimal parameters determined from GATE simulation. When comparing the mean spatial error of the reconstructed trajectories, it was found that dynamic parameter optimisation results in mean spatial errors in trajectory reconstruction decreasing by 76.5% over the default parameters and by 4.03% over the best constant parameters [47]. This study makes a significant contribution to the field of PEPT research by showing that the optimal parameter values for the Birmingham Method can be predicted with GATE simulations. Further, the Birmingham Method was extended to use locally optimal parameters based on an estimate of the tracer position, resulting in the best possible trajectory [100]. This research demonstrates that GATE simulations, when conducted in conjunction with PEPT experiments, can be used to eliminate the need for expert knowledge in choosing the parameters f and N_{LoRs}, making the Birmingham Method PEPT algorithm more rigorous and more accessible.

The work of Chapter 7 sought to demonstrate how not only can simulations of PEPT experiments be used to improve the current focus of PEPT research, but also be used to develop entirely new methodologies. To this end, the development and demonstration of a novel PEPT technique for measuring tracer size was presented. In many types of systems of industrial relevance, it is expected that particles will lose material through attrition, dissolution, or other similar mechanisms [130]. Measuring the rate at which material is lost and where this loss occurs is crucial to designing better systems to encourage or discourage

this happening [22]. Therefore it would be beneficial to use PEPT to not only record the trajectory of a tracer but also measure its size.

It was hypothesised the optimal value for the f parameter of the Birmingham Method PEPT algorithm, which has been shown to be directly correlated with the fraction of LoRs which originate from the tracer, could be used to measure the amount of positron-emitting material contained within the PEPT tracer [47]. As a tracer loses positron-emitting material to the experimental system in which it is contained, it was expected that the optimal values for f would also decrease. Thus by measuring this change, the loss of tracer volume and also its size could possibly be determined. To test this, GATE simulations were designed to record LoRs from a spherical, volumetrically activated graphite tracer placed in the centre of a water cylinder for which in successive simulations, a fraction of the tracer activity was transferred to the water cylinder and the tracer radius decreased.

The results of the GATE simulation showed a remarkable ability to reconstruct the tracer size using the determined values for f. This type of information has never before been able to be generated from PEPT measurements. Additionally, this technique can be used with the same PEPT algorithm, the Birmingham Method, which has been used in decades of PEPT research but was never intended to be used for this purpose [100, 149]. The findings of this study represent a crucial initial step towards achieving the proposed capability. It is worth noting that this proof-of-concept experiment was conducted under idealised conditions, such as uniform tracer activation, a spherical static tracer, and instantaneous and uniform dispersion of tracer activity. Despite these simplifications, the demonstration of the feasibility of measuring tracer size opens the door for future research to extend the technique to more realistic scenarios.

8.2 Context of Research

The PEPT technique is useful for a variety of scientific and engineering fields because it can provide fully three-dimensional trajectories of positron-emitting tracers which are contained within opaque media and cannot adequately be imaged through conventional optically-based techniques [100, 147, 148]. To perform PEPT measurements, three things are required: a typically small positron-emitting tracer which matches the properties of the media studied, (2) a position-sensitive radiation detector capable of recording the coincident 511 keV gamma rays formed by positron annihilation with an electron, and (3) an algorithm which can convert samples of coincident gamma rays into an estimate of the location over time [96, 97, 149]. These three components of PEPT experiments are called PEPT tracers, PEPT detector systems, and PEPT algorithms. Regarding PEPT tracers, a great deal of research has been conducted to determine the best methods of producing positron-emitting isotopes using cyclotrons and radioisotope generators, finding methods of labelling tracers of positron-emitting material, and altering the properties of tracers to match specific aspects of the media being studied [98, 96, 32, 107, 17, 18]. Additionally, much work has gone into acquiring, maintaining, and building sensitive PEPT detectors which can record lines-of-response (LoRs) formed from coincident gamma rays at a high rate and spatial resolution [43, 103, 72, 39, 97]. Further, there has been extensive development of PEPT algorithms capable of locating PEPT tracers accurately and often [100, 139, 73, 138, 94, 91, 82].

Using these tools, PEPT has been used successfully to interrogate granular media, complex fluid flow, and many types of industrially relevant systems like rotating drums and fluidised beds [147, 137, 87, 146]. It is the ability of PEPT to reveal tracer motion inside dense systems which makes it so useful and led to the discovery of previously unknown phenomena such as the existence of convection rolls in vibro-fluidised beds [141, 143]. However, because the motion of tracers inside opaque media cannot often be determined through

complementary measurements, the response of PEPT detectors is a complex relationship between the source activity, attenuating media, and detector settings, and there has not been a method of producing synthetics PEPT data, the ability to fully understand PEPT algorithm performance, optimise data acquisition, and investigate the feasibility of current of future PEPT experiments has not been possible [45].

The work presented in this thesis demonstrates a method of simulating all aspects of PEPT experiments, introducing several robust tools and methodologies. The GATE v9.1 simulation software is used to perform Monte Carlo simulations of radiation transport and convert the interactions of these particles with a model of a PEPT detector into synthetic data identical to that output by a real system [56]. However, the validation of these GATE models is crucial to generating realistic detector responses [90, 62, 36]. While there exist several validated GATE models of detectors, none of these, with the exception of the Siemens Inveon and Siemens ECAT, have been used at PEPT facilities [115, 74, 57].

Therefore, in this work, two new GATE models of PEPT detector systems were developed and validated through industry-standard performance characterisation experiments such that realistic synthetic data could be generated using the same systems on which PEPT are currently carried out [90, 45, 97]. These validated models have been made available and free to use through a GitHub repository (GATE model repository). Additionally, to aid in the calibration of the GATE models, a method of using an evolutionary algorithm to perform this task is developed using the Autonomous Characterisation and Calibration using Evolutionary Simulation (ACCES) software which eliminates the need for manual tuning of GATE models. The new method ultimately resulted in the ability to calibrate GATE models to match experimental results which outperform manual tuning in terms of accuracy in matching experimental performance characteristics and also in terms of the time needed to produce a calibrated model [92, 45]. Similar to before, the software for calibrating GATE models is freely available and an example of how to use this method has been provided in a

freely available GitHub repository (optimisation example).

These GATE models are used in this thesis to generate synthetic data through which methods of evaluating PEPT algorithm performance, optimising PEPT, and developing new PEPT methodologies are conducted [149, 46, 47]. The comparative review of the seven most used PEPT algorithms using a GATE simulated data set where the position of the PEPT tracer was prescribed produced several interesting and important results [149]. To begin with, this was the first comparative review of PEPT algorithms ever conducted, bringing together PEPT users from across the globe. Further, while no PEPT algorithm performed superior to all others, the relative strengths and weaknesses of each algorithm were able to be highlighted. For example, while Birmingham Method is computationally efficient and produced the best spatial accuracy when the tracer was static. When the tracer was moving the Birmingham Method's accuracy decreased and was surpassed by newer, but more computationally intensive algorithms, like PEPT-ML and PEPT-EM [100, 91, 82]. When two tracers separated by a small distance were simulated, it was found that several of the algorithms were able to resolve the two tracers up to and including the point-of-contact between them [149]. This potentially opens up a new avenue of PEPT research which could examine tracer-tracer collisions. Moreover, as a result of this international collaboration, the FPI and Birmingham Method algorithms were incorporated into the 'pept' Python package, joining the existing PEPT-ML algorithms, and advancing community collaboration and future development within the PEPT research community. Researchers can now access all of these tools by simply installing one package and can easily switch between them. Moreover, the assessment of the performance of each algorithm was conducted transparently with comparison functions, available and free to use through a GitHub repository (comparison functions).

Additionally, included in this thesis are frameworks for optimising experimental tracer activity and head separation, a method of using DEM simulated movement of PEPT tracers

to determine the optimal length of time for experimental data acquisitions as well as the development to correcting factors for PEPT reconstructed field data, and novel method of selecting the optimal PEPT algorithm parameters to most accurately reconstruct a trajectory [46, 47]. The development of the technique would not have been possible without GATE simulations allowing the evaluation of more conditions that could be tested in reality and the ability to compare a PEPT reconstructed trajectory to its known prescribed position. For the first example, optimising tracer activity and PEPT detector system head separation is an important development in PEPT research because prior to this work an exhaustive set of experiments would be needed to determine these optimal experimental parameters. However, by using GATE simulations to produce realistic synthetic LoRs, a range of different parameter combinations can quickly be tested. This creates an opportunity for users to identify the set of parameter values which will produce the best PEPT trajectory reconstruction in simulation, then use those values when running a real experiment. This represents a significant development in tailoring experimental parameters.

Further, by combining DEM simulation of a proposed PEPT experiment with GATE simulation, the relationship between the accuracy of reconstructed system behaviour can be determined such as a user can identify the amount of time needed to run an experiment to reach a desired level of agreement of the bulk steady-state behaviour [46]. This is important because when PEPT experiments are conducted, often a range of experiments testing a range of operating conditions of a system are tested [148]. For example, this would be normal if the PEPT data was being used to validate a DEM simulation of the same system [63, 145, 15]. However, the time and resources of PEPT facilities are limited [148]. Therefore, finding the amount of time needed to run an experiment to reach a desired level of accuracy is beneficial because it results in better use of time and resources. Moreover, the comparison between PEPT-reconstructed and GATE-prescribed trajectories has led to the development of a new method to correct measurements which rely on fluctuating velocity terms like

granular temperature [142]. In this work, the errors between the DEM simulation and PEPT reconstruction were plotted through a bin-by-bin comparison, and it was found that at low granular temperatures, PEPT over-predicts the values and at high granular temperatures, PEPT measurements under-predict the values which was caused by the uncertainty in PEPT measurements inducing velocity fluctuations when in reality the tracer velocity was not fluctuating and also caused real fluctuations to be smoothed. This work marks the first attempt to correlate PEPT-measured quantities to known DEM quantities and shows that it is possible to correct PEPT measurements, opening up a new avenue of research [46]. To aid the use of the described methods of using DEM simulated results in GATE situations, the software tools which made this work possible have been made available and free to download through a GitHub repository (DEM2GATE).

With regard to the optimisation of PEPT algorithm parameter values, this work makes significant progress in estimating the optimal parameters and in understanding why certain values were selected. Of note, it was found that for the Birmingham Method and at least for static tracers, the value of the f parameter which minimises the tracer position uncertainty is approximately the fraction of true LoRs in the sample [47]. While this had long been assumed there has never before been direct experimental evidence for this [100, 103, 97]. Further, it was found the optimal values for the f and N_{LoRs} parameters for a moving tracer are highly position dependent. By using GATE simulations, it was shown that these values could be determined for all tested tracer positions and be used to dynamically update the parameter used for reconstructing the tracer using an iterative method which begins by using the non-optimal parameters to first estimate the tracer location [47]. This method is an important development in PEPT research because it represents the most sophisticated method of selecting optimal algorithm parameters currently developed, but for a justifiable reason, since it results in an improvement in spatial accuracy of 76.5% over the default parameters and by 4.03% over the best constant parameters [47].

In addition to these previously mentioned uses of the simulation of PEPT experiments, one further remains to be discussed. This is the use of simulated data to develop entirely new PEPT methodologies which could be applied to real experiments in the future. This work develops a new PEPT methodology which could be used to measure tracer size. Being able to measure the size of the tracer and, in particular, changes in the size of the tracer is desirable because, in many types of systems used in industry, it is expected that the media will lose some of its volume through attrition, dissolution, or other mechanisms [130]. Therefore, being able to measure the rate of material loss and where this occurs is crucial to designing better systems to encourage or discourage this behaviour [22]. Using a simulated data set based on a spherical graphite tracer in a water cylinder, as some of its activity was transferred to the cylinder and the tracer decreased in size, it was shown that measuring changes in the optimal f parameter could reliably be used to infer the tracer size. While these simulations were idealised, it is nevertheless a useful development in PEPT research, demonstrating this type of capability is within the realm of possibility.

8.3 Future Outlook

While the work in this thesis attempted to comprehensively provide an answer to many of the research questions involving the simulation of PEPT experiments, there remain many open questions and opportunities for further development. For example, the use of ACCES to autonomously calibrate GATE models was shown to be a useful method which could be applied to any future GATE model of a PEPT detector system [92, 45]. If a user were to create a geometric model of a PEPT detector system, the same methodology described in Chapter 2 of conducting performance characterisation experiments and defining a cost function to compare the results of the simulation could be used to efficiently calibrate free parameters of the digitizer model. However, one addition that would make this method

even more rigorous and useful would be to improve ACCES such that not only could the parameter values of free parameters of the digitizer be optimised, but also the structure of the digitizer model itself.

Additionally, the combination of DEM and GATE simulations presents several opportunities for further use of the technique to better understand distortions in PEPT measurements and methods of correcting measurements in post-processing. The work presented here only examined two-dimensional fields. In the future, this could readily be extended to fully three-dimensional fields. Moreover, measurements of other types could be analysed to determine if there are correctable systematic errors caused by PEPT.

Concerning the optimisation of PEPT algorithm parameters in Chapter 6, it should be noted that this approach may also apply to other PEPT algorithms with free parameters, such as the FPI PEPT algorithm [138]. It is currently unknown if this approach using the f and N_{LoRs} parameters of the Birmingham Method only worked as well as it did because these parameters are directly tied to physical reality, in this case, the fraction of LoRs which emanate from the PEPT tracer and the detection rate of LoRs, or if this approach could similarly be extended to the voxel thresholding parameter of the FPI algorithm.

With respect to the new PEPT methodology discussed in Chapter 7, it was shown how PEPT can be extended to infer a tracer's size. However, the GATE simulations used to test these methods were simplistic. This was done intentionally to provide the best chance of success. Further developments of these techniques could be made in the future to tackle more realistic scenarios. There are several potential avenues for future development, including conducting additional GATE simulations that examine the technique's performance under various conditions, such as tracers of different shapes, moving tracers, and non-uniform activity distributions. Ultimately, it would be best to test the techniques using a real data set to fully validate the methodologies. For the tracer size measurement technique, this data

could be generated from a well-designed PEPT experiment in which a graphite tracer is blended to exfoliate layers of graphene. The use of graphite in the GATE-simulated example was chosen because carbon-11, which decays through positron-emission, can be produced from irradiation of the natural carbon-12 in graphite by using a cyclotron [13, 98].

In summary, there is no shortage of paths for the future development of the techniques presented in this work and many of the methods developed here are ready to be used to inform PEPT experiments. In fact, identifying the optimal tracer activity and head separation for a proposed experiment which would use the ADAC Forte PEPT detector system is recommended. This would ensure that the highest spatial accuracy and temporal resolution of PEPT trajectories are produced.

References

[1] J. Allison et al. "Recent developments in Geant4". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 835 (2016), pp. 186–225. ISSN: 0168-9002. DOI: 10.1016/j.nima.2016.06.125. URL: https://www.sciencedirect.com/science/article/pii/S0168900216306957.

[2] K. Alsabti, S. Ranka, and V. Singh. "An efficient k-means clustering algorithm". In: *Electrical Engineering and Computer Science - All Scholarship* (1997). URL: https://surface.syr.edu/eecs/43.

[3] C. D. Anderson. "The Positive Electron". In: *Physical Review* 43.6 (1933), pp. 491–494. DOI: 10.1103/PhysRev.43.491. URL: https://link.aps.org/doi/10.1103/PhysRev.43.491.

[4] T. W. Anderson and D. A. Darling. "Asymptotic Theory of Certain "Goodness of Fit" Criteria Based on Stochastic Processes". In: *The Annals of Mathematical Statistics* 23.2 (1952), pp. 193–212. ISSN: 0003-4851. DOI: 10.1214/aoms/1177729437. URL: http://projecteuclid.org/euclid.aoms/1177729437.

[5] H. Anger and D. Rosenthal. *Scintillation camera and positron camera*. INIS Reference Number: 43112886. International Atomic Energy Agency (IAEA): IAEA, 1959, pp. 59–75. URL: https://inis.iaea.org/search/search.aspx?orig_q=RN:43112891.

[6] G. Baldi, R. Conti, and E. Alaria. "Complete suspension of particles in mechanically agitated vessels". In: *Chemical Engineering Science* 33.1 (1978), pp. 21–25. ISSN: 0009-2509. DOI: 10.1016/0009-2509(78)85063-5. URL: https://www.sciencedirect.com/science/article/pii/0009250978850635.

[7] C. R. Bemrose et al. "Application of positron emission tomography to particulate flow measurement in chemical engineering processes". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 273.2-3 (1988), pp. 874–880. ISSN: 01689002. DOI: 10.1016/0168-9002(88)90111-8. URL: https://linkinghub.elsevier.com/retrieve/pii/0168900288901118 (visited on 05/14/2020).

[8] D. Brasse et al. "Correction Methods for Random Coincidences in Fully 3D Whole-Body PET: Impact on Data and Image Quality". In: *Journal of Nuclear Medicine* 46.5 (2005), pp. 859–867. ISSN: 0161-5505, 2159-662X. URL: https://jnm.snmjournals.org/content/46/5/859.

[9] A. Buffler et al. "A new positron emission particle tracking facility at iThemba LABS". In: *Proceedings of International Topical Meeting on Nuclear Research Applications and Utilization of Accelerators* (2010). URL: http://inis.iaea.org/Search/search.aspx?orig_q=RN:41127761.

[10] D. Calloway. "Beer-Lambert Law". In: *Journal of Chemical Education* 74.7 (1997), p. 744. ISSN: 0021-9584, 1938-1328. DOI: 10.1021/ed074p744.3. URL: https://pubs.acs.org/doi/abs/10.1021/ed074p744.3.

[11] R. Campello, D. Moulavi, and J. Sander. "Density-Based Clustering Based on Hierarchical Density Estimates". In: *Advances in Knowledge Discovery and Data Mining*. Ed. by J. Pei et al. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, 2013, pp. 160–172. ISBN: 9783642374562. DOI: 10.1007/978-3-642-37456-2_14.

[12] J. W. Carson, L. Cid, and J. McInerney. "Flow Measurement of Powders and Bulk Solids". In: *Kirk-Othmer Encyclopedia of Chemical Technology*. Ed. by Inc John Wiley & Sons. 1st ed. Wiley, 2020, pp. 1–23. ISBN: 9780471484943. DOI: 10.1002/0471238961.koe00053. URL: https://onlinelibrary.wiley.com/doi/10.1002/0471238961.koe00053.

[13] A. R. Chalise et al. "Carbon-11 and Carbon-12 beam range verifications through prompt gamma and annihilation gamma measurements: Monte Carlo simulations". In: *Biomedical physics and engineering express* 6.6 (2020), p. 065013. ISSN: 2057-1976. DOI: 10.1088/2057-1976/abb8b6. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8148632/.

[14] Y. F. Chang and A. C. Hoffmann. "A Lagrangian study of liquid flow in a reverse-flow hydrocyclone using positron emission particle tracking". In: *Experiments in Fluids* 56.1 (2015), p. 4. ISSN: 0723-4864, 1432-1114. DOI: 10.1007/s00348-014-1875-5.

[15] H. Che et al. "PEPT validated CFD-DEM model of aspherical particle motion in a spouted bed". In: *Chemical Engineering Journal* 453 (2023), p. 139689. ISSN: 13858947. DOI: 10.1016/j.cej.2022.139689. URL: https://linkinghub.elsevier.com/retrieve/pii/S1385894722051683.

[16] F. Chiti et al. "Using positron emission particle tracking (PEPT) to study the turbulent flow in a baffled vessel agitated by a Rushton turbine: Improving data treatment and validation". In: *Chemical Engineering Research and Design* 89.10 (2011), pp. 1947–1960. ISSN: 0263-8762. DOI: 10.1016/j.cherd.2011.01.015. URL: https://www.sciencedirect.com/science/article/pii/S0263876211000414.

[17] K. Cole et al. "A surface coating method to modify tracers for positron emission particle tracking (PEPT) measurements of froth flotation". In: *Powder Technology* 263 (2014), pp. 26–30. ISSN: 0032-5910. DOI: 10.1016/j.powtec.2014.04.083. URL: https://www.sciencedirect.com/science/article/pii/S0032591014004136.

[18] K. Cole et al. “Standard method for performing positron emission particle tracking (PEPT) measurements of froth flotation at PEPT Cape Town”. In: *MethodsX* 9 (2022), p. 101680. ISSN: 2215-0161. DOI: 10.1016/j.mex.2022.101680. URL: https://www.sciencedirect.com/science/article/pii/S2215016122000644.

[19] A. H. Compton. “A Quantum Theory of the Scattering of X-rays by Light Elements”. In: *Physical Review* 21.5 (1923), pp. 483–502. DOI: 10.1103/PhysRev.21.483.

[20] M. Conti and B. Bendriem. “The new opportunities for high time resolution clinical TOF PET”. In: *Clinical and Translational Imaging* 7.2 (2019), pp. 139–147. ISSN: 2281-7565. DOI: 10.1007/s40336-019-00316-5. URL: https://doi.org/10.1007/s40336-019-00316-5.

[21] M. Conti and L. Eriksson. “Physics of pure and non-pure positron emitters for PET: a review and a discussion”. In: *EJNMMI Physics* 3.1 (2016), p. 8. ISSN: 2197-7364. DOI: 10.1186/s40658-016-0144-5. URL: http://ejnmmiphys.springeropen.com/articles/10.1186/s40658-016-0144-5.

[22] R. Crawford and J. Ralston. “The influence of particle size and contact angle in mineral flotation”. In: *International Journal of Mineral Processing* 23.1 (1988), pp. 1–24. ISSN: 0301-7516. DOI: 10.1016/0301-7516(88)90002-6. URL: https://www.sciencedirect.com/science/article/pii/0301751688900026.

[23] J. C. Crocker and D. G. Grier. “Methods of Digital Video Microscopy for Colloidal Studies”. In: *Journal of Colloid and Interface Science* 179.1 (1996), pp. 298–310. ISSN: 00219797. DOI: 10.1006/jcis.1996.0217. URL: https://linkinghub.elsevier.com/retrieve/pii/S0021979796902179.

[24] D. E. Cullen, J. H. Hubbell, and L. Kissel. “EPDL97: the evaluated photo data library ‘97 version”. In: *United States Department of Energy Office of Scientific and Technical Information* (1997). DOI: 10.2172/295438. URL: https://www.osti.gov/biblio/295438.

[25] P. A. Cundall. *The measurement and analysis of accelerations in rock slopes*. University of London, 1971. URL: https://books.google.co.uk/books?id=N-CGHAAACAAJ.

[26] D. Daraio et al. "Validation of a Discrete Element Method (DEM) Model of the Grinding Media Dynamics within an Attritor Mill Using Positron Emission Particle Tracking (PEPT) Measurements". In: *Applied Sciences* 9.22 (2019), p. 4816. DOI: 10.3390/app9224816.

[27] C. Darwin. *On the origin of species, 1859*. Routledge, 2004.

[28] Y. L. Ding et al. "Scaling relationships for rotating drums". en. In: *Chemical Engineering Science* 56.12 (2001), pp. 3737–3750. ISSN: 00092509. DOI: 10.1016/S0009-2509(01)00092-6. (Visited on 11/03/2021).

[29] A. Einstein. "Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt". In: *Annalen der Physik* 322.6 (1905), pp. 132–148. ISSN: 00033804, 15213889. DOI: 10.1002/andp.19053220607.

[30] E. R. L. Espiritu et al. "Investigation of the rotating drum technique to characterise powder flow in controlled and low pressure environments". In: *Powder Technology* 366 (2020), pp. 925–937. ISSN: 0032-5910. DOI: 10.1016/j.powtec.2020.03.029.

[31] B. Faddegon et al. "The TOPAS tool for particle simulation, a Monte Carlo simulation tool for physics, biology and clinical research". In: *Physica Medica* 72 (2020), pp. 114–121. ISSN: 11201797. DOI: 10.1016/j.ejmp.2020.03.019. URL: https://linkinghub.elsevier.com/retrieve/pii/S1120179720300715.

[32] X. Fan, D. J. Parker, and M. D. Smith. "Enhancing 18F uptake in a single particle for positron emission particle tracking through modification of solid surface chemistry". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 558.2 (2006), pp. 542–546. ISSN:

0168-9002. DOI: 10.1016/j.nima.2005.12.186. URL: https://www.sciencedirect.com/science/article/pii/S0168900205026094.

[33] F. Fulchini et al. "Development of a methodology for predicting particle attrition in a cyclone by CFD-DEM". In: *Powder Technology*. Expanding Boundaries in Particle Technology: A select collision of papers from the World Congress in Particle Technology VIII 357 (2019), pp. 21–32. ISSN: 0032-5910. DOI: 10.1016/j.powtec.2019.08.101. URL: https://www.sciencedirect.com/science/article/pii/S0032591019307004.

[34] N. Ghods et al. "CFD-DEM modelling of particles attrition in jet-in-fluidized beds". In: *Chemical Engineering Research and Design* 148 (2019), pp. 336–348. ISSN: 0263-8762. DOI: 10.1016/j.cherd.2019.06.015. URL: https://www.sciencedirect.com/science/article/pii/S0263876219302977.

[35] W. D. Griffiths et al. "Tracking inclusions in aluminium alloy castings using positron emission particle tracking (PEPT)". In: *Materials Science and Technology* 26.5 (2010), pp. 528–533. ISSN: 0267-0836. DOI: 10.1179/174328409X405643. URL: https://doi.org/10.1179/174328409X405643.

[36] D. Guez et al. "Counting Rates Modeling for PET Scanners With GATE". In: *IEEE Transactions on Nuclear Science* 55.1 (2008), pp. 516–523. ISSN: 1558-1578. DOI: 10.1109/TNS.2007.910880.

[37] A. Guida, A. W. Nienow, and M. Barigou. "PEPT measurements of solid–liquid flow field and spatial phase distribution in concentrated monodisperse stirred suspensions". In: *Chemical Engineering Science* 65.6 (2010), pp. 1905–1914. ISSN: 0009-2509. DOI: 10.1016/j.ces.2009.11.005. URL: https://www.sciencedirect.com/science/article/pii/S0009250909007908.

[38] G. Hamerly and C. Elkan. "Learning the k in k-means". In: *Advances in Neural Information Processing Systems*. Ed. by S. Thrun, L. Saul, and B. Schölkopf. Vol. 16.

NIPS'03. Cambridge, MA, USA: MIT Press, 2003. URL: https://proceedings.neurips.cc/paper/2003/file/234833147b97bb6aed53a8f4f1c7a7d8-Paper.pdf.

[39] D. M. Hampel et al. "SuperPEPT: A new tool for positron emission particle tracking; first results". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1028 (2022), p. 166254. ISSN: 0168-9002. DOI: 10.1016/j.nima.2021.166254. URL: https://www.sciencedirect.com/science/article/pii/S0168900221010858.

[40] T. Han, A. Levy, and H. Kalman. "DEM simulation for attrition of salt during dilute-phase pneumatic conveying". In: *Powder Technology* 129.1–3 (2003), pp. 92–100. ISSN: 00325910. DOI: 10.1016/S0032-5910(02)00252-8. URL: https://linkinghub.elsevier.com/retrieve/pii/S0032591002002528.

[41] N. Hansen. "The CMA Evolution Strategy: A Comparing Review". In: *Towards a New Evolutionary Computation: Advances in the Estimation of Distribution Algorithms*. Studies in Fuzziness and Soft Computing. Berlin, Heidelberg: Springer, 2006, pp. 75–102. ISBN: 9783540324942. DOI: 10.1007/3-540-32494-1_4. URL: https://doi.org/10.1007/3-540-32494-1_4.

[42] N. Hansen et al. *CMA-ES/pycma: r3.2.2*. Zenodo, 2022. DOI: 10.5281/ZENODO.2559634. URL: https://zenodo.org/record/2559634.

[43] M. R. Hawkesworth et al. "A positron camera for industrial application". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 253.1 (1986), pp. 145–157. ISSN: 01689002. DOI: 10.1016/0168-9002(86)91138-1. URL: https://linkinghub.elsevier.com/retrieve/pii/0168900286911381 (visited on 05/14/2020).

[44] M. Herald, T. Kokalova-Wheldon, and C. R. K. Windows-Yule. "Monte Carlo model validation of a detector system used for Positron Emission Particle Tracking". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spec-*

trometers, Detectors and Associated Equipment 993 (2021), p. 165073. ISSN: 0168-9002. DOI: 10.1016/j.nima.2021.165073. URL: https://www.sciencedirect.com/science/article/pii/S0168900221000577.

[45] M. Herald et al. "Autonomous digitizer calibration of a Monte Carlo detector model through evolutionary simulation". In: *Scientific Reports* 12.1 (2022), p. 19535. ISSN: 2045-2322. DOI: 10.1038/s41598-022-24022-x. URL: https://www.nature.com/articles/s41598-022-24022-x.

[46] M. Herald et al. "DEM2GATE: Combining discrete element method simulation with virtual positron emission particle tracking experiments". In: *Powder Technology* 401 (2022), p. 117302. ISSN: 00325910. DOI: 10.1016/j.powtec.2022.117302. URL: https://linkinghub.elsevier.com/retrieve/pii/S0032591022001966.

[47] M. Herald et al. "Improving the accuracy of PEPT algorithms through dynamic parameter optimisation". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1047 (2023), p. 167831. ISSN: 01689002. DOI: 10.1016/j.nima.2022.167831. URL: https://linkinghub.elsevier.com/retrieve/pii/S0168900222011238.

[48] M. Herald et al. "Monte Carlo Model of the Large Modular Array for Positron Emission Particle Tracking". In: *IEEE Access* 11 (2023), pp. 25982–25990. ISSN: 2169-3536. DOI: 10.1109/ACCESS.2023.3255505.

[49] M. Herald et al. "Simulated time-dependent data to estimate uncertainty in fluid flow measurements". In: *Nuclear Engineering and Design* 337 (2018), pp. 221–227. ISSN: 00295493. DOI: 10.1016/j.nucengdes.2018.07.005. URL: https://linkinghub.elsevier.com/retrieve/pii/S0029549318303844.

[50] A. C. Hoffmann et al. "PET investigation of a fluidized particle: spatial and temporal resolution and short term motion". In: *Measurement Science and Technology* 16.3 (2005), pp. 851–858. ISSN: 0957-0233, 1361-6501. DOI: 10.1088/0957-0233/16/3/029.

[51] W. J. Hong, P. Yang, and K. Tang. "Evolutionary Computation for Large-scale Multi-objective Optimization: A Decade of Progresses". In: *International Journal of Automation and Computing* 18.2 (2021), pp. 155–169. ISSN: 1751-8520. DOI: 10.1007/s11633-020-1253-0. URL: https://doi.org/10.1007/s11633-020-1253-0.

[52] F. Hossein et al. "Application of acoustic techniques to fluid-particle systems – A review". In: *Chemical Engineering Research and Design* 176 (2021), pp. 180–193. ISSN: 0263-8762. DOI: 10.1016/j.cherd.2021.09.031. URL: https://www.sciencedirect.com/science/article/pii/S0263876221003993.

[53] A. N. Huang et al. "DEM study of particle segregation in a rotating drum with internal diameter variations". en. In: *Powder Technology* 378 (2021), pp. 430–440. ISSN: 00325910. DOI: 10.1016/j.powtec.2020.10.019. URL: https://linkinghub.elsevier.com/retrieve/pii/S0032591020309621 (visited on 10/22/2021).

[54] H. Iams and B. Salzberg. "The Secondary Emission Phototube". In: *Proceedings of the Institute of Radio Engineers* 23.1 (1935), pp. 55–64. ISSN: 2162-6626. DOI: 10.1109/JRPROC.1935.227243.

[55] S. Incerti, J. M. C. Brown, and S. Guatelli. "Advances in Geant4 applications in medicine". In: *Physica medica* 70 (2020), pp. 224–227. ISSN: 1724-191X. DOI: 10.1016/j.ejmp.2020.01.019.

[56] S. Jan et al. "GATE: a simulation toolkit for PET and SPECT". In: *Physics in Medicine and Biology* 49.19 (2004), pp. 4543–4561. ISSN: 0031-9155. DOI: 10.1088/0031-9155/49/19/007.

[57] S. Jan et al. "Monte Carlo Simulation for the ECAT EXACT HR+ system using GATE". In: *IEEE Transactions on Nuclear Science* 52.3 (2005), pp. 627–633. ISSN: 0018-9499. DOI: 10.1109/TNS.2005.851461. URL: http://ieeexplore.ieee.org/document/1487692/.

[58] C. R. Jones et al. "Dynamics of fabric and dryer sheet motion in domestic clothes dryers". In: *Drying Technology* 40.10 (2022), pp. 2087–2104. ISSN: 0737-3937. DOI: 10.1080/07373937.2021.1918706. URL: https://doi.org/10.1080/07373937.2021.1918706.

[59] H. W. de Jong et al. "A method to measure PET scatter fractions for daily quality control: SF measurement method for daily QC". In: *Medical Physics* 36.10 (2009), pp. 4609–4615. ISSN: 00942405. DOI: 10.1118/1.3213096. URL: http://doi.wiley.com/10.1118/1.3213096.

[60] K. O. Jung et al. "Whole-body tracking of single cells via positron emission tomography". In: *Nature Biomedical Engineering* 4.8 (2020), pp. 835–844. ISSN: 2157-846X. DOI: 10.1038/s41551-020-0570-5. URL: https://www.nature.com/articles/s41551-020-0570-5.

[61] M. H. Kalos and P. A. Whitlock. *Monte carlo methods*. John Wiley & Sons, 2009. URL: https://onlinelibrary.wiley.com/doi/10.1002/bbpc.198800128.

[62] S. Kerhoas-Cavata and D. Guez. "Modeling electronic processing in GATE". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*. Proceedings of the 3rd International Conference on Imaging Technologies in Biomedical Sciences 569.2 (2006), pp. 330–334. ISSN: 0168-9002. DOI: 10.1016/j.nima.2006.08.107. URL: https://www.sciencedirect.com/science/article/pii/S0168900206014690.

[63] W. Ketterhagen and C. Wassgren. "A perspective on calibration and application of DEM models for simulation of industrial bulk powder processes". In: *Powder Technology* 402 (2022), p. 117301. ISSN: 0032-5910. DOI: 10.1016/j.powtec.2022.117301. URL: https://www.sciencedirect.com/science/article/pii/S0032591022001954.

[64] C. Kloss et al. "Models, algorithms and validation for opensource DEM and CFD-DEM". In: *Progress in Computational Fluid Dynamics, An International Journal*

12.2/3 (2012), p. 140. ISSN: 1468-4349, 1741-5233. DOI: 10.1504/PCFD.2012.047457. URL: http://www.inderscience.com/link.php?id=47457.

[65] G. F. Knoll. *Radiation detection and measurement*. 4th ed. Hoboken, N.J: John Wiley, 2010. ISBN: 9780470131480.

[66] G. D. Knott. *Interpolating Cubic Splines*. Boston, MA: Birkhäuser, 2000. DOI: 10.1007/978-1-4612-1320-8. URL: http://link.springer.com/10.1007/978-1-4612-1320-8.

[67] K. Kouris et al. "Coincidence time window, ring sampling and attenuation problem in positron emission tomography". In: *Nuclear Instruments and Methods in Physics Research* 193.1 (1982), pp. 215–222. ISSN: 0167-5087. DOI: 10.1016/0029-554X(82)90699-1. URL: https://www.sciencedirect.com/science/article/pii/0029554X82906991.

[68] W. D. Kunze, M. Baehre, and E. Richter. "PET with a dual-head coincidence camera: spatial resolution, scatter fraction, and sensitivity". In: *Journal of Nuclear Medicine: Official Publication, Society of Nuclear Medicine* 41.6 (2000), pp. 1067–1074. ISSN: 0161-5505. URL: https://pubmed.ncbi.nlm.nih.gov/10855637/.

[69] S. Langford et al. "Positron Emission Particle Tracking (PEPT) Validation for Jet Flow". In: *Volume 5: Student Paper Competition*. Charlotte, North Carolina, USA: American Society of Mechanical Engineers, 2016, V005T15A065. ISBN: 9780791850053. DOI: 10.1115/ICONE24-60880. URL: https://asmedigitalcollection.asme.org/ICONE/proceedings/ICONE24/50053/Charlotte,%5C%20North%5C%20Carolina,%5C%20USA/252136.

[70] S. Langford et al. "Three-dimensional spatiotemporal tracking of fluorine-18 radiolabeled yeast cells via positron emission particle tracking". In: *PLOS ONE* 12.7 (2017). Ed. by Chin-Tu Chen, e0180503. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0180503. URL: https://dx.plos.org/10.1371/journal.pone.0180503.

[71] T. Leadbeater. "The development of positron imaging systems for applications in industrial process tomography". PhD thesis. University of Birmingham, 2009. URL: https://etheses.bham.ac.uk//id/eprint/521/.

[72] T. Leadbeater, D. J. Parker, and J. Gargiuli. "Characterization of the latest Birmingham modular positron camera". In: *Measurement Science and Technology* 22.10 (2011), p. 104017. ISSN: 0957-0233, 1361-6501. DOI: 10.1088/0957-0233/22/10/104017. URL: https://iopscience.iop.org/article/10.1088/0957-0233/22/10/104017.

[73] K. S. Lee, T. J. Kim, and G. Pratx. "Single-Cell Tracking With PET Using a Novel Trajectory Reconstruction Algorithm". In: *IEEE Transactions on Medical Imaging* 34.4 (2015), pp. 994–1003. ISSN: 1558-254X. DOI: 10.1109/TMI.2014.2373351.

[74] S. Lee, J. Gregor, and D. Osborne. "Development and validation of a complete GATE model of the Siemens Inveon trimodal imaging platform". In: *Molecular Imaging* 12.7 (2013), pp. 1–13. ISSN: 1536-0121. URL: https://pubmed.ncbi.nlm.nih.gov/23962650/.

[75] T. K. Lewellen. "The Challenge of Detector Designs for PET". In: *American Journal of Roentgenology* 195.2 (2010), pp. 301–309. ISSN: 0361-803X. DOI: 10.2214/AJR.10.4741. URL: https://www.ajronline.org/doi/10.2214/AJR.10.4741.

[76] T. Lindmo et al. "Accuracy and noise in optical Doppler tomography studied by Monte Carlo simulation". In: *Physics in Medicine and Biology* 43.10 (1998), pp. 3045–3064. ISSN: 0031-9155, 1361-6560. DOI: 10.1088/0031-9155/43/10/025. URL: https://iopscience.iop.org/article/10.1088/0031-9155/43/10/025.

[77] Y. Liu, M. Gonzalez, and C. Wassgren. "Modeling granular material blending in a rotating drum using a finite element method and advection-diffusion equation multiscale model". en. In: *AIChE Journal* 64.9 (2018), pp. 3277–3292. ISSN: 00011541. DOI: 10.1002/aic.16179. URL: https://onlinelibrary.wiley.com/doi/10.1002/aic.16179 (visited on 10/22/2021).

[78] G. Lumay et al. "Measuring the flowing properties of powders and grains". In: *Powder Technology* 224 (2012), pp. 19–27. ISSN: 0032-5910. DOI: 10.1016/j.powtec.2012.02.015.

[79] A. Lyon. "Why are Normal Distributions Normal?" In: *The British Journal for the Philosophy of Science* 65.3 (2014), pp. 621–649. ISSN: 0007-0882. DOI: 10.1093/bjps/axs046. URL: https://www.journals.uchicago.edu/doi/abs/10.1093/bjps/axs046.

[80] S. Maaß et al. "Automated drop detection using image analysis for online particle size monitoring in multiphase systems". In: *Computers and Chemical Engineering* 45 (2012), pp. 27–37. ISSN: 00981354. DOI: 10.1016/j.compchemeng.2012.05.014. URL: https://linkinghub.elsevier.com/retrieve/pii/S0098135412001603.

[81] K. Magota et al. "Performance characterization of the Inveon preclinical small-animal PET/SPECT/CT system for multimodality imaging". In: *European Journal of Nuclear Medicine and Molecular Imaging* 38.4 (2011), pp. 742–752. ISSN: 1619-7070, 1619-7089. DOI: 10.1007/s00259-010-1683-y. URL: http://link.springer.com/10.1007/s00259-010-1683-y.

[82] S. Manger, A. Renaud, and J. Vanneste. "An expectation–maximization algorithm for positron emission particle tracking". In: *Review of Scientific Instruments* 92.8 (2021), p. 085102. ISSN: 0034-6748. DOI: 10.1063/5.0053545. URL: https://aip.scitation.org/doi/10.1063/5.0053545.

[83] M. Marigo et al. "Application of Positron Emission Particle Tracking (PEPT) to validate a Discrete Element Method (DEM) model of granular flow and mixing in the Turbula mixer". In: *International Journal of Pharmaceutics* 446.1–2 (2013), pp. 46–58. ISSN: 03785173. DOI: 10.1016/j.ijpharm.2013.01.030. URL: https://linkinghub.elsevier.com/retrieve/pii/S0378517313000616.

[84] L. McInnes, J. Healy, and S. Astels. "hdbscan: Hierarchical density based clustering". In: *Journal of Open Source Software* 2.11 (2017), p. 205. ISSN: 2475-9066. DOI: 10.21105/joss.00205. URL: https://joss.theoj.org/papers/10.21105/joss.00205.

[85] C. L. Melcher. "Scintillation crystals for PET". In: *Journal of Nuclear Medicine: Official Publication, Society of Nuclear Medicine* 41.6 (2000), pp. 1051–1055. ISSN: 0161-5505. URL: https://pubmed.ncbi.nlm.nih.gov/10855634/.

[86] N. Metropolis et al. "Equation of State Calculations by Fast Computing Machines". In: *The Journal of Chemical Physics* 21.6 (1953), pp. 1087–1092. ISSN: 0021-9606. DOI: 10.1063/1.1699114. URL: https://aip.scitation.org/doi/10.1063/1.1699114.

[87] A. J. Morrison et al. "The shape and behaviour of a granular bed in a rotating drum using Eulerian flow fields obtained from PEPT". In: *Chemical Engineering Science* 152 (2016), pp. 186–198. ISSN: 0009-2509. DOI: 10.1016/j.ces.2016.06.022. URL: https://www.sciencedirect.com/science/article/pii/S0009250916303189.

[88] W. W. Moses. "Fundamental limits of spatial resolution in PET". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 648 (2011), S236–S240. ISSN: 01689002. DOI: 10.1016/j.nima.2010.11.092. URL: https://linkinghub.elsevier.com/retrieve/pii/S0168900210026276.

[89] M. Moszyński et al. "Energy resolution of scintillation detectors". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*. Special Issue in memory of Glenn F. Knoll 805 (2016), pp. 25–35. ISSN: 0168-9002. DOI: 10.1016/j.nima.2015.07.059. URL: https://www.sciencedirect.com/science/article/pii/S0168900215008979.

[90] NEMA. *Performance Measurements of Positron Emission Tomographs (PET) 2018*. Standard. National Electrical Manufacturers Association, 2018. URL: https://ww

w.nema.org/standards/view/Performance-Measurements-of-Positron-Emission-Tomographs.

[91] A. Nicuşan and C. R. K. Windows-Yule. "Positron emission particle tracking using machine learning". In: *Review of Scientific Instruments* 91.1 (2020), p. 013329. ISSN: 0034-6748. DOI: 10.1063/1.5129251. URL: https://aip.scitation.org/doi/abs/10.1063/1.5129251.

[92] A. Nicuşan et al. *ACCES: Autonomous Characterisation and Calibration via Evolutionary Simulation*. 2022. URL: https://github.com/uob-positron-imaging-centre/Coexist.

[93] G. T. Nolan and P. E. Kavanagh. "Computer simulation of random packing of hard spheres". In: *Powder Technology* 72.2 (1992), pp. 149–155. ISSN: 0032-5910. DOI: 10.1016/0032-5910(92)88021-9. URL: https://www.sciencedirect.com/science/article/pii/0032591092880219.

[94] A. Odo et al. "A PEPT algorithm for predefined positions of radioisotopes relative to the tracer particle". In: *Applied Radiation and Isotopes* 151 (2019), pp. 299–309. ISSN: 0969-8043. DOI: 10.1016/j.apradiso.2019.06.011. URL: https://www.sciencedirect.com/science/article/pii/S0969804319303884.

[95] D. J. Parker. "Positron emission particle tracking and its application to granular media". In: *Review of Scientific Instruments* 88.5 (2017), p. 051803. ISSN: 0034-6748, 1089-7623. DOI: 10.1063/1.4983046. URL: http://aip.scitation.org/doi/10.1063/1.4983046.

[96] D. J. Parker and X. Fan. "Positron emission particle tracking—Application and labelling techniques". In: *Particuology*. Selected papers from 1st UK-China Particle Technology Forum 6.1 (2008), pp. 16–23. ISSN: 1674-2001. DOI: 10.1016/j.cpart.2007.10.004. URL: https://www.sciencedirect.com/science/article/pii/S1672251507001455.

[97] D. J. Parker, D. M. Hampel, and T. Kokalova-Wheldon. "Performance Evaluation of the Current Birmingham PEPT Cameras". In: *Applied Sciences* 12.14 (2022), p. 6833. ISSN: 2076-3417. DOI: 10.3390/app12146833. URL: https://www.mdpi.com/2076-3417/12/14/6833.

[98] D. J. Parker and C. Wheldon. "The Birmingham MC40 Cyclotron Facility". In: *Nuclear Physics News* 28.4 (2018), pp. 15–20. ISSN: 1061-9127, 1931-7336. DOI: 10.1080/10619127.2018.1463021. URL: https://www.tandfonline.com/doi/full/10.1080/10619127.2018.1463021.

[99] D. J. Parker et al. "Developments in particle tracking using the Birmingham Positron Camera". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 392.1–3 (1997), pp. 421–426. ISSN: 01689002. DOI: 10.1016/S0168-9002(97)00301-X. URL: https://linkinghub.elsevier.com/retrieve/pii/S016890029700301X.

[100] D. J. Parker et al. "Positron emission particle tracking - a technique for studying flow within engineering equipment". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 326.3 (1993), pp. 592–607. ISSN: 0168-9002. DOI: 10.1016/0168-9002(93)90864-E. URL: https://www.sciencedirect.com/science/article/pii/016890029390864E.

[101] D. J. Parker et al. "Positron emission particle tracking studies of spherical particle motion in rotating drums". en. In: *Chemical Engineering Science* 52.13 (1997), pp. 2011–2022. ISSN: 00092509. DOI: 10.1016/S0009-2509(97)00030-4. URL: https://linkinghub.elsevier.com/retrieve/pii/S0009250997000304 (visited on 10/22/2021).

[102] D. J. Parker et al. "Positron emission particle tracking using a modular positron camera". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 604.1–2 (2009), pp. 339–

342. ISSN: 01689002. DOI: 10.1016/j.nima.2009.01.085. URL: https://linkinghub.elsevier.com/retrieve/pii/S0168900209001880.

[103] D. J. Parker et al. "Positron emission particle tracking using the new Birmingham positron camera". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 477.1–3 (2002), pp. 540–545. ISSN: 01689002. DOI: 10.1016/S0168-9002(01)01919-2. URL: https://linkinghub.elsevier.com/retrieve/pii/S0168900201019192.

[104] N. Patel, C. Wiggins, and A. E. Ruggles. "Positron emission particle tracking in pulsatile flow". In: *Experiments in Fluids* 58.5 (2017), p. 42. ISSN: 0723-4864, 1432-1114. DOI: 10.1007/s00348-017-2330-1. URL: http://link.springer.com/10.1007/s00348-017-2330-1.

[105] K. R. Paton et al. "Scalable production of large quantities of defect-free few-layer graphene by shear exfoliation in liquids". In: *Nature Materials* 13.6 (2014), pp. 624–630. ISSN: 1476-4660. DOI: 10.1038/nmat3944. URL: https://www.nature.com/articles/nmat3944.

[106] M. Pawlik-Niedźwiecka et al. "Preliminary Studies of J-PET Detector Spatial Resolution". In: *Acta Physica Polonica A* 132.5 (2017), pp. 1645–1649. ISSN: 0587-4246, 1898-794X. DOI: 10.12693/APhysPolA.132.1645. URL: http://przyrbwn.icm.edu.pl/APP/PDF/132/app132z5p47.pdf.

[107] J. Pellico et al. "Synthesis and 68Ga radiolabelling of calcium alginate beads for positron emission particle tracking (PEPT) applications". In: *Chemical Engineering Science* 264 (2022), p. 118159. ISSN: 00092509. DOI: 10.1016/j.ces.2022.118159. URL: https://linkinghub.elsevier.com/retrieve/pii/S0009250922007436.

[108] S. T. Perkins, D. E. Cullen, and S. M. Seltzer. *EEDL Evaluated Electron Data Library of the Lawrence Livermore National Laboratory, USA Summary documenta-*

tion. IAEA-NDS–157 INIS Reference Number: 26038451. International Atomic Energy Agency (IAEA), 1994, p. 2. URL: https://www.osti.gov/etdeweb/biblio/31490.

[109] E. Rault et al. "Accurate Monte Carlo modelling of the back compartments of SPECT cameras". In: *Physics in Medicine and Biology* 56.1 (2011), pp. 87–104. ISSN: 0031-9155, 1361-6560. DOI: 10.1088/0031-9155/56/1/006. URL: https://iopscience.iop.org/article/10.1088/0031-9155/56/1/006 (visited on 05/03/2020).

[110] V. Ravisankar et al. "Studying particle attrition in a solid-liquid agitated vessel using focused beam reflectance measurement (FBRM)". In: *Chemical Engineering and Processing - Process Intensification* 183 (2023), p. 109256. ISSN: 0255-2701. DOI: 10.1016/j.cep.2022.109256. URL: https://www.sciencedirect.com/science/article/pii/S0255270122004597.

[111] L. M. Rios and N. V. Sahinidis. "Derivative-free optimization: a review of algorithms and comparison of software implementations". In: *Journal of Global Optimization* 56.3 (2013), pp. 1247–1293. ISSN: 0925-5001, 1573-2916. DOI: 10.1007/s10898-012-9951-y. URL: https://link.springer.com/10.1007/s10898-012-9951-y.

[112] P. K. Romano and B. Forget. "The OpenMC Monte Carlo particle transport code". In: *Annals of Nuclear Energy* 51 (2013), pp. 274–281. ISSN: 0306-4549. DOI: 10.1016/j.anucene.2012.06.040. URL: https://www.sciencedirect.com/science/article/pii/S0306454912003283.

[113] J. Salvadori et al. "Monte Carlo simulation of digital photon counting PET". In: *EJNMMI physics* 7.1 (2020), pp. 1–19. ISSN: 2197-7364. DOI: 10.1186/s40658-020-00288-w. URL: https://doi.org/10.1186/s40658-020-00288-w.

[114] A. Sanchez-Crespo. "Comparison of Gallium-68 and Fluorine-18 imaging characteristics in positron emission tomography". In: *Applied Radiation and Isotopes* 76 (2013), pp. 55–62. ISSN: 09698043. DOI: 10.1016/j.apradiso.2012.06.034. URL: https://linkinghub.elsevier.com/retrieve/pii/S0969804312004708.

[115] D. Sarrut et al. "Advanced Monte Carlo simulations of emission tomography imaging systems with GATE". In: *Physics in Medicine and Biology* 66.10 (2021), 10TR03. ISSN: 0031-9155, 1361-6560. DOI: 10.1088/1361-6560/abf276. URL: https://iopscience.iop.org/article/10.1088/1361-6560/abf276.

[116] A. Savitzky and M. J. E. Golay. "Smoothing and Differentiation of Data by Simplified Least Squares Procedures." In: *Analytical Chemistry* 36.8 (1964), pp. 1627–1639. ISSN: 0003-2700, 1520-6882. DOI: 10.1021/ac60214a047. URL: https://pubs.acs.org/doi/abs/10.1021/ac60214a047.

[117] I. F. Sbalzarini and P. Koumoutsakos. "Feature point tracking and trajectory analysis for video imaging in cell biology". In: *Journal of Structural Biology* 151.2 (2005), pp. 182–195. ISSN: 10478477. DOI: 10.1016/j.jsb.2005.06.002. URL: https://linkinghub.elsevier.com/retrieve/pii/S1047847705001267.

[118] C. R. Schmidtlein et al. "Validation of GATE Monte Carlo simulations of the GE Advance/Discovery LS PET scanners: Validation of GATE simulations of GE PET scanners". In: *Medical Physics* 33.1 (2005), pp. 198–208. ISSN: 00942405. DOI: 10.1118/1.2089447. URL: http://doi.wiley.com/10.1118/1.2089447.

[119] M. Al-Shemmeri et al. "Coffee bean particle motion in a spouted bed measured using Positron Emission Particle Tracking (PEPT)". In: *Journal of Food Engineering* 311 (2021), p. 110709. ISSN: 0260-8774. DOI: 10.1016/j.jfoodeng.2021.110709. URL: https://www.sciencedirect.com/science/article/pii/S026087742100234X.

[120] H. Shi, G. Lumay, and S. Luding. "Stretching the limits of dynamic and quasi-static flow testing on cohesive limestone powders". In: *Powder Technology* 367 (2020), pp. 183–191. ISSN: 0032-5910. DOI: 10.1016/j.powtec.2020.03.036.

[121] J. W. Son et al. "SimPET: a Preclinical PET Insert for Simultaneous PET/MR Imaging". In: *Molecular Imaging and Biology* 22.5 (2020), pp. 1208–1217. ISSN: 1536-

1632, 1860-2002. DOI: 10.1007/s11307-020-01491-y. URL: https://link.springer.com/10.1007/s11307-020-01491-y.

[122] V. Spanoudaki and C. S. Levin. "Photo-Detectors for Time of Flight Positron Emission Tomography (ToF-PET)". In: *Sensors* 10.11 (2010), pp. 10484–10505. DOI: 10.3390/s101110484.

[123] S. Staelens et al. "Monte Carlo simulations of a scintillation camera using GATE: validation and application modelling". In: *Physics in Medicine & Biology* 48.18 (2003), p. 3021. ISSN: 0031-9155, 1361-6560. DOI: 10.1088/0031-9155/48/18/305. URL: https://iopscience.iop.org/article/10.1088/0031-9155/48/18/305.

[124] J. Stafford et al. "Real-time monitoring and hydrodynamic scaling of shear exfoliated graphene". In: *2D Materials* 8.2 (2021), p. 025029. ISSN: 2053-1583. DOI: 10.1088/2053-1583/abdf2f. URL: https://iopscience.iop.org/article/10.1088/2053-1583/abdf2f.

[125] C. W. Stearns et al. "Random coincidence estimation from single event rates on the Discovery ST PET/CT scanner". In: *2003 IEEE Nuclear Science Symposium. Conference Record (IEEE Cat. No.03CH37515)*. Portland, OR, USA: IEEE, 2004, pp. 3067–3069. ISBN: 9780780382572. DOI: 10.1109/NSSMIC.2003.1352545. URL: http://ieeexplore.ieee.org/document/1352545/ (visited on 05/14/2020).

[126] R. Stine et al. "Fabrication, Optimization, and Use of Graphene Field Effect Sensors". In: *Analytical Chemistry* 85.2 (2013), pp. 509–521. ISSN: 0003-2700, 1520-6882. DOI: 10.1021/ac303190w. URL: https://pubs.acs.org/doi/10.1021/ac303190w.

[127] M. Strugari et al. "NEMA NU 1-2018 performance characterization and Monte Carlo model validation of the Cubresa Spark SiPM-based preclinical SPECT scanner". In: *EJNMMI Physics* (2022). DOI: 10.21203/rs.3.rs-1946160/v1. URL: https://www.researchsquare.com/article/rs-1946160/v1.

[128] C. C. Sun. "True Density of Microcrystalline Cellulose". en. In: *Journal of Pharmaceutical Sciences* 94.10 (2005), pp. 2132–2134. ISSN: 00223549. DOI: 10.1002/jps.20459. URL: https://linkinghub.elsevier.com/retrieve/pii/S0022354916318676 (visited on 10/22/2021).

[129] C. Suplee. *XCOM: Photon Cross Sections Database*. text. 2009. URL: https://www.nist.gov/pml/xcom-photon-cross-sections-database (visited on 05/14/2020).

[130] J. P. M. Syvitski. *Principles, Methods and Application of Particle Size Analysis*. 1st ed. Cambridge University Press, 1991. ISBN: 9780521364720. DOI: 10.1017/CBO9780511626142.

[131] A. Taheri, S. Heidary, and H. Shahrabi. "Monte Carlo simulation of a SPECT system: GATE, MCNPX or SIMIND? (a comparative study)". In: *Journal of Instrumentation* 12.12 (2017), P12022–P12022. ISSN: 1748-0221. DOI: 10.1088/1748-0221/12/12/P12022. URL: https://iopscience.iop.org/article/10.1088/1748-0221/12/12/P12022.

[132] A. Tiwari et al. "Monte Carlo evaluation of hypothetical long axial field-of-view PET scanner using GE Discovery MI PET front-end architecture". In: *Medical Physics* 49.2 (2022), pp. 1139–1152. ISSN: 0094-2405, 2473-4209. DOI: 10.1002/mp.15422. URL: https://onlinelibrary.wiley.com/doi/10.1002/mp.15422.

[133] *University of Birmingham BlueBEAR*. Accessed March 01, 2022. 2022. URL: https://intranet.birmingham.ac.uk/it/teams/infrastructure/research/bear/index.aspx.

[134] S. Usman and A. Patil. "Radiation detector deadtime and pile up: A review of the status of science". In: *Nuclear Engineering and Technology* 50.7 (2018), pp. 1006–1016. ISSN: 17385733. DOI: 10.1016/j.net.2018.06.014. URL: https://linkinghub.elsevier.com/retrieve/pii/S1738573318302596 (visited on 05/03/2020).

[135] G. Venter. *Review of Optimization Techniques*. Chichester, UK: John Wiley and Sons, Ltd, 2010. ISBN: 9780470754405. DOI: 10.1002/9780470686652.eae495.

[136] T. S. Volkwyn et al. "Studies of the effect of tracer activity on time-averaged positron emission particle tracking measurements on tumbling mills at PEPT Cape Town". In: *Minerals Engineering* 24.3–4 (2011), pp. 261–266. ISSN: 08926875. DOI: 10.1016/j.mineng.2010.08.020.

[137] C. Wiggins, L. Carasik, and A. E. Ruggles. "Noninvasive interrogation of local flow phenomena in twisted tape swirled flow via positron emission particle tracking (PEPT)". In: *Nuclear Engineering and Design* 387 (2022), p. 111601. ISSN: 0029-5493. DOI: 10.1016/j.nucengdes.2021.111601. URL: https://www.sciencedirect.com/science/article/pii/S0029549321005537.

[138] C. Wiggins, R. Santos, and A. E. Ruggles. "A feature point identification method for positron emission particle tracking with multiple tracers". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 843 (2017), pp. 22–28. ISSN: 01689002. DOI: 10.1016/j.nima.2016.10.057. URL: https://linkinghub.elsevier.com/retrieve/pii/S0168900216311184.

[139] C. Wiggins, R. Santos, and A. E. Ruggles. "A novel clustering approach to positron emission particle tracking". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 811 (2016), pp. 18–24. ISSN: 01689002. DOI: 10.1016/j.nima.2015.11.136. URL: https://linkinghub.elsevier.com/retrieve/pii/S0168900215015533 (visited on 05/14/2020).

[140] C. Wiggins et al. "Qualification of multiple-particle positron emission particle tracking (M-PEPT) technique for measurements in turbulent wall-bounded flow". In: *Chemical Engineering Science* 204 (2019), pp. 246–256. ISSN: 0009-2509. DOI: 10.1016/j.ces.2019.04.030. URL: https://www.sciencedirect.com/science/article/pii/S0009250919303951.

[141] R. D. Wildman, J. M. Huntley, and D. J. Parker. "Convection in Highly Fluidized Three-Dimensional Granular Beds". In: *Physical Review Letters* 86.15 (2001), pp. 3304–

3307. DOI: 10.1103/PhysRevLett.86.3304. URL: https://link.aps.org/doi/10.1103/PhysRevLett.86.3304.

[142] R. D. Wildman and D. J. Parker. "Coexistence of Two Granular Temperatures in Binary Vibrofluidized Beds". In: *Physical Review Letters* 88.6 (2002), p. 064301. DOI: 10.1103/PhysRevLett.88.064301. URL: https://link.aps.org/doi/10.1103/PhysRevLett.88.064301.

[143] R. D. Wildman et al. "Single-particle motion in three-dimensional vibrofluidized granular beds". In: *Physical Review E* 62.3 (2000), pp. 3826–3835. DOI: 10.1103/PhysRevE.62.3826. URL: https://link.aps.org/doi/10.1103/PhysRevE.62.3826.

[144] C. R. K. Windows-Yule. "Ensuring adequate statistics in particle tracking experiments". In: *Particuology* 59 (2021), pp. 43–54. ISSN: 16742001. DOI: 10.1016/j.partic.2020.10.013.

[145] C. R. K. Windows-Yule and A. Neveu. "Calibration of DEM simulations for dynamic particulate systems". In: *Papers in Physics* 14 (2022), p. 140010. ISSN: 1852-4249. DOI: 10.4279/pip.140010. URL: https://www.papersinphysics.org/papersinphysics/article/view/795.

[146] C. R. K. Windows-Yule et al. "Effect of distributor design on particle distribution in a binary fluidised bed". In: *Powder Technology* 367 (2020), pp. 1–9. ISSN: 0032-5910. DOI: 10.1016/j.powtec.2020.03.034. URL: https://www.sciencedirect.com/science/article/pii/S003259102030231X.

[147] C. R. K. Windows-Yule et al. "Positron Emission Particle Tracking of Granular Flows". In: *Annual Review of Chemical and Biomolecular Engineering* 11.1 (2020), pp. 367–396. ISSN: 1947-5438, 1947-5446. DOI: 10.1146/annurev-chembioeng-011620-120633. URL: https://www.annualreviews.org/doi/10.1146/annurev-chembioeng-011620-120633.

[148] C. R. K. Windows-Yule et al. *Positron Emission Particle Tracking: A comprehensive guide*. IOP Publishing, 2022. ISBN: 9780750330718. URL: https://iopscience.iop.org/book/mono/978-0-7503-3071-8.

[149] C. R. K. Windows-Yule et al. "Recent advances in positron emission particle tracking: a comparative review". In: *Reports on Progress in Physics* 85.1 (2022), p. 016101. ISSN: 0034-4885, 1361-6633. DOI: 10.1088/1361-6633/ac3c4c. URL: https://iopscience.iop.org/article/10.1088/1361-6633/ac3c4c.

[150] C. R. K. Windows-Yule et al. "Understanding and exploiting competing segregation mechanisms in horizontally rotated granular media". In: *New Journal of Physics* 18.2 (2016), p. 023013. ISSN: 1367-2630. DOI: 10.1088/1367-2630/18/2/023013. URL: https://iopscience.iop.org/article/10.1088/1367-2630/18/2/023013.

[151] H. Yang et al. "Granular dynamics of cohesive powders in a rotating drum as revealed by speckle visibility spectroscopy and synchronous measurement of forces due to avalanching". In: *Chemical Engineering Science* 146 (2016), pp. 1–9. ISSN: 0009-2509. DOI: 10.1016/j.ces.2016.02.023. URL: https://www.sciencedirect.com/science/article/pii/S0009250916300690.

[152] R. Y. Yang et al. "Numerical simulation of particle dynamics in different flow regimes in a rotating drum". en. In: *Powder Technology* 188.2 (2008), pp. 170–177. ISSN: 00325910. DOI: 10.1016/j.powtec.2008.04.081. URL: https://linkinghub.elsevier.com/retrieve/pii/S0032591008002167 (visited on 10/22/2021).

[153] Y. Yang et al. "Effect of operating parameters on fine particle grinding in a vertically stirred media mill". In: *Separation Science and Technology* 52.6 (2017), pp. 1143–1152. ISSN: 0149-6395. DOI: 10.1080/01496395.2016.1276931. URL: https://doi.org/10.1080/01496395.2016.1276931.

[154] B. M. Yoo et al. "Graphene and graphene oxide and their uses in barrier polymers: Review". In: *Journal of Applied Polymer Science* 131.1 (2014), n/a–n/a. ISSN: 00218995.

DOI: 10.1002/app.39628. URL: https://onlinelibrary.wiley.com/doi/10.1002/app.39628.

[155] S. Yu et al. "The use of positron emission particle tracking (PEPT) to study milling of roll-compacted microcystalline cellulose ribbons". In: *Powder Technology*. Pharmaceutical Particle Technology 285 (2015), pp. 74–79. ISSN: 0032-5910. DOI: 10.1016/j.powtec.2015.06.051.

[156] C. Zheng et al. "Investigation of granular dynamics in a continuous blender using the GPU-enhanced discrete element method". In: *Powder Technology* 412 (2022), p. 117968. ISSN: 00325910. DOI: 10.1016/j.powtec.2022.117968. URL: https://linkinghub.elsevier.com/retrieve/pii/S003259102200849X.

Appendix Zero

GATE Model of the ADAC Forte

Listing 1: GATE Script for the validated model of the ADAC Forte PEPT detector. This PEPT detector model is described in Chapter 3. GitHub repository: https://github.com/uob-positron-imaging-centre/GATE_Models.

```
# GATE Model of the ADAC Forte
# Created by Matthew Herald
# Univeristy of Birmingham
# mxh1092@bham.ac.uk
# 2023

#------------------------ooooooOOOOOO000000OOOOOOoooooo------------------------#
#
#          D E F I N E   Y O U R   S C A N N E R   A N D
#
#          Y O U R   C O M P L E T E   S I M U L A T I O N
#
#          F O R   I M A G I N G   A P P L I CA T I O N S
#
#------------------------ooooooOOOOOO000000OOOOOOoooooo------------------------#

#============================================================
# VISUALISATION
#============================================================

#/vis/disable
#/vis/open                                   OGL
```

```
#/vis/viewer/reset
#/vis/viewer/set/viewpointThetaPhi    120 30
#/vis/viewer/set/viewpointThetaPhi 30 30
#/vis/viewer/set/viewpointThetaPhi    0 0
#/vis/viewer/zoom                     1.5
#/vis/viewer/set/style surface
#/vis/drawVolume
#/tracking/storeTrajectory            1
#/vis/scene/add/trajectories
#/vis/scene/endOfEventAction accumulate
#/vis/scene/add/axes
#/vis/viewer/update

#=============================================================
# GEOMETRY
#=============================================================

/gate/geometry/setMaterialDatabase    data/GateMaterials.db
/gate/world/geometry/setXLength       300. cm
/gate/world/geometry/setYLength       300. cm
/gate/world/geometry/setZLength       300. cm
/gate/world/setMaterial               Air
#/gate/world/vis/setColor              white
#/gate/world/vis/forceWireframe
#/gate/world/vis/setVisible            1

#-------------------ooooooOOOOOO000000OOOOOOoooooo-------------------#
#                                                                   #
#     D E F I N I T I O N   A N D   D E S C R I T I O N             #
#        O F   Y O U R   P E T   D E V I C E                        #
#                                                                   #
#-------------------ooooooOOOOOO000000OOOOOOoooooo-------------------#


#       ADAC_Forte
/gate/world/daughters/name                    ADAC
/gate/world/daughters/systemType              cylindricalPET
/gate/world/daughters/insert                  box
/gate/ADAC/placement/setTranslation       0.0 0.0 0.0 cm
/gate/ADAC/geometry/setXLength        200. cm
```

```
/gate/ADAC/geometry/setYLength          200. cm
/gate/ADAC/geometry/setZLength          200. cm
/gate/ADAC/placement/setRotationAxis        0 0 1
/gate/ADAC/placement/setRotationAngle       0 deg
/gate/ADAC/setMaterial                      Air
#/gate/ADAC/vis/setColor                    white
#/gate/ADAC/vis/forceWireframe
#/gate/ADAC/vis/setVisible                  1

#    RING
/gate/world/daughters/name          ring
/gate/world/daughters/insert        tessellated
/gate/ring/placement/setTranslation         0.0 0.0 -750.0 mm
/gate/ring/geometry/setPathToSTLFile        data/ring.stl
/gate/ring/setMaterial                      Aluminium
#/gate/ring/vis/forceWireframe
#/gate/ring/vis/setColor                    gray
#/gate/ring/vis/setVisible                  1

#    RING COVER
/gate/world/daughters/name                      plasticRing
/gate/world/daughters/insert                    cylinder
/gate/plasticRing/placement/setTranslation      0.0 0.0 -350.0 mm
/gate/plasticRing/geometry/setRmax              800 mm
/gate/plasticRing/geometry/setRmin              350 mm
/gate/plasticRing/geometry/setHeight            5.0 mm
/gate/plasticRing/placement/setRotationAngle    90 deg
/gate/plasticRing/setMaterial                   Plastic
#/gate/plasticRing/vis/forceWireframe
#/gate/plasticRing/vis/setColor                 gray
#/gate/plasticRing/vis/setVisible               1

#       HEAD (Move this volume to control the head separation)
/gate/ADAC/daughters/name           head
/gate/ADAC/daughters/insert         box
/gate/head/placement/setTranslation         337.0 0.0 0.0 mm
# Head distance from centre of system. User must add 87 mm to this value.
/gate/head/placement/setRotationAxis        0 1 0
/gate/head/placement/setRotationAngle       90 deg
/gate/head/geometry/setXLength              600 mm
```

```
/gate/head/geometry/setYLength                700 mm
/gate/head/geometry/setZLength                255 mm
/gate/head/setMaterial                        Air
#/gate/head/vis/setColor                       red
#/gate/head/vis/forceWireframe
#/gate/head/vis/setVisible                     1

/gate/head/daughters/name                      lightguide
/gate/head/daughters/insert                    tessellated
/gate/lightguide/placement/setTranslation       -190 -255 -95 mm
/gate/lightguide/geometry/setPathToSTLFile      data/lightguide.stl
/gate/lightguide/setMaterial                    Glass
#/gate/lightguide/vis/forceSolid
#/gate/lightguide/vis/setColor                    blue
#/gate/lightguide/vis/setVisible                  1

/gate/head/daughters/name                      PMTs
/gate/head/daughters/insert                    tessellated
/gate/PMTs/placement/setTranslation             -190 -255 -95 mm
/gate/PMTs/geometry/setPathToSTLFile            data/PMTs.stl
/gate/PMTs/setMaterial                          PMT
#/gate/PMTs/vis/forceSolid
#/gate/PMTs/vis/setColor                          yellow
#/gate/PMTs/vis/setVisible                        1

/gate/head/daughters/name                      electronics
/gate/head/daughters/insert                    tessellated
/gate/electronics/placement/setTranslation      -190 -255 -95 mm
/gate/electronics/geometry/setPathToSTLFile     data/electronics.stl
/gate/electronics/setMaterial                   Elec
#/gate/electronics/vis/forceSolid
#/gate/electronics/vis/setColor                   grey
#/gate/electronics/vis/setVisible                 0

/gate/head/daughters/name                      shielding
/gate/head/daughters/insert                    tessellated
/gate/shielding/placement/setTranslation        -190 -255 -95 mm
/gate/shielding/geometry/setPathToSTLFile       data/shielding.stl
/gate/shielding/setMaterial                     Lead
#/gate/shielding/vis/forceWireframe
```

```
#/gate/shielding/vis/setColor                    gray
#/gate/shielding/vis/setVisible                  0

/gate/head/daughters/name                        casing
/gate/head/daughters/insert                      tessellated
/gate/casing/placement/setTranslation            -190 -255 -95 mm
/gate/casing/geometry/setPathToSTLFile           data/casing.stl
/gate/casing/setMaterial                         Aluminium
#/gate/casing/vis/forceWireframe
#/gate/casing/vis/setColor                       gray
#/gate/casing/vis/setVisible                     0

/gate/head/daughters/name                        frontAl
/gate/head/daughters/insert                      box
/gate/frontAl/placement/setTranslation           0. 0.0 -96 mm
/gate/frontAl/geometry/setXLength                500 mm
/gate/frontAl/geometry/setYLength                600 mm
/gate/frontAl/geometry/setZLength                1.5 mm
/gate/frontAl/setMaterial                        Aluminium
#/gate/frontAl/vis/forceWireframe
#/gate/frontAl/vis/setColor                      white
#/gate/frontAl/vis/setVisible                    1

/gate/head/daughters/name                        frontPlastic
/gate/head/daughters/insert                      box
/gate/frontPlastic/placement/setTranslation      0. 0.0 -100 mm
/gate/frontPlastic/geometry/setXLength           500 mm
/gate/frontPlastic/geometry/setYLength           600 mm
/gate/frontPlastic/geometry/setZLength           10 mm
/gate/frontPlastic/setMaterial                   Perspex
#/gate/frontPlastic/vis/forceWireframe
#/gate/frontPlastic/vis/setColor                 white
#/gate/frontPlastic/vis/setVisible               1

/gate/head/daughters/name                        crystal
/gate/head/daughters/insert                      box
/gate/crystal/placement/setTranslation           0.0 0.0 -87 mm
# Crystal center is 13 mm from the front of the detector head
/gate/crystal/geometry/setXLength                16.5 cm
/gate/crystal/geometry/setYLength                19.5 cm
```

```
/gate/crystal/geometry/setZLength                1.6 cm
/gate/crystal/setMaterial                        NaI
#/gate/crystal/vis/forceWireframe
#/gate/crystal/vis/setColor                       red
#/gate/crystal/vis/setVisible                     1

/gate/crystal/repeaters/insert                   cubicArray
/gate/crystal/cubicArray/setRepeatNumberX        3
/gate/crystal/cubicArray/setRepeatNumberY        3
/gate/crystal/cubicArray/setRepeatNumberZ        1
/gate/crystal/cubicArray/setRepeatVector         16.5 19.5 0 cm

### M O D U L E ###

/gate/crystal/daughters/name                     active
/gate/crystal/daughters/insert                   box
/gate/active/placement/setTranslation            0.0 0.0 0.0 cm
/gate/active/geometry/setXLength                 16.5 cm
/gate/active/geometry/setYLength                 19.5 cm
/gate/active/geometry/setZLength                 1.6 cm
/gate/active/setMaterial                         NaI
#/gate/active/vis/forceSolid
#/gate/active/vis/setColor                        green
#/gate/active/vis/setVisible                      1


###     R E P E A T     H E A D ###
/gate/head/repeaters/insert                      ring
/gate/head/ring/setPoint1 0. 1. 0. mm
/gate/head/ring/setPoint2 0. 0. 0. mm
/gate/head/ring/setRepeatNumber                  2

/gate/geometry/rebuild
#/vis/viewer/update

/gate/systems/list

#       A T T A C H     S Y S T E M
/gate/systems/ADAC/rsector/attach      head
/gate/systems/ADAC/module/attach       crystal
```

```
/gate/systems/ADAC/crystal/attach      active

#       A T T A C H    C R Y S T A L  SD
/gate/systems/list
/gate/active/attachCrystalSD

#       A T T A C H     P H A N T O M  SD

#/gate/world/attachPhantomSD
#/gate/plasticRing/attachPhantomSD
#/gate/ring/attachPhantomSD
#/gate/lightguide/attachPhantomSD
#/gate/PMTs/attachPhantomSD
#/gate/electronics/attachPhantomSD
#/gate/frontPlastic/attachPhantomSD
#/gate/frontAl/attachPhantomSD
#/gate/shielding/attachPhantomSD
#/gate/casing/attachPhantomSD
#/gate/crystal/attachPhantomSD

/gate/physics/addPhysicsList emstandard

#———————————————ooooooOOOOOO000000OOOOOOoooooo———————————————#
#                                                               #
# D E F I N I T I O N   O F   Y O U R   A C Q U I S I T I O N   #
# D I G I T I Z E R   &   C O I N C I D E N C E   S H O R T E R #
#                                                               #
#———————————————ooooooOOOOOO000000OOOOOOoooooo———————————————#

#          A D D E R
/gate/digitizer/Singles/insert adder

#       E N E R G Y   B L U R R I N G
/gate/digitizer/Singles/insert                                  blurring
/gate/digitizer/Singles/blurring/setResolution                  0.15
/gate/digitizer/Singles/blurring/setEnergyOfReference           511. keV

#       P I L E -  U P
/gate/digitizer/Singles/insert pileup
/gate/digitizer/Singles/pileup/setDepth 4
```

```
/gate/digitizer/Singles/pileup/setPileup  498 ns #300 ns

#      T I M E      R E S O L U T I O N
/gate/digitizer/Singles/insert                                  timeResolution
/gate/digitizer/Singles/timeResolution/setTimeResolution    17 ns

#      S P A T I A L    B L U R R I N G
/gate/digitizer/Singles/insert spblurring
/gate/digitizer/Singles/spblurring/setSpresolution 5.0 mm
#/gate/digitizer/Singles/spblurring/verbose 1

#       E N E R G Y   C U T
/gate/digitizer/Singles/insert                                  thresholder
/gate/digitizer/Singles/thresholder/setThreshold            285 keV #200 keV
/gate/digitizer/Singles/insert                                  upholder
/gate/digitizer/Singles/upholder/setUphold                  1020 keV #1200 keV

#       D E A D    T I M E
/gate/digitizer/Singles/insert                                  deadtime
/gate/digitizer/Singles/deadtime/setDeadTime                1.07 us # 1.2  us
/gate/digitizer/Singles/deadtime/setMode                    paralysable
/gate/digitizer/Singles/deadtime/chooseDTVolume             crystal

#       E N E R G Y     W I N D O W
/gate/digitizer/name                                    EWindow
/gate/digitizer/insert                                  singleChain
/gate/digitizer/EWindow/setInputName                    Singles
/gate/digitizer/EWindow/insert                          thresholder
# 50 % energy window
/gate/digitizer/EWindow/thresholder/setThreshold    360 keV
/gate/digitizer/EWindow/insert                          upholder
/gate/digitizer/EWindow/upholder/setUphold          640 keV

#           C O I N C I D E N C E    S O R T E R
/gate/digitizer/Coincidences/setInputName           EWindow
/gate/digitizer/Coincidences/setWindow              15 ns
/gate/digitizer/Coincidences/minSectorDifference    1
/gate/digitizer/Coincidences/describe
/gate/digitizer/Coincidences/MultiplesPolicy takeAllGoods
```

```
###  P H A N T O M S ###
/gate/world/daughters/name                      phantom
/gate/world/daughters/insert                    sphere
/gate/phantom/setMaterial                       Plastic
#/gate/phantom/vis/forceWireframe
#/gate/phantom/vis/setColor                     white
#/gate/phantom/vis/setVisible                   0

/gate/phantom/geometry/setRmin                  0 mm
/gate/phantom/geometry/setRmax                  10 mm
/gate/phantom/geometry/setPhiStart              0 deg
/gate/phantom/geometry/setDeltaPhi              360 deg
/gate/phantom/geometry/setThetaStart            0 deg
/gate/phantom/geometry/setDeltaTheta            360 deg
/gate/phantom/placement/setTranslation          0 0 0 mm
/gate/phantom/moves/insert                      translation
/gate/phantom/translation/setSpeed              0 0 0 mm/s
#/gate/phantom/attachPhantomSD
/gate/geometry/rebuild
#/vis/viewer/update

### S O U R C E     D E F I N I T I O N ###

/gate/phantom/daughters/name                        Source1
/gate/phantom/daughters/insert                      sphere
/gate/Source1/setMaterial                   Plastic
#/gate/Source1/vis/forceWireframe
#/gate/Source1/vis/setColor                 white
#/gate/Source1/vis/setVisible               1

/gate/Source1/geometry/setRmin              0 mm
/gate/Source1/geometry/setRmax              0.1 mm
/gate/Source1/geometry/setPhiStart          0 deg
/gate/Source1/geometry/setDeltaPhi          360 deg
/gate/Source1/geometry/setThetaStart        0 deg
/gate/Source1/geometry/setDeltaTheta        360 deg
/gate/Source1/placement/setTranslation      0 0 0 mm

#/gate/Source1/attachPhantomSD
/gate/geometry/rebuild
```

```
#/vis/viewer/update

#=====================================================
# INITIALISATION
#=====================================================

/gate/run/initialize

###SOURCE DEFINITION###

#SOURCE
/gate/source/addSource twogamma
/gate/source/twogamma/setActivity 10 MBq
/gate/source/twogamma/setType backtoback

# POSITION
/gate/source/twogamma/gps/centre 0. 0. 0. mm

# PARTICLE
/gate/source/twogamma/gps/particle gamma
/gate/source/twogamma/gps/energytype Mono
/gate/source/twogamma/gps/monoenergy 0.511 MeV

# TYPE = Volume or Surface or Point

/gate/source/twogamma/gps/type Volume
/gate/source/twogamma/gps/shape Sphere
/gate/source/twogamma/gps/radius 0.1 mm
/gate/source/twogamma/gps/centre 0.0 0.0 0.0 mm

# SET THE ANGULAR DISTRIBUTION OF EMISSION
/gate/source/twogamma/gps/angtype iso

# SET MIN AND MAX EMISSION ANGLES
/gate/source/twogamma/gps/mintheta 0. deg
/gate/source/twogamma/gps/maxtheta 180. deg
/gate/source/twogamma/gps/minphi 0. deg
/gate/source/twogamma/gps/maxphi 360. deg

/gate/source/twogamma/attachTo Source1
```

```
/gate/source/list

#————————————————ooooooOOOOOO000000OOOOOOoooooo————————————————#
#                                                                 #
#                      D E F I N I T I O N   O F                  #
#                    Y O U R   O U T P U T   F I L E              #
#                                                                 #
#————————————————ooooooOOOOOO000000OOOOOOoooooo————————————————#

#=================================================
#  D A T A   O U T P U T
#=================================================

/gate/random/setEngineSeed auto

#/gate/output/root/enable
#/gate/output/root/setFileName                output/test_
#/gate/output/root/setRootNtupleFlag          0
#/gate/output/root/setRootHitFlag             0
#/gate/output/root/setRootSinglesFlag         0
#/gate/output/root/setRootCoincidencesFlag    1
#/gate/output/root/setRootdelayFlag           0

/gate/output/ascii/enable
/gate/output/ascii/setFileName                  output/test_
/gate/output/ascii/setOutFileHitsFlag           0
/gate/output/ascii/setOutFileSinglesFlag        0
/gate/output/ascii/setOutFileEWindowFlag        0
/gate/output/ascii/setOutFileCoincidencesFlag   1
/gate/output/ascii/setCoincidenceMask 0 0 0 0 0 0 1 0 1 1 1 0 0 0 0 0 0 0 0 0 0 0 0 1 1 1 0 0 0 0 0 0

#=================================================
#   M E A S U R E M E N T   S E T T I N G S
#=================================================

#/gate/application/setTotalNumberOfPrimaries 50

/gate/application/setTimeSlice      0.01   s
/gate/application/setTimeStart      0.0    s
```

```
/gate/application/setTimeStop       0.01   s

/gate/application/startDAQ

exit
```

Appendix Zero

GATE Model of Large Modular Array

Listing 2: GATE Script for the validated model of the Large Modular Array PEPT detector. This PEPT detector model is described in Chapter 3. GitHub repository: https://github.com/uob-positron-imaging-centre/GATE_Models.

```
# GATE Model of the Large Modular Array
# Created by Matthew Herald
# Univeristy of Birmingham
# mxh1092@bham.ac.uk
# 2023

#------------------------ooooooOOOOOO000000OOOOOOoooooo------------------------#
#
#          D E F I N E   Y O U R   S C A N N E R   A N D
#
#          Y O U R   C O M P L E T E   S I M U L A T I O N
#
#          F O R   I M A G I N G   A P P L I CA T I O N S
#
#------------------------ooooooOOOOOO000000OOOOOOoooooo------------------------#

#=============================================================
# VISUALISATION
#=============================================================

#/vis/disable
#/vis/open                                   OGL
```

```
#/vis/viewer/reset
#/vis/viewer/set/viewpointThetaPhi    120 30
#/vis/viewer/set/viewpointThetaPhi 30 30
#/vis/viewer/set/viewpointThetaPhi    0 0
#/vis/viewer/zoom                     1
#/vis/viewer/set/style surface
#/vis/drawVolume
#/tracking/storeTrajectory            1
#/vis/scene/add/trajectories
#/vis/scene/endOfEventAction accumulate
#/vis/scene/add/axes
#/vis/viewer/update

#=============================================================
# GEOMETRY
#=============================================================

/gate/geometry/setMaterialDatabase     data/GateMaterials.db

#
#     W O R L D
#
/gate/world/geometry/setXLength        200. cm
/gate/world/geometry/setYLength        200. cm
/gate/world/geometry/setZLength        200. cm
/gate/world/setMaterial                Air
#/gate/world/vis/setColor               white
#/gate/world/vis/forceWireframe
#/gate/world/vis/setVisible             0
#       L A R G E       M O D U L A R   A R R A Y
/gate/world/daughters/name                          Camera
/gate/world/daughters/systemType                    cylindricalPET
/gate/world/daughters/insert                        box
/gate/Camera/placement/setTranslation       0.0 0.0 0.0 cm
/gate/Camera/geometry/setXLength       200. cm
/gate/Camera/geometry/setYLength       200. cm
/gate/Camera/geometry/setZLength       200. cm
/gate/Camera/setMaterial                    Air
#/gate/Camera/vis/setColor                  white
#/gate/Camera/vis/forceWireframe
```

```
#/gate/Camera/vis/setColor white
#/gate/Camera/vis/setVisible                    0

#      M O D U L E
/gate/Camera/daughters/name      module
/gate/Camera/daughters/insert    box
/gate/module/setMaterial Air
/gate/module/geometry/setXLength 355 mm
/gate/module/geometry/setYLength 80 mm
/gate/module/geometry/setZLength 450 mm
/gate/module/placement/setTranslation   0. 0. 0. mm  # 475 mm = 250 mm separation
#/gate/module/vis/forceWireframe
#/gate/module/vis/setColor          red
#/gate/module/vis/setVisible                    1

#   F R A M E   T O P
/gate/module/daughters/name  frametop
/gate/module/daughters/insert    box
/gate/frametop/placement/setTranslation 0 39 0 mm
/gate/frametop/geometry/setXLength  355 mm
/gate/frametop/geometry/setYLength  2 mm
/gate/frametop/geometry/setZLength  450 mm
/gate/frametop/setMaterial  Aluminium
#/gate/frametop/vis/forceWireframe
#/gate/frametop/vis/setColor grey
#/gate/frametop/vis/setVisible 1

#   F R A M E   B O T T O M
/gate/module/daughters/name  framebottom
/gate/module/daughters/insert    box
/gate/framebottom/placement/setTranslation 0 -39 0 mm
/gate/framebottom/geometry/setXLength  355 mm
/gate/framebottom/geometry/setYLength  2 mm
/gate/framebottom/geometry/setZLength  450 mm
/gate/framebottom/setMaterial  Aluminium
#/gate/framebottom/vis/forceWireframe
#/gate/framebottom/vis/setColor grey
#/gate/framebottom/vis/setVisible 1

#   F R A M E   S I D E 1
```

```
/gate/module/daughters/name  frameside1
/gate/module/daughters/insert    box
/gate/frameside1/placement/setTranslation 176.5 0 0 mm
/gate/frameside1/geometry/setXLength  2 mm
/gate/frameside1/geometry/setYLength  74 mm
/gate/frameside1/geometry/setZLength  450 mm
/gate/frameside1/setMaterial  Aluminium
#/gate/frameside1/vis/forceWireframe
#/gate/frameside1/vis/setColor grey
#/gate/frameside1/vis/setVisible 1

#   F R A M E   S I D E 2
/gate/module/daughters/name  frameside2
/gate/module/daughters/insert    box
/gate/frameside2/placement/setTranslation -176.5 0 0 mm
/gate/frameside2/geometry/setXLength  2 mm
/gate/frameside2/geometry/setYLength  74 mm
/gate/frameside2/geometry/setZLength  450 mm
/gate/frameside2/setMaterial  Aluminium
#/gate/frameside2/vis/forceWireframe
#/gate/frameside2/vis/setColor grey
#/gate/frameside2/vis/setVisible 1

#   F R A M E   B A C K
/gate/module/daughters/name  frameback
/gate/module/daughters/insert    box
/gate/frameback/placement/setTranslation 0 0 -224 mm
/gate/frameback/geometry/setXLength  351 mm
/gate/frameback/geometry/setYLength  74 mm
/gate/frameback/geometry/setZLength  2 mm
/gate/frameback/setMaterial  Aluminium
#/gate/frameback/vis/forceWireframe
#/gate/frameback/vis/setColor grey
#/gate/frameback/vis/setVisible 1

#   F R A M E   F R O N T
/gate/module/daughters/name  framefront
/gate/module/daughters/insert    box
/gate/framefront/placement/setTranslation 0 0 224 mm
/gate/framefront/geometry/setXLength  351 mm
```

```
/gate/framefront/geometry/setYLength   74 mm
/gate/framefront/geometry/setZLength   2 mm
/gate/framefront/setMaterial   Aluminium
#/gate/framefront/vis/forceWireframe
#/gate/framefront/vis/setColor grey
#/gate/framefront/vis/setVisible 1

#     W I N D O W
/gate/framefront/daughters/name window
/gate/framefront/daughters/insert box
/gate/window/placement/setTranslation -136.5 0. 0  mm
/gate/window/geometry/setXLength 47 mm
/gate/window/geometry/setYLength 56 mm
/gate/window/geometry/setZLength 2 mm
/gate/window/setMaterial Air
#/gate/window/vis/setColor white
#/gate/window/vis/forceWireframe
#/gate/window/vis/setVisible                  1

#     B L O C K
/gate/module/daughters/name block
/gate/module/daughters/insert box
/gate/block/placement/setTranslation -136.5 0.0 147.5 mm
/gate/block/geometry/setXLength 55 mm
/gate/block/geometry/setYLength 60 mm
/gate/block/geometry/setZLength 155.0 mm
/gate/block/setMaterial Air
#/gate/block/vis/setColor blue
#/gate/block/vis/forceWireframe
#/gate/block/vis/setVisible                  1

#   E L E C T R O N I C S
/gate/module/daughters/name electronics
/gate/module/daughters/insert box
/gate/electronics/placement/setTranslation 0 0 -50 mm
/gate/electronics/geometry/setXLength 300 mm
/gate/electronics/geometry/setYLength 40 mm
/gate/electronics/geometry/setZLength 150 mm
/gate/electronics/setMaterial Elec
#/gate/electronics/vis/setColor yellow
```

```
#/gate/electronics/vis/forceWireframe
#/gate/electronics/vis/setVisible          1

#    L I G H T G U I D E
/gate/block/daughters/name lightguide
/gate/block/daughters/insert box
/gate/lightguide/placement/setTranslation 0 0 41 mm
/gate/lightguide/geometry/setXLength 50 mm
/gate/lightguide/geometry/setYLength 54 mm
/gate/lightguide/geometry/setZLength 5 mm
/gate/lightguide/setMaterial Glass
#/gate/lightguide/vis/setColor white
#/gate/lightguide/vis/forceWireframe
#/gate/lightguide/vis/setVisible      1

#    P M T
/gate/block/daughters/name PMTs
/gate/block/daughters/insert box
/gate/PMTs/placement/setTranslation 0 0 -11.5 mm
/gate/PMTs/geometry/setXLength 50 mm
/gate/PMTs/geometry/setYLength 54 mm
/gate/PMTs/geometry/setZLength 100 mm
/gate/PMTs/setMaterial PMT
#/gate/PMTs/vis/setColor green
#/gate/PMTs/vis/forceWireframe
#/gate/PMTs/vis/setVisible 1

#        C R Y S T A L
/gate/block/daughters/name crystal
/gate/block/daughters/insert box
/gate/crystal/placement/setTranslation 0 0 58.5 mm
/gate/crystal/geometry/setXLength 6.25 mm
/gate/crystal/geometry/setYLength 6.75 mm
/gate/crystal/geometry/setZLength 30.0 mm
/gate/crystal/setMaterial BGO
#/gate/crystal/vis/setColor yellow
#/gate/crystal/vis/setVisible                  1

#        B G O
/gate/crystal/daughters/name BGO
```

```
/gate/crystal/daughters/insert box
/gate/BGO/geometry/setXLength 6.25 mm
/gate/BGO/geometry/setYLength 6.75 mm
/gate/BGO/geometry/setZLength 30.0 mm
/gate/BGO/setMaterial BGO
#/gate/BGO/vis/setColor yellow
#/gate/BGO/vis/setVisible

#       R E P E A T    C R Y S T A L
/gate/crystal/repeaters/insert cubicArray
/gate/crystal/cubicArray/setRepeatNumberX 8
/gate/crystal/cubicArray/setRepeatNumberY 8
/gate/crystal/cubicArray/setRepeatNumberZ 1
/gate/crystal/cubicArray/setRepeatVector 6.25 6.75  0 mm

#       R E P E A T    BLOCK
/gate/block/repeaters/insert linear
/gate/block/linear/setRepeatNumber 4
/gate/block/linear/setRepeatVector 91. 0. 0. mm
/gate/block/linear/autoCenter false

#       R E P E A T    W I N D O W
/gate/window/repeaters/insert linear
/gate/window/linear/setRepeatNumber 4
/gate/window/linear/setRepeatVector 91. 0. 0. mm
/gate/window/linear/autoCenter false


#   R E P E A T    M O D U L E
/gate/module/repeaters/insert genericRepeater
#/gate/module/genericRepeater/useRelativeTranslation 1
/gate/module/genericRepeater/setPlacementsFilename data/count-rate-calibration.placements

#       A T T A C H    C R Y S T A L  SD
/gate/BGO/attachCrystalSD
/gate/systems/list

#   M A K E    A     S Y S T E M
/gate/systems/Camera/rsector/attach module
/gate/systems/Camera/module/attach block
```

```
/gate/systems/Camera/crystal/attach crystal
/gate/systems/Camera/layer0/attach BGO

/gate/geometry/rebuild
#/vis/viewer/update

/gate/physics/addPhysicsList emstandard

#————————————————ooooooOOOOOO000000OOOOOOoooooo————————————————#
#                                                                 #
# D E F I N I T I O N   O F   Y O U R   A C Q U I S I T I O N   #
# D I G I T I Z E R   &   C O I N C I D E N C E   S H O R T E R #
#                                                                 #
#————————————————ooooooOOOOOO000000OOOOOOoooooo————————————————#

#       A D D E R
/gate/digitizer/Singles/insert adder

/gate/distributions/name energy_distrib
/gate/distributions/insert Gaussian
/gate/distributions/energy_distrib/setMean 511 keV
/gate/distributions/energy_distrib/setSigma 1 keV

/gate/distributions/name dt_distrib
/gate/distributions/insert Exponential
/gate/distributions/dt_distrib/setLambda 1.97 us

/gate/digitizer/Singles/insert noise
/gate/digitizer/Singles/noise/setDeltaTDistribution dt_distrib
/gate/digitizer/Singles/noise/setEnergyDistribution energy_distrib

#       R E A D O U T
#/gate/digitizer/Singles/insert readout
#/gate/digitizer/Singles/readout/setDepth        2
/gate/digitizer/Singles/insert readout
/gate/digitizer/Singles/readout/setPolicy TakeEnergyCentroid
/gate/digitizer/Singles/readout/setDepth 4

#       T I M E   R E S O L U T I O N
/gate/digitizer/Singles/insert timeResolution
```

```
/gate/digitizer/Singles/timeResolution/setTimeResolution 13.72 ns

#       P I L E U P
/gate/digitizer/Singles/insert pileup
/gate/digitizer/Singles/pileup/setDepth 2
/gate/digitizer/Singles/pileup/setPileup 637.4 ns

#      E N E R G Y   B L U R R I N G
/gate/digitizer/Singles/insert blurring
/gate/digitizer/Singles/blurring/setResolution 0.5
/gate/digitizer/Singles/blurring/setEnergyOfReference 511. keV

#       E N E R G Y   C U T
/gate/digitizer/Singles/insert                          thresholder
/gate/digitizer/Singles/thresholder/setThreshold        324 keV
/gate/digitizer/Singles/insert                          upholder
/gate/digitizer/Singles/upholder/setUphold              2000 keV

#       D E A D    T I M E
/gate/digitizer/Singles/insert deadtime
/gate/digitizer/Singles/deadtime/setDeadTime 6.63 us
/gate/digitizer/Singles/deadtime/setMode nonparalysable
/gate/digitizer/Singles/deadtime/chooseDTVolume block

#       E N E R G Y     W I N D O W
/gate/digitizer/name                                    EWindow
/gate/digitizer/insert                                  singleChain
/gate/digitizer/EWindow/setInputName                    Singles
# 50 % energy window
/gate/digitizer/EWindow/insert                          thresholder
/gate/digitizer/EWindow/thresholder/setThreshold        400 keV
/gate/digitizer/EWindow/insert                          upholder
/gate/digitizer/EWindow/upholder/setUphold              700 keV

#           C O I N C I D E N C E    S O R T E R
/gate/digitizer/Coincidences/setInputName               EWindow
/gate/digitizer/Coincidences/setWindow                  12 ns
/gate/digitizer/Coincidences/MultiplesPolicy            takeAllGoods
/gate/digitizer/Coincidences/minSectorDifference        1
/gate/digitizer/Coincidences/describe
```

```
/gate/digitizer/name delayed
/gate/digitizer/insert coincidenceSorter
/gate/digitizer/delayed/setInputName EWindow
/gate/digitizer/delayed/setOffset 256 ns
/gate/digitizer/delayed/setWindow 12 ns
/gate/digitizer/delayed/MultiplesPolicy          takeAllGoods
/gate/digitizer/delayed/minSectorDifference 1
/gate/digitizer/delayed/describe

/gate/digitizer/name finalCoinc
/gate/digitizer/insert coincidenceChain
/gate/digitizer/finalCoinc/addInputName delayed
/gate/digitizer/finalCoinc/addInputName Coincidences
#/gate/digitizer/finalCoinc/usePriority False
/gate/digitizer/finalCoinc/describe

###  P H A N T O M S ###
/gate/world/daughters/name                       phantom
/gate/world/daughters/insert                     sphere
/gate/phantom/setMaterial                        Plastic
#/gate/phantom/vis/forceWireframe
#/gate/phantom/vis/setColor                       white
#/gate/phantom/vis/setVisible                     0

/gate/phantom/geometry/setRmin                   0 mm
/gate/phantom/geometry/setRmax                   10 mm
/gate/phantom/geometry/setPhiStart               0 deg
/gate/phantom/geometry/setDeltaPhi               360 deg
/gate/phantom/geometry/setThetaStart             0 deg
/gate/phantom/geometry/setDeltaTheta             360 deg
/gate/phantom/placement/setTranslation           0 0 0 mm
/gate/phantom/moves/insert                       translation
/gate/phantom/translation/setSpeed               0 0 0 mm/s
#/gate/phantom/attachPhantomSD
/gate/geometry/rebuild
#/vis/viewer/update

### S O U R C E     D E F I N I T I O N ###
/gate/phantom/daughters/name                       Source1
```

```
/gate/phantom/daughters/insert                    sphere
/gate/Source1/setMaterial                 Plastic
#/gate/Source1/vis/forceWireframe
#/gate/Source1/vis/setColor                white
#/gate/Source1/vis/setVisible              1

/gate/Source1/geometry/setRmin            0 mm
/gate/Source1/geometry/setRmax            0.1 mm
/gate/Source1/geometry/setPhiStart        0 deg
/gate/Source1/geometry/setDeltaPhi        360 deg
/gate/Source1/geometry/setThetaStart      0 deg
/gate/Source1/geometry/setDeltaTheta      360 deg
/gate/Source1/placement/setTranslation    0 0 0 mm

#/gate/Source1/attachPhantomSD
/gate/geometry/rebuild
#/vis/viewer/update

#============================================================
# INITIALISATION
#============================================================

/gate/run/initialize

### S O U R C E     D E F I N I T I O N ###

# S O U R C E
/gate/source/addSource twogamma
/gate/source/twogamma/setActivity 10 MBq
/gate/source/twogamma/setType backtoback

# POSITION
/gate/source/twogamma/gps/centre 0. 0. 0. mm

# PARTICLE
/gate/source/twogamma/gps/particle gamma
/gate/source/twogamma/gps/energytype Mono
/gate/source/twogamma/gps/monoenergy 0.511 MeV

# TYPE = Volume or Surface or Point
```

```
/gate/source/twogamma/gps/type Volume
/gate/source/twogamma/gps/shape Sphere
/gate/source/twogamma/gps/radius 0.1 mm
/gate/source/twogamma/gps/centre 0.0 0.0 0.0 mm

# SET THE ANGULAR DISTRIBUTION OF EMISSION
/gate/source/twogamma/gps/angtype iso

# SET MIN AND MAX EMISSION ANGLES
/gate/source/twogamma/gps/mintheta 0. deg
/gate/source/twogamma/gps/maxtheta 180. deg
/gate/source/twogamma/gps/minphi    0. deg
/gate/source/twogamma/gps/maxphi 360. deg

/gate/source/twogamma/attachTo      Source1

/gate/source/list

#--------------------ooooooOOOOOO000000OOOOOOoooooo--------------------#
#                                                                       #
#                    D E F I N I T I O N   O F                          #
#                 Y O U R   O U T P U T   F I L E                       #
#                                                                       #
#--------------------ooooooOOOOOO000000OOOOOOoooooo--------------------#

#=======================================================
# D A T A   O U T P U T
#=======================================================

/gate/random/setEngineSeed auto

#/gate/output/root/enable
#/gate/output/root/setFileName                output/test_
#/gate/output/root/setRootNtupleFlag          0
#/gate/output/root/setRootHitFlag             0
#/gate/output/root/setRootSinglesFlag         0
#/gate/output/root/setRootCoincidencesFlag    1
#/gate/output/root/setRootdelayFlag           0

/gate/output/ascii/enable
```

```
/gate/output/ascii/setFileName                    output/test_
/gate/output/ascii/setOutFileHitsFlag             0
/gate/output/ascii/setOutFileSinglesFlag          0
/gate/output/ascii/setOutFileEWindowFlag          0
/gate/output/ascii/setOutFileCoincidencesFlag     1
/gate/output/ascii/setCoincidenceMask 0 0 0 0 0 0 1 0 1 1 1 0 0 0 0 0 0 0 0 0 0 0 1 1 1 0 0 0 0 0 0

#=====================================================
#   M E A S U R E M E N T   S E T T I N G S
#=====================================================

#/gate/application/setTotalNumberOfPrimaries 50

/gate/application/setTimeSlice      0.01   s
/gate/application/setTimeStart      0.0    s
/gate/application/setTimeStop       0.01   s

/gate/application/startDAQ

exit
```

Appendix Zero

Functions for Comparing GATE Prescribed and PEPT Reconstructed Tracer Locations

Listing 3: Comparsison functions written in Python to compare trajectories between GATE prescribed locations and PEPT reconstructed locations. These functions are used in Chapter 4 to benchmark the performance of PEPT algorithms on a GATE-simulated data set. GitHub repository: https://github.com/mxh1092/RoPP-Comparison-Functions.

```
import numpy as np
from scipy import interpolate

def checkDims(exp):

    if exp.ndim < 2:
        exp = exp

    else:
        xmean = np.mean(exp[:, 0])
        ymean = np.mean(exp[:, 1])
        zmean = np.mean(exp[:, 2])

        exp = np.array([xmean, ymean, zmean])
```

```
    return exp

def stationaryCompare(exp):

    """
    Compares the positions returned by PEPT
    algorithms from a single particle tracking
    test. A single stationary
    tracer is placed in the field-of-view. Exp is
    an (N, 3) np.array of x, y, z, positions. The
    position of the real
    position for comparison may be left as is,
    edited to a new position, or changed to read
    in a new position from a
    file.
    """

    #real = np.loadtxt('Keys/Key_Stationary.csv')

    # Real position of the tracer
    x_real = 0
    y_real = 10
    z_real = 0

    # If multiple positions are given, compute the mean
    exp = checkDims(exp)

    # Break exp into individual positional components
    x_exp = exp[0]
    y_exp = exp[1]
    z_exp = exp[2]

    # Compute individual error components
    x_error = abs(x_real - x_exp)
    y_error = abs(y_real - y_exp)
    z_error = abs(z_real - z_exp)

    # Compute 3D error
    error = np.sqrt(x_error ** 2 + y_error ** 2 + z_error ** 2)
```

```
    return error


def scatterCompare(exp):


    """
    Compares the positions returned by PEPT
    algorithms from a single particle tracking
    test. A single stationary
    tracer is placed in the field-of-view and
    placed inside a sphere of material to induce
    scattering. Exp is a (N, 3)
    np.array of x, y, z, positions. The position
    of the real position for comparison may be
    left as is, edited to a new
    positions, or changed to read in a new
    position from a file.
    """
    #real = np.loadtxt('Keys/Key_Scatter.csv')

    # Real particle position
    x_real = 0
    y_real = 20
    z_real = 0

    # If multiple positions are given, compute the mean
    exp = checkDims(exp)

    # Break exp in to individual positional components
    x_exp = exp[0]
    y_exp = exp[1]
    z_exp = exp[2]

    # Compute individual error components
    x_error = abs(x_real - x_exp)
    y_error = abs(y_real - y_exp)
    z_error = abs(z_real - z_exp)
```

```
    # Compute 3D error
    error = np.sqrt(x_error ** 2 + y_error ** 2 + z_error ** 2)

    return error


def fieldOfViewCompare(exp):

    """
    Compares the positions returned by PEPT
    algorithms from a single particle tracking
    test. A single tracer enters one
    side of the field-of-view then exits out the
    other. Exp is an (N, 4) np.array of t, x, y, z,
    positions. The position
    of the real tracer is described by an equation
    as a function of time. The position to compare
    is computed for the
    timestep associated with the PEPT detection.
    """

    # Read in PEPT detected trajectory
    t_exp = exp[:, 0]
    x_exp = exp[:, 1]
    y_exp = exp[:, 2]
    z_exp = exp[:, 3]

    # A tracer moves in a straight line at a velocity of 250 mm/s in the z direction
    v = 250  # mm/s
    x_real = 0
    y_real = 0
    z_real = -250 + (t_exp / 1000) * v

    # Computer individual error components
    x_error = abs(x_real - x_exp)
    y_error = abs(y_real - y_exp)
    z_error = abs(z_real - z_exp)

    # Compute instantaneous errors
    errors = np.sqrt(x_error**2+y_error**2+z_error**2)
```

```
    # Compute mean error
    error = np.mean(errors)

    # Compute standard deviation of the individual error components
    x_std = np.std(x_error)
    y_std = np.std(y_error)
    z_std = np.std(z_error)

    # Compute overall standard deviation of the 3D particle position error
    std = np.mean(np.sqrt(x_std ** 2 + y_std ** 2 + z_std ** 2))

    return error, std, errors


def velocityCompare(exp, real):

    """
    Compares the positions returned by PEPT
    algorithms from a single particle tracking
    test. A single tracer moves
    between two positions at constant velocity.
    Exp is a (N, 4) np.array of t, x, y, z,
    positions. The position of the
    real tracer is read in from a file and passed
    to this function for each of the different
    tests. The file contains a
    series of timesteps and positions for the
    prescribed movements. To compare to the
    associated PEPT detections, the
    PEPT detected timestep is used as the basis
    for linearly interpolating to the expected
    real position.
    """

    # Break up exp into time and position components
    t_exp = exp[:, 0]
    x_exp = exp[:, 1]
    y_exp = exp[:, 2]
    z_exp = exp[:, 3]
```

```
    # Break of real into individual time an position components
    t_real = real[:, 0]
    x_real = real[:, 1]
    y_real = real[:, 2]
    z_real = real[:, 3]

    # Create linear interpolation functions
    fx = interpolate.interp1d(t_real, x_real)
    fy = interpolate.interp1d(t_real, y_real)
    fz = interpolate.interp1d(t_real, z_real)

    # Use the PEPT detected timesteps to calculate the real particle position
    x_real = fx(t_exp)
    y_real = fy(t_exp)
    z_real = fz(t_exp)

    # Compute instantaneous error for the individual position components
    x_error = abs(x_real - x_exp)
    y_error = abs(y_real - y_exp)
    z_error = abs(z_real - z_exp)

    # Compute the mean 3D error
    error = np.mean(np.sqrt(x_error ** 2 + y_error ** 2 + z_error ** 2))

    # Compute the 1D standard deviation
    x_std = np.std(x_error)
    y_std = np.std(y_error)
    z_std = np.std(z_error)

    # Compute the 3D standard deviation
    std = np.mean(np.sqrt(x_std ** 2 + y_std ** 2 + z_std ** 2))

    return error, std


def seperationCompare(exp1, exp2, real):

    """
    Compares the positions returned by PEPT
```

```
    algorithms from a two-particle tracking test.
    Two tracers are separated by
    small distances, testing the ability to
    differentiate increasing small tracer
    separations. Exp is a (N, 3) np.array
    x, y, z, positions. The position of the real
    tracer is passed to this function for each of
    the different tests. The
    centre-to-centre separations used in this test
    are 2, 3, 4, 6, 8, 10, 12, 16, 20, and 30 mm.
    Tests are conducted in
    both the transaxial (x) and axial (z)
    orientations of the detector, with the other
    position components left as 0 mm.
    """

    # Particle 1 position components
    x1_exp = exp1[:, 0]
    y1_exp = exp1[:, 1]
    z1_exp = exp1[:, 2]

    # Particle 2 position components
    x2_exp = exp2[:, 0]
    y2_exp = exp2[:, 1]
    z2_exp = exp2[:, 2]

    # Real particle 1 position components
    x1_real = -real[0]
    y1_real = -real[1]
    z1_real = -real[2]

    # Real particle 2 position components
    x2_real = real[0]
    y2_real = real[1]
    z2_real = real[2]

    # Compute the instantaneous 1D error for particle 1
    x1_error = abs(x1_real - x1_exp)
    y1_error = abs(y1_real - y1_exp)
    z1_error = abs(z1_real - z1_exp)
```

```
    # Compute mean 3D error for particle 1
    error1 = np.mean(np.sqrt(x1_error ** 2 + y1_error ** 2 + z1_error ** 2))

    # Compute 1D standard deviation for particle 1 positions
    x1_std = np.std(x1_error)
    y1_std = np.std(y1_error)
    z1_std = np.std(z1_error)

    # Compute 3D standard deviation for particle 1
    std1 = np.sqrt(x1_std ** 2 + y1_std ** 2 + z1_std ** 2)

    # Compute the instantaneous 1D error for particle 2
    x2_error = abs(x2_real - x2_exp)
    y2_error = abs(y2_real - y2_exp)
    z2_error = abs(z2_real - z2_exp)

    # Compute mean 3D error for particle 2
    error2 = np.mean(np.sqrt(x2_error ** 2 + y2_error ** 2 + z2_error ** 2))

    # Compute 1D standard deviation for particle 2 positions
    x2_std = np.std(x2_error)
    y2_std = np.std(y2_error)
    z2_std = np.std(z2_error)

    # Compute 3D standard deviation for particle 2
    std2 = np.mean(np.sqrt(x2_std ** 2 + y2_std ** 2 + z2_std ** 2))

    # Mean error for both particles
    error = (error1 + error2) / 2

    # Mean standard deviation for both particles
    std = (std1 + std2) / 2

    return error, std

def linkingCompare(exp, real):

    """
    Compares the positions returned by PEPT
```

```
    algorithms from a three-particle tracking
    test. Three tracers are separated
    by a constant separation and continuously
    tracked as they move about the surface of the
    sphere, testing the ability
    link positions into trajectories for multiple
    particles simultaneously. Exp is a (N, 4)
    np.array t, x, y, z,
    positions. The position of the real tracer is
    passed to this function for each of the
    different tests. The PEPT
    trajectories must be correctly associated with
    the right real trajectory before passing to
    this function.
    """

    # Break up exp into individual time and position components
    t_exp = exp[:, 0]
    x_exp = exp[:, 1]
    y_exp = exp[:, 2]
    z_exp = exp[:, 3]

    # Break up the real trajectory into individual time and position components
    t_real = real[:, 0]
    x_real = real[:, 1]
    y_real = real[:, 2]
    z_real = real[:, 3]

    # Create interpolation functions for the real trajectory
    fx = interpolate.interp1d(t_real, x_real)
    fy = interpolate.interp1d(t_real, y_real)
    fz = interpolate.interp1d(t_real, z_real)

    # Calculate the expected real tracer position for each detected time
    x_real = fx(t_exp)
    y_real = fy(t_exp)
    z_real = fz(t_exp)

    # Compute the instantaneous 1D errors
    x_error = abs(x_real - x_exp)
```

```
    y_error = abs(y_real - y_exp)
    z_error = abs(z_real - z_exp)

    # Compute the mean 3D error
    error = np.mean(np.sqrt(x_error ** 2 + y_error ** 2 + z_error ** 2))

    # Compute the 1D standard deviation
    x_std = np.std(x_error)
    y_std = np.std(y_error)
    z_std = np.std(z_error)

    # Compute the 3D standard deviation
    std = np.sqrt(x_std ** 2 + y_std ** 2 + z_std ** 2)

    # Record how many timesteps are given for the PEPT detected trajectories
    time = len(t_exp)

    return error, std, time


def multipleCompare(exp, real):

    """
    Compares the positions returned by PEPT
    algorithms from a multiple particle tracking
    test. A random number of
    static tracers are placed with the detection
    volume of the detectors. An array of the PEPT-detected positions for
    each test are compared to the real position. A
    tracer is considered found if it is the
    nearest tracer to the PEPT
    returned positions. A list of the found
    tracers is scanned for unique IDs and returned
    as the number of particles
    found. Additionally, the mean 3D error and the
    3D standard deviation is returned. This is
    meant to test the
    ability to track high numbers of particles
    simultaneously. Exp is an (N, 3) np.array x, y,
    z, positions. The position
```

```
    of the real tracers is passed to this function
    for each of the different tests.
    """

    # Break exp into position components
    x_exp = exp[:, 0]
    y_exp = exp[:, 1]
    z_exp = exp[:, 2]

    # Break up real into position components
    x_real = real[:, 0]
    y_real = real[:, 1]
    z_real = real[:, 2]

    # Create mean 3D error array
    errors_mean = []
    # Allocate space for instantaneous 3D errors
    errors = np.zeros(len(x_exp))
    # Create a 3D standard deviation array
    stds = []
    # Allocate space for the nearest particle IDs
    IDs = np.zeros(len(x_exp))

    # Loop over all detected positions
    for j in range(len(x_exp)):

        # Allocate space for instananeous errors over all real particles
        errors_ind = np.zeros(len(x_real))

        # Loop over all real particles
        for k in range(len(x_real)):

            # Compute 1D error between exp and all real particles
            x_error = abs(x_exp[j] - x_real[k])
            y_error = abs(y_exp[j] - y_real[k])
            z_error = abs(z_exp[j] - z_real[k])

            # Compute 3D error between exp and all real particles
            error = np.sqrt(x_error ** 2 + y_error ** 2 + z_error ** 2)
```

```
            # Fill in instanteous error array
            errors_ind[k] = error

        # Find the closest real particle and record the error
        IDs[j] = np.argmin(errors_ind)
        errors[j] = np.min(errors_ind)

    # The length of unique particle IDs is the number of particles found
    N_unique = len(np.unique(IDs))

    # Record mean paritcle 3D error
    errors_mean.append(np.mean(errors))

    # Record the 3D standard deviation
    stds.append(np.std(errors))

    return errors_mean, stds, N_unique


def closepackedCompare(exp, real):

    """
    Compares the positions returned by PEPT
    algorithms from a multiple particle tracking
    test. Arrangements of static
    tracers are placed with the detection volume
    of the detectors. An array of the PEPT-
    detected positions for
    each test are compared to the real position. A
    tracer is considered found if it is the
    nearest tracer to the PEPT
    returned positions. A list of the found
    tracers is scanned for unique IDs and returned
    as the number of particles
    found. Additionally, the mean 3D error and the
    3D standard deviation is returned. This is
    meant to test the
    ability to track high numbers of particles in
    close proximity to each other with high
    degrees of symmetry
```

```
    simultaneously. Exp is a (N, 3) np.array x, y,
    z, positions. The position of the real tracers
    is passed to this
    function for each of the different tests.
    """

    # Break exp into components
    x_exp = exp[:, 0]
    y_exp = exp[:, 1]
    z_exp = exp[:, 2]

    # Break real into components
    x_real = real[:, 0]
    y_real = real[:, 1]
    z_real = real[:, 2]

    # Create lists for the overall mean 3D errors and 3D standard deviations for each test
    errors = []
    stds = []

    # Loop over all real positions
    for k in range(len(x_real)):

        # Compute the 1D error compared to each real position
        x_error = abs(x_exp - x_real[k])
        y_error = abs(y_exp - y_real[k])
        z_error = abs(z_exp - z_real[k])

        # Compute the 3D error
        error = np.mean(np.sqrt(x_error ** 2 + y_error ** 2 + z_error ** 2))

        # Compute the 1D standard deviations
        x_std = np.std(x_error)
        y_std = np.std(y_error)
        z_std = np.std(z_error)

        # Compute the 3D standard deviation
        std = np.mean(np.sqrt(x_std ** 2 + y_std ** 2 + z_std ** 2))

        # Append to previous arrays
```

```
        errors.append(error)
        stds.append(std)

    # Find the closest real particle
    error = np.min(errors)
    std = stds[np.where(errors==np.min(errors))[0][0]]

    return error, std
```

Appendix Zero

LIGGGHTS Script for DEM Simulation of the GranuTools GranuDrum

Listing 4: LIGGGHTS script for the DEM simulation of the GranuTools GranuDrum rotating at 45 RPM and with monodisperse spheres of microcrystalline cellulose. The results of this simulation are described in Chapter 5.

```
# LIGGGHTS simulation of the GranuDrum rotating at 45 RPM filled with monodisperse MCC spheres
processors  * 2 *
variable timestep equal 1e-5
variable rotationPeriod equal 60/45     #s per revolution
variable N equal 10000                          #Particle number

variable fricPW      equal 0.319687387                  # friction particle wall
variable fricPSW     equal 0.319687387                  # friction particle sidewall
variable cohPW       equal 0
variable fric        equal 1.0          # friction value, does not matter it defines the wall-wall frictio

variable corPP       equal 0.3                  #Coeff of restitution Particle particle
variable fricPP      equal 0.319687387                  # friction particle particle
variable fricRoll    equal 0.00247952321            #Rolling friction
variable youngmodP   equal 5e6
variable poissP      equal 0.3
variable dens        equal 1580.0
variable cohPP       equal 0
```

```
#Radii for the Particles
variable r1          equal 0.0018/2

#MAIN CODE rotating drum

atom_style       granular
atom_modify      map array

boundary             f f f           #non-periodic
newton               off
#Sets newtons 3rd law on or off (Effect on computation time due to less computation but higher communi
communicate      single vel yes
units                si

region       domain block -0.048 0.048 -0.002 0.0221 -0.048 0.048 units box #creates a region
create_box       3 domain
#read_restart    base.start

# creates the neighbor-list (due to only contact models this is very short
neighbor             0.0018 bin
neigh_modify     delay 0

##############Particle/Wall properties#################

fix  m1 all property/global youngsModulus peratomtype ${youngmodP} ${youngmodP} ${youngmodP}
fix  m2 all property/global poissonsRatio peratomtype ${poissP} ${poissP} ${poissP}
fix  m3 all property/global coefficientRestitution peratomtypepair 3 &
${corPP}      ${corPP}      ${corPP}  &
${corPP}      ${corPP}      ${corPP}  &
${corPP}      ${corPP}      ${corPP}  &

fix  m4 all property/global coefficientFriction peratomtypepair 3 &
${fricPP}           ${fricPW}           ${fricPSW}  &
${fricPW}           ${fric}             ${fric}     &
${fricPSW}          ${fric}             ${fric}     &

fix  m6 all property/global cohesionEnergyDensity peratomtypepair 3 &
 ${cohPP}        ${cohPW}        ${cohPW}    &
 ${cohPW}        ${cohPW}        ${cohPW}    &
 ${cohPW}        ${cohPW}        ${cohPW}    &
```

```
fix  m7 all property/global coefficientRollingFriction peratomtypepair 3 &
${fricRoll}         ${fricRoll}         ${fricRoll}    &
${fricRoll}         ${fricRoll}         ${fricRoll}    &
${fricRoll}         ${fricRoll}         ${fricRoll}    &

fix  m5 all property/global k_finnie peratomtypepair 3  1.0 1.0 1.0 1.0 1.0 1.0 1.0 1.0 1.0

#New pair style
pair_style  gran model hertz tangential history rolling_friction cdt
pair_coeff  * *

timestep    ${timestep}

fix  gravi all gravity 9.81 vector 0.0 0.0 -1.0

#the Drum/faces
fix  cad all mesh/surface/stress file mesh/rotating_drum_mm.stl type 2 scale 0.001 wear finnie
fix  side all mesh/surface/stress file mesh/rotating_drum_side_mm.stl type 3 scale 0.001
fix  inface all mesh/surface file mesh/inface_mm.stl type 3 scale 0.001
fix  granwalls all wall/gran model hertz tangential history rolling_friction cdt  mesh &
    n_meshes  2 meshes cad side

#distributions for insertion
fix  pts1 all particletemplate/sphere 15485863 atom_type 1 density constant ${dens} radius &
    constant ${r1}

fix  pdd1 all particledistribution/discrete/numberbased  32452843 1  pts1 1.0

#no need for particle insertion as the system is read in by restart file
#particle insertion
fix    ins all insert/stream seed 32452867 distributiontemplate pdd1 &
       nparticles ${N} particlerate 1000000 overlapcheck yes all_in no vel constant &
       0.0 0.0 -0.1 insertion_face inface extrude_length 0.03

#apply nve integration to all particles that are inserted as single particles
fix    integr all nve/sphere

#output settings, include total thermal energy
```

```
thermo          1000

thermo_style  custom step atoms ke vol time
thermo_modify lost ignore norm no

#insert the first particles so that dump is not empty
run     1000
dump    dmp all custom/vtk 1000 post/drum_*.vtk id type type x y z ix iy iz vx vy vz fx fy &
    fz omegax omegay omegaz radius
# save every milisecond

modify_timing on

fix move all move/mesh mesh cad rotate origin 0. 0. 0. axis  0. 1. 0. period ${rotationPeriod}

reset_timestep 0
restart 100000 restart.drum restart.drum

run     180000000 # 30 minutes
```

Appendix Zero

First Author Journal Papers

.1 Monte Carlo model validation of a detector system used for positron emission particle tracking

Nuclear Inst. and Methods in Physics Research, A 993 (2021) 165073

Contents lists available at ScienceDirect

Nuclear Inst. and Methods in Physics Research, A

journal homepage: www.elsevier.com/locate/nima

Monte Carlo model validation of a detector system used for Positron Emission Particle Tracking

Matthew Herald [a,*], Tzany Wheldon [b,c], Christopher Windows-Yule [a,c]

[a] School of Chemical Engineering, University of Birmingham, Birmingham, United Kingdom
[b] School of Physics and Astronomy, University of Birmingham, Birmingham, United Kingdom
[c] Positron Imaging Centre, University of Birmingham, Birmingham, United Kingdom

ARTICLE INFO

Keywords:
PEPT
PET
Monte Carlo
Experimental validation
GATE
Radiation detection

ABSTRACT

The spatiotemporal resolution of Lagrangian particle trajectories captured using Positron Emission Particle Tracking (PEPT) is difficult to predict prior to experimentation, since this relies on the detector systems, source activity distribution, and experimental apparatus. However, understanding the limitations of an experiment is crucial to quantifying error and ensuring that the captured trajectories reveal phenomena of interest in enough detail for meaningful analysis. These factors are especially important in PEPT experiments since this technique is applied to image opaque systems lacking optical access for complementary measurement techniques, such as Particle Image Velocimetry. Using the Monte Carlo simulator Geant4 Application for Tomographic Emission (GATE), a computational model of the ADAC/Phillips Forte, a detector system used at the Positron Imaging Centre (PIC) for PEPT studies, is created and validated against experiments testing the spatial resolution, sensitivity, scatter fraction, and count-rates following National Electronic Manufactures Association standards. In this work, fluorine-18 sources and experimental geometries are recreated, generating synthetic data analogous to experimentally acquired data. Over all experiments and activities tested, this GATE model reports agreement to within 1%–10% of experiments. In the future, this model is expected to be used by the PIC to conduct feasibility studies of potential experiments. Further, optimization of experiments can now be conducted without expending the considerable time and resources required for physical experimentation, representing a major improvement of the PIC's PEPT modeling capabilities.

1. Introduction

Positron Emission Particle Tracking (PEPT) is a technique similar to Positron Emission Tomography (PET) that has been developed to study opaque systems such as rotating drums, fluidized beds, and other industrial equipment without optical access [1]. A positron-emitting substance is bound chemically or physically to a discrete particle, then tracked as a Lagrangian tracer as it moves within a system [2,3]. This technique relies on detecting pairs of annihilation photons emitted by the tracer particle, using these Lines-of-Response (LORs) to successively calculate the position of the tracer [4]. Analysis of trajectories reveals the velocity flow fields, particle residence times, recirculation periods and other time-averaged behavior used in engineering analysis of industrial equipment [5–8] (see Fig. 1).

Using Monte Carlo simulators such as Geant4 Application for Tomographic Emission (GATE), it is possible to model the PET systems on which PEPT experiments are conducted, generating synthetic data analogous to real LORs [9,10] [11]. GATE modeled PET systems such as the Siemens Inveon, Philips Allegro, and the General Electric Discovery report agreement with validation experiments to within 1%–10% in spatial resolution, sensitivity, count-rates, and scatter fraction [12–14]. Additionally, several single-head gamma-cameras such as the Philips AXIS, GEMS DST Xli, and Millennium VG Hawk-Eye have been modeled and validated using GATE [15,16]. These models are useful for iterating through variations of experiments for optimization or testing image reconstruction algorithms without expending the considerable time and resources required for physical experimentation.

In this work, GATE is used to model the ADAC/Phillps Forte, a PET system used for PEPT research at the University of Birmingham Positron Imaging Centre (PIC). The Monte Carlo model is then validated by experiments. While this is an older PET system, no validated GATE models exist to the authors knowledge. Previous to this work, the performance of real experiments has been difficult to predict since there is often a complex experimental geometry, a moving source within the FOV, and a nonlinear response of the detector to the source activity and energy. For future PEPT experiments at the PIC, this model is expected to be used for feasibility studies of proposed PEPT experiments by providing an estimate of the Forte output. Recently, an approach for

* Corresponding author.
E-mail address: mxh1092@student.bham.ac.uk (M. Herald).

https://doi.org/10.1016/j.nima.2021.165073
Received 4 May 2020; Received in revised form 4 January 2021; Accepted 16 January 2021
Available online 21 January 2021
0168-9002/© 2021 Elsevier B.V. All rights reserved.

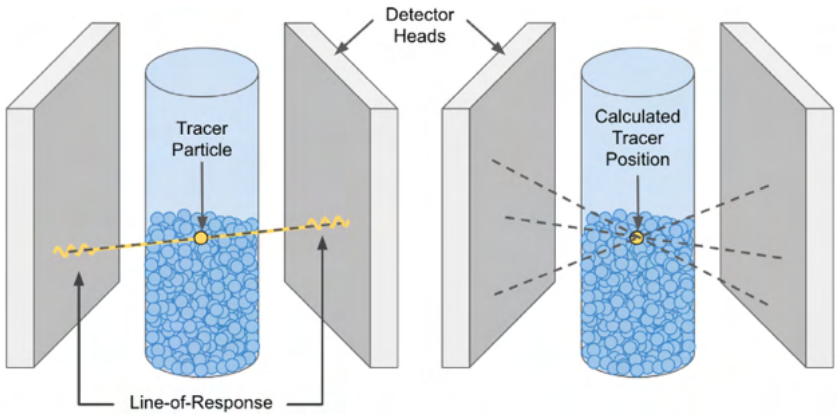

Fig. 1. PEPT locates a positron-emitting tracer in a dense granular system by calculating the intersection of LORs, generated by detecting pairs of annihilation photons.

Fig. 2. The Forte can accommodate large experiments for investigation with PEPT due to its dual-headed geometry. The coordinate system in this work begins in the center of the FOV.

creating virtual PEPT experiments in GATE has been used which integrates tracer trajectory data from simulations with pre-existing GATE models of PET systems [17]. GATE simulations coupled to simulated trajectories produced by multiphysics modeling such as Computational Fluid Dynamics (CFD) or the Discrete Element Method (DEM) may prove especially useful, representing a major improvement in the PIC's PEPT modeling capabilities. The development of an accurate simulation model of the ADAC Forte camera will also provide an invaluable testing ground for the development of new PEPT algorithms (see Fig. 2).

2. Methods

For validating GATE models, a number of approaches have been used in the past, generally relying on standards published by the National Electronics Manufacturers Association (NEMA) [18]. The NEMA standards describe methods for calculating the true, scattered, and random coincidences count-rates, as well as the spatial resolution and sensitivity of PET systems using standard procedures and phantoms. The performance of PET systems following NEMA standards are usually provided by the manufacturer, providing a way for customers to compare the performance of systems in such a way as to not bias one system over the other, providing a guaranteed level of performance.

In this work, experiments testing the spatial resolution, sensitivity, scatter fraction, and count-rates of the Forte are conducted following the NEMA standards for data analysis. All experimental phantoms are replicated in GATE simulations to provide a basis for validation. The Forte has also been characterized in other work to demonstrate improved capabilities, replacing a predecessor system used for PEPT studies at the PIC [19]. This previous characterization provides an important comparison for the results of this work.

Table 1
Positions of the spatial resolution tracer in the Forte FOV.

X-Position (mm)	Y-Position (mm)	Z-Position (mm)
0	0	0
0	127.5	0
0	0	95
150	0	0
150	127.5	0
150	0	95

2.1. Spatial resolution

Spatial resolution is the ability to distinguish the position of a source within the system field of view (FOV). It is reported as the full-width half-maximum (FWHM) of a source projection created by a back-projection of LORs taken at the source's position. In this work, the axial and transaxial spatial resolution is reported for a small spherical source in air measured in the FOV at locations prescribed by the NEMA protocol. For a coordinate system starting at the central point between the two detector heads separated by 600 mm, the positions of the source used in this experiment are recorded in Table 1.

When calculating the FWHM, 1,000,000 LORs are used to reduce statistical uncertainty. The FWHM of each acquisition is determined by linear interpolation of the projection profile created by single-slice rebinning of the LORs using 1 mm bins.

$$Res_{Transaxial_{Center}} = (Res_{y_{x=0,y=0,z=0}} + Res_{z_{x=0,y=0,z=0}} + Res_{y_{x=0,y=0,z=95}} + Res_{z_{x=0,y=0,z=95}})/4 \quad (1)$$

$$Res_{Transaxial_{1/2FOV}} = (Res_{y_{x=150,y=127.5,z=0}} + Res_{z_{x=150,y=127.5,z=0}} + Res_{y_{x=0,y=127.5,z=0}} + Res_{z_{x=0,y=127.5,z=0}})/4 \quad (2)$$

$$Res_{Axial_{Center}} = (Res_{x_{x=0,y=0,z=0}} + Res_{x_{x=0,y=0,z=95}})/2 \quad (3)$$

$$Res_{Axial_{1/2FOV}} = (Res_{x_{x=150,y=0,z=0}} + Res_{x_{x=150,y=0,z=95}})/2 \quad (4)$$

The phantom and source used in this experiment is a 0.2 mm diameter anionic exchange resin bead. Anionic exchange beads selectively absorb fluorine-18 ions from a solution of water and fluorine-18, volumetrically activating the bead with a high activity concentration [20]. After soaking in the solution for approximately 30 min, a resin bead was extracted and its activity measured with a well-counter to be

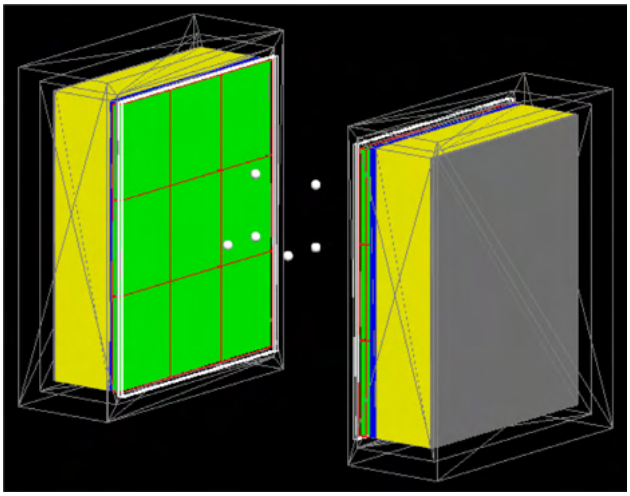

Fig. 3. Locations of the spatial resolution phantom positions, with exaggerated phantom size.

approximately 27.8 MBq ± 0.1. The resin bead was attached to the end of a 1.5 mm diameter stainless-steel syringe using a UV activated epoxy. The syringe was taped to a 200 mm long and 5 mm diameter wooden dowel, then placed in the FOV of the detectors.

A GATE model of the spatial resolution phantom and source has been designed to recreate the experimental geometry. The exact composition of the resin bead is unknown but was modeled as a 0.2 mm diameter spherical plastic source, encapsulated in a spherical plastic phantom 2 mm in diameter. This outer sphere represents the UV activated epoxy and provides material for positrons to annihilate near the source. By the end of the experiments, the source activity decayed to approximately 19 MBq (see Fig. 3).

The source projections are created using the 'Voxels' base-class of the PEPT Library, an open-source Python library of functions useful for analyzing PEPT data [21]. Voxels are the 3D equivalent of 2D pixels. Similar to an image, voxels can be arranged into 3D arrays. In the context of PEPT experiment, the detected LORs are converted into voxel images by finding and increasing the scalar value of the voxels which are crossed by the line. This process can be repeated with successive LORs, creating a array corresponding to the density of LORs and can be visualized by looking at a 2D slice of the 3D voxel array. The line-density is the back-projected image of the source used to fit the FWHM (see Fig. 4).

2.2. Sensitivity

Sensitivity is the rate of detected coincidences compared to the rate of positron-annihilation produced by a source. Geometric and intrinsic factors such as the detector and source positions as well the scintillation crystal material determine the sensitivity of PET systems, however, at high count-rates, the dead-time from electronic pulse processing can restrict the number of events able to be recorded. For low activity sources the rate of detected coincidences, R, is approximately the source activity, A, multiplied by the geometric efficiency, ε_g, and multiplied by the square of the intrinsic efficiency, ε_i, since both coincident photons must be detected. The intrinsic efficiency should remain relatively constant for these experiments but has some dependency on the head separation and source position due to the depth of photon interaction [22].

$$R = A\varepsilon_g\varepsilon_i^2 \tag{5}$$

For a centrally located point source in between the two detector heads separated by a distance S, the geometric efficiency can be defined as the solid angle between the point source and limits of the detector, length L and height H. The geometric efficiency of the Forte was studied in previous work and an analytical expression derived to calculate the geometric efficiency for an ideal point source anywhere in the FOV [23]. This expression transforms the Cartesian coordinates into spherical coordinates and integrates across the zenith, ϕ, and azimuth, θ. Using trigonometric identities, the geometric efficiency can be expressed solely by integrating the azimuth angle.

$$d\varepsilon_g = \frac{\sin\phi d\theta d\phi}{2\pi} \tag{6}$$

$$\theta_{min} = \arctan\left(\max\left\{\frac{S-2x}{L-2z}; \frac{S+2x}{L+2z}\right\}\right) \tag{7}$$

$$\theta_{max} = \pi - \arctan\left(\max\left\{\frac{S-2x}{L+2z}; \frac{S+2x}{L-2z}\right\}\right) \tag{8}$$

$$\varepsilon_g = \frac{1}{2\pi}\sum_{q=1}^{2}\int_{\theta_{min}}^{\theta_{max}}\left(\max\left\{\frac{S-2x}{H+(-1)^q 2y}; \frac{S+2x}{H-(-1)^q 2y}\right\}\right)^2 \csc^2\theta + 1)^{-\frac{1}{2}} d\theta \tag{9}$$

For the Forte, the detector area is 380 mm in length and 510 mm in height, resulting in a geometric efficiency of approximately 12% for a head separation of 600 mm. Sensitivity is highest in the center of the field of view and decreases rapidly as the source moves off-axis. Additionally, by sampling the geometric efficiency at several points within the FOV and multiplying by the square of the intrinsic efficiency, a sensitivity field can be created and visualized (see Fig. 5).

In this work, a directly activated 2 mm glass bead attached to the end of a long stainless-steel syringe is placed in the center of the FOV with a detector separation of 600 mm. Since positrons emitted from fluorine-18 can have a considerable range in air, the peak sensitivity is found by removing successive layers of shielding from the point source. The shielding provides material for positrons to annihilate with an electron, then the expected counts for a bare source in air are found by extrapolating the shielded count-rates. The parameters C_1 and C_2 are fitted, so as long as the shielding is of uniform material the attenuation coefficient is not needed.

$$R = C_1 e^{-C_2 t} \tag{10}$$

The activity of the source in these experiments is measured by a well-counter to be 10.5 MBq ± 0.1, decaying to approximately 8 MBq when the last shield is placed over the source. The shields are four 150 mm long, 0.75 mm thick stainless-steel tubes of increasing diameter placed over the point source. For each measurement 1,000,000 LORs are used to reduce statistical uncertainty. Using a value of 24% for the intrinsic efficiency the expected sensitivity of the Forte at 600 mm is 6.49 kHz/MBq.

2.3. Scatter fraction and count-rates

All LORs acquired by PET systems are either true, scattered or random. The scatter fraction is the fraction of detected coincidences in which one or both rays undergo scattering. This corrupts the LORs, since they no longer pass through the point of positron-annihilation. As the experimental system of interest becomes denser, more rays are scattered, reducing the maximum achievable true LOR count-rate. In addition to scattered coincidences, two unrelated events can be detected and associated as a coincidence if they occur within the coincidence window. At high count-rates, random LORs can constitute a large fraction of the total LORs since they scale with the square of the singles rate, S [24].

$$R_{Total} = R_{True} + R_{Scattered} + R_{Random} \tag{11}$$

$$R_{Random} = 2\tau S^2 \tag{12}$$

To measure the scatter fraction, the scattered LORs must be extracted from the true and random LORs. NEMA describes an experiment using a line source inserted into a high-density polyethylene (HDPE) phantom as the basis for calculating the scatter fraction and count-rates.

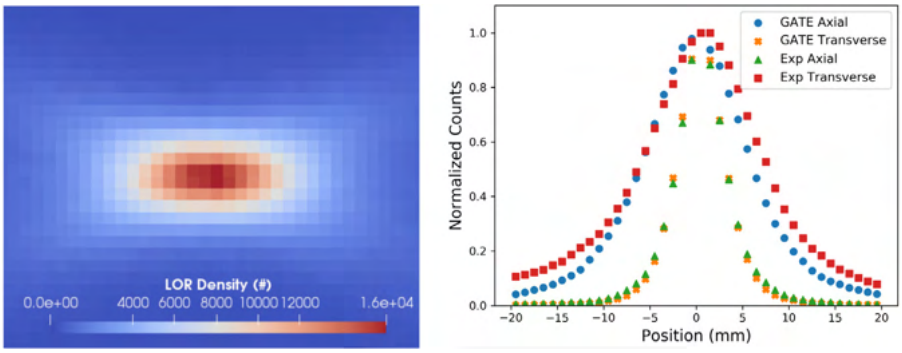

Fig. 4. Left: 2D slice of the 3D voxel array for a source near the center of the FOV. Both the Axial and Transaxial spatial resolutions are visible. Each voxel edge is 1 mm. Right: 1D back-projection of the experimental and GATE simulated spatial resolution phantom at the central position of the Forte FOV.

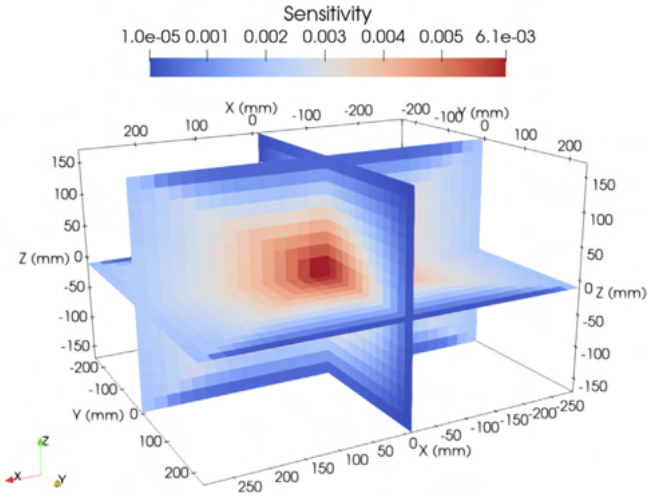

Fig. 5. 3D Total Sensitivity field of the Forte at a head separation of 600 mm.

However, due to the unavailability of specialized NEMA compliant phantoms, other phantoms are used in their place. As reported by others, the non-standard phantoms can be used in place of NEMA phantoms with little effect on the overall results [25,26]. For the purposes of this work, any phantom could be used, since the objective is to create a validated GATE model, as opposed to a performing quality assessment of the Forte. A smaller cylindrical HDPE phantom is used in this work and the NEMA methods applied. This phantom has an outside diameter of 50 mm and length 120 mm, with an internal cylindrical cavity of diameter 10 mm and length 100 mm. The phantom is filled with approximately 100 MBq of fluorine-18 well-mixed with water at the start of the experiment then imaged over several half-lives. The head separation is set to 445 mm.

When imaging this source using single-slice rebinning, the true, scattered, and random LORs can be isolated by following the NEMA protocol. The true LORs are all contained within the 40 mm strip centered on the maximum pixel with all other counts outside this strip being scattered or random. To remove scattered and random counts within the 40 mm strip, the pixel intensities on the edges of this strip are linearly interpolated and subtracted, leaving only the true LORs. The random LORs can be further extracted using a random coincidence estimate, however, for simplicity, the scatter fraction in this work is reported as the ratio of scattered and random LORs acquired during the lowest activity run. In this way, the random contribution is minimal.

$$R_{Scatter} = R_{Total} - R_{True} \tag{13}$$

$$SF = R_{Scatter}/R_{Total} \tag{14}$$

In addition to the true, scattered, and random count-rates, the Noise Equivalent count-rate (NEC) is computed as the square of the true counts divided by the total number of counts. The NEC has little relationship to actual noise, but rather provides a measure of relative count-rate contributions from true, scattered, and random coincidences. It is commonly used as a figure of merit for comparing PET systems [27].

$$NEC = R^2_{True}/R_{Total} \tag{15}$$

2.4. GATE model

GATE is a Monte Carlo simulator based on Geant4 libraries, designed for the simulation of PET scanners and other medical imaging devices [28]. The geometry of detectors and experiments can be created using simple shapes or geometries imported from mesh files, meaning a wide range of systems can be modeled. Further, the affects of signal-processing can be emulated using GATE, including the energy resolution of scintillation crystals, energy windows for accepting events, coincidence timing windows, time-resolution of recorded events, and dead-time of the PMTs. Accurate modeling of signal-processing is crucial for replicating realistic detector response over a wide-range of activities since the energy response and count-rates in real systems are greatly affected by these processes [29].

In this work, the geometric model of the Forte is informed by design drawings provided by the manufacturer, measurements taken at the PIC. The Forte possesses a wide-area NaI scintillation crystal of dimensions 590 x 470 x 16 mm, optically coupled to a glass light-guide. This component is considered the sensitive detector in the GATE simulation, the volume in which events are detected. Additionally, an aluminum gantry ring, aluminum casing around the detector heads, and lead shielding are added since these components will contribute significantly to the overall scatter. Further add to the geometric model, a generalized back-compartment corresponding to the ‘intermediate complexity’ back-compartment described by Rault et al. is added to account for scattering within the PMTs and internal electronics [30]. An accurate back-compartment is needed when modeling photons over 300 keV since these rays can pass through the relatively thin scintillation crystal and back-scatter, contaminating the energy spectrum (see Fig. 6)

The signal-processing of the Forte, referred to as the digitizer, is informed by previously published work using the Forte and experiments conducted in this work [19]. The energy resolution is reported to be approximately 14% at 511 keV for moderate count-rates. When used for PEPT, a 50% energy window of 350–650 keV is applied to discard the majority of scattered rays, while maintaining high coincidence rates. The detector is operated in coincidence with a window of 15 ns and reported time resolution of 15 ns. Additionally, only pulses from the central 510 x 380 mm area of the detector crystal are accepted. When the pulses falling outside this range are rejected for a central point

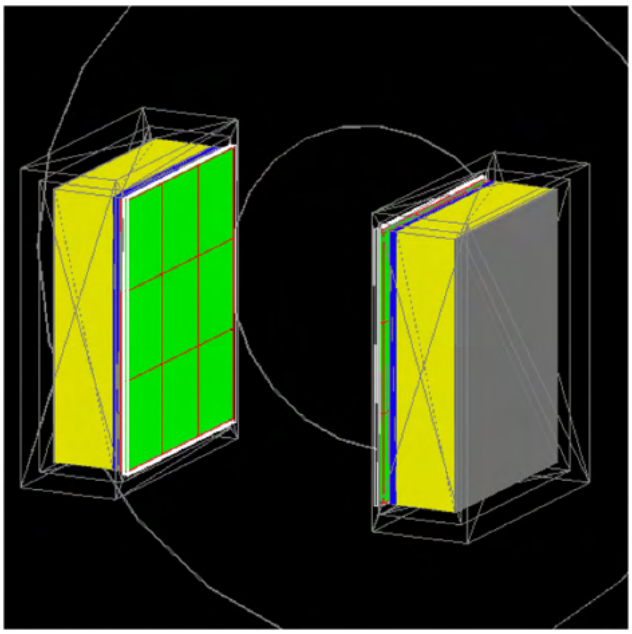

Fig. 6. GATE model of the Forte. NaI crystal (green, red outline), glass lightguide (blue), PMTs (yellow), electronics (gray), Al casing and Pb shielding (gray box outline), Al ring (gray ring outline), plastic front cover (white outline). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

source, the count-rate is reduced by 20%–30%, determined by the geometric efficiency. This effect is added in post-processing of the GATE generated LORs by rejecting LORs falling outside the active area.

Another consideration for the GATE model is the amount of spatial blurring to apply to the detected LORs. This blurring refers to an applied random movement in each position component of the detected LOR. The fluctuation in position follows a normal distribution with an FWHM set by the user. GATE simulation results will generally report a better spatial resolution than experimental measurements if spatial blurring is not simulated [31]. To investigate this effect, an additional series of simulations are conducted with the spatial blurring module activated in the digitizer. It is unknown how much additional spatial blurring is needed to recreate experimental results, so 3, 5, and 7 mm blurring are investigated and the best fit to experimental data is selected.

Additionally, the digitizer settings are crucial for accurate dead-time emulation. When a photon interaction triggers detector to convert the light generated by the scintillation crystal into an electronic pulse, dead-time occurs during which the detector is unable to accept a new pulse and other pulses may 'pile-up', being detected as a single pulse or simply lost [32]. A wide range of models have been proposed to account for this effect in PET systems including paralyzable, non-paralyzable, and hybrid models [33]. These models refer to the behavior of the dead-time, describing whether new pulses occurring during the dead-time extend the dead-time (paralyzable), do not extend the dead-time (non-paralyzable), or cause some mixture of the two behaviors (hybrid). The dead-time of the Forte is considered to have a hybrid dead-time, since each pixel is paralyzable, however, the recording of LORs is a non-paralyzable process.

Complicating the addition of dead-time in this GATE model is the need to account for the pseudo-independent zones of the scintillation crystal. When a scintillation event occurs, approximately only 6–7 individual photomultiplier tubes are activated out of an array of 55 on each head. The position of photon interaction is then calculated as the energy-centroid of the light collected by these PMTs, referred to as Anger logic [34]. To model these pseudo-independent zones, in lieu of not applying Anger logic directly, 9 pixels are used, arranged in a 3 × 3 array, since this the closest approximation of the area covering 6–7 PMTs. The dead-time per pixel is determined to be 1.15 μs by fitting the GATE produced count-rate curves to previously published data and an independent count-rate experiment conducted in this work [19]. This value disagrees with the published value of 170 ns for the non-paralyzable model described by Parker et al. but is closer to the manufacturer specification 1.3 μs, though it is not known how the manufacturer's value was derived. The non-paralyzable dead-time is also set to 1.15 μs (see Tables 2 and 3).

Next, the phantom and source geometries must be described. Three geometries are used for these experiments: an activated resin bead, coated in epoxy on the end of a hollow stainless-steel syringe, stainless steel cylinders placed over the point source, and a high-density polyethylene (HDPE) cylinder filled with water well-mixed with fluorine-18. These are the spatial resolution phantom, sensitivity phantom, and scatter fraction phantom, respectively. These phantoms are described in detail in the sections above (see Fig. 7, Table 4).

Table 2
Forte geometry description.

Geometric Model	Values
Detector dimensions	590 × 470 × 16 mm
Useful detector area	510 × 380 mm
Dual-Head separation	250–800 mm
Detector crystal	Sodium Iodide
Number of pixels	9

Table 3
Forte digitizer description.

Digitizer model	Values
Coincidence window	15 ns
Time resolution	15 ns
Paralyzable Dead-Time per pixel	1.15 μs
Non-Paralyzable Dead-Time	1.15 μs
Energy resolution at 511 keV	14%
Energy window	350–650 keV (50% Photopeak)
Intrinsic efficiency	24%

Table 4
Experimental and simulated phantom characteristics.

Characteristic	Spatial resolution	Sensitivity	Scatter and Count-Rates
Material	Resin, Plastic, Stainless-steel, Wood	Glass, Stainless-steel	HDPE, Water
Dimensions	0.2 mm diameter resin bead, 1.5 mm diameter syringe, dowel 200 mm length	2 mm diameter bead, 0.75 mm thickness shield, 150 mm length	120 mm height, 50 mm diameter, 12 mm inner diameter, 100 mm inner height
Head Separation	600 mm	600 mm	445 mm
Initial Activity	27.8 MBq	10 MBq	100 MBq

3. Results

3.1. Spatial resolution

The experimental spatial resolution phantom in the center of the FOV of the Forte produces a central axial FWHM of 16.28 mm and transaxial FWHM of 5.39 mm. The spatial resolution was also measured at 1/2 FOV. The 1/2 FOV axial and transaxial FWHM are measured to be 19.62 mm and 5.13, respectively. Compared to the experiment, the GATE simulation without spatial blurring underpredicts the FWHM

Fig. 7. The side view of the spatial resolution phantom (left), top down view of the sensitivity phantom (center), and cut-away view of the scatter fraction and count-rates phantom (right).

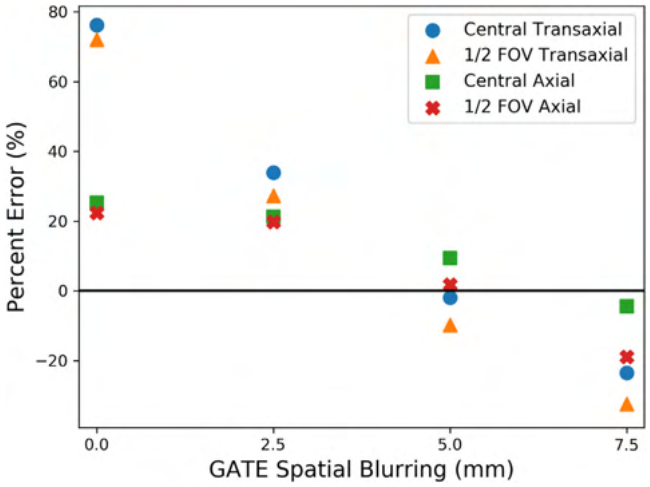

Fig. 8. The amount of spatial blurring applied to the GATE simulated LORs was varied and compared to the experimentally determined spatial resolution of the Forte.

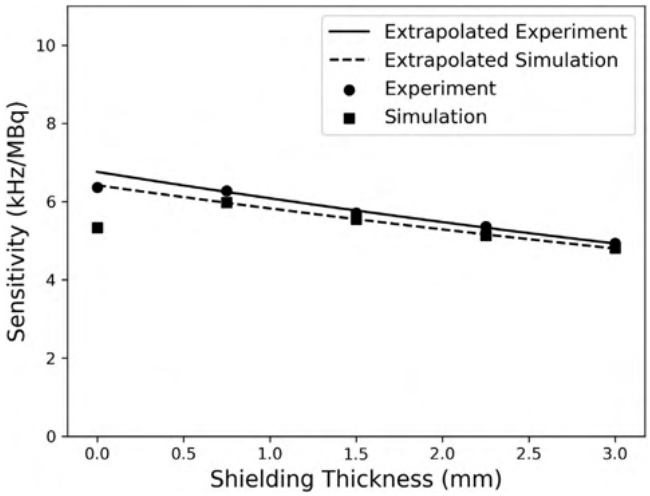

Fig. 9. Measured and extrapolated counts of the sensitivity experiment, showing close agreement when shielding is present and poor agreement for a bare source.

in both the axial and axial directions, indicating that spatial blurring is required. The 2.5, 5, and 7.5 mm spatial blurring simulations were compared, with the 5 mm blurring GATE simulation agreeing best with the experiment, though still under-predicting the transaxial FWHM. In previous work, the Forte was measured to have an transaxial spatial resolution of approximately 6 mm by fitting 1D a Gaussian function to a point source profile in air [19]. The measurements in this work vary, but are consistent with this value (see Fig. 8, Table 5).

3.2. Sensitivity

Comparing the results of the sensitivity experiment, the experimental sensitivity of the Forte is measured to be approximately 6.76 kHz/MBq, whereas the simulated sensitivity is reported to be 6.42 kHz/MBq. This translates into a −5.03% error between the experiment and simulation. The predicted theoretical sensitivity based on the geometric and intrinsic efficiencies, is approximately 6.49 kHz/MBq, showing agreement between theory, experiment, and simulation (see Fig. 9).

Further, the linear attenuation coefficient of the shielding in the experiment, calculated by fitting the count-rates, is 0.105 mm^{-1} compared to simulated value of 0.097 mm^{-1}, representing a −7.62% error. When the shielding is present the experiment and simulation agree well, however, when the source is imaged without shielding, the count-rates between the experiment and simulate vary significantly. This discrepancy is addressed in the discussion section.

Table 5
Experimental and simulated spatial resolution results.

Spatial resolution	Experiment	0 mm Blurring	5 mm Blurring
Central axial (mm)	14.680	10.965	13.294
Central transaxial (mm)	5.567	3.159	5.675
1/2 FOV axial (mm)	17.217	13.368	16.904
1/2 FOV transaxial (mm)	4.941	2.871	5.477

3.3. Scatter fraction and count-rates

The scatter fraction of the experiment and simulation in the lowest activity acquisition show relatively good agreement, however the experiment reports a slightly higher scatter fraction than the simulation. The experimental measured scatter fraction is approximately 0.167, compared to the simulated value of approximately 0.155, representing a −7.19% error. Over 1–100 MBq, simulated scatter and random fraction continue to agree, though at high source activities the scattered plus random and true fraction begin to diverge, with a maximum error of −14.21% at 100 MBq, possible reasons for this are covered in the discussion (see Fig. 10).

Further, the results of the count-rate experiment and simulation are in good agreement, with the total, true, scattered plus random, and NEC count-rates differing by less than 10% over the range of 1–100 MBq. The total and true count-rates perform particularly well, with an average error of 2.87% and 6.04%, respectively over the range of tested activities. The peak true coincidence rate of the experiment is

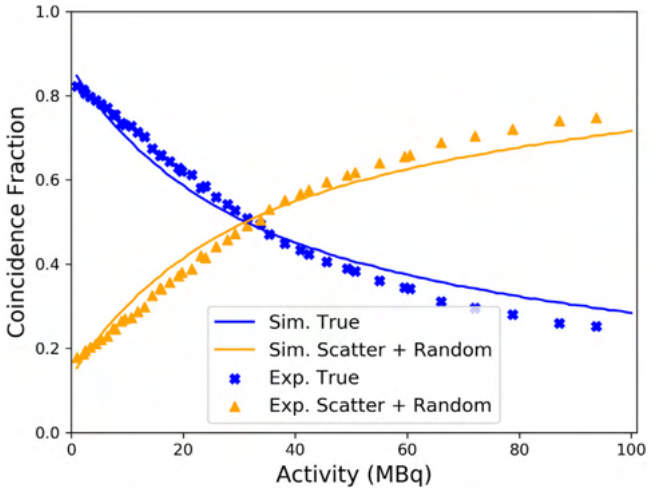

Fig. 10. The true LORs and scattered plus random LORs for the simulation and experiments.

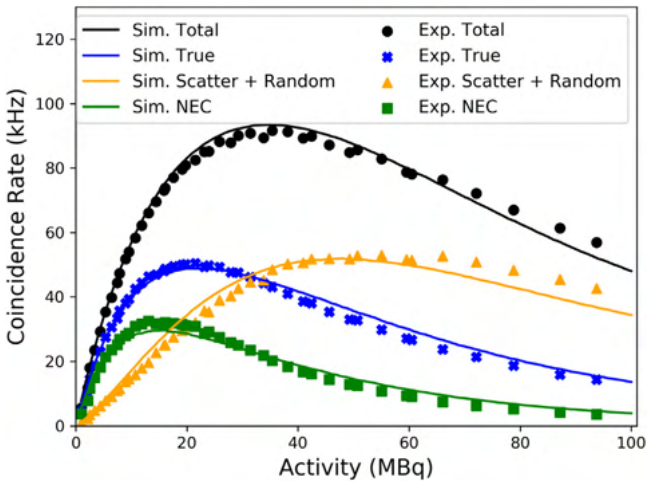

Fig. 11. Results of the count-rate experiment with the total (black), true (blue), scattered plus random (orange), and NEC (green) rates for both the experiment and simulation. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

50.3 kHz at 21.4 MBq compared to the simulated values of 48.5 kHz at 21.4 MBq. Additionally, the peak experimental NEC rate of 32.1 kHz occurs at 13.2 MBq compared to the simulated values of 29.6 kHz at 14.2 MBq, showing close agreement (see Fig. 11).

4. Discussion

The spatial resolution of the GATE simulation without added spatial blurring is considerably higher than what was measured in the experiment, demonstrating that the spatial blurring module is needed to degrade the resolution of detected LORs in order to match the experimental data. Accurate replication of the spatial resolution is important for simulating PEPT experiments since it will affect the fidelity of reconstructed trajectories, distorting the tracked position of source and can introduce spurious tracer velocities. For the Forte, approximately 5 mm of spatial blurring is required at activities around 20 MBq. It is expected that at lower activities the amount of spatial blurring can be decreased, but more experiments are needed to confirm this.

The sensitivity simulation agrees with experiment in both the overall sensitivity of Forte, as well as the calculated linear attenuation coefficient of the stainless-steel shielding. In previous work using the Forte, intrinsic efficiency is reported to be 23%, though a value of 24% appears to match this work better. While the theoretical, experimental, and simulated sensitivity are in good agreement, the counts for the unshielded source in the simulation do not agree, introducing a discrepancy. One reason for this could be differences in the air definition of the simulation compared to the actual conditions on the day of the experiment. The range of positrons emitted in diffuse media, such as air, can travel considerable distances before annihilating with an electron [35]. The density of air in GATE is 1.29 kg/m^3, however, the density of air is sensitive to both temperature and humidity which can affect the positron range. This is likely the source of error for this case, but differences between the source geometry or activity in the simulation could also be responsible.

The scatter fraction and count-rate experiment and simulation at low count-rates are at low and moderate activities are in good agreement, demonstrating that the source and geometric definitions are accurate since these are main contributors to the count-rates in these cases. At high activities the scattered and random counts diverge, with a maximum error of −14.21% at 100 MBq, which could indicate inaccuracies in the digitizer model or the pixelated detector crystal. Since crosstalk between the crystals or afterglow caused by defective crystals are not explicitly simulated, at high count rates these effects may introduce excessive error in positioning the point of interaction of photons [36]. This could explain why the total count-rate remains accurate, while the scatter and random count-rate is higher in the experiment than the simulation. However, this discrepancy is not significant since PEPT experiments using the Forte rarely use source activities where this would become an issue.

Table 6
Results from the validation experiments conducted in this work.

Validation characteristic	Percent error (%)	Count-Rate	Absolute percent error (%)
Central axial	9.44	–	–
Central transaxial	−1.90	–	–
1/2 FOV axial	1.82	Total	2.87
1/2 FOV transaxial	−9.79	True	4.59
Sensitivity	−5.03	Corrupted	8.08
Scatter fraction	−7.19	NEC	10.18

5. Conclusions

In this work, three experiments were conducted to validate a GATE model of the Forte, testing the spatial resolution, sensitivity, the scatter fraction, and count-rates. Overall, this GATE model agrees with experiment to within approximately 10% over all experiments and activities, consistent with values reported of other systems modeled using GATE (see Table 6).

Spatial resolution, sensitivity, scatter fraction, and count-rates are characteristics of interest in predicting the performance of PEPT experiments. By using realistic phantoms descriptions, this model has been shown to generate synthetic LORs from replicated experimental conditions. This is useful since a balance between the quantity and quality of LORs is desired in PEPT experiments. High count-rates are needed to create sufficient temporal resolution or tracers along their trajectory, however, increasing the source activity can result in poor count-rates due to excessive dead-time. Additionally, the fraction of corrupted events also increases at high count-rates, degrading spatial resolution of reconstructed trajectories.

Combined with the DEM simulation capabilities of the PIC, this GATE model has the potential to resolve several outstanding questions pertaining to the optimization of PEPT experiments. What is the source activity for a given experiment that will return the highest true LOR count-rates? How much material can annihilation photons

penetrate without being excessively corrupted by scattering? Can the spatiotemporal resolution of trajectories be improved by changing the detector geometry or digitizer settings? These questions can now be systematically investigated using Monte Carlo simulations without expending considerable time using the Forte, which is often constantly in use by researchers and companies conducting PEPT experiments. For proposed experiments, this model can also be used for feasibility studies, ensuring that the expected tracer velocities can be tracked or that the experimental apparatus will not introduce excessive scattering, representing a major improvement in PEPT modeling by the PIC.

CRediT authorship contribution statement

Matthew Herald: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation management, Writing - original draft, Visualization. **Christopher Windows-Yule:** Conceptualization, Writing - review & editing, Supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

Work at the Positron Imaging Centre is supported in part by a grant from the Engineering and Physical Science Research Council EP/R045046/1, Probing Multiscale Complex Multiphase Flows with Positrons for Engineering and Biomedical Applications.

References

[1] D.J. Parker, Positron emission particle tracking and its application to granular media, Rev. Sci. Instrum. 88 (5) (2017) 051803, http://dx.doi.org/10.1063/1.4983046, URL: http://aip.scitation.org/doi/10.1063/1.4983046.

[2] R. Shaw, Signalling particles for introduction into blood flowing through a vessel of interest, 1978, URL: https://patents.google.com/patent/US4224303A/en.

[3] M. Hawkesworth, D. Parker, P. Fowles, J. Crilly, N. Jefferies, G. Jonkers, Nonmedical applications of a positron camera, Nucl. Instrum. Methods Phys. Res. A 310 (1–2) (1991) 423–434, http://dx.doi.org/10.1016/0168-9002(91)91073-5, URL: https://linkinghub.elsevier.com/retrieve/pii/0168900291910735.

[4] D. Parker, C. Broadbent, P. Fowles, M. Hawkesworth, P. McNeil, Positron emission particle tracking - a technique for studying flow within engineering equipment, Nucl. Instrum. Methods Phys. Res. A 326 (3) (1993) 592–607, http://dx.doi.org/10.1016/0168-9002(93)90864-E, URL: https://linkinghub.elsevier.com/retrieve/pii/016890029390864E.

[5] R.D. Wildman, J.M. Huntley, D.J. Parker, Convection in highly fluidized three-dimensional granular beds, Phys. Rev. Lett. 86 (15) (2001) 3304–3307, http://dx.doi.org/10.1103/PhysRevLett.86.3304, URL: https://link.aps.org/doi/10.1103/PhysRevLett.86.3304.

[6] R. Pérez-Mohedano, N. Letzelter, C. Amador, C. VanderRoest, S. Bakalis, Positron Emission Particle Tracking (PEPT) for the analysis of water motion in a domestic dishwasher, Chem. Eng. J. 259 (2015) 724–736, http://dx.doi.org/10.1016/j.cej.2014.08.033, URL: https://linkinghub.elsevier.com/retrieve/pii/S1385894714010870.

[7] S. Langford, C. Wiggins, D. Tenpenny, A. Ruggles, Positron emission particle tracking (PEPT) for fluid flow measurements, Nucl. Eng. Des. 302 (2016) 81–89, http://dx.doi.org/10.1016/j.nucengdes.2016.01.017, URL: https://linkinghub.elsevier.com/retrieve/pii/S0029549316000273.

[8] C. Windows-Yule, J. Seville, A. Ingram, D. Parker, Positron emission particle tracking of granular flows, Annu. Rev. Chem. Biomol. Eng. 11 (1) (2020) annurev–chembioeng–011620–120633, http://dx.doi.org/10.1146/annurev-chembioeng-011620-120633, URL: https://www.annualreviews.org/doi/10.1146/annurev-chembioeng-011620-120633.

[9] G. Santin, D. Strul, D. Lazaro, L. Simon, M. Krieguer, M. Martins, V. Breton, C. Morel, GATE: A Geant4-based simulation platform for PET and SPECT integrating movement and time management, IEEE Trans. Nucl. Sci. 50 (5) (2003) 1516–1521, http://dx.doi.org/10.1109/TNS.2003.817974, URL: http://ieeexplore.ieee.org/document/1236960/.

[10] S. Jan, G. Santin, D. Strul, S. Staelens, K. Assié, D. Autret, S. Avner, R. Barbier, M. Bardiès, P.M. Bloomfield, D. Brasse, V. Breton, P. Bruyndonckx, I. Buvat, A.F. Chatziioannou, Y. Choi, Y.H. Chung, C. Comtat, D. Donnarieix, L. Ferrer, S.J. Glick, C.J. Groiselle, D. Guez, P.-F. Honore, S. Kerhoas-Cavata, A.S. Kirov, V. Kohli, M. Koole, M. Krieguer, D.J. van der Laan, F. Lamare, G. Largeron, C. Lartizien, D. Lazaro, M.C. Maas, L. Maigne, F. Mayet, F. Melot, C. Merheb, E. Pennacchio, J. Perez, U. Pietrzyk, F.R. Rannou, M. Rey, D.R. Schaart, C.R. Schmidtlein, L. Simon, T.Y. Song, J.-M. Vieira, D. Visvikis, R. Van de Walle, E. Wieërs, C. Morel, GATE - Geant4 Application for Tomographic Emission: a simulation toolkit for PET and SPECT, Phys. Med. Biol. 49 (19) (2004) 4543–4561, URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3267383/.

[11] S. Jan, G. Santin, D. Strul, S. Staelens, K. Assié, D. Autret, S. Avner, R. Barbier, M. Bardiès, P.M. Bloomfield, D. Brasse, V. Breton, P. Bruyndonckx, I. Buvat, A.F. Chatziioannou, Y. Choi, Y.H. Chung, C. Comtat, D. Donnarieix, L. Ferrer, S.J. Glick, C.J. Groiselle, D. Guez, P.-F. Honore, S. Kerhoas-Cavata, A.S. Kirov, V. Kohli, M. Koole, M. Krieguer, D.J.v.d. Laan, F. Lamare, G. Largeron, C. Lartizien, D. Lazaro, M.C. Maas, L. Maigne, F. Mayet, F. Melot, C. Merheb, E. Pennacchio, J. Perez, U. Pietrzyk, F.R. Rannou, M. Rey, D.R. Schaart, C.R. Schmidtlein, L. Simon, T.Y. Song, J.-M. Vieira, D. Visvikis, R.V.d. Walle, E. Wieërs, C. Morel, GATE: a simulation toolkit for PET and SPECT, Phys. Med. Biol. 49 (19) (2004) 4543–4561, http://dx.doi.org/10.1088/0031-9155/49/19/007, URL: https://iopscience.iop.org/article/10.1088/0031-9155/49/19/007.

[12] S. Lee, J. Gregor, D. Osborne, Development and validation of a complete GATE model of the Siemens Inveon trimodal imaging platform, Mol. Imag. 12 (7) (2013) http://dx.doi.org/10.2310/7290.2013.00058, URL: http://journals.sagepub.com/doi/10.2310/7290.2013.00058, 7290.2013.00058.

[13] F. Lamare, A. Turzo, Y. Bizais, C.C.L. Rest, D. Visvikis, Validation of a Monte Carlo simulation of the Philips Allegro/GEMINI PET systems using GATE, Phys. Med. Biol. 51 (4) (2006) 943–962, http://dx.doi.org/10.1088/0031-9155/51/4/013, URL: https://iopscience.iop.org/article/10.1088/0031-9155/51/4/013.

[14] C.R. Schmidtlein, A.S. Kirov, S.A. Nehmeh, Y.E. Erdi, J.L. Humm, H.I. Amols, L.M. Bidaut, A. Ganin, C.W. Stearns, D.L. McDaniel, K.A. Hamacher, Validation of GATE Monte Carlo simulations of the GE Advance/Discovery LS PET scanners: Validation of GATE simulations of GE PET scanners, Med. Phys. 33 (1) (2005) 198–208, http://dx.doi.org/10.1118/1.2089447, URL: http://doi.wiley.com/10.1118/1.2089447.

[15] S. Staelens, D. Strul, G. Santin, S. Vandenberghe, M. Koole, Y.D. Asseler, I. Lemahieu, R.V.d. Walle, Monte Carlo simulations of a scintillation camera using GATE: validation and application modelling, Phys. Med. Biol. 48 (18) (2003) 3021–3042, http://dx.doi.org/10.1088/0031-9155/48/18/305, URL: https://iopscience.iop.org/article/10.1088/0031-9155/48/18/305.

[16] D. Autret, A. Bitar, L. Ferrer, A. Lisbona, M. Bardiès, Monte Carlo modeling of gamma cameras for I-131 imaging in targeted radiotherapy, Cancer Biother. Radiopharmaceuticals 20 (1) (2005) 77–84, http://dx.doi.org/10.1089/cbr.2005.20.77, URL: http://www.liebertpub.com/doi/10.1089/cbr.2005.20.77.

[17] M. Herald, Z. Bingham, R. Santos, A. Ruggles, Simulated time-dependent data to estimate uncertainty in fluid flow measurements, Nucl. Eng. Des. 337 (2018) 221–227, http://dx.doi.org/10.1016/j.nucengdes.2018.07.005, URL: https://linkinghub.elsevier.com/retrieve/pii/S0029549318303844.

[18] Performance Measurements of Positron Emission Tomographs, Technical Report NU-2-2007, National Electronics Manufacturer Association, 2007.

[19] D. Parker, R. Forster, P. Fowles, P. Takhar, Positron emission particle tracking using the new Birmingham positron camera, Nucl. Instrum. Methods Phys. Res. A 477 (1–3) (2002) 540–545, http://dx.doi.org/10.1016/S0168-9002(01)01919-2, URL: https://linkinghub.elsevier.com/retrieve/pii/S0168900201019192.

[20] X. Fan, D. Parker, M. Smith, Labelling a single particle for positron emission particle tracking using direct activation and ion-exchange techniques, Nucl. Instrum. Methods Phys. Res. A 562 (1) (2006) 345–350, http://dx.doi.org/10.1016/j.nima.2006.03.015.

[21] A.L. Nicuşan, C.R.K. Windows-Yule, Positron emission particle tracking using machine learning, Rev. Sci. Instrum. 91 (1) (2020) 013329, http://dx.doi.org/10.1063/1.5129251.

[22] S. Ahmadi, S. Ashrafi, F. Yazdansetad, A method to calculate the gamma ray detection efficiency of a cylindrical NaI (Tl) crystal, J. Instrum. 13 (05) (2018) P05019, http://dx.doi.org/10.1088/1748-0221/13/05/P05019, URL: https://iopscience.iop.org/article/10.1088/1748-0221/13/05/P05019.

[23] A. Guida, Positron Emission Particle Tracking Applied to Solid-Liquid Mixing in Mechanically Agitated Vessles (Ph.D. thesis), University of Birmingham, 2010, URL: https://etheses.bham.ac.uk/id/eprint/935/.

[24] C. Stearns, D. McDaniel, S. Kohlmyer, P. Arul, B. Geiser, V. Shanmugam, Random coincidence estimation from single event rates on the Discovery ST PET/CT scanner, in: 2003 IEEE Nuclear Science Symposium. Conference Record (IEEE Cat. No.03CH37515), IEEE, Portland, OR, USA, 2004, pp. 3067–3069, http://dx.doi.org/10.1109/NSSMIC.2003.1352545, URL: http://ieeexplore.ieee.org/document/1352545/.

[25] H.W. de Jong, M. Lubberink, H. Watabe, H. Iida, A.A. Lammertsma, A method to measure PET scatter fractions for daily quality control: SF measurement method for daily QC, Med. Phys. 36 (10) (2009) 4609–4615, http://dx.doi.org/10.1118/1.3213096, URL: http://doi.wiley.com/10.1118/1.3213096.

[26] W. Kunze, M. Baehre, E. Richter, PET with a dual-head coincidence camera: spatial resolution, scatter fraction, and sensitivity., J. Nucl. Med. 41 (6) (2000) 1067–1074.

[27] S. Strother, M. Casey, E. Hoffman, Measuring PET scanner sensitivity: relating countrates to image signal-to-noise ratios using noise equivalents counts, IEEE Trans. Nucl. Sci. 37 (2) (1990) 783–788, http://dx.doi.org/10.1109/23.106715, URL: http://ieeexplore.ieee.org/document/106715/.

[28] S. Agostinelli, J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce, M. Asai, D. Axen, S. Banerjee, G. Barrand, F. Behner, L. Bellagamba, J. Boudreau, L. Broglia, A. Brunengo, H. Burkhardt, S. Chauvie, J. Chuma, R. Chytracek, G. Cooperman, G. Cosmo, P. Degtyarenko, A. Dell'Acqua, G. Depaola, D. Dietrich, R. Enami, A. Feliciello, C. Ferguson, H. Fesefeldt, G. Folger, F. Foppiano, A. Forti, S. Garelli, S. Giani, R. Giannitrapani, D. Gibin, J. Gómez Cadenas, I. González, G. Gracia Abril, G. Greeniaus, W. Greiner, V. Grichine, A. Grossheim, S. Guatelli, P. Gumplinger, R. Hamatsu, K. Hashimoto, H. Hasui, A. Heikkinen, A. Howard, V. Ivanchenko, A. Johnson, F. Jones, J. Kallenbach, N. Kanaya, M. Kawabata, Y. Kawabata, M. Kawaguti, S. Kelner, P. Kent, A. Kimura, T. Kodama, R. Kokoulin, M. Kossov, H. Kurashige, E. Lamanna, T. Lampén, V. Lara, V. Lefebure, F. Lei, M. Liendl, W. Lockman, F. Longo, S. Magni, M. Maire, E. Medernach, K. Minamimoto, P. Mora de Freitas, Y. Morita, K. Murakami, M. Nagamatu, R. Nartallo, P. Nieminen, T. Nishimura, K. Ohtsubo, M. Okamura, S. O'Neale, Y. Oohata, K. Paech, J. Perl, A. Pfeiffer, M. Pia, F. Ranjard, A. Rybin, S. Sadilov, E. Di Salvo, G. Santin, T. Sasaki, N. Savvas, Y. Sawada, S. Scherer, S. Sei, V. Sirotenko, D. Smith, N. Starkov, H. Stoecker, J. Sulkimo, M. Takahata, S. Tanaka, E. Tcherniaev, E. Safai Tehrani, M. Tropeano, P. Truscott, H. Uno, L. Urban, P. Urban, M. Verderi, A. Walkden, W. Wander, H. Weber, J. Wellisch, T. Wenaus, D. Williams, D. Wright, T. Yamada, H. Yoshida, D. Zschiesche, Geant4—a simulation toolkit, Nucl. Instrum. Methods Phys. Res. A 506 (3) (2003) 250–303, http://dx.doi.org/10.1016/S0168-9002(03)01368-8, URL: https://linkinghub.elsevier.com/retrieve/pii/S0168900203013688.

[29] S. Kerhoas-Cavata, D. Guez, Modeling electronic processing in GATE, Nucl. Instrum. Methods Phys. Res. A 569 (2) (2006) 330–334, http://dx.doi.org/10.1016/j.nima.2006.08.107.

[30] E. Rault, S. Staelens, R. Van Holen, J. De Beenhouwer, S. Vandenberghe, Accurate Monte Carlo modelling of the back compartments of SPECT cameras, Phys. Med. Biol. 56 (1) (2011) 87–104, http://dx.doi.org/10.1088/0031-9155/56/1/006, URL: https://iopscience.iop.org/article/10.1088/0031-9155/56/1/006.

[31] F. Bataille, C. Comtat, S. Jan, R. Trebossen, Monte Carlo simulation for the ECAT HRRT using GATE, in: IEEE Symposium Conference Record Nuclear Science 2004, Vol. 4, 2004, pp. 2570–2574 Vol. 4, http://dx.doi.org/10.1109/NSSMIC.2004.1462778.

[32] S. Usman, A. Patil, Radiation detector deadtime and pile up: A review of the status of science, Nucl. Eng. Technol. 50 (7) (2018) 1006–1016, http://dx.doi.org/10.1016/j.net.2018.06.014, URL: https://linkinghub.elsevier.com/retrieve/pii/S1738573318302596.

[33] M. Yousaf, T. Akyurek, S. Usman, A comparison of traditional and hybrid radiation detector dead-time models and detector behavior, Prog. Nucl. Energy 83 (2015) 177–185, http://dx.doi.org/10.1016/j.pnucene.2015.03.018, URL: https://linkinghub.elsevier.com/retrieve/pii/S0149197015000839.

[34] W.-H. Wong, H. Li, J. Uribe, A high count rate position decoding and energy measuring method for nuclear cameras using Anger logic detectors, IEEE Trans. Nucl. Sci. 45 (3) (1998) 1122–1127, http://dx.doi.org/10.1109/23.681989, URL: http://ieeexplore.ieee.org/document/681989/.

[35] M. Conti, L. Eriksson, Physics of pure and non-pure positron emitters for PET: a review and a discussion, EJNMMI Phys. 3 (1) (2016) 8, http://dx.doi.org/10.1186/s40658-016-0144-5, URL: http://ejnmmiphys.springeropen.com/articles/10.1186/s40658-016-0144-5.

[36] A.N. Blacklocks, A.V. Chadwick, R.A. Jackson, K.B. Hutton, Investigation into thallium sites and defects in doped scintillation crystals, Phys. Status Solidi (C) 4 (3) (2007) 1008–1011, http://dx.doi.org/10.1002/pssc.200673704, URL: http://doi.wiley.com/10.1002/pssc.200673704.

.2 DEM2GATE: Combining discrete element method simulation with virtual positron emission particle tracking experiments

Powder Technology 401 (2022) 117302

Contents lists available at ScienceDirect

Powder Technology

journal homepage: www.elsevier.com/locate/powtec

DEM2GATE: Combining discrete element method simulation with virtual positron emission particle tracking experiments

M.T. Herald [a], J.A. Sykes [a,b,*], D. Werner [a], J.P.K. Seville [a], C.R.K. Windows-Yule [a]

[a] School of Chemical Engineering, the University of Birmingham, Edgbaston, Birmingham B15 2TT, UK
[b] School of Physics and Astronomy, the University of Birmingham, Edgbaston, Birmingham B15 2TT, UK

ARTICLE INFO

Article history:
Received 8 November 2021
Received in revised form 7 March 2022
Accepted 17 March 2022
Available online 23 March 2022

Keywords:
DEM
PEPT
GATE
Validation
Simulation

ABSTRACT

Positron Emission Particle Tracking (PEPT) is a Lagrangian particle tracking technique useful for imaging flow in opaque granular media. The scientific and industrial systems studied using PEPT and the PEPT technique can also be simulated: Discrete Element Method (DEM) for granular systems and Monte Carlo for PEPT. Using the open-source DEM2GATE library, we quantify the reconstruction of particle trajectories and overall system behaviour by integrating particle trajectories from a rotating drum simulation with a PEPT detector simulation. The results show the main drivers of PEPT spatial error are particle acceleration and gamma-ray scattering. The accuracy of reconstructed flow fields increases logarithmically with measurement time. Measurements of higher-order, derived quantities require more time to reach a desired level of accuracy. Surprisingly, our results imply PEPT measurements relying on fluctuating velocity deviate substantially from the system's true behaviour. However, using DEM2GATE, these errors are predictable and can be corrected.

© 2022 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

1. Introduction

1.1. Background and motivation

Granular materials are used in a wide range of contexts in industry, but are often dense and opaque, making them difficult to study using optically-based particle tracking techniques like Particle Tracking Velocimetry (PTV) or Particle Image Velocimetry (PIV) [1]. In response to this, more penetrating imaging methods like x-ray tomography and refractive index matching tomography have been developed [2,3]. Positron Emission Particle Tracking (PEPT) is a widely used Lagrangian particle tracking technique that relies on the detection of coincident 511 keV gamma-rays to locate a positron-emitting source [4–6]. Even through dense optically inaccessible vessels, PEPT can be used to study granular systems [7].

However, it is necessary to establish that the PEPT result is a faithful representation of real trajectories and that the behaviour of the particle is representative of all particles in the system. There are a number of ways that Lagrangian trajectories can misrepresent the real behaviour of a system, such as not having enough temporal resolution to capture high-frequency motions of the tracer or introducing noise to a slow-moving particle [8]. In the first case when a trajectory does not have enough resolution to resolve rapid fluctuations in velocity, the tracer will appear to be moving more slowly than it actually is, since this motion will be smoothed out. On the other hand, if a particle is moving very slowly, the noise in the position of the tracer will make it appear to be 'jiggling' in place. Further, segregation caused by differences in particle size or density can lead to particles having very different behaviour [9]. This will necessarily distort the PEPT measurement as these particles may not explore all areas or phase states of the system.

To validate PEPT measurements, complementary particle tracking experiments such as Particle Image Velocimetry (PIV) can be conducted simultaneously, but particles must be imaged in optically transparent systems [10]. In general, these experiments have shown that PEPT locates the tracer accurately, but complementary experiments cannot be carried out for all cases due to the opaque nature of many granular systems. Furthermore, PIV cannot guarantee that the true data are completely representative of the real particle motion either. This necessitates a study of the extent to which PEPT trajectories represent particle motion in a system where the true trajectories and system behaviour are known a priori.

In order to fulfill the above stated, in this work, we introduce a new method of using highly-detailed Lagrangian trajectories of particles from a DEM simulation as the basis for virtual PEPT experiments allowing, for the first time, a meaningful and direct comparison between PEPT and DEM data. The prescribed particle position provides a means of direct comparison between a known and highly specified (virtual) trajectory and its (virtual) measurement using PEPT. First, a DEM

* Corresponding author at: School of Physics and Astronomy, the University of Birmingham, Edgbaston, Birmingham B15 2TT, UK.
E-mail address: jas653@student.bham.ac.uk (J.A. Sykes).

https://doi.org/10.1016/j.powtec.2022.117302
0032-5910/© 2022 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

simulation of a system representative of a typical PEPT experiment is created and a single particle trajectory extracted. Then this trajectory is prescribed as the motion of a positron-emitting tracer in a virtual PEPT experiment using calibrated Monte Carlo simulations. The virtual PEPT experiment which reproduces the behaviour of the full PEPT camera system, produces outputs quantitatively similar to that of real detectors and can be processed using a PEPT algorithm to locate the tracer. To assess the difference between the DEM trajectory and the PEPT reconstructed trajectory, the trajectories are first compared point-by-point to compute an average 2-dimensional spatial error. Then, both trajectories are post-processed to reconstruct Eulerian fields and scalar indicators which describe the system behaviour. Since values generated from DEM are the model data, PEPT reconstructed values can be compared to assess to what extent the modelled system's behaviour is captured - and crucially to what extent, and in what manner, the PEPT-reconstructed trajectories differ from the model DEM.

This work sets out to describe how to take a DEM simulation of a granular system and use extracted particle trajectories to create a realistic virtual PEPT experiment, introducing a publicly available repository of useful functions for this task called 'DEM2GATE'. Along with this workflow, methods to compare the DEM trajectories with the PEPT trajectories are demonstrated, providing a quantifiable and reproducible comparison. To aid the reader, an overview of the paper's structure and where to find key information is included below:

- Rotating drum DEM simulation (Section 2.1)
  - Mono-disperse GranuDrum (Section 2.1.1)
  - Poly-Disperse GranuDrum (Section 2.1.2)
- DEM2GATE Data Extraction (Section 2.2)
  - Tracer trajectory (Section 2.2.1)
  - Packing-density voxel array (Section 2.2.2)
- PEPT detector simulation (Section 2.3)
- Post-Processing DEM and PEPT data (Section 2.5)
  - Eulerian Fields (Section 2.5.1)
  - Scalar Indicators (Section 2.5.3)
- Results and Key Findings (Section 3)
  - Eulerian field accuracy increases logarithmically with measurement time (Fig. 17)
  - Higher-order derived quantities require longer measurement times (Fig. 17)
  - Particles of mean diameter best recreate occupancy (Fig. 22)
  - PEPT spatial error is caused by particle acceleration and gamma-ray scattering (Fig. 27)
  - PEPT measurements of fluctuating velocity can deviate substantially from the prescribed behaviour (Fig. 28)
  - Errors in PEPT are predictable and can be corrected (Fig. 29)

1.2. Positron emission particle tracking

PEPT is a fully 3-dimensional Lagrangian particle tracking technique capable of capturing high spatial and temporal resolution trajectories of a single particle in a granular system [4]. To conduct a PEPT experiment, a particle from a granular system is labelled with a positron-emitting isotope, usually fluorine-18, then placed back into the system [11]. The annihilation photons of a positron-electron pair are recorded by a radiation detector and used to reconstruct the position of the tracer over time [4]. Positrons annihilate with electrons close to their point-of-emission producing coincident, back-to-back 511 keV gamma-rays. The gamma-rays can be detected using gamma cameras operated in coincidence mode. The line between the coincident 511 keV gamma-rays intersects the tracer location and is called a line-of-response (LoR). In principle, two LoRs define the position of the tracer by triangulation. In practice, to account for scattering and other uncertainties, a larger sample of LoRs is used. In a typical PEPT experiment, approximately 50–500 LoRs are used [6,12], with the optimal number of LoRs being a complex function of the scattering environment, tracer velocity, and source activity. Higher event count rates of LoRs are needed to locate the tracer accurately in environments with a high degree of scattering or when the tracer is moving quickly, but the source activity is a limiting factor in the rate at which LoRs are detected. 200 LoRs has been shown to give a satisfactory compromise between spatial and temporal resolution in a number of systems, and was recently used in an extensive comparative study of the performance of PEPT algorithms by Windows-Yule et al. 2021, RoPP [12]. A sample of LoRs and segment of a trajectory from a PEPT tracer are shown in Fig. 1. The LoRs which do not intersect the tracer position are scattered or random LoRs.

PEPT trajectories are post-processed to calculate properties of granular systems. In other work, PEPT has been used to study properties of industrial relevance in systems such as rotating drums, fluidised beds, and pipe flows, among others [4,13–18].

While the Birmingham PEPT algorithm is the most widely used PEPT algorithm, newer methods such as the Positron Emission Particle Tracking using Machine Learning (PEPT-ML) algorithm have been developed which provide more robust tools for locating, identifying, and tracking particles [6,19]. PEPT-ML is used in this work; however, other algorithms such as the voxel-based Feature Point Identification (FPI) method or the Stanford B-Spline Reconstruction (SBSR) method [20,21] could be used. The FPI method offers advantages in multiple particle tracking capabilities whereas the SBSR method is capable of using fewer LoRs [12].

1.3. Simulating granular media

The Discrete Element Method (DEM) is a simulation technique for numerical modelling granular systems by solving Newton's equations of motion in discrete time-steps. Developed by Cundall in 1971 [22], this method resolves the Lagrangian motion of the particles, detects contacts, and calculates contact forces between particles. DEM is therefore a powerful tool to predict granular behaviour [23]. However, often when modelling granular materials consisting of very small particles, like fine powders, a larger 'coarse-grained' DEM particle is used. This is done to reduce the amount of computation needed in the simulation since fine powder may consist of billions or trillions of individual particles. When this is done, some forces which may be particularly relevant to small particles, for example, the fluid drag force and Van Der Waals forces, are not modelled, thus the particle behaviour is determined only by inertial and contact forces.

Trajectories of a single particle can be derived from DEM calculations; however, DEM simulations may produce unrealistic results if not correctly calibrated [24]. This is where PEPT proves itself as a useful method for calibrating and validating DEM [25] and can be made more rigorous and effective by using the DEM2GATE methods proposed in this paper, as DEM2GATE can be used to determine the degree to which if the particle motion captured by PEPT accurately represents the model particle trajectory, and use this information to map the PEPT data so as to provide more precise model DEM data for validation.

In this work, the DEM software LAMMPS Improved for General Granular and Granular Heat Transfer Simulations (LIGGGHTS) is used to simulate a granular system which is then used as the basis for a virtual PEPT experiment. LIGGGHTS is an open-source DEM software used in industry and academia to model granular systems [26]. A small rotating drum, based on the commercially available powder test apparatus GranuDrum (GranuTools, Awans, Belgium) is chosen as the model system [27]. The GranuDrum is a system with rich phenomenology whose measurements are sensitive to the cohesive and frictional properties of particles. In practice, the GranuDrum is used to characterise powder flow, which can then be used to model the flowability of

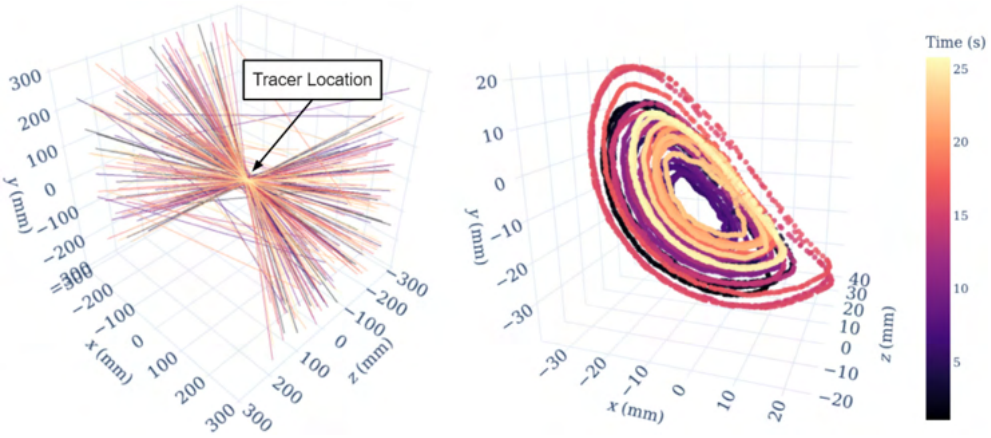

Fig. 1. A sample of 200 LoRs from a PEPT tracer (left) and a trajectory from clustered samples of LoRs, coloured by time (right).

materials in larger industrial systems [28,29]. Trajectories extracted from this simulation are the *model* basis on which to compare the PEPT reconstructed trajectories (see Section 3.2). The GranuDrum is shown in Fig. 2.

Previous work has shown that the time needed to run an experiment largely depends on the characteristic length of the vessel and tracer velocity [30]. In this case, a 30-min long trajectory ensures that the tracer explores all areas of the rotating drum and goes to these areas often enough such that accurate statistics of the calculated Eulerian fields and scalar indicators can be ensured. Other systems may require longer or shorter trajectory times to acquire adequate statistics, but 30 min is suitable in this instance, as will be seen in Section 3.1.1.

1.4. Creating virtual PEPT experiments

Virtual PEPT experiments are Monte Carlo simulations of positron-emitting sources (or back-to-back 511 keV gamma-rays) inside validated models of radiation detectors designed to replicate the conditions of a real PEPT experiment. The virtual experiments in this work are made using the Geant4 Application for Tomographic Emission (GATE). GATE is an open-source software that models the transport of radioactive particles through matter and is designed specifically for imaging applications such as Positron Emission Tomography, Single Photon Emission Computed Tomography, and PEPT [31,32]. Using GATE, complex scattering from objects in the field-of-view (FOV) and the behaviour of multi-stage pulse-processing chains in detectors can be recreated [33].

Prior to this work, GATE has been used to estimate the uncertainty of PEPT trajectories extracted from CFD simulations of pipe flow using a model of the Siemens Inveon at the University of Tennessee [34,35]. Recently, a GATE model of the ADAC Forte, a detector system used for PEPT experiments at the University of Birmingham (shown in Fig. 3) has been developed and validated against experimental measurements [36]. The model of the ADAC Forte includes the energy-resolution, time-resolution, and a pulse-processing stage that mimics the data rates recorded by the real detector, replicating experimental data to within 10% error. However, some aspects of the real detector are not included in this model such as an Anger logic algorithm used to determine the position of a gamma detection in the crystal [37]. Instead, the exact position of interaction is recorded in GATE and then blurred using a Gaussian filter which has been calibrated to yield the closest approximation to the spatial resolution of the detector. This creates an opportunity to develop virtual PEPT experiments which can be used prior to, and to optimise real PEPT experiments.

In this work, the DEM simulation of the GranuDrum, detailed in Section 2.1.1, is combined with a virtual PEPT experiment by using extracted DEM trajectories to move a virtual 10 MBq positron-emitting tracer. Two versions of the virtual PEPT experiment are created, one to examine a PEPT experiment without the influence of LoRs scattering within the GranuDrum and then another to include this scattering. In the first virtual experiment, the tracer is placed only in a volume of air in the centre of the FOV, minimising the scattering of LoRs, thereby creating an idealised experiment. Then in the second experiment, the tracer is placed inside a structural model of the GranuDrum which

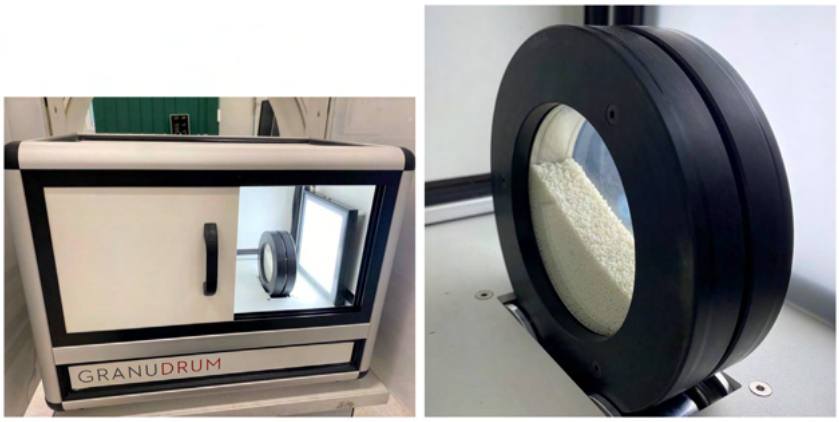

Fig. 2. A picture of the GranuDrum in between the ADAC Forte (left) and a close-up of the drum with micro-crystalline cellulose spheres (right).

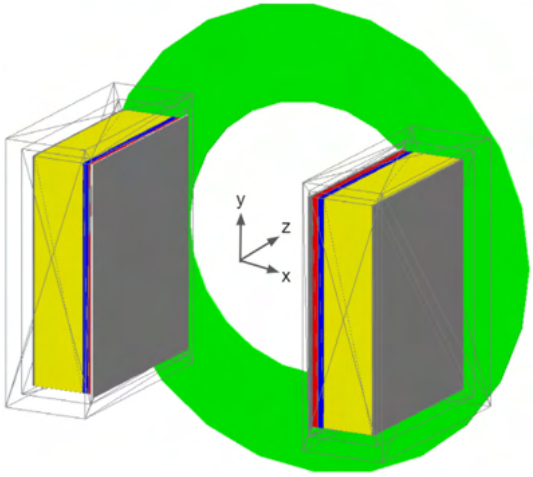

Fig. 3. ADAC Forte GATE geometry.

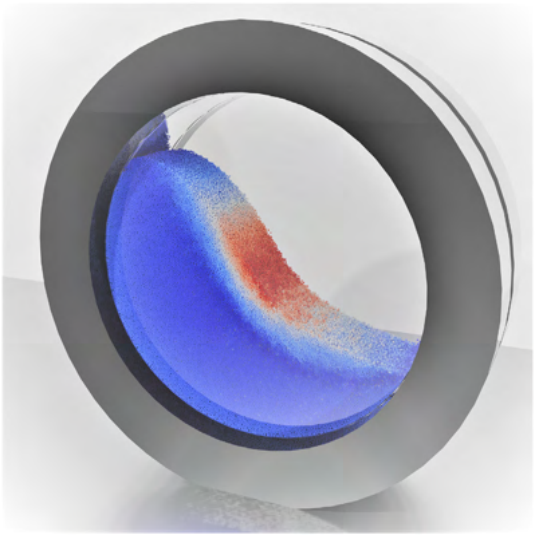

Fig. 4. The LIGGGHTS simulation of the GranuDrum viewed in ParaView and particles coloured by velocity.

includes the 3-dimensional packing density of the powder flow, a realistic representation of a PEPT experiment. From these virtual experiments, LoRs are recorded in the same manner as they would be in a real experiment, with the trajectories extracted using the PEPT-ML algorithm. The trajectories are post-processed to generate a series of 2-dimensional Eulerian fields and scalar indicators which are compared to the model DEM data to assess PEPT reconstruction accuracy based on the methods outlined in Section 2.5.

2. Methods

2.1. GranuDrum DEM simulation

2.1.1. Mono-disperse GranuDrum

The GranuDrum is a type of rheometer, or flow characterisation equipment, for measuring properties of granular media [29]. It consists of a small, thin, rotating drum inside a box with a camera to measure the dynamic free surface which allows for the calculation of the angle-of-repose [38]. For our work, a LIGGGHTS simulation is created with a hollow rotating cylinder of the same dimensions as the GranuDrum and filled with simulated mono-disperse spheres of microcrystalline cellulose (MCC). MCC is a common material in pharmaceutical industries which has been used in PEPT experiments in previous work [39]. A geometric description of the GranuDrum can be found in Table 1 and the simulated GranuDrum is shown in Fig. 4.

The drum is filled with 10,000 mono-disperse MCC spheres of diameter 1.8 mm and rotated at 45 rotations per minute (RPM) for 30 min. At this RPM, the MCC particles are within the cascading flow regime, where mixing is expected to be strong. Higher regimes such as cataracting and centrifuging can result in the crushing of particles or particles centrifuging around the edge of the drum, whereas lower speed regimes such as slipping and slumping result in the particles sliding over each other rather than properly mixing [40,41]. The positions of all particles are recorded and used to extract trajectories for comparison with PEPT detected trajectories of a single tracer. At the end of the simulation, the trajectories of all 10,000 particles have been recorded at a frequency of 1 ms, as this keeps the distance travelled by the fastest particles (moving at approximately 1 ms^{-1}) between successive locations smaller than the radius of the MCC sphere itself. This frequency is also representative of the achievable time resolution of the ADAC we are modelling.

The material properties of the DEM particles used in the simulation are shown in Table 2. The first four parameter values in the table were determined using the Autonomous Characterisation and Calibration using Evolutionary Simulation Software, ACCES [42] (see ACCES GitHub and ACCES Conference Talk). The remaining values were obtained from prior work as given in the references.

2.1.2. Poly-disperse GranuDrum

In addition to the mono-disperse simulation of the GranuDrum, another DEM simulation is generated with a poly-disperse particle size distribution. These particles have the same properties as the mono-disperse particles described previously, with only their diameter changing. The size distribution chosen is determined by a particle size distribution measurement performed on MCC particles using a Canty Lab Solidsizer [43]. The mean particle diameter of the measurement was approximately 1.2 mm. For the purpose of this simulation and to make the mean particle size equal to that of the mono-disperse system, the particle sizes were scaled to make the mean particle diameter 1.8 mm. Ten evenly spaced particle sizes were extracted from the data using a log-normal fit of the particle sizes. The particle size distribution and the extracted sizes are shown in Fig. 5.

As demonstrated in other work, particle size and density in a poly-disperse system are drivers of segregation [9]. PEPT measurements, due to tracking only one particle at a time, could be distorted if the

Table 1
GranuDrum description.

Characteristic	Values
Drum Diameter	84 mm
Drum Thickness	20 mm
Glass Wall Thickness	3 mm
Box Length	560 mm
Box Height	360 mm
Box Depth	410 mm

Table 2
Material properties of MCC particles in the LIGGGHTS simulation.

Simulated property	Values
Particle-Wall Friction	0.32
Particle-Particle Friction	0.32
Coefficient of Restitution	0.3
Rolling Friction	0.0025
Young's Modulus	5×10^6 Nm^{-2} [57]
Poisson's Ratio	0.3 [58]
Density of MCC	1580 kgm^{-3} [59]

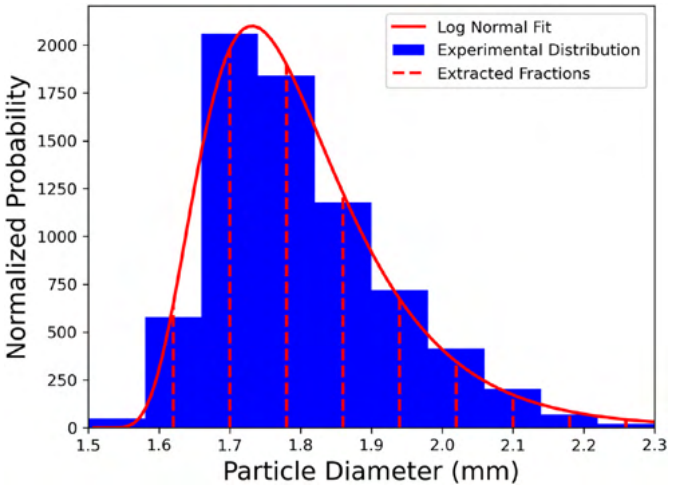

Fig. 5. The scaled particle size distribution for the poly-disperse simulation of the GranuDrum with a mean particle size of 1.8 mm. The data are fitted to a log-normal distribution and ten particle sizes are extracted.

particle being tracked is segregated. The evidence of segregation would be most pronounced in the occupancy distribution, thus in Section 3 the occupancy distribution for the whole poly-disperse system and the particles representing each size fraction are compared in Section 3.1.2.

2.2. DEM2GATE data extraction

Following the simulation of the DEM model, the raw DEM data were converted into a format that can be used in a GATE simulation. Two things are desired from this data: 1) individual trajectories of particles, and 2) a 3-dimensional array representing the packing density of the flowing MCC powder which can be used to approximate the scattering of gamma-rays in the drum. To perform both of these tasks, a Python library - DEM2GATE - was created. GATE already has capabilities to read from files containing trajectory information and can import complex geometries from 3-dimensional image data. The novel work described here is the process of extracting this information from a LIGGGHTS simulation and converting it to a GATE readable format. The workflow for this task is as follows:

1. Tracer Trajectories
   (a) Extract particle position from LIGGGHTS;
   (b) Convert trajectories to GATE Placements Format;
   (c) Prescribe trajectory to GATE source.

2. Voxel Array
   (a) Calculate a 3-dimensional array of packing density;
   (b) Export array as a MetaImage;
   (c) Define GATE material definitions;
   (d) Create a GATE range translator;
   (e) Place voxel array in GATE simulation.

The software developed for this paper is free and open-sources, and so can be used by the reader at will. The software can be found in the Positron Imaging Centre's DEM2GATE GitHub Repository.

2.2.1. Tracer trajectories

The DEM simulation of the GranuDrum tracks the position of all 10,000 simulated particles simultaneously. For every time-step, the positions are updated based on the current forces acting on the particles. At regular intervals of 1 ms, the positions of all tracers are written to the VTK file. To extract a DEM trajectory, the VTK files are read into Python using a VTK reader, allowing us to view and select the data we want. Next, the time and x, y, z positions for a specified particle are extracted and written into the GATE placements file format. The Placements file is a text file consisting of a header describing the format of the data accompanied by the data columns. GATE uses this file to identify where a source should be placed inside the simulation during a particular time-step by selecting the time and corresponding positions that are closest to the current GATE time-step. Finally, the particle trajectory is prescribed in the GATE simulation by declaring that the source position should be read from this placement file via the Generic Move function.

2.2.2. Packing density voxel array

The second set of data to extract from the simulated rotating drum is a 3-dimensional array of the packing density, i.e. a voxel array. In this representation of the GranuDrum, each voxel has a value corresponding to the number of particles it contains; when all voxels are the same size, this provides a measure of packing density. To generate a voxel array of the flowing particles in the GranuDrum, an empty array is first created which breaks up the drum volume into a $50 \times 50 \times 50$ grid over the simulation space of the LIGGGHTS simulation, -0.048 to 0.048 m in the x and z, then -0.002 to 0.0221 m in the y-direction. This array is created in Python using NumPy arrays binning the VTK particle data. Then, for each time-step, the number of particles in each voxel is summed, creating a time-averaged packing density. This voxel array is normalised and stored in a GATE-readable MetaImage format.

To map the normalised voxel value to the real corresponding bulk density, the maximum bulk density of the MCC particles must be calculated by multiplying the particle density, $\rho = 1580\ kgm^{-3}$, with a loose random packing fraction of 0.6, yielding a bulk density of $\rho = 948\ kgm^{-3}$ [44]. This is then equated to the maximum voxel value. All other voxels are mapped to 50 discrete combinations in between the bulk density and the density of air, e.g. 20% of the bulk particle density and 80% air density would lead to $\rho_{20/80} = 189.9\ kgm^{-3}$. This is accomplished using a GATE range translator which uses the voxel value to choose a material from the list of predefined material mixtures.

Now that the voxel array - shown in Fig. 6 - has been created, this must be converted to a GATE readable format. This is accomplished by saving the voxel image as the medical imaging format (MetaImage) which GATE can read, followed by creating a range translator to map the voxel values to predefined materials in the GATE material database. The 3-dimensional NumPy array of packing density is converted to a MetaImage by using the SimpleITK Python package and specifically

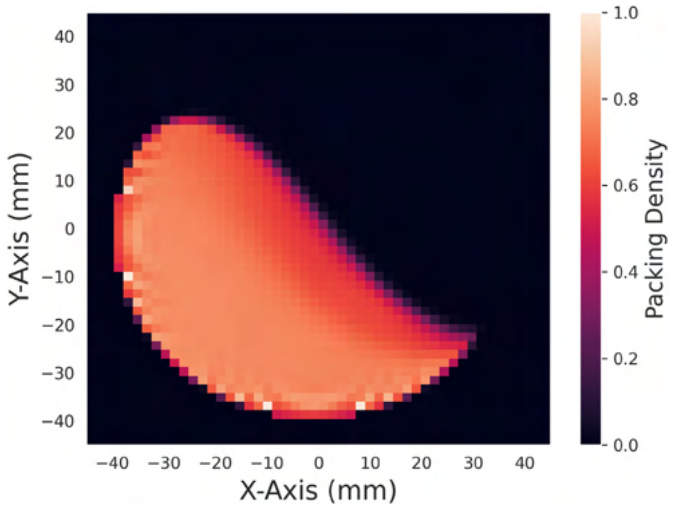

Fig. 6. The packing density of the GranuDrum from a slice of the 3-dimensional voxel array.

the GetImageFromArray function. Finally, a way for GATE to map the voxel values to the predefined materials is needed. A GATE range translator is applied which uses the voxel value to choose a material ranging from all solid to no solid. To ensure that a smooth transition from all bulk particles to all air occurs, 50 material combinations between bulk particles and air are generated and given the appropriate fractions of materials and density.

2.3. GATE simulation

A GATE simulation emulating a virtual PEPT experiment is created by starting with a validated model of the ADAC Forte detector [36]. This model has been validated against experiments following the National Electronics Manufacturers Association (NEMA) protocol for assessing the performance of Positron Emission Tomographs [5,45]. Thus, this model can reliably be used to generate synthetic LoRs for virtual PEPT experiments as previous work on other GATE modelled PEPT experiments have demonstrated [12,35]. The geometric model is shown in Fig. 7. Two scenarios are simulated with the data extracted from DEM2GATE: a single DEM trajectory in air and a single DEM trajectory in the GranuDrum geometry. These represent a best-case (minimal scattering) and a realistic case (considerable scattering) for a PEPT experiment, respectively. It is worth noting that the velocity of the tracer in both cases remains the same, and the tracer follows the same trajectory.

The geometric model of the GranuDrum follows the description given in Section 2.1.1. To be consistent with the real GranuDrum, the rotating drum inside the GranuDrum structure is translated 145 mm closer to one side of the ADAC Forte detector. This has the effect of reducing the count-rate detected from the tracer due to the lower geometric efficiency. The trajectory is co-registered with the voxel array such that the DEM trajectory is fully within the voxels at all times. This case will have considerably more scattered LoRs due to the added material between the tracer and detector. In turn, it will result in both a lower detection rate for the tracer and less accurate detection, caused by the incorporation of scattered LoRs in PEPT algorithms, thus providing a full, true-to-life representation of a real PEPT experiment.

Once the detector model is applied and set to a detector separation of 600 mm, the next step is to define the source. The DEM particle is a solid 1.8 mm diameter sphere of MCC with density $\rho = 1580\ kgm^{-3}$. This material definition is added to GATE and modelled as a volumetrically activated, positron-emitting source of 10 MBq activity. A 10 MBq tracer has been shown to result in count-rate near the peak true LoR count-rate for the ADAC Forte at 600 mm separation as demonstrated in Herald et al. 2021, NIMA [36]. The positrons are prescribed an energy spectrum corresponding to that of fluorine-18. The GATE modelled tracer and a positron-emission are shown in Fig. 8 (see Section 1.2 for positron-emission and annihilation photons explanation). Then, using the Generic Move function, the DEM trajectory which has been converted to a GATE placements file is prescribed as the successive positions of the source. This allows the GATE modelled tracer to take the same path as the DEM tracer, with the addition of emitting positrons in the GATE simulation and the annihilation photons detected by the ADAC Forte detector model.

To run GATE simulations efficiently, the inherent parallelisation of GATE is leveraged by splitting the simulation into smaller jobs, then recombining the results of these simulations later. The simulations are run on BlueBEAR, the University of Birmingham's High-Performance Computer [46]. The 30-min simulation is broken into 200 smaller jobs, each running for 9 s of the simulation. The output of the simulation is a file containing detected coincidences which can be processed with a PEPT algorithm to extract a trajectory of the tracer.

2.4. PEPT-ML clustering

Several PEPT algorithms exist, as has been mentioned in Section 1.2, and their performance has been assessed in other work [12]. The PEPT-ML algorithm used in this work and implemented in the PEPT Python library is described in the following steps [19]. The LoRs collected during a PEPT experiment are split into samples containing N_{LoRs}. For every possible pair of LoRs in the sample, the point that minimises the distance between pairs of lines is calculated, called a 'cutpoint'. Cutpoints that fall within a user-defined distance from the LoR pairs are recorded, then clustered using a hierarchical, density-based clustering algorithm [47]. For a sample of LoRs, producing a number of cutpoints, one cluster centre is calculated. Additionally, samples of cluster centres can be re-clustered in a process called 'second-pass clustering', resulting in a smoothed trajectory [19]. This process can be seen visually in Fig. 9.

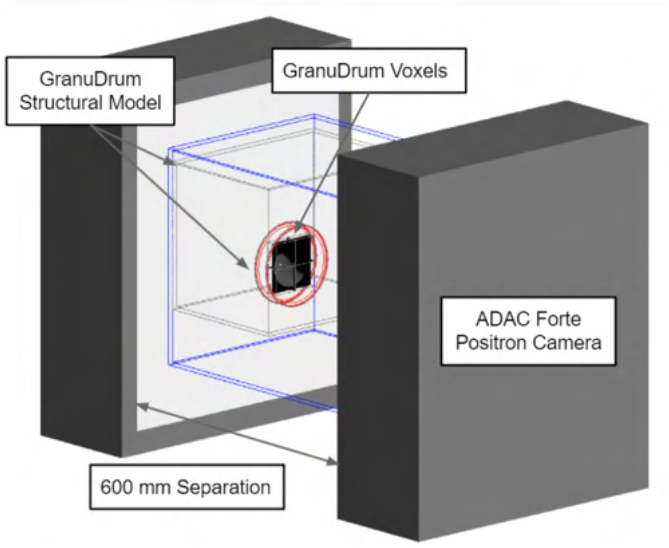

Fig. 7. The GATE modelled GranuDrum in the ADAC Forte Geometry with a positron-emitting source.

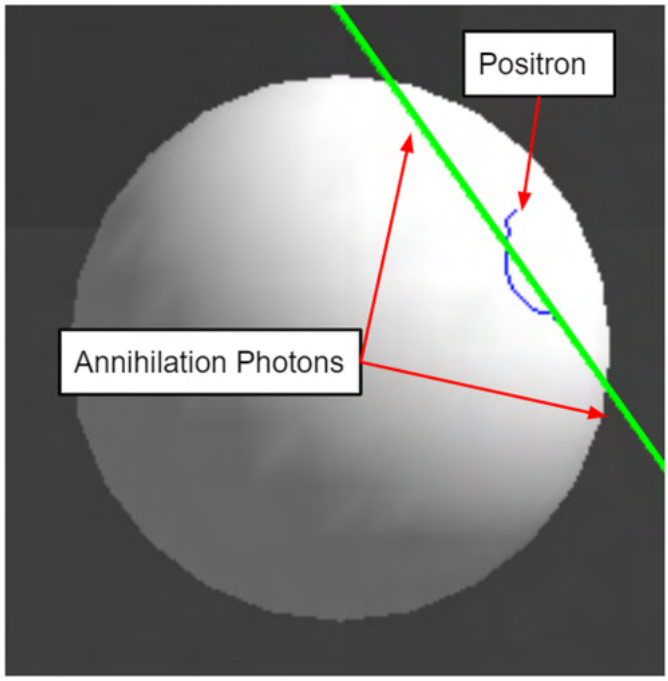

Fig. 8. GATE modelled MCC sphere with a positron-emission and annihilation present.

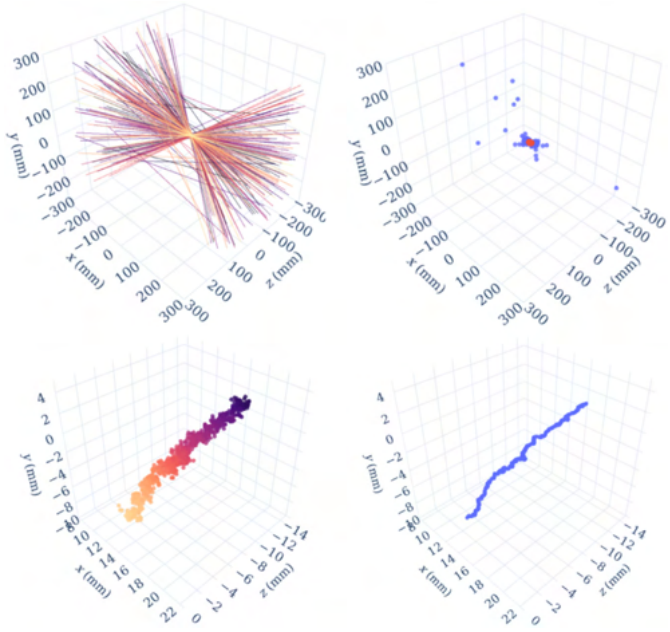

Fig. 9. (Top left) A sample of 200 LoRs is transformed into cutpoints and clustered using a hierarchical density-based algorithm (top right). Blue points represent the cutpoints calculated from the LoRs, with the red points being the cluster centre of these cutpoints. (Bottom left) The 1st pass cluster centres for a segment of the trajectory are coloured by time then the second pass of clustering tightens these points (bottom right).

In this work, PEPT 0.4.1 (PEPT-GitHub) is installed on the University of Birmingham's BlueBEAR High-Performance Computing facility [46] and is used to calculate tracer trajectories from GATE data. Each 9-s GATE simulation is initially clustered using only one round of clustering. Then these points are combined into a single file and second-pass clustered using tighter parameters. This process results in a continuously tracked particle, with outlier points removed by second-pass clustering. The tracer trajectory can then be directly compared to the prescribed DEM trajectory or used to compute Eulerian fields or scalar indicators.

The sample size of LoRs is prescribed as $N_{LoRs} = 200$; the choice of number was discussed earlier in Section 1.2. A moving window of LoRs is used to provide more samples of LoRs, with the latter 50% of LoRs in a given sample used in the next sample. All cutpoints falling within 0.15 mm of the LoR pair are saved for clustering. For each sample of cutpoints, the densest 15% of points are considered core points of the cluster and their centroid is calculated. The result is the 1st-pass position of the tracer. After this is completed, the 1st-pass positions are re-clustered with a sample size of 30 points, with 70% of the densest points being the core points and their centroid calculated. To produce nearly as many 2nd-pass cluster centres as the 1st-pass method, a moving window of 1 minus the sample size of centres is applied, in this case a sample size of 29.

2.5. Post-processing

Once a PEPT trajectory has been produced, there are two ways to assess the accuracy of the PEPT detected trajectory. The first method is to compare the point data of the trajectories, averaging the mean error between their recorded positions at the same time-step. This has been used in previous work with static tracers and those moving in simple trajectories to compare PEPT algorithm performance [12].

$$\varepsilon_y = y_{PEPT} - y_{DEM} \tag{1}$$

$$\varepsilon_{2D} = \sqrt{\varepsilon_x^2 + \varepsilon_y^2} \tag{2}$$

However, while looking at the point data of a trajectory is useful for characterising the uncertainty in the position of an individual tracer, an individual tracer does not necessarily explore the whole phase-space of particles within the system, particularly if it has not been exploring the system for a long enough time [4]. The decision to run the simulations for 30 min overcomes this issue for this specific system, but with DEM2GATE a second, more robust, method of comparison becomes available, which is comparing a reconstructed field of the system. For example, an Eulerian flow field generated from virtual PEPT measurement, such as the velocity field, can be compared to the field reconstructed using the model DEM data. Comparison of macroscopic fields shows not only to what extent PEPT trajectories follow an individual tracer but also how well the bulk behaviour of the system is reconstructed from a single tracer. This type of comparison is more significant for PEPT users because it is a measure of how the uncertainty in PEPT trajectories actually affects the measurement of a bulk system behaviour.

In this work, both the overall spatial error of the PEPT trajectory and the effect of this on the measurement of bulk system behaviour is examined. The generation of PEPT reconstructed fields also offers an opportunity to take advantage of the recent developments in the PEPT-ML library which uses time-step interpolation and polynomial fitting of positions in a Savitzky-Golay filter to derive tracer velocity [48]. The reconstructed Eulerian fields for the PEPT trajectory in air and PEPT trajectory inside the model of the GranuDrum are compared against fields generated by averaging all DEM trajectories.

All Eulerian fields are generated using the same procedure, whereby for each time-step of a trajectory, a value corresponding to an instantaneous measurement of a system field can be assigned to bins spanning the drum volume. For simplicity and because the GranuDrum has a small drum thickness in comparison to its diameter, only 2-dimensional fields are computed in this work. To perform the binning over the trajectory history, the SciPy function scipy.stats.binned_statistic_dd is used. The number of bins in each dimension as well as the value range over which the function is applied is provided by the user, then the mean, standard deviation, or counting is performed on the values falling within each bin. The fields calculated in this work are velocity components, velocity magnitude, acceleration components, acceleration magnitude, velocity relative standard deviation (RSD), granular temperature, and occupancy.

In addition to Eulerian fields, scalar indicators of the system behaviour are calculated. For a rotating drum, four scalar indicators that are important for characterising the behaviour of the drum are the locations of the so-called bulk 'shoulder' and bulk 'toe' regions, the centre-of-circulation (CoC), and also the dynamic angle-of-repose (AoR) [13]. A full description of these indicators and how they are calculated is described in Section 2.5.3.

2.5.1. Generating Eulerian fields

All fields are computed on the same 30×30 grid spanning ± 45 mm in the x and y directions centred on the centre of the GranuDrum, producing a field with 3×3 mm resolution. Velocity is calculated using a Savitzky-Golay filter with a window size of y points as a step within the PEPT-ML algorithm and is written into the trajectory as an appended column of data, thus the velocity components and magnitude are pre-computed and only need to be binned. The velocity fields are calculated by finding the mean values within each bin. Acceleration is calculated by numerically differentiating the velocity by time according to Eqs. 3 and 4, then binning. The acceleration fields are calculated by finding the mean values within each bin.

$$a_x = \Delta v_x / \Delta t \tag{3}$$

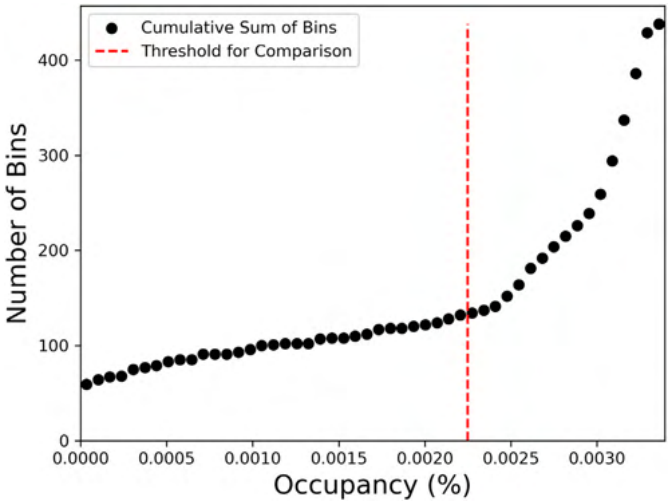

Fig. 10. The cumulative probability of occupancy is used to determine which bins to consider for comparison.

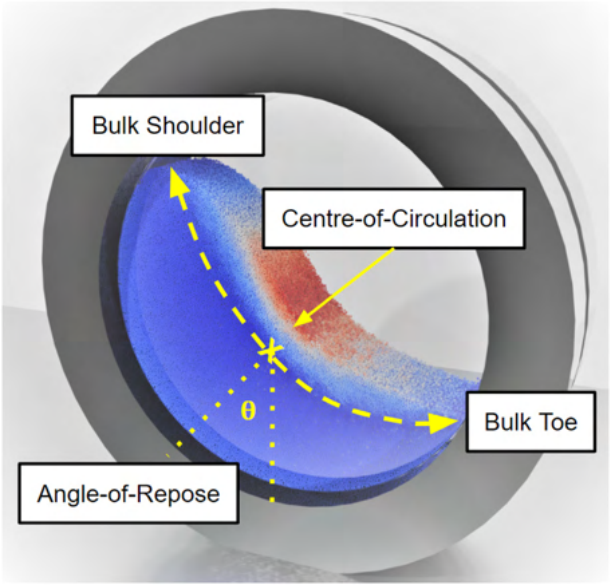

Fig. 12. Diagram of the bulk shoulder, bulk toe, and centre-of-circulation extracted from the equilibrium surface.

$$a = \sqrt{a_x^2 + a_y^2} \tag{4}$$

The velocity RSD and granular temperature are two statistics that use fluctuations from the mean velocity to measure the bulk behaviour of the system. Velocity RSD is a measurement of how the standard deviation in velocity of an area of the system differs from the mean velocity in that area, measured as a percent difference. Higher velocity RSD values show the tracer trajectory is deviating more substantially from the mean flow behaviour. To compute the velocity RSD, first, the standard deviation in all velocity components within a bin are calculated, then the magnitude of the velocity standard deviation is divided by the mean velocity in the bin. The velocity standard deviation and RSD are shown in Eqs. (5) and (6).

$$\sigma_{v_x} = \sqrt{\frac{\Sigma(v_{xi} - \overline{v}_x)^2}{N}} \tag{5}$$

$$\sigma_{RSD} = 100\frac{\sqrt{\sigma_{v_x}^2 + \sigma_{v_y}^2}}{\overline{v}} \tag{6}$$

On the other hand, the granular temperature is a measure of the squared velocity fluctuations and is not scaled by the mean velocity [49]. Thus, larger magnitude velocity fluctuations will always produce higher granular temperatures. This is shown in Eqs. (7) and (8), through which the 1-dimensional and 2-dimensional granular temperatures are calculated. To compute the granular temperature field, first, the x and y velocity fields are calculated and the bin number for each instantaneous velocity is recorded. Using this information, the mean velocities are subtracted from the instantaneous velocity to produce the fluctuating velocity according to the tracer location within the system. Then the granular temperature can be calculated using the squared velocity fluctuations and multiplying by the mass of the particle [49].

$$T_x = m(v_{xi} - \overline{v}_x)^2 \tag{7}$$

$$T = \frac{1}{2}(T_x + T_y) \tag{8}$$

The occupancy measures how long a particle spends within a region of the system. This is useful since occupancy is an indirect measurement of particle density in different areas of the systems, which can be used to infer particle dynamics in industrial systems as has recently been demonstrated in a coffee-roaster [50]. In this work, the DEM and PEPT trajectory positions are recorded in constant intervals of time and continuously tracked, thus the occupancy percentage for a bin can be calculated by counting how many times a particle has been found in that bin, multiplying by the time-step, then dividing by the difference in time between the last and first particle position in the trajectory. The occupancy

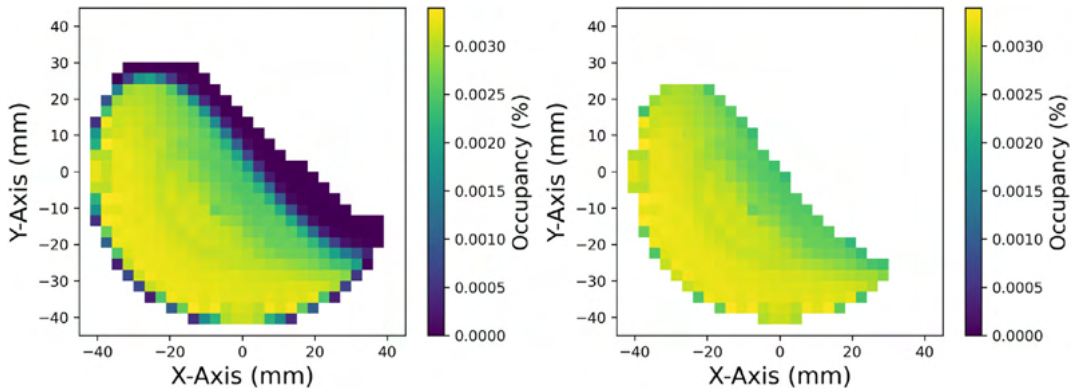

Fig. 11. (Left) Apply a threshold to removes bins which the tracer explores only a few times. (Right) Bins subject to the threshold are the only bins used for comparison in later analysis.

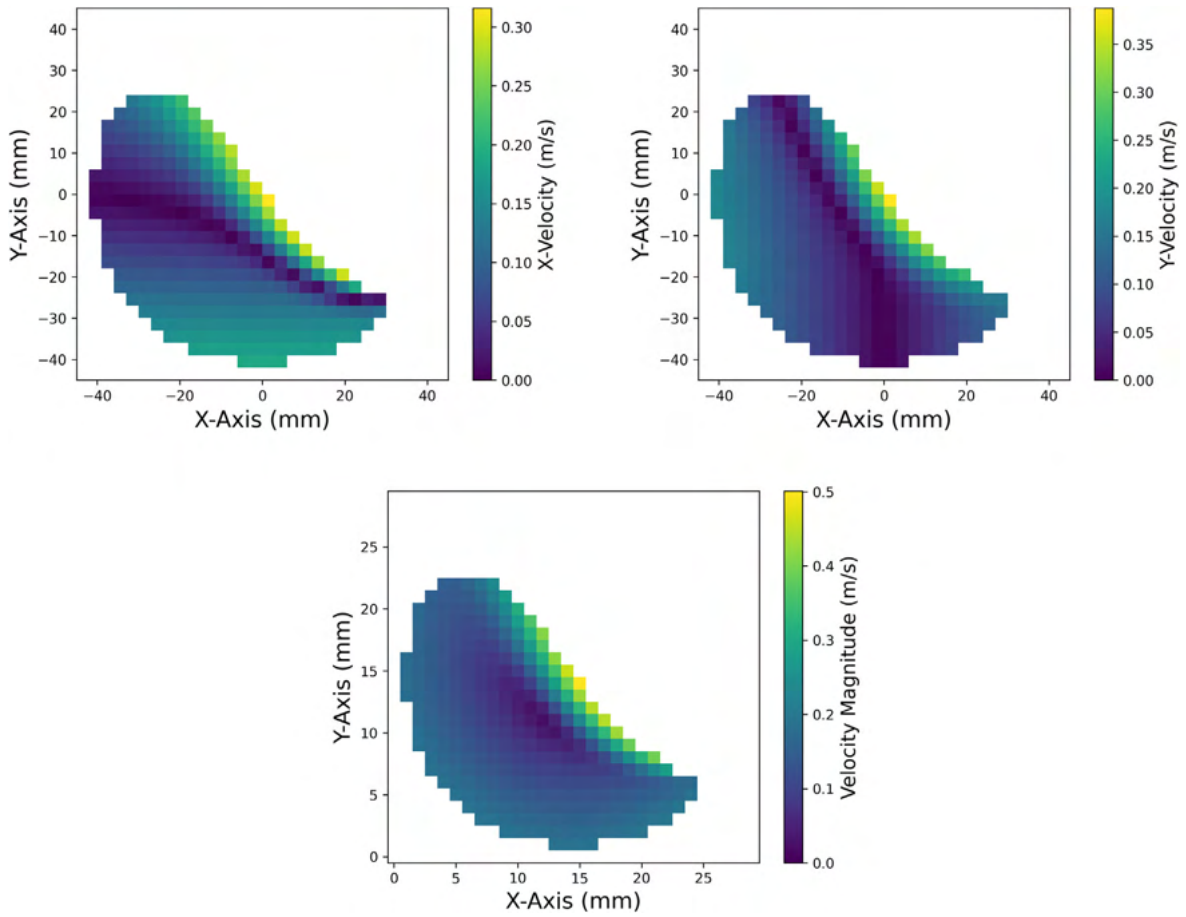

Fig. 13. (Top left) x-velocity and y-velocity (top right) of the full DEM data. (Bottom) x- and y-velocity fields combined to form the magnitude velocity field of the full DEM data.

is calculated using Eq. 9. The calculation of most of these fields is described in greater detail in 'Positron Emission Particle Tracking: A Comprehensive Guide' by Windows-Yule et al. 2022, IOP [51].

$$O = 100\frac{N_{points}t_{step}}{t_f - t_0} \tag{9}$$

2.5.2. Comparing Eulerian fields

All Eulerian fields for DEM and PEPT are binned in 3×3 mm bins, computed over the same range, and using the same number of bins. To compare the fields, the mean bin percent difference is calculated. The percent difference is calculated by finding the absolute difference between the PEPT reconstructed fields and baseline DEM model fields,

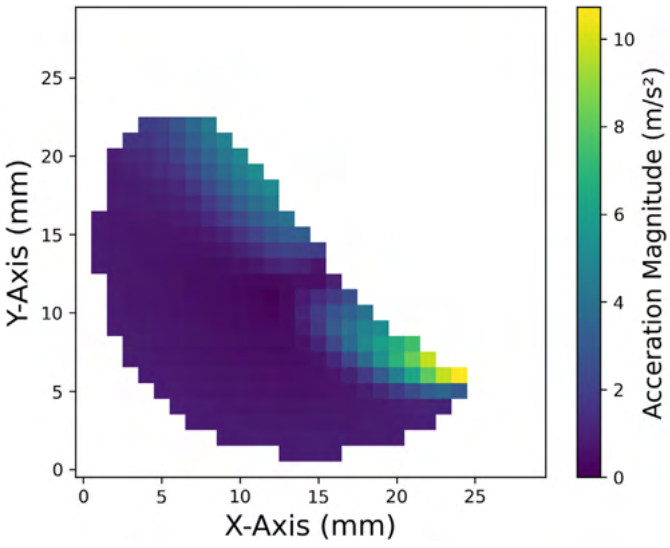

Fig. 14. The acceleration magnitude field using all DEM data.

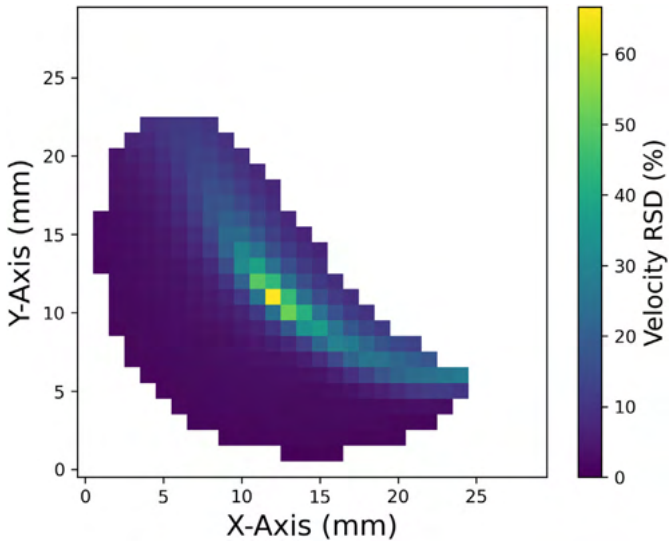

Fig. 15. The velocity RSD field is generated from all model DEM trajectories in the monodisperse GranuDrum.

then dividing by the DEM model value and multiplying by 100. The individual bin percent differences are then averaged as shown in Eq. (10).

$$\varepsilon_{Mean} = \frac{100}{N_{bins}} \sum_{N=1} N_{bins} \frac{|X_{Experiment} - X_{Model}|}{X_{Model}} \quad (10)$$

However, to avoid outlier bin errors distorting the mean percent difference due to a low number of tracer passes, a threshold is set using the occupancy so that only bins which have an occupancy higher than 0.00225% are considered. This value was chosen by looking at the cumulative probability distribution of the bin occupancy, as shown in Fig. 10. An occupancy of 0.00225% immediately precedes a sharp rise in the plot, thereby excluding only the bins which are not adequately explored and do not contribute significantly to the overall system behaviour. The results of this threshold applied to the occupancy are shown in Fig. 11.

2.5.3. Scalar indicators

Scalar indicators such as the bulk shoulder, bulk toe, CoC, and the AoR can be used to characterise the behaviour of a rotating drum by classifying the flow regime [13]. All of these values can be derived by calculating the equilibrium surface which is created by combining the horizontal (x-velocity) and vertical (y-velocity) equilibrium surfaces. These are the surfaces where the net flow in each respective direction equals zero and are computed in Section 2.5.1 [52]. The point where the two surfaces intersect is the CoC. The points where the equilibrium surface intersects the drum wall are the bulk shoulder and bulk toe; the shoulder is the upper intersection and the toe is the lower one. The equilibrium surface is the continuous surface that connects these three points. The AoR is defined as the angle between the tangent to the equilibrium surface at the CoC and the downward direction. These scalar indicators including the equilibrium surface are shown in Fig. 12.

3. Results and discussion

Our results can be summarised according to the three main components of this work: 1) Analysis of the DEM simulations, 2) Comparison of the PEPT reconstruction with the DEM model, and 3) Demonstration of a new method to correct PEPT reconstructed values based on the model DEM.

3.1. Analysis of DEM simulations

3.1.1. Mono-disperse analysis

Using the full 30-min trajectories extracted from the DEM simulation of the mono-disperse GranuDrum, the Eulerian fields and scalar indicators described in Sections 2.5.1 and 2.5.3 were generated. In Fig. 13, the velocity magnitude field is shown, revealing an active, fast region on the free surface whilst having a passive and slow-moving region in the bulk; both are well-known features of rotating drums operating at the Froude number of 0.1 [53].

Numerically differentiating the tracer's velocity with respect to time yields the acceleration. The acceleration magnitude field is produced and shown in Fig. 14. The shoulders and toes experience the highest accelerations since particles in the shoulder falling down the cascade are accelerated by gravity and particles in the toe are abruptly decelerated by crashing into the drum wall.

The velocity RSD field is shown in Fig. 15. A band of high velocity RSD values reaches from the shoulder to the toe of the drum. These are likely caused by cascade fluctuations which cause velocity to change drastically in both magnitude and direction, leading to a large standard deviation in velocity. The measurement of cascade fluctuations is an important part of the GranuDrum, however, in PEPT measurements only a single (or small handful of) particle's behaviour can be seen. Thus PEPT lends itself better to quantifying the time-averaged behaviour of the system, yet this shows that single particle measurements can indirectly capture cascade fluctuations, at least in this system.

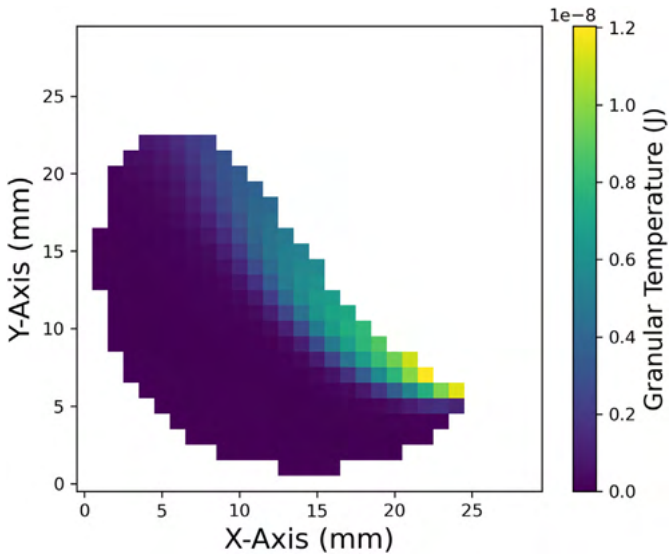

Fig. 16. Granular Temperature in the mono-disperse GranuDrum DEM simulation.

The granular temperature is shown in Fig. 16. The granular temperature is highest in the active region of the drum and especially in the toe. This is because collisions in this area are frequent and energetic, leading to large fluctuations in velocity [54].

These fields were then generated for every minute of the simulation time for individual DEM particles. These are compared to the previously shown fields produced after 30 min using all of the particles which represent the steady-state behaviour of the drum. This is done to investigate how the measured system behaviour reconstructed from a single particle begins to break down as the measurement is reduced. The results of this comparison are shown in Fig. 17.

Across all measured fields, the differences from steady-state fields are found to logarithmically decrease with the measurement time. Derived measurements of lower-order than acceleration, such as velocity, are found to reach a high level of accuracy relatively quickly; 1% mean difference after 10 min of measurement time. Higher-order derived measurements like acceleration, which rely on the velocity measurement, take longer to reach the same level of accuracy. For all fields, the mean percent differences are under 10% by 30 min. These results

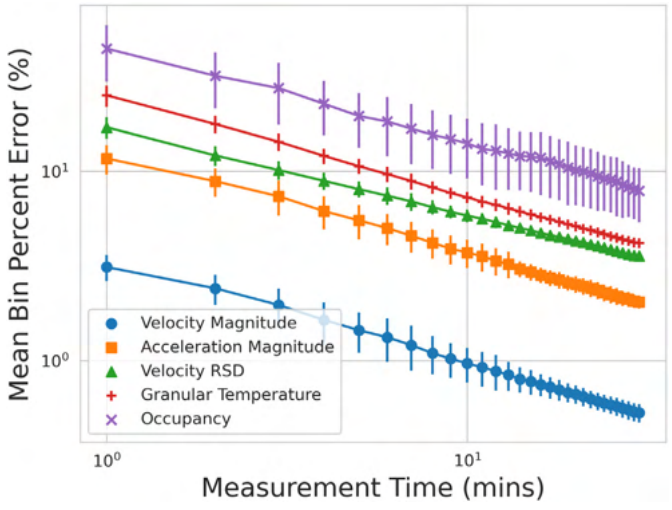

Fig. 17. Estimating the time needed to ensure adequate statistics using results from multiple DEM tracers for different Eulerian fields. Error bars are the standard deviation of the percent differences.

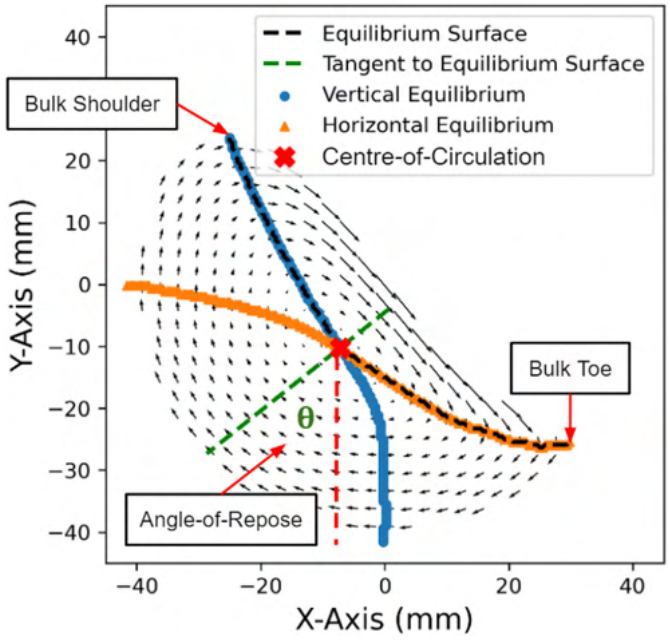

Fig. 18. A diagram showing the equilibrium surface and extracted scalar indicators using all mono-disperse DEM trajectories.

show for this particular system not only can system behaviour be reconstructed from single particles, but also the length of measurement time needed to reach a desired level of accuracy can be estimated.

In addition to the Eulerian fields, scalar indicators are calculated after 30 min. These indicators are the bulk shoulder, bulk toe, CoC, and the AoR shown in Fig. 18. Each scalar indicator is extracted at every minute for each trajectory, then compared to the value produced after 30 min. The bulk shoulder, bulk toe, and CoC are compared by the location of the indicator while the AoR is compared by the difference in degrees. The results are shown in Fig. 19, revealing these values converge constant values within 5 min for the bulk shoulder and bulk toe and within 20 min for the CoC and AoR. Scalar indicators are thus found to need shorter times to get accurate measurements than 2-dimensional Eulerian fields. This further supports the argument that measurements that rely on higher-order, derived values (e.g. velocity derived from position, and acceleration derived from velocity) the uncertainty in the original measurement propagates upward.

An unexpected result is the equilibrium surface, and thus the bulk shoulder, bulk toe, and CoC, intersect the bins with the high velocity RSD. The equilibrium surface is overlaid onto the velocity RSD field in Fig. 20to demonstrate this. While the relationship between velocity RSD and equilibrium surface is not explored in this work, nevertheless Fig. 20 suggests velocity RSD is a measurement that can reveal useful information about rotating drums.

3.1.2. Poly-disperse analysis

In a poly-disperse system, the assumption that a single particle is representative of all the particles of the system no longer holds. When particles have different sizes or densities, segregation can occur, leading to different particle behaviours depending on the species being tracked. This has been demonstrated numerous times using the PEPT technique across a wide range of systems [4,13,25,55,56]. Fig. 21 shows the occupancy distribution for the largest, smallest, and all the DEM particles in the poly-disperse GranuDrum simulation described in Section 2.1.2. A PEPT measurement conducted with a tracer taken from the extremes of the particle distribution would not be reflective of the occupancy of all particles. When the mean percent differences in the occupancy are plotted against the particle size in Fig. 22, this shows, as might be expected, that particles most similar to the mean particle diameter best replicate the system occupancy. However, while occupancy, in this case, is best reconstructed with the mean particle size, this does not guarantee other types of measurements will also be reconstructed to the same degree. Therefore, when conducting a PEPT experiment attention should be paid to whether the particle being tracked actually mimics the bulk system behaviour. Using the tools provided with DEM2GATE, a poly-disperse simulation can help users determine which particles are good candidates as a PEPT tracer prior to experimentation.

3.2. PEPT reconstruction

After running the GATE simulations of a particle in air and of a particle in the simulated GranuDrum, the reconstructed PEPT trajectories are compared to the prescribed DEM trajectory. For a 2 s segment of the simulation, the positions for the DEM tracer and the two PEPT reconstructed trajectories are plotted in Fig. 23, then compared point-by-point to compute a mean spatial error. Further, the instantaneous

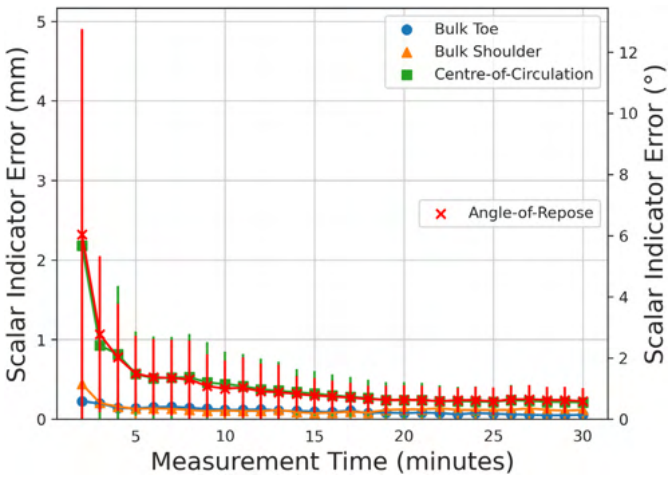

Fig. 19. The difference in location and angle of the four scalar indicators in the mono-disperse GranuDrum simulation over time compared to their location using all the trajectories after 30 min. The error bars represent the standard deviation of the individual trajectories.

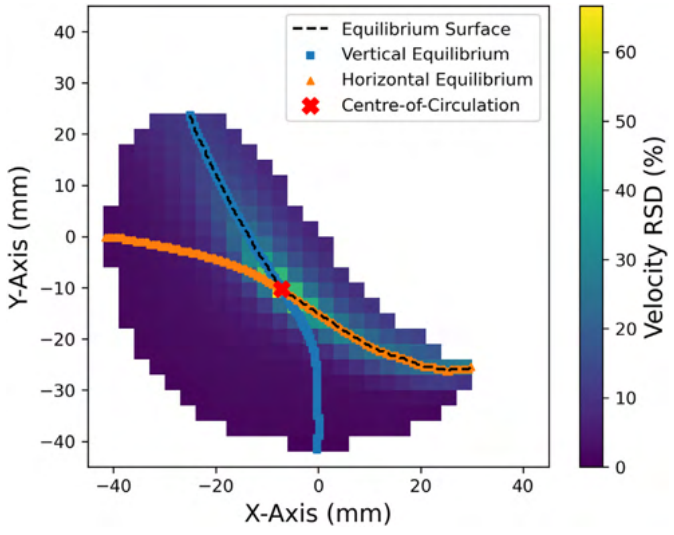

Fig. 20. The equilibrium surface intersects the highest velocity RSD values and the bin with the high velocity RSD contains the CoC.

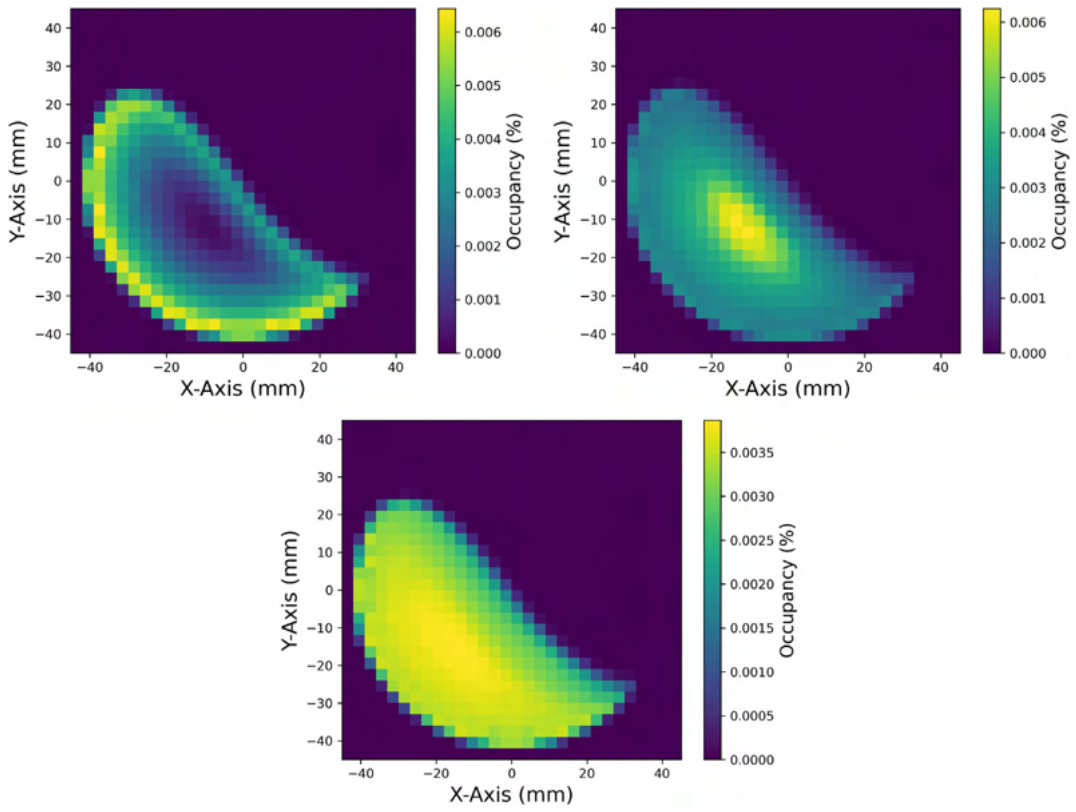

Fig. 21. Poly-dispersed occupancy shows large particle size segregation (top left), small particle size segregation (top right) and the occupancy for all particles (bottom).

spatial errors are plotted according to the tracer's positions in Fig. 24. The tracer in air and tracer in the GranuDrum voxels over their 30 min trajectory are tracked to a mean spatial accuracy of 0.175 mm and 0.412 mm, respectively. Additionally, the detection rates of the two tracers are 889 Hz and 355 Hz, respectively. This demonstrates that DEM2GATE can be used to reconstruct particle behaviour and then quantify the spatial and temporal resolution of the reconstruction. It is worth noting that, in PEPT, spatial accuracy is always in terms of length and this spatial accuracy does not directly depend on the size of the system, though larger systems typically cause more scattering to occur.

The spatial errors can be broken down further into the mean spatial error in each bin. The spatial errors for the two GATE simulations are shown in Fig. 25. These plots show a more complete story about what is happening in Fig. 24, showing higher errors when the tracer is falling from the shoulder of the drum or crashing into the toe region. These are the two areas that experience the highest acceleration, as can be seen in Fig. 26. When a bin-by-bin comparison between the spatial error in PEPT measurement and the model DEM particle acceleration is plotted in Fig. 27, the relationship between the two is linearly increasing spatial error with increasing tracer acceleration. When a straight line is fitted to the data, the y-intercept is the mean error that is to be expected for a

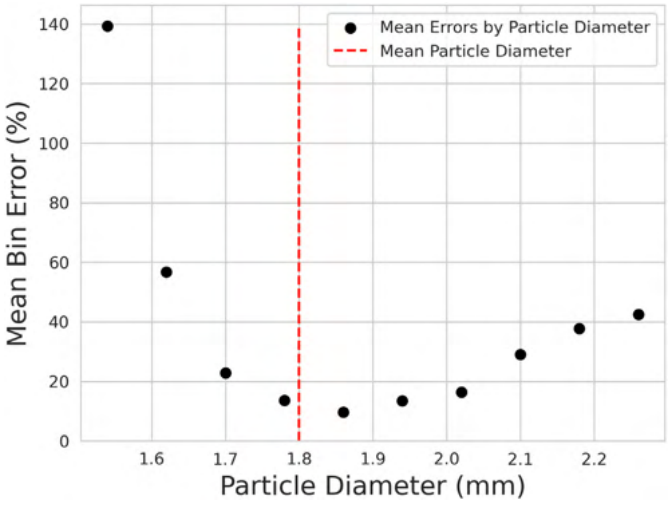

Fig. 22. The mean bin error in the occupancy across all particle sizes in the poly-disperse DEM simulation.

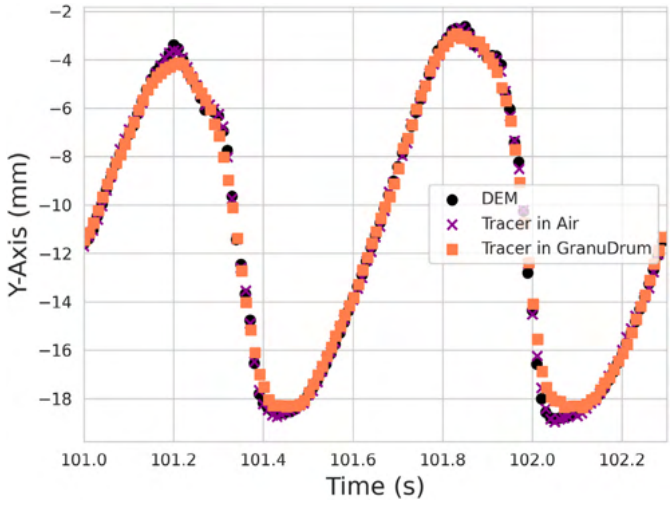

Fig. 23. Comparison of a DEM trajectory to the PEPT detected trajectory of the tracer in air and tracer in the GranuDrum.

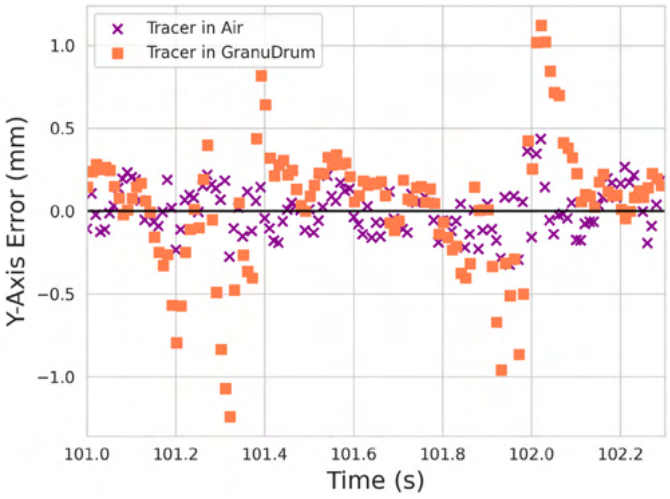

Fig. 24. Errors in the vertical y-axis for tracer in air and tracer in GranuDrum over a 1-s time-step.

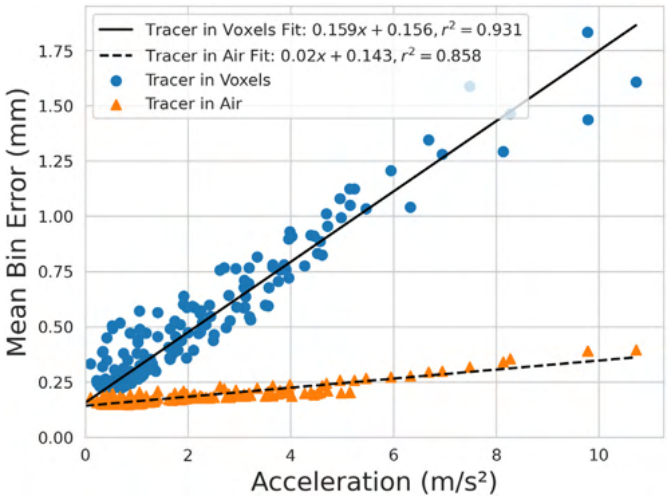

Fig. 27. The PEPT spatial error increases linearly with acceleration, with the y-intercept being the mean error for a constant velocity tracer; the y-intercept is higher for the PEPT tracer in voxels because of the increased scattering of LoRs.

tracer at rest or at constant velocity. The y-intercept value is higher for the tracer in the GranuDrum voxels, most likely due to the increase in scattered LoRs as this has been shown to increase PEPT error in other work and thus demonstrates that this is an important relation with real physical significance [12].

Additional analysis of the PEPT trajectories is performed by reconstructing the Eulerian fields and scalar indicators. These are calculated in the same manner as the DEM fields. These values are calculated every minute and compared to the DEM values calculated using all the model DEM trajectories after 30 min of the simulation. The results of comparing the Eulerian fields are shown in Fig. 28. Fig. 28 shows that for the velocity, acceleration, and occupancy fields the PEPT reconstruction aligns with the model DEM fields, but for velocity RSD and granular temperature, the fields do not align and do not get more accurate with longer measurement times. The reason for this is explained in Section 3.3 and a novel method to correct these errors demonstrated.

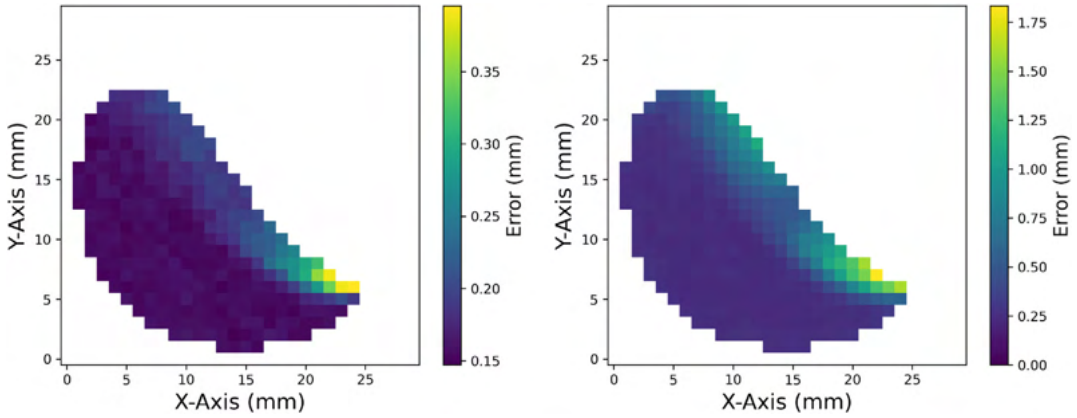

Fig. 25. (Left) Error field for the PEPT tracer in air. (Right) Error field for the PEPT tracer in the GranuDrum.

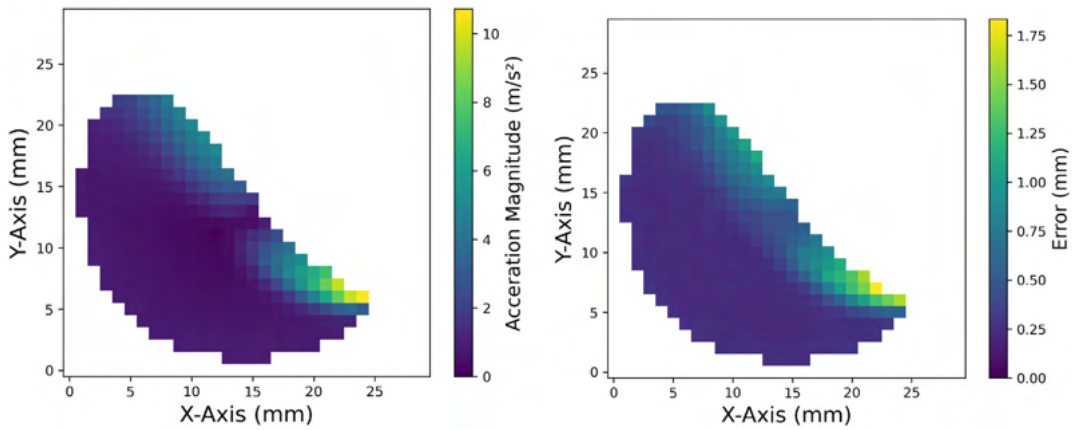

Fig. 26. (Left) Acceleration field using all DEM data and (right) error in the PEPT data with the tracer in the GranuDrum.

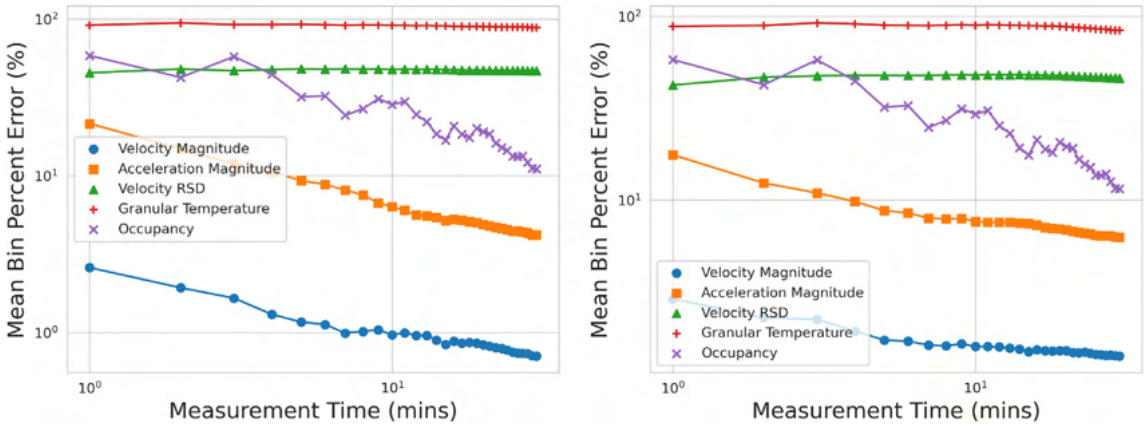

Fig. 28. Estimating the time needed to ensure adequate statistics using results from a PEPT tracer in air (left) and in the GranuDrum (right). Notice the velocity RSD and Granular Temperature are not improving for either case.

Table 3
Scalar indicators of the mono-disperse GranuDrum DEM simulation and from the reconstructed virtual PEPT experiments of the tracer in air and the tracer in the GranuDrum voxels after 30 min of measurement time.

Scalar indicator	Model DEM	Tracer in air	Tracer in GranuDrum voxels
Bulk Shoulder (x and y)	−25.01, 23.71 mm	−25.01 23.71 mm	−25.45, 23.71 mm
Bulk Toe (x and y)	29.80, −25.45 mm	29.80–25.45 mm	29.80, −25.45 mm
Centre-of-Circulation (x and y)	−7.18, −10.22	−7.18 -10.22	−6.96, −10.44 mm
Angle-of-Repose	38.47°	38.87	40.15°

The scalar indicators can be analysed in a similar fashion. As before, the bulk shoulder, bulk toe, and CoC are compared by the location of the indicator while the AoR is compared by their difference in degrees. The results for the tracer in air and the tracer in the GranuDrum voxels resulted in nearly constant reconstructed values which agree with the model DEM values throughout the 30 min of the measurement time. The values reconstructed after 30 min are given in Table 3. However, for the tracer in the GranuDrum GATE simulation, the values of the CoC and AoR are consistently different by approximately 0.5 mm and 1.5°, respectively. This shows that as more scattering occurs and the quality of the PEPT trajectory is degraded this also degrades other PEPT reconstructed quantities as well.

3.3. Using DEM2GATE to correct PEPT measurements

As described in Section 3.2 the PEPT reconstructed velocity RSD and granular temperature fields (see Fig. 28) did not agree with the model DEM fields. These two fields rely on the fluctuating velocity components and, as such, they are particularly prone to errors caused by both low temporal resolution and the inherent uncertainty in PEPT measurements. The low temporal resolution of a PEPT measurement in comparison with the behaviour of the DEM tracer smooths out fluctuations in velocity while uncertainty in the PEPT measurement induces fluctuations in velocity. These two factors come together to produce large errors in the PEPT reconstructed fields.

However, when the PEPT and model DEM fields are compared bin-by-bin there exists a correlation between them, shown in Fig. 29. When the velocity RSD and square root of the granular temperature (square root since this is calculated using a squared velocity term) for PEPT reconstructed and model DEM are compared, a straight line can be fitted to the data. The fitted equations show that bins with low velocity RSD and granular temperature in the model DEM fields are over-predicted by PEPT, as evidenced by a positive y-intercept. Further, at high velocity RSD and granular temperatures in the model DEM, PEPT

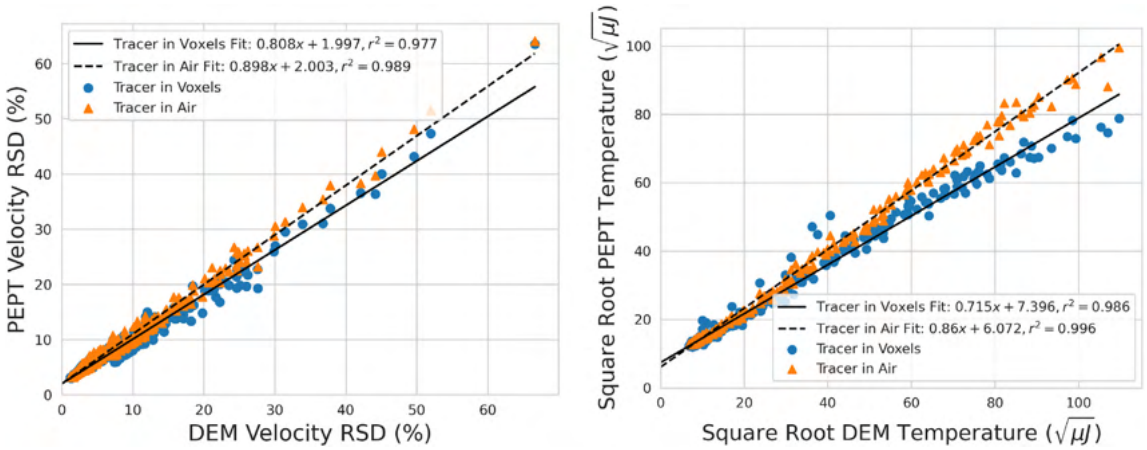

Fig. 29. Linear correlations between the PEPT reconstructed and model DEM fields can be used to correct the PEPT velocity RSD and granular temperature.

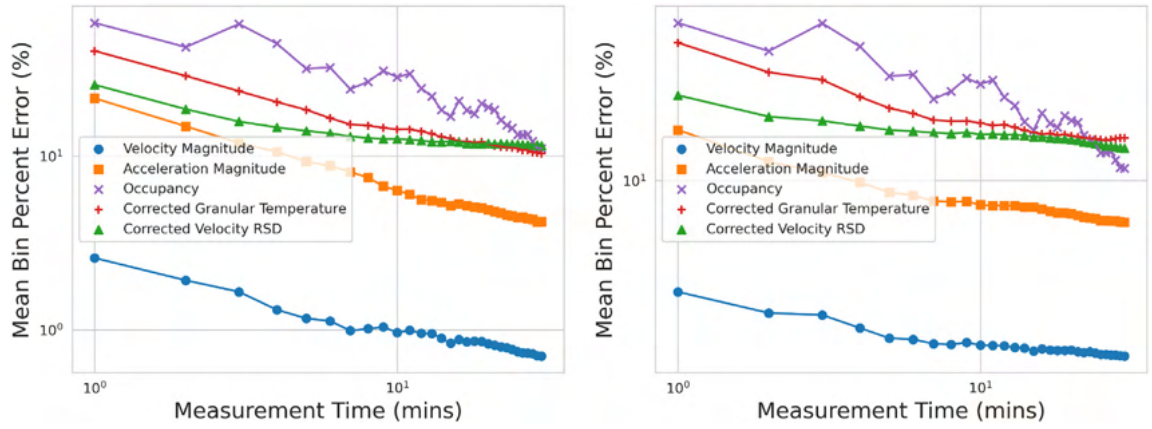

Fig. 30. The mean bin percent differences for the PEPT reconstructed tracer in air (left) and the tracer in the GranuDrum (right) over the measurement time. When the correction for velocity RSD and granular temperature is applied, the mean bin percent differences for these fields align are improved, reaching approximately 10% after 30-min.

measurements under-predict the values, as evidenced by a slope gradient less the 1.

Not only does this result provide insight into a possible source of error when comparing PEPT measurements with DEM models, but also shows that the PEPT measurements *have predictable errors which can be corrected*. To correct the PEPT measurement, the terms of the linear model can be rearranged and applied to the PEPT reconstructed fields. When this is done, the mean bin percent difference over the measurement time aligns more closely to what is expected from the model DEM simulation, as shown in Fig. 30. This correction method is robust because it is not dependent on the overall behaviour of the particle within the system, only its velocity fluctuations. Thus, in order to produce a correlation between the PEPT reconstructed values and model DEM values, a fully calibrated DEM model is not required, only a model which produces velocity fluctuations across the range which might be expected in the experiment.

4. Conclusion

In this work, we demonstrate the use of novel tools in the open-source DEM2GATE library to combine DEM simulation with virtual PEPT experiments. A LIGGGHTS DEM simulation of mono-disperse and poly-disperse MCC particles in a model of the GranuTools GranuDrum was simulated and the mono-disperse simulation was used as the basis of a virtual PEPT experiment. From this DEM simulation, Lagrangian trajectories of individual particles were extracted at regular time steps of 1 ms. Along with these trajectories, a volume representing the time-averaged particle density throughout the rotating drum was produced.

A GATE simulation was created using the extracted DEM trajectories to move a 10 MBq positron-emitting source. In a separate GATE simulation, the volume of particle density and a structural model of the GranuDrum was added to replicate gamma-ray scattering. The synthetic LoRs collected over the 30 min-long trajectory were processed using the PEPT-ML algorithm, yielding a PEPT tracked version of the trajectory in both virtual experiments. The reconstructed trajectories were compared with the prescribed model DEM trajectories through a point-by-point comparison, as well as through 2-dimensional Eulerian fields and scalar indicators. This analysis showed that the PEPT reconstructed trajectories were found to have a mean 2-dimensional spatial error of 0.175 mm and 0.412 mm for the tracer in air and tracer in the simulated GranuDrum, respectively. Importantly, two drivers of spatial error were found and quantified for the PEPT measurement, particle acceleration and gamma-ray scattering.

Further, the accuracy of reconstructed fields and scalar indicators are shown to depend on the length of measurement time. The DEM data show that accuracy improves logarithmically for the reconstructed fields and that higher-order derived measurements such as acceleration take longer to reach a desired level of accuracy compared to lower-order derived measurements like velocity. After 30-min of a single-particle virtual PEPT measurement velocity, acceleration, and occupancy fields were reconstructed to within 10%. However, the velocity RSD and granular temperature PEPT measurements produced mean bin percent errors of approximately 50% and above. A correction for the PEPT data was developed by comparing the PEPT and DEM values for these two measurements on a bin-by-bin basis, finding they were linearly correlated. Low model DEM values in these fields were over-predicted in the PEPT measurement, then under-predicted for higher values. When the PEPT measurement correction was applied, the velocity RSD and granular temperature for both virtual PEPT experiments after 30 min of measurement time reached approximately a 10% mean bin error.

This work shows, to the extent that DEM simulations replicate the real behaviour of particles in granular media, that virtual PEPT experiments faithfully reproduce the PEPT measured Lagrangian particle trajectories with high spatio-temporal resolution. The Eulerian fields and scalar indicators reconstructed from the single-particle virtual PEPT experiments are capturing the dynamics produced from full-field DEM data, thus suggesting that the PEPT experiments are capturing real system behaviour. The methods introduced in this work can be similarly applied to any DEM simulated granular system, provided that the DEM output has been verified to represent realistic flow behaviour of the system. A trajectory extracted from the DEM simulation can be used to assess whether the expected trajectories from a real PEPT experiment are of good enough resolution to track the real granular motion, and thus would be sufficient to calibrate a DEM model.

CRediT authorship contribution statement

M.T. Herald: Conceptualization, Methodology, Visualization, Writing – original draft. **J.A. Sykes:** Software, Writing – original draft, Writing – review & editing. **D. Werner:** Resources, Software, Validation. **J.P.K. Seville:** Supervision, Writing – review & editing. **C.R.K. Windows-Yule:** Supervision, Writing – review & editing.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors would like to recognise the valuable contributions made by Andrei Leonard Nicuşan for development of the pept Python package

on GitHub. The computations described in this paper were performed using the University of Birmingham's BlueBEAR HPC service, which provides a High Performance Computing service to the University's research community. See http://www.birmingham.ac.uk/bear for more details [46]. This work was supported by the Engineering and Physical Sciences Research Council (grant number EP/T034327/1); and the Engineering and Physical Sciences Research Council Centre for Doctoral Training in Topological Design (grant number EP/S02297X/1).

References

[1] S. Bhattacharya, P.P. Vlachos, Volumetric particle tracking velocimetry (ptv) uncertainty quantification, Exp. Fluids 61 (9) (2020) 197, https://doi.org/10.1007/s00348-020-03021-6.
[2] L. Farber, G. Tardos, J.N. Michaels, Use of x-ray tomography to study the porosity and morphology of granules, Powder Technol. 132 (1) (2003) 57–63, https://doi.org/10.1016/S0032-5910(03)00043-3.
[3] S. Wiederseiner, N. Andreini, G. Epely-Chauvin, C. Ancey, Refractive-index and density matching in concentrated particle suspensions: a review, Exp. Fluids 50 (5) (2011) 1183–1206, https://doi.org/10.1007/s00348-010-0996-8.
[4] C.R.K. Windows-Yule, J.P.K. Seville, A. Ingram, D.J. Parker, Positron emission particle tracking of granular flows, Annu. Rev. Chem. Biomol. Eng. 11 (1) (2020) 367–396, https://doi.org/10.1146/annurev-chembioeng-011620-120633, URL https://www.annualreviews.org/doi/10.1146/annurev-chembioeng-011620-120633.
[5] D.J. Parker, R.N. Forster, P. Fowles, P.S. Takhar, Positron emission particle tracking using the new Birmingham positron camera, Nucl. Instrum. Meth. Phys. Res. Sect. A 477 (1–3) (2002) 540–545, https://doi.org/10.1016/S0168-9002(01)01919-2, URL https://linkinghub.elsevier.com/retrieve/pii/S0168900201019192.
[6] D.J. Parker, C.J. Broadbent, P. Fowles, M.R. Hawkesworth, P. McNeil, Positron emission particle tracking - a technique for studying flow within engineering equipment, Nucl. Instrum. Meth. Phys. Res. Sect. A 326 (3) (1993) 592–607, https://doi.org/10.1016/0168-9002(93)90864-E, URL https://linkinghub.elsevier.com/retrieve/pii/016890029390864E.
[7] D.J. Parker, Positron emission particle tracking and its application to granular media, Rev. Sci. Instrum. 88 (5) (2017) 051803 https://doi.org/10.1063/1.4983046.
[8] A.K. Prasad, R.J. Adrian, C.C. Landreth, P.W. Offutt, Effect of resolution on the speed and accuracy of particle image velocimetry interrogation, Exp. Fluids 13 (2) (1992) 105–116, https://doi.org/10.1007/BF00218156.
[9] C.R.K. Windows-Yule, B.J. Scheper, A.J.V.-D. Horn, N. Hainsworth, J. Saunders, D.J. Parker, A.R. Thornton, Understanding and exploiting competing segregation mechanisms in horizontally rotated granular media, New J. Phys. 18 (2) (2016), 023013https://doi.org/10.1088/1367-2630/18/2/023013, URL https://iopscience.iop.org/article/10.1088/1367-2630/18/2/023013.
[10] S. Langford, Jet Flow Validation of Positron Emission Particle Tracking Utilizing High Speed Video, Masters Theses. URL https://trace.tennessee.edu/utk_gradthes/3782 May 2016.
[11] X. Fan, D.J. Parker, M.D. Smith, Labelling a single particle for positron emission particle tracking using direct activation and ion-exchange techniques, Nucl. Instrum. Meth. Phys. Res. Sect. A 562 (1) (2006) 345–350, https://doi.org/10.1016/j.nima.2006.03.015.
[12] C.R.K. Windows-Yule, M.T. Herald, A.L. Nicuşan, C.S. Wiggins, G. Pratx, S. Manger, A.E. Odo, T. Leadbeater, J. Pellico, R.T.M. de Rosales, A. Renaud, L. Govender, L.B. Carasik, A.E. Ruggles, T. Kokalova-Wheldon, J.P.K. Seville, D.J. Parker, Recent Advances in Positron Emission Particle Tracking: A Comprehensive Review and Benchmarking, Reports on Progress in Physics, In review 2021.
[13] A.J. Morrison, I. Govender, A.N. Mainza, D.J. Parker, The shape and behaviour of a granular bed in a rotating drum using eulerian flow fields obtained from pept, Chem. Eng. Sci. 152 (2016) 186–198, https://doi.org/10.1016/j.ces.2016.06.022.
[14] E.M.E. Ahmed, I. Govender, A. Mainza, The geometric axial surface profiles of granular flows in rotating drums, J. South. Afr. Inst. Min. Metall. 121 (5) (2021) https://doi.org/10.17159/2411-9717/1228/2021, URL http://ref.scielo.org/mqb2xx.
[15] T.M. Povall, I. Govender, A.T. McBride, Dense granular flows in rotating drums: a computational investigation of constitutive equations, Powder Technol. 393 (2021) 238–249, https://doi.org/10.1016/j.powtec.2021.07.051, URL https://linkinghub.elsevier.com/retrieve/pii/S0032591021006410.
[16] M. Stein, Y.L. Ding, J.P.K. Seville, D.J. Parker, Solids motion in bubbling gas fluidised beds, Chem. Eng. Sci. 55 (22) (2000) 5291–5300, https://doi.org/10.1016/S0009-2509(00)00177-9, URL https://linkinghub.elsevier.com/retrieve/pii/S0009250900001779.
[17] M. Marigo, M. Davies, T. Leadbeater, D.L. Cairns, A. Ingram, E.H. Stitt, Application of positron emission particle tracking (pept) to validate a discrete element method (dem) model of granular flow and mixing in the turbula mixer, Int. J. Pharm. 446 (1–2) (2013) 46–58, https://doi.org/10.1016/j.ijpharm.2013.01.030 , URL https://linkinghub.elsevier.com/retrieve/pii/S0378517313000616.
[18] A. Ingram, J.P.K. Seville, D.J. Parker, X. Fan, R.G. Forster, Axial and radial dispersion in rolling mode rotating drums, Powder Technol. 158 (1–3) (2005) 76–91, https://doi.org/10.1016/j.powtec.2005.04.030, URL https://linkinghub.elsevier.com/retrieve/pii/S0032591005001592.
[19] A.L. Nicuşan, C.R.K. Windows-Yule, Positron emission particle tracking using machine learning, Rev. Sci. Instrum. 91 (1) (2020) 013329 https://doi.org/10.1063/1.5129251, URL http://aip.scitation.org/doi/10.1063/1.5129251.
[20] C. Wiggins, R. Santos, A. Ruggles, A feature point identification method for positron emission particle tracking with multiple tracers, Nucl. Instrum. Meth. Phys. Res. Sect. A 843 (2017) 22–28, https://doi.org/10.1016/j.nima.2016.10.057 , URL https://linkinghub.elsevier.com/retrieve/pii/S0168900216311184.
[21] K.S. Lee, T.J. Kim, G. Pratx, Single-cell tracking with pet using a novel trajectory reconstruction algorithm, IEEE Trans. Med. Imaging 34 (4) (2015) 994–1003, https://doi.org/10.1109/TMI.2014.2373351.
[22] P.A. Cundall, The Measurement and Analysis of Accelerations in Rock Slopes, 1971.
[23] R.P. Jensen, P.J. Bosscher, M.E. Plesha, T.B. Edil, DEM simulation of granular media—structure interface: effects of surface roughness and particle shape, Int. J. Numer. Anal. Methods Geomech. 23 (6) (1999) 531–547, https://doi.org/10.1002/(SICI)1096-9853(199905)23:6<531::AID-NAG980>3.0.CO;2-V.
[24] C.R.K. Windows-Yule, D.R. Tunuguntla, D.J. Parker, Numerical modelling of granular flows: a reality check, Comput. Part. Mech. 3 (3) (2016) 311–332, https://doi.org/10.1007/s40571-015-0083-2, URL http://link.springer.com/10.1007/s40571-015-0083-2.
[25] C.R.K. Windows-Yule, T. Weinhart, D.J. Parker, A.R. Thornton, Effects of packing density on the segregative behaviors of granular systems, Phys. Rev. Lett. 112 (9) (2014) 098001 https://doi.org/10.1103/PhysRevLett.112.098001.
[26] C. Kloss, C. Goniva, A. Hager, S. Amberger, S. Pirker, Models, algorithms and validation for opensource dem and cfd–dem, Progr. Comput. Fluid Dyn. Int. J. 12 (2–3) (2012) 140–152, https://doi.org/10.1504/PCFD.2012.047457.
[27] Granutools, Granudrum: Our Powder Cohesion Measurement Method, accessed October 22 vol. 2021, 2021, URL https://www.granutools.com/en/granudrum.
[28] E.R.L. Espiritu, A. Kumar, A. Nommeots-Nomm, J.A.M. Lerma, M. Brochu, Investigation of the rotating drum technique to characterise powder flow in controlled and low pressure environments, Powder Technol. 366 (2020) 925–937, https://doi.org/10.1016/j.powtec.2020.03.029.
[29] G. Lumay, F. Boschini, K. Traina, S. Bontempi, J.-C. Remy, R. Cloots, N. Vandewalle, Measuring the flowing properties of powders and grains, Powder Technol. 224 (2012) 19–27, https://doi.org/10.1016/j.powtec.2012.02.015.
[30] C.R.K. Windows-Yule, Ensuring adequate statistics in particle tracking experiments, Particuology (Dec 2020) https://doi.org/10.1016/j.partic.2020.10.013, URL https://www.sciencedirect.com/science/article/pii/S1674200120301267.
[31] S. Jan, G. Santin, D. Strul, S. Staelens, K. Assié, D. Autret, S. Avner, R. Barbier, M. Bardiès, P.M. Bloomfield, D. Brasse, V. Breton, P. Bruyndonckx, I. Buvat, A.F. Chatziioannou, Y. Choi, Y.H. Chung, C. Comtat, D. Donnarieix, L. Ferrer, S.J. Glick, C.J. Groiselle, D. Guez, P.-F. Honore, S. Kerhoas-Cavata, A.S. Kirov, V. Kohli, M. Koole, M. Krieguer, D.J.V.D. Laan, F. Lamare, G. Largeron, C. Lartizien, D. Lazaro, M.C. Maas, L. Maigne, F. Mayet, F. Melot, C. Merheb, E. Pennacchio, J. Perez, U. Pietrzyk, F.R. Rannou, M. Rey, D.R. Schaart, C.R. Schmidtlein, L. Simon, T.Y. Song, J.-M. Vieira, D. Visvikis, R.V.-D. Walle, E. Wieërs, C. Morel, GATE: a simulation toolkit for PET and SPECT, Phys. Med. Biol. 49 (19) (2004) 4543–4561, https://doi.org/10.1088/0031-9155/49/19/007, URL https://iopscience.iop.org/article/10.1088/0031-9155/49/19/007.
[32] G. Santin, D. Strul, D. Lazaro, L. Simon, M. Krieguer, M.V. Martins, V. Breton, C. Morel, GATE: a Geant4-based simulation platform for PET and SPECT integrating movement and time management, IEEE Trans. Nucl. Sci. 50 (5) (2003) 1516–1521, https://doi.org/10.1109/TNS.2003.817974, URL https://ieeexplore.ieee.org/document/1236960/.
[33] S. Raychaudhuri, Introduction to Monte Carlo simulation, 2008 Winter Simulation Conference, IEEE, Miami, FL, USA 2008, pp. 91–100, https://doi.org/10.1109/WSC.2008.4736059[, URL https://ieeexplore.ieee.org/document/4736059/.
[34] S. Lee, J. Gregor, D. Osborne, Development and validation of a complete gate model of the siemens inveon trimodal imaging platform, Mol. Imaging 12 (7) (2013) https://doi.org/10.2310/7290.2013.00058.
[35] M. Herald, Z. Bingham, R. Santos, A. Ruggles, Simulated time-dependent data to estimate uncertainty in fluid flow measurements, Nucl. Eng. Des. 337 (2018) 221–227, https://doi.org/10.1016/j.nucengdes.2018.07.005.
[36] M. Herald, T. Wheldon, C.R.K. Windows-Yule, Monte Carlo model validation of a detector system used for positron emission particle tracking, Nucl. Instrum. Meth. Phys. Res. Sect. A 993 (2021) 165073, https://doi.org/10.1016/j.nima.2021.165073, URL https://linkinghub.elsevier.com/retrieve/pii/S0168900221000577.
[37] T.E. Peterson, L.R. Furenlid, Spect detectors: the anger camera and beyond, Phys. Med. Biol. 56 (17) (2011) R145–R182, https://doi.org/10.1088/0031-9155/56/17/R01, URL https://iopscience.iop.org/article/10.1088/0031-9155/56/17/R01.
[38] H. Shi, G. Lumay, S. Luding, Stretching the limits of dynamic and quasi-static flow testing on cohesive limestone powders, Powder Technol. 367 (2020) 183–191, https://doi.org/10.1016/j.powtec.2020.03.036.
[39] S. Yu, C.-Y. Wu, M.J. Adams, G. Reynolds, B. Gururajan, J. Gargiuli, T. Leadbeater, R. Roberts, D.J. Parker, The use of positron emission particle tracking (pept) to study milling of roll-compacted microcystalline cellulose ribbons, Powder Technol. 285 (2015) 74–79, https://doi.org/10.1016/j.powtec.2015.06.051.
[40] D.J. Parker, A.E. Dijkstra, T.W. Martin, J.P.K. Seville, Positron emission particle tracking studies of spherical particle motion in rotating drums, Chem. Eng. Sci. 52 (13) (1997) 2011–2022, https://doi.org/10.1016/S0009-2509(97)00030-4, URL https://linkinghub.elsevier.com/retrieve/pii/S0009250997000304.
[41] R.Y. Yang, A.B. Yu, L. McElroy, J. Bao, Numerical simulation of particle dynamics in different flow regimes in a rotating drum, Powder Technol. 188 (2) (2008) 170–177, https://doi.org/10.1016/j.powtec.2008.04.081, URL https://linkinghub.elsevier.com/retrieve/pii/S0032591008002167.
[42] A.L. Nicuşan, Autonomous Characterisation and Calibration via Evolutionary Simulation, To be published. URL https://github.com/uob-positron-imaging-centre/Coexist 2022.
[43] Canty Lab Solidsizer, accessed March 01, 2022. URL https://www.jmcanty.com/product/lab-solidsizer/ 2022.
[44] G.T. Nolan, P.E. Kavanagh, Computer simulation of random packing of hard spheres, Powder Technol. 72 (2) (1992) 149–155, https://doi.org/10.1016/0032-5910(92)88021-9, URL https://www.sciencedirect.com/science/article/pii/0032591092880219.

[45] National Electronics Manufacturer Association, Performance Measurements of Positron Emission Tomographs, Tech. Rep. NU-2-2007 2007.

[46] University of birmingham bluebear, accessed March 01, 2022. URL https://intranet.birmingham.ac.uk/it/teams/infrastructure/research/bear/index.aspx 2022.

[47] R.J.G.B. Campello, D. Moulavi, J. Sander, Density-based clustering based on hierarchical density estimates, in: J. Pei, V.S. Tseng, L. Cao, H. Motoda, G. Xu (Eds.), Advances in Knowledge Discovery and Data Mining, Lecture Notes in Computer Science, Springer, Berlin, Heidelberg 2013, pp. 160–172, https://doi.org/10.1007/978-3-642-37456-2_14.

[48] A. Savitzky, M.J.E. Golay, Smoothing and differentiation of data by simplified least squares procedures, Anal. Chem. 36 (8) (1964) 1627–1639, https://doi.org/10.1021/ac60214a047, URL https://pubs.acs.org/doi/abs/10.1021/ac60214a047.

[49] R.D. Wildman, J.M. Huntley, J.-P. Hansen, D.J. Parker, D.A. Allen, Single-particle motion in three-dimensional vibrofluidized granular beds, Phys. Rev. E 62 (3) (2000) 3826–3835, https://doi.org/10.1103/PhysRevE.62.3826, URL https://link.aps.org/doi/10.1103/PhysRevE.62.3826.

[50] M. Al-Shemmeri, C.R.K. Windows-Yule, E. Lopez-Quiroga, P.J. Fryer, Coffee bean particle motion in a spouted bed measured using positron emission particle tracking (pept), J. Food Eng. 311 (2021) 110709, https://doi.org/10.1016/j.jfoodeng.2021.110709, URL https://linkinghub.elsevier.com/retrieve/pii/S026087742100234X.

[51] C.R.K. Windows-Yule, M.T. Nicuşan, A.L. Herald, S. Manger, D.J. Parker, Positron Emission Particle Tracking: A Comprehensive Guide, Institute of Physics, Temple Circus Temple Way, Bristol, BS1 6HG, UK, 2022.

[52] D.N. de Klerk, I. Govender, A.N. Mainza, Geometric features of tumbling mill flows: a positron emission particle tracking investigation, Chem. Eng. Sci. 206 (2019) 41–49, https://doi.org/10.1016/j.ces.2019.05.020, URL https://linkinghub.elsevier.com/retrieve/pii/S000925091930452X.

[53] Y.L. Ding, R.N. Forster, J.P.K. Seville, D.J. Parker, Scaling relationships for rotating drums, Chem. Eng. Sci. 56 (12) (2001) 3737–3750, https://doi.org/10.1016/S0009-2509(01)00092-6, URL https://linkinghub.elsevier.com/retrieve/pii/S0009250901000926.

[54] H. Yang, G.L. Jiang, H.Y. Saw, C. Davies, M.J. Biggs, V. Zivkovic, Granular dynamics of cohesive powders in a rotating drum as revealed by speckle visibility spectroscopy and synchronous measurement of forces due to avalanching, Chem. Eng. Sci. 146 (2016) 1–9, https://doi.org/10.1016/j.ces.2016.02.023, URL https://www.sciencedirect.com/science/article/pii/S0009250916300690.

[55] D. Daraio, J. Villoria, A. Ingram, A. Alexiadis, E. Hugh Stitt, M. Marigo, Validation of a discrete element method (dem) model of the grinding media dynamics within an attritor mill using positron emission particle tracking (pept) measurements, Appl. Sci. 9 (22) (2019) 4816, https://doi.org/10.3390/app9224816.

[56] Y.L. Ding, R. Forster, J.P.K. Seville, D.J. Parker, Segregation of granular flow in the transverse plane of a rolling mode rotating drum, Int. J. Multiphase Flow 28 (4) (2002) 635–663, https://doi.org/10.1016/S0301-9322(01)00081-7, URL https://www.sciencedirect.com/science/article/pii/S0301932201000817.

[57] A.N. Huang, T.H. Cheng, W.Y. Hsu, C.C. Huang, H.P. Kuo, DEM study of particle segregation in a rotating drum with internal diameter variations, Powder Technol. 378 (2021) 430–440, https://doi.org/10.1016/j.powtec.2020.10.019, URL https://linkinghub.elsevier.com/retrieve/pii/S0032591020309621.

[58] Y. Liu, M. Gonzalez, C. Wassgren, Modeling granular material blending in a rotating drum using a finite element method and advection-diffusion equation multiscale model, AICHE J. 64 (9) (2018) 3277–3292, https://doi.org/10.1002/aic.16179, URL https://onlinelibrary.wiley.com/doi/10.1002/aic.16179.

[59] C.C. Sun, True density of microcrystalline cellulose, J. Pharm. Sci. 94 (10) (2005) 2132–2134, https://doi.org/10.1002/jps.20459, URL https://linkinghub.elsevier.com/retrieve/pii/S0022354916318676.

.3 Autonomous digitizer calibration of a Monte Carlo detector model through evolutionary simulation

scientific reports

Check for updates

OPEN

Autonomous digitizer calibration of a Monte Carlo detector model through evolutionary simulation

Matthew Herald[1]✉, Andrei Nicuşan[1], Tzany Kokalova Wheldon[2,3], Jonathan Seville[1,3] & Christopher Windows-Yule[1,3]

Simulating the response of a radiation detector is a modelling challenge due to the stochastic nature of radiation, often complex geometries, and multi-stage signal processing. While sophisticated tools for Monte Carlo simulation have been developed for radiation transport, emulating signal processing and data loss must be accomplished using a simplified model of the electronics called the digitizer. Due to a large number of free parameters, calibrating a digitizer quickly becomes an optimisation problem. To address this, we propose a novel technique by which evolutionary algorithms calibrate a digitizer autonomously. We demonstrate this by calibrating six free parameters in a digitizer model for the ADAC Forte. The accuracy of solutions is quantified via a cost function measuring the absolute percent difference between simulated and experimental coincidence count rates across a robust characterisation data set, including three detector configurations and a range of source activities. Ultimately, this calibration produces a count rate response with 5.8% mean difference to the experiment, improving from 18.3% difference when manually calibrated. Using evolutionary algorithms for model calibration is a notable advancement because this method is novel, autonomous, fault-tolerant, and achieved through a direct comparison of simulation to reality. The software used in this work has been made freely available through a GitHub repository.

Simulating the response of detectors to radiation is an important aspect in a variety of physics and medical fields because this allows users to test imaging algorithms, optimise experiments, and design new detectors[1–3]. This is typically achieved by using Monte Carlo radiation transport codes to simulate the interactions of a radiation field with a geometric model of the detector and then applying a pulse-processing chain to the recorded events to emulate the detector's response[4]. Software such as the Geant4 Application for Tomographic Emission (GATE) has been developed specifically for the purpose of running Monte Carlo simulations and emulating detector responses[5,6]. In GATE, the 'digitizer' determines how the timing, energy, and position of interactions with the detector geometry are recorded, how events are grouped and implements the pulse-processing logic of the system[7]. However, digitizer models must be precisely tuned to replicate the behaviour of a real detector.

Several detectors have been modelled using GATE and validated against experimental measurements such as the ADAC Forte, Siemens Inveon, and Phillips Vereos Positron Emission Tomography (PET) scanners[8–10]. For PET systems, performance characterisation experiments are described by the National Electronics Manufacturers Associated (NEMA) which test the spatial resolution, sensitivity, and count-rate response[11]. The GATE model's digitizer is then calibrated to achieve the closest agreement with these experiments. GATE models which do not have well-calibrated digitizers may produce an unrealistic simulated detector response.

State-of-the-art. Current methods of calibrating GATE models, as demonstrated in other work, are achieved by using known properties of the detector or by fitting models to count-rate experiments[7,10,12,13]. Many steps in the digitizer model correspond directly to measurable properties of the detector, such as the energy resolution, dead-time, or time resolution[10]. Values for these properties are often provided by the manufacturer and this can serve as a reliable starting point, but manual tuning is still needed to match the simulated and experimental response of the detector due to variation between each detector[8]. Conversely, with a model-fitting approach, manual tuning can be avoided, but other challenges arise. For example, when fitting a dead-time model to the count rates or fitting a Gaussian function to the 511 keV photo-peak to determine the energy reso-

[1]School of Chemical Engineering, University of Birmingham, Birmingham, UK. [2]School of Physics and Astronomy, University of Birmingham, Birmingham, UK. [3]Positron Imaging Centre, University of Birmingham, Birmingham, UK. ✉email: mxh1092@bham.ac.uk

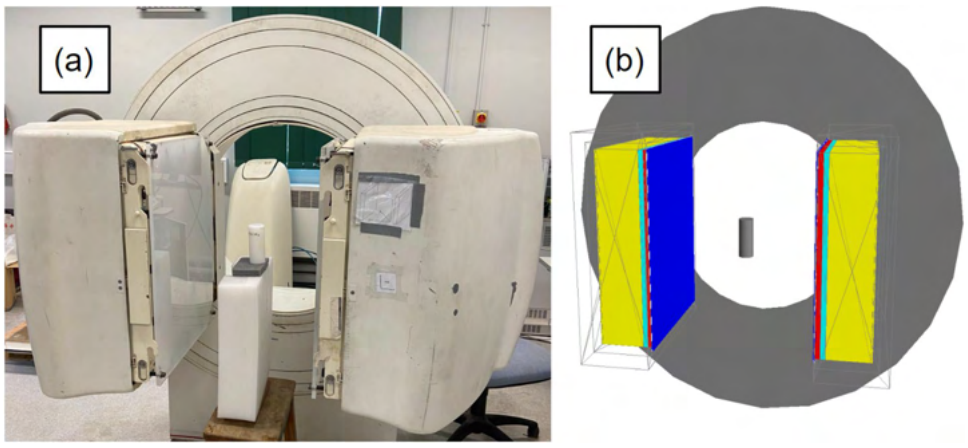

Figure 1. The ADAC Forte at the Positron Imaging Centre during the count-rate experiment (**a**) and the GATE model of the detector and replicated experiment (**b**).

lution, this relies on having both the singles and coincidence count rates, which may not both be available, and also involves fitting simplified models to the detector response, which may not capture the complexity of a real system[14]. In summary, manual tuning of GATE models can produce a good agreement between simulation and experiment, but at the expense of time, resources, and objectivity, whereas fitting simplified models to determine the digitizer parameter values is a quicker, more objective, method but the information is not always available and can still produce inaccurate simulations.

Proposed methodology. In this work, we propose a new procedure which leverages recent advances in metaheuristics to perform an efficient optimisation of parameter values in a detector digitizer model created using GATE v9.1. The goal of the optimisation is to produce a set of parameters which can replicate the count-rate response of the detector across varied source activities and detector separations.

To do this we use an evolutionary algorithm to modify the free parameters of the digitizer , resembling Darwinian evolution, and directly compare the simulated results of candidate solutions to the experimental data. The evolutionary algorithm chosen for this approach is the Covariance Matrix Adaptation Evolutionary Strategy (CMA-ES), which is a stochastic optimiser for robust non-linear non-convex numerical optimisation[15,16]. Parameter combinations are generated following a multivariate normal distribution; in our case, finding the optimum digitizer parameters is equivalent to "evolving" the mean and covariance matrix of this distribution. A particular advantage of this setup is that the underlying optimisation function – i.e. the digitizer response – does not need to have a continuous response. The addition of stochastic "mutations" to the inputs tried, so as to mimic the injection of new genetic material in the biological population, allows CMA-ES to escape local, false minima, which gradient-based optimisers are prone to falling into[17]. We demonstrate this procedure by calibrating the GATE digitizer model of the ADAC Forte, a dual-headed positron camera operated in coincidence mode[8]. The Forte and its digital-twin GATE model are shown in Fig. 1. Six free parameters in the model are calibrated simultaneously by CMA-ES.

In order to interface with the existing CMA-ES optimiser and extend the types of problems it can be used with, we have developed a Python library called the Autonomous Calibration and Characterisation via Evolutionary Software (ACCES) v0.2.2. The purpose of ACCES is to use meta-programming in conjunction with an arbitrary Python script defining the simulation to populate the user-defined free parameters with candidate solutions generated by CMA-ES, then autonomously re-launch the simulation, analyse the results, and use CMA-ES to generate candidate solutions in a cycle until a termination criterion is met[16,18]. The absolute percent difference between the total, true, and scattered plus random coincidence count rates are optimised using a multi-objective cost function to combine their differences into a single value. This method offers improvements over previous calibration procedures since the optimal parameters are chosen by directly comparing the performance of the optimised digitizer to count-rate experiments and multiple experiments are optimised simultaneously.

Methods

Count rate experiment. Characterisation experiments are conducted that measure the coincidence count rates of the Forte as a function of source activity. The experimental coincidence count rates are chosen to be compared to the simulated count rates to assess the optimisation of the model since this is the observable output from the detector in experiments and simulations. Additionally, the coincidence count rates have a complex relationship to the digitizer parameters, source activity, and detector configuration making this an ideal metric for comparison. Three detector separations representing the closest, median, and furthest separations possible for the detector are tested. The initial source activities for each separation are selected to test both the high-activity range where the effect of detector dead-time induces count-losses and, as the source decays over several half-lives, testing the low-activity range where count-rates are linearly proportional to the source activity. The optimisation of the digitizer seeks to find a common set of parameter values to replicate the behaviour of the

Experiment	Head separation (mm)	Initial activity (MBq)
Experiment 1	800	75
Experiment 2	525	60
Experiment 3	250	40

Table 1. Head separations and initial activities for each calibration experiment.

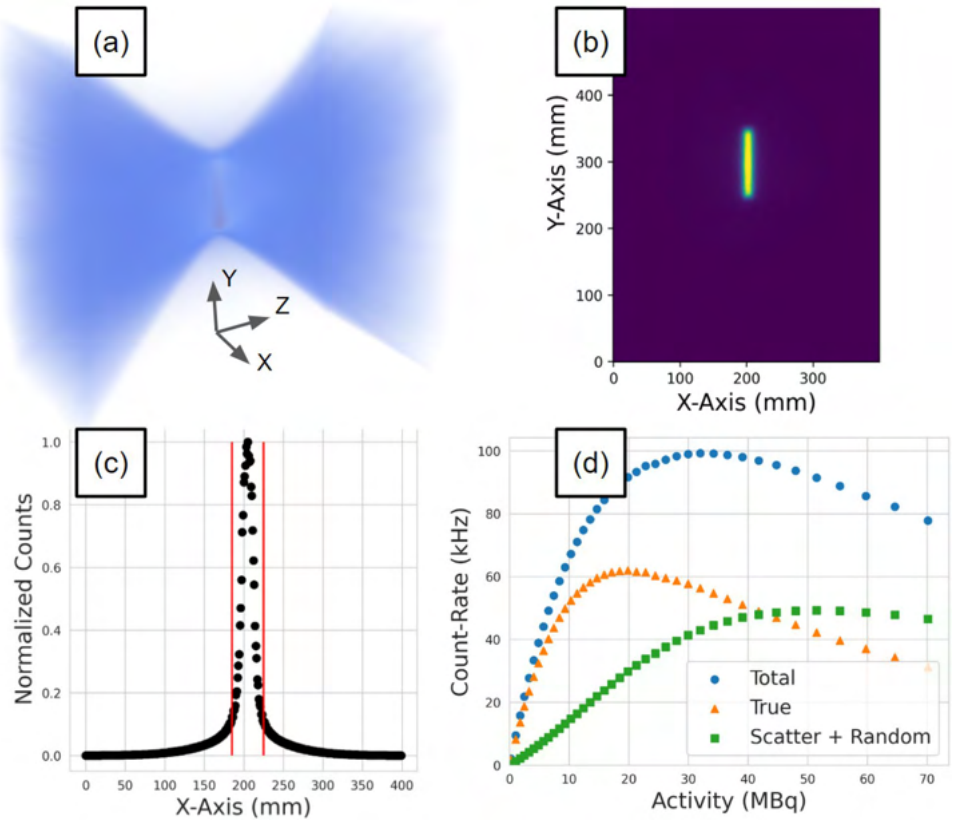

Figure 2. A demonstration of the protocol for extracting count rates from a sample of LoRs: (**a**) a sample of LoRs collected during the experiment is converted into voxels, (**b**) the slice containing the maximum number of LoRs is extracted, (**c**) the slice is collapsed into a line profile and the counts in the central 40 mm strip are summed and background counts subtracted to yield a total, true, and scattered + random count-rate (**d**). Steps a-c are repeated for multiple samples to generate the count-rate response as a function of activity.

detector across all of these conditions. Since there are six free parameters and 45 individual data points for each parameter combination (three detector separations, five activities per separation, and three coincidence count rates per activity), the optimisation problem is considered to be well-constrained.

For these experiments, the source consists of a high-density polyethene (HDPE) cylindrical phantom filled with a solution of water and fluorine-18. The phantom measures 120 mm long and 50 mm in diameter. The inner cylinder in which the water and fluorine-18 solution is filled measures 100 mm long and 12 mm in diameter. The phantom is filled with an initial activity, then placed in the centre of the field-of-view of the Forte and imaged over several half-lives until the activity is below 1 MBq. The three head separations and initial activities for each experiment are found in Table 1.

For each experiment, the total, true, and scattered plus random (corrupted) count rates are extracted as a function of the source activity. This is achieved by applying the NEMA protocol to projection images of the source[11]. A demonstration of the workflow for extracting count rates from the acquisition is shown in Fig. 2. First, samples of a minimum of 500,000 lines-of-response (LoRs) are used to create a three-dimensional voxelised representation of the FOV with a 1 mm voxel size. At this stage, the source activity is calculated using exponential decay equations. From the voxels, a two-dimensional slice is extracted which is both parallel with the detector face and contains the voxel with the maximum number of LoRs. The slice is then collapsed into a line profile of the pixel intensities. All points within ± 20 mm of the maximum pixel are summed. To subtract the background counts, the values at both ends of the ± 20 mm are averaged, multiplied by the size of the window, and subtracted from the counts under the peak leaving only the true counts. The total counts are the sum of all LoRs passing through

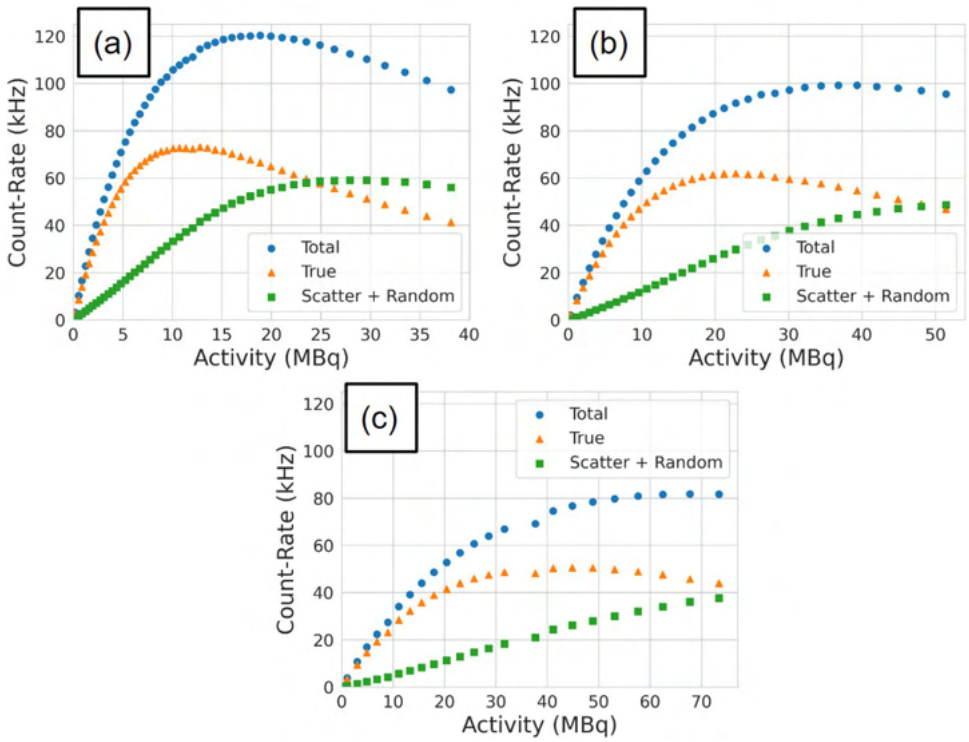

Figure 3. The total, true, and scatter + random coincidence count rates as a function of source activity for (**a**) 250 mm, (**b**) 525 mm, and (**c**) 800 mm head separation.

the slice and the scattered plus random coincidence count rate is the total counts subtracted by the true counts. The extracted coincidence count rates are shown in Fig. 3.

GATE model. GATE v9.1 is an extension of Geant4 v10.7.3 designed for the simulation of radiation detectors used in physics, medicine, and engineering applications[5,6]. GATE uses Geant4 to run Monte Carlo radiation transport simulations, generating a history of interactions of the detector with a radiation field, and then mimicking how the detector would respond[19]. Using GATE consists of 6 steps: defining the geometries (detector and experiment), adding radioactive sources, describing the detector pulse processing chain (digitizer), including physics processes, specifying data output format, and prescribing acquisition settings (run time and time slice)[5].

In this work, we use a detector geometry and GATE model for the ADAC Forte previously developed by the authors. The ADAC Forte is a dual-headed positron camera used at the Positron Imaging Centre[8]. A full description of this model and its original calibration can be found in Herald et al.[8]. The experiment geometry is the same HDPE cylindrical phantom as described in "Count rate experiment" section. The radioactive source is a solution of water and fluorine-18 prescribed as emitting back-to-back 511 keV gamma rays. Since the mean positron range in water and HDPE it can be assumed that all positrons annihilate before leaving the phantom, thus making a back-to-back gamma source a reasonable approximation that decreases the time needed to run the simulations. The detector model's digitizer structure follows the same as described in Herald et al., (2021). Six key parameters of the digitizer will be calibrated. These are the singles dead-time, coincidence dead-time, pileup, lower energy discriminator, upper energy discriminator, and the time resolutions as will be discussed is 2.2. Physics processes are imported through the GATE's 'emstandard' physics list, which includes the Livermore model for photon interactions and is based on the Evaluated Photon Data Library, 1997 (EPDL1997)[20]. The output format is coincidence data saved as a text file. The acquisition was prescribed as a 10 second simulation with the time slice saving data every 10 ms of simulated time.

Once the simulation begins the source activity determines the decay rate and individual decays are modelled on a Poisson distribution. Each event (two back-to-back 511 keV gamma rays) is initialised randomly within the source volume and prescribed a direction isotropically. As the gamma rays pass through the geometry, they have a stochastic chance of interacting with the materials following Beer-Lambert's Law and using attenuation coefficients generated from material composition, density, and cross-sections from EPDL1997. Interactions which occur within the 'Sensitive Detector', in this case, the scintillation crystals, are termed 'hits'. From the list of hits, which contains information about the type of interaction, time, position, and energy, the GATE digitizer converts hits into 'pulses'. A pulse is the response of the detector element that is analogous to a signal which can be processed, eventually producing an output of what a real detector would record. The digitizer model for the ADAC Forte is shown in Fig. 4.

In this work, we demonstrate a novel application of evolutionary algorithms to calibrate the digitizer for a GATE model of the ADAC Forte, a dual-headed positron camera used at the Positron Imaging Centre[8]. The

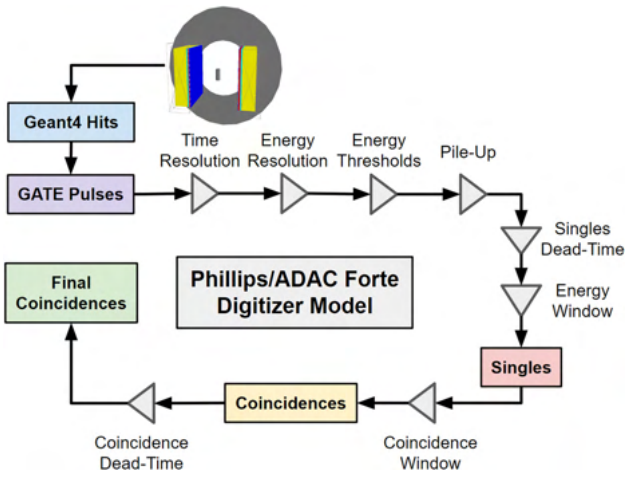

Figure 4. The digitizer model of the ADAC Forte.

primary use of this GATE model is to emulate experiments using positron emission particle tracking, a radio-imaging technique, in order to estimate the spatial and temporal resolution of tracer trajectories and to optimise experiment design[21,22]. The detector consists of two wide-area sodium-iodide crystals measuring 590 mm × 470 mm and 16 mm thick[23]. The active area for recording coincidences measures 510 mm × 380 mm. Additionally, the two detector heads can be moved between 250 mm and 800 mm of separation which can accommodate a variety of experiments[24,25]. The Forte and the GATE model of the Forte are shown in Fig. 1.

The six digitizer parameters chosen to be optimised are the singles dead-time, coincidence dead-time, pileup, lower energy discriminator, upper energy discriminator, and the time resolutions which are explained below. These parameters were chosen because they have not been measured directly through a characterisation experiment meaning there is uncertainty in the optimal values. The singles dead-time is a paralysable dead-time which affects each pulse, rendering the detector unable to record another pulse until the dead-time has ended[14]. If another gamma ray enters the detector before the singles dead-time is completed, the dead-time is reset and the gamma-ray is not recorded. Paralysable dead-time results in count losses and at high source activities can cause the count rate to decrease. Coincidence dead-time is a separate, non-paralysable dead-time affecting the recording of a coincidence[26]. Unlike a paralysable model, a non-paralysable dead-time does not get reset with additional events. Pile-up time is the time between the detection of a single gamma-ray triggering the recording of the pulse and the time at which other events can 'pile-up' onto the same pulse[14]. Pile-up has the effect of creating count-losses at high source activities. The lower and upper energy discriminators are the minimum and maximum energies of events which can trigger the singles dead-time[27]. The time resolution is the uncertainty in the timing of precision of the detector, defined by a Gaussian blurring with a full-width half-maximum[28]. If two gamma rays interact with the detector within the coincidence window of 15 ns, they are not guaranteed to be detected in coincidence due to the timing uncertainty. This has the effect of disregarding some real coincidences and accepting more random coincidences.

ACCES. When trying to calibrate a simulation's free parameters so that an experimental measurement can be replicated, it is often useful to test a range of conditions and assess how the tested parameter values replicate the measurement. In the simplest case with only one free parameter, the value that minimises the error to the measurement can be easily found and visualised by plotting parameter values and the error as a two-dimensional plot. This can also be extended to two free parameters by plotting the error as a third dimension on the plot. Beyond three dimensions, the number of parameter values needed to explore the solution space increases exponentially and the relationship between the parameters becomes non-intuitive. For these problems, an optimiser is needed to efficiently test a range of parameter values and converge to a set of optimal parameter values. However, in simulations and experiments, there often exist noisy measurements, thus a function defining the difference between experiment and simulation will be non-smooth and potentially have many false local minima. This means that gradient-based optimisers are ill-suited for calibrating simulations.

In these difficult optimisations, evolutionary algorithms excel[29]. Evolutionary algorithms are a type of bio-inspired computing which mimics natural selection. For example, in a population where individuals have a randomised set of genes and selective pressure is exerted, only the individuals which have genes that enable them to survive will reproduce. Due to this, the next generation of individuals will be more adapted to selective pressure. Similarly, when an evolutionary algorithm is applied to a model function with quantitative free parameters which can be tuned, the parameter values act as genes, a model with a specific set of parameters is an individual, and a group of individual simulations is a generation[29]. For each generation, a cost function determines an individual's fitness and acts as a selective pressure. Using this method, parameter value combinations which result in a low cost function are prioritised until the solutions converge to a set of optimal values. A flow diagram of how an evolutionary algorithm can be applied to digitizer calibration is shown in Fig. 5.

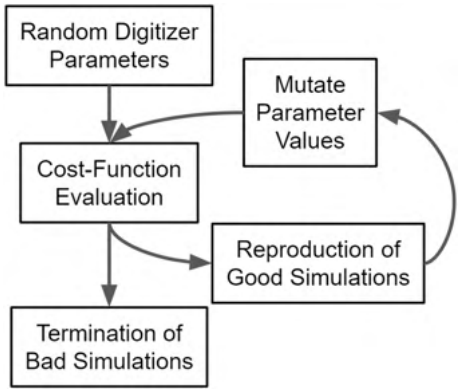

Figure 5. The flow diagram for ACCES is applied to optimising free parameters in a digitizer of a GATE detector model.

While there exist several types of evolutionary algorithms the CMA-ES algorithm is used in this work since it performed well in a comparative review of optimisation algorithms and there is a well-documented Python implementation CMA v3.0.3[16,30]. To use CMA-ES, the ACCES Python library v0.2.2 is employed to interface with CMA-ES and edit an arbitrary script for updating free parameter values in the simulation[18]. Using code inspection and meta-programming, simulation scripts are parallelised by ACCES allowing them to be launched locally or on a high-performance computer. The difference between the simulated system and experimental reality can then be quantified by a cost function so that CMA-ES can determine the next generation of solutions. ACCES offers improvements over other interfaces to optimisers in that it is fault-tolerant and designed for high-performance computing.

ACCES needs only the bounds of the search parameters, and the number of individuals in a population, and stores the results after each generation, or 'epoch', so that the optimisation state can be restored at any point. The default implementation of CMA-ES requires the use of a single initial standard deviation for all parameters - i.e. assuming that all parameters have comparable value ranges and sensitivities. ACCES scales the parameter values by 40 % of each parameter's allowed range, such that parameters of vastly different scales can be optimised together - e.g. singles dead times in the range [0, 2] and pile up between [0, 600]. As parameter combinations are drawn from normal distributions, an initial standard deviation of 40 % naturally covers the entire parameter range.

In order to allow the use of complex, potentially thread-unsafe simulations written in different programming languages, ACCES launches each simulation as a completely separate OS process, which is either scheduled by the kernel to be run locally on a shared-memory machine (e.g. a laptop) or using an external workload manager to launch jobs on multi-node clusters; in this study, ACCES automatically sets up and launches batch jobs for each parameter combination to be evaluated using GATE. To summarise, the two critical CMA-ES configuration parameters are automatically determined by the computing resources available and the possible parameter ranges, such that no manual adjustments of optimiser settings for a given problem is necessary.

Digitizer calibration. We use ACCES in this work to optimise the six free parameters within the digitizer of the Forte GATE model described in "GATE model" section. The experiments described in "Count rate experiment" section are used to determine the fitness of parameter combinations. Specifically, a cost function is applied which measures the percent difference between the experimentally observed and simulated count rates for the total, true, and scatter plus random count-rates across all three head separations and activities. The sum for each of these percent differences is denoted as ε_R, ε_T, and ε_{SR} respectively and computed using Eq. (1). Each type of count rate is treated as an objective to optimise and combined into a multi-objective optimisation by multiplying them together using Eq. (2). In this case, each type of count rate is treated as equally important; this could be changed by adding weights to each percent difference.

$$\varepsilon_R = \sum 100 \frac{|R_{exp} - R_{GATE}|}{R_{exp}} \tag{1}$$

$$\varepsilon = \varepsilon_R \varepsilon_T \varepsilon_{SR} \tag{2}$$

To run ACCES, three things must be prescribed: the number of simulations per epoch, the bounds of the parameter guesses, and the terminating criterion. The number of simulations per epoch should be large enough that sufficient learning can occur and the bounds of parameters must be set so as to keep guesses within a realistic range. We have tested ACCES using a simple analytical cost function, the Ackley function, which is widely used for testing optimization algorithms[31]. This function, described in Eq. (3), has many local minima and one global minimum. The number of epochs needed to find the global minimum as well as the total number of cost functions evaluated can be studied as a function of the number of solutions per epoch. The results from this study are shown in Fig. 6. We used a two-dimensional (d = 2) Ackley function with the parameters a = 20, b = 0.2 and c = 2π.

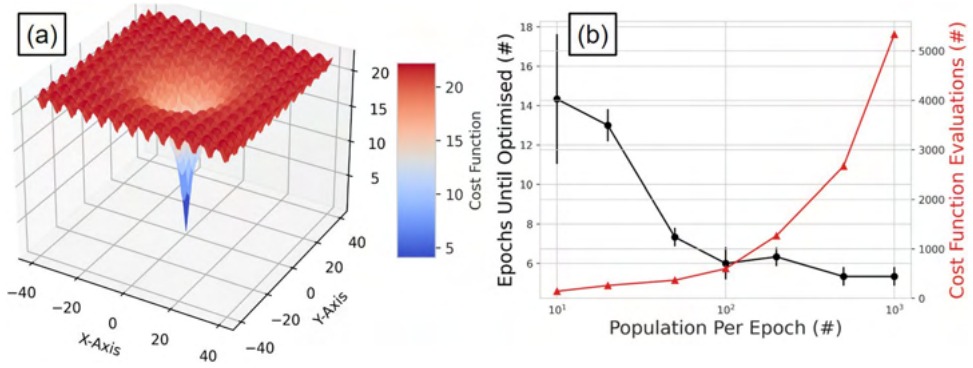

Figure 6. (**a**) The two-dimensional Ackley function with the with evaluation as the third dimension. (**b**) The behaviour of the number of epochs needed to find the global minimum as well the total number of cost functions evaluated as function of the number of solutions per epoch.

Parameter	Lower bound	Upper bound	Initial guess
Singles dead-time (ns)	0	2	1
Coincidence dead-time (ns)	0	2	1
Pile-up (ns)	0	600	300
Lower energy discriminator (keV)	0	360	180
Upper energy discriminator (keV)	640	1200	920
Time resolution (ns)	10	20	15

Table 2. Digitizer parameter bounds and initial guesses.

$$\varepsilon(x_i) = -a\exp\left(-b\sqrt{\sum_{i=1}^{d}\frac{1}{d}x_i^2}\right) - \exp\sum_{i=1}^{d}\frac{1}{d}cos(cx_i) + a + \exp(1) \tag{3}$$

The results from this show that larger numbers of solutions optimise the parameters in fewer epochs, but at a cost of increasing the number of cost function evaluations. When optimising the Monte Carlo model's digitizer, in order to be computationally efficient, the lowest reasonable number of simulations per epoch should be run. The number of simulations is set to 150 so that a wide range of different parameter combinations can be tried and the simulations can all be launched in parallel while not affecting the shared usage of the BlueBEAR high-performance computing (HPC) system. In this work, each set of simulations with a common set of parameter solutions is run on a single Intel Icelake core of the BlueBEAR HPC with 8 GB of memory allocated. The maximum run-time is set to 4 hours and 30 minutes, which is approximately twice as long as the mean run-time expected. In the event that a set of simulations takes longer than 4 hours and 30 minutes, the job is terminated and the results are not in the solution space for the next generation of parameter solutions.

Additionally, the bounds of the parameter guesses are set to only explore solutions which make physical sense, excluding options like a negative dead-time or upper energy level being below the upper energy window. The bounds are also limited where needed such that the solution space is finite, yet spanning a range likely to contain the optimal value based on an estimate from a previously calibrated system[8]. A list of the bounds and the initial guesses are shown in Table 2.

The termination criterion for the optimisation is the standard deviation for each parameter reaching 10% of the initial standard deviation. This range is chosen such that variation in the parameter values will not significantly affect the accuracy of the model. The initial standard deviation is equal to the range of the bounds at the beginning of the optimisation and the scaled standard deviation is defined as unity. Once the optimal values are identified, they are input to the digitizer model and a coincidence count-rate response is generated to compare with the experimental data. These simulations are run at 2 MBq intervals starting at 1 MBq and reaching into the upper activities for each experiment. A study of the accuracy of extracted count-rates for the simulation with the highest separation and lowest activity (800 mm and 2 MBq) at different numbers of LoRs used to produce projection images showed that at least 10,000 events are needed to ensure that variance in the extracted count-rates is well below 10%. The results of this study are shown in Fig. 7. The lowest count rate that would be expected in an experiment is approximately 1 kHz. As a result, we determined that simulations should be run for 10 seconds of simulated time at each activity in order to ensure that 10,000 events are captured. As more events are recorded, the covariance of the extracted count rates decreases exponentially.

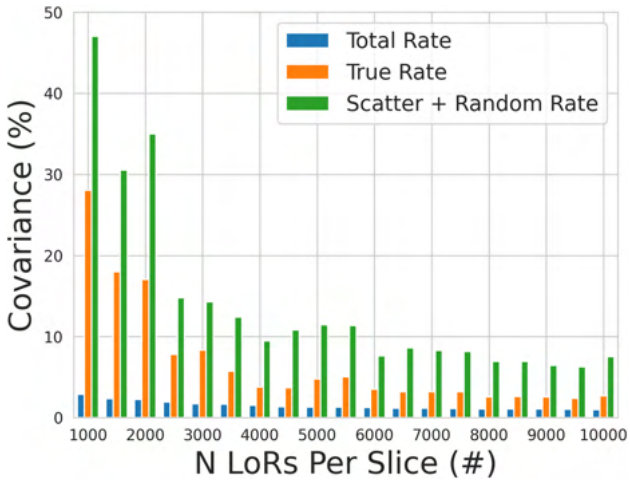

Figure 7. The covariance of the total, true, and scattered plus random count rates for the 800 mm and 2 MBq simulation as a function of the number of LoRs used to generate a projection image.

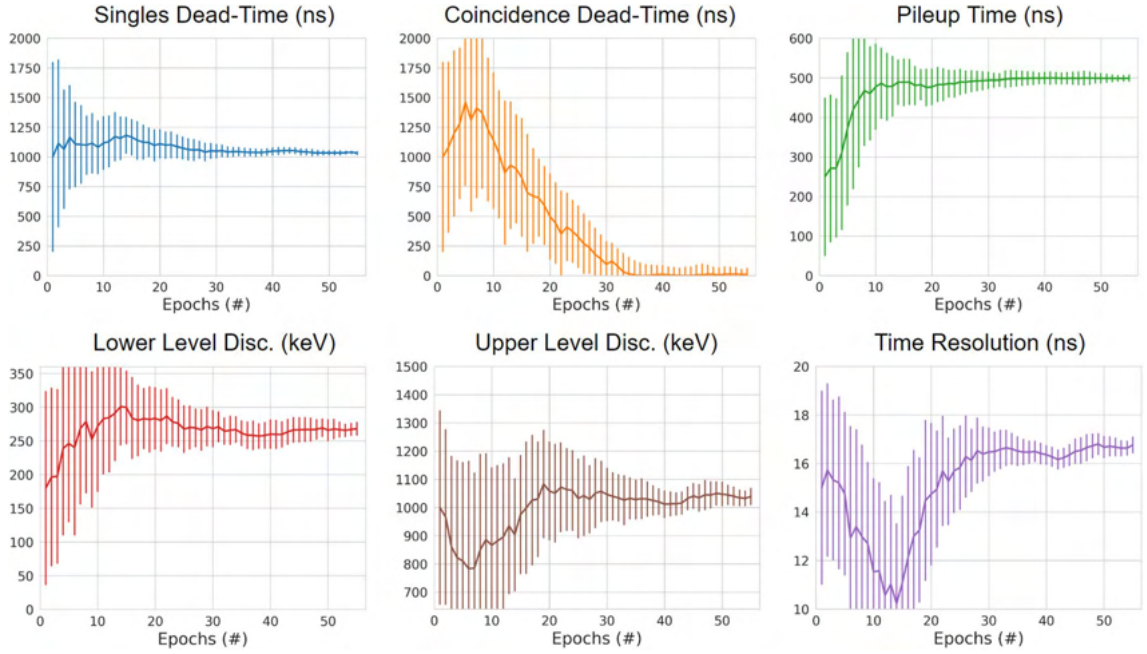

Figure 8. The mean parameter value guesses for each of the six free parameters with the standard deviation of the guesses are plotted as error bars. After 56 epochs all parameters are below 10% standard deviation and the optimisation is completed.

In order to contextualise the proposed for calibrating Monte Carlo detector models through evolutionary simulation to the existing methods, the ACCES-calibrated model is compared to the existing model described in Herald et al., (2021) through the ability to reproduce the real count-rate response of the ADAC Forte. This previous model was calibrated by using parameter values for the digitizer which were determined from the manufacturer's characterisation and by manual calibration, taking a considerable amount of time and computational resources to achieve. The main advantage of using evolutionary simulation is the ability to achieve similar or, in this case, better results than manual calibration without spending the time and resources needed to run simulations, compare results, and update parameter values through iteration.

Results and discussion

In total, the ACCES optimisation took 56 epochs, 8400 cost function evaluations, and approximately 4 days to complete. At the beginning of the ACCES optimisation, the guesses for the six free parameters are broad so as to explore the solution space. After this initial period, the guesses begin to converge to their optimal values as shown in Fig. 8 where the mean solution values and their standard deviations are plotted for the parameters

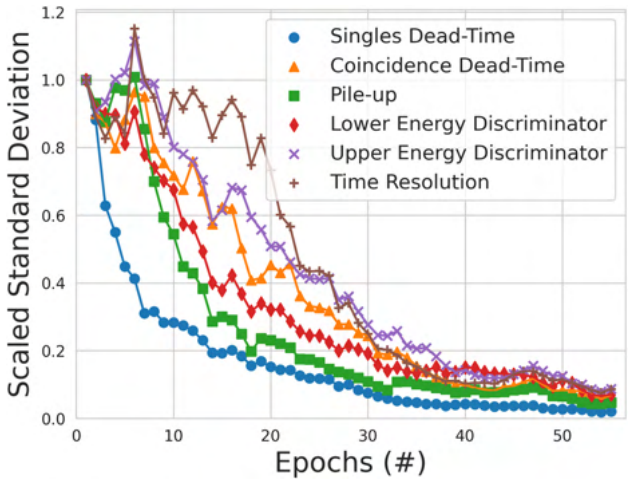

Figure 9. The standard deviation of the parameter value guesses tried by ACCES. A lower standard deviation represents lower uncertainty in discovering the optimal parameter values.

Parameter	Calibrated value	Uncertainty
Singles dead-time (ns)	1070	± 16.7
Coincidence dead-time (ns)	10	± 54.7
Pile-up (ns)	498	± 9.31
Lower energy discriminator (keV)	284	± 10.1
Upper energy discriminator (keV)	1020	± 30.2
Time resolution (ns)	17	± 0.347

Table 3. Calibrated digitizer parameter values.

over the optimisation. The scaled standard deviations are also shown in Fig. 9 to depict how the uncertainty in the optimisation decreases as the optimisation progresses.

Once the optimisation reached 10% uncertainty for each parameter, 0.1 scaled standard deviation, the parameters are considered calibrated. The final calibrated values are presented in Table 3. For all parameters except for the coincidence dead-time, the optimal solutions are well within their bounds, suggesting an optimal calibration that would not change with different bounds. However, the optimal coincidence dead-time is found to be approximately 0 ns. While this could be due to the bounds being ill-suited to the problem, in this case, we believe this demonstrates that coincidence dead-time is insignificant to the digitizer model. Further, support for this is that the model under-predicts the peak count rates. The opposite would be expected if coincidence dead-time was important.

To assess the ability of the ACCES-calibrated digitizer model to replicate the experimental data, a new set of simulations is run for each head separation using the optimised values. After the simulations are finished, the results were plotted against the experimental data in Fig. 10. Visually, the count-rate response of the GATE model matches the general form of the real experiment. To quantify the accuracy, a mean absolute percent difference is calculated for each head separation and each type of count rate and presented in Table 4. Additionally, the results for the manually calibrated digitizer model are presented in Table 5.

The average mean absolute percent differences for the 250 mm, 525 mm, and 800 mm are 7.55%, 4.30%, and 5.48%, respectively. The separation which was closest to the phantom experienced the highest error between the simulation and experiment. This could be caused by the closer separation amplifying differences between the phantom's position in the simulation versus the experiment. In addition to this, the ACCES-calibrated model improves the match between simulation and experiment compared to a manually calibrated digitizer model which produced a mean absolute percent difference in the count rate response of 17.78%, 15.42%, and 21.75%. This represents ACCES producing a calibration which achieves a nearly three times better agreement with the experiments. This is a significant improvement and one accomplished without guiding the optimiser to these solutions. Overall, this calibration represents an agreement with the experiment that would be sufficient for the GATE model to be used as a predictive tool to generate data representative of real experiments.

To assist users in developing their own optimisations using ACCES, we have included an example within the GitHub repository found here. This example uses a simulated count-rate response of the ADAC Forte GATE model with prescribed parameter values in the digitizer as the ground truth response, then uses ACCES to calibrate two parameters, the singles dead-time and time resolution, to match the ground truth response. Two parameters were chosen because this is a more complex optimisation than a single parameter, yet easier to

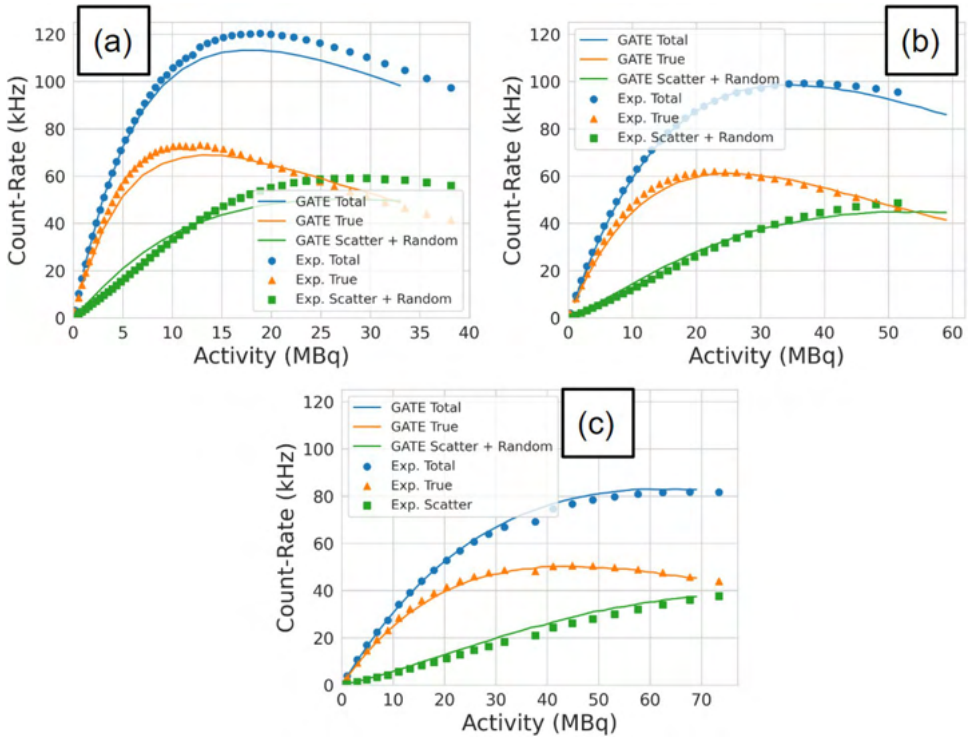

Figure 10. The optimised GATE model count rates are plotted against the experimental data for the (**a**) 250 mm experiment, (**b**) 525 mm experiment, (**c**) and the 800 mm experiment.

ACCES-calibrated results			
Head separation	**250 mm**	**525 mm**	**800 mm**
Total count-rate	4.76%	1.71%	2.43%
True count-rate	4.55%	2.85%	2.25%
Scatter + random	13.33%	8.33%	11.77%
Average error	7.55%	4.30%	5.48%

Table 4. Mean absolute percent differences in the count rate of the ACCES-calibrated digitizer model.

Manually calibrated results			
Head separation	**250 mm**	**525 mm**	**800 mm**
Total count-rate	11.50%	4.48%	6.25%
True count-rate	17.86%	15.68%	10.85%
Scatter + random	23.98%	26.10%	48.15%
Average error	17.78%	15.42%	21.75%

Table 5. Mean absolute percent differences in the count rate of the manually calibrated digitizer model.

visualise than an optimisation with three or more parameters. The prescribed values for the singles dead-time and the time resolution are 1000 ns and 15 ns, respectively. The methodology in this simple example follows the same as that described in "Digitizer calibration" section. The results from this optimisation in Fig. 11 show the optimal parameter was determined to be 995.016 ns for the singles dead-time and 15.022 ns for the time resolution, which matches the prescribed parameters.

Conclusions

In this work, we have demonstrated the calibration of a GATE digitizer model using an evolutionary algorithm. The model's accuracy was quantified by a direct comparison of the ability of different parameter value combinations to replicate the count rate response of the detector across a diverse set of experiments. Importantly, the

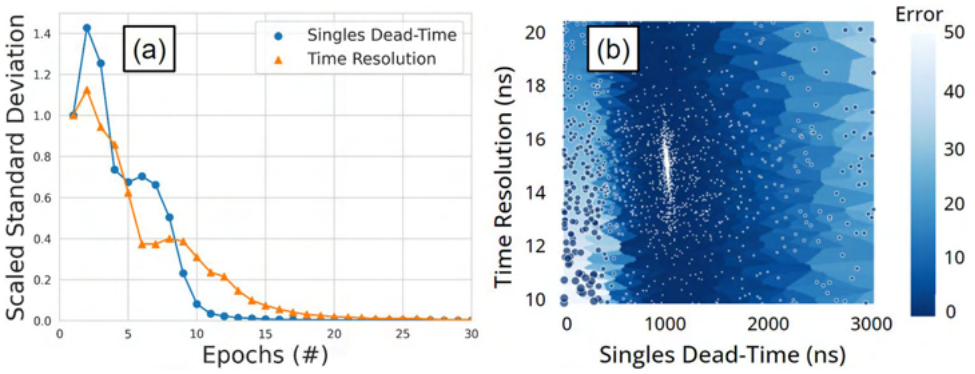

Figure 11. An example of an optimised GATE model with two free parameters. (**a**) The scaled standard deviation for each parameter over several epochs of simulations (**b**) a Voronoi diagram of the parameter combinations shows the solution converges to the optimal parameters. Each point is a candidate solution and the larger the point or lighter blue the Voronoi plot the higher the error in the cost-function evaluation.

calibration was completed autonomously, needing only the number of simulations desired, the bounds of the search parameters, and the user-defined stopping criterion. This represents an advancement which brings simulations closer to reality. By employing the ACCES software available from our GitHub repository to perform this calibration, the need for users to perform a calibration through trial-and-error is eliminated. Even though this method needs a relatively long time and a large number of computational resources, the ability for ACCES to run on a high-performance computing system and periodically save the optimisation state makes this method useful and practical for users who have these resources at their disposal. While this workflow was demonstrated through the calibration of a specific detector model following the NEMA protocol, this same type of method can be applied to other models and also expanded to cover other types of measurements such as spatial resolution and sensitivity. Additionally, this method of optimisation can be improved in the future by including a strategy for calibrating the structure of the digitizer by including or excluding pulse-processing stages and by adjusting parameters which are categorical instead of quantitative, such as the type of dead-time model (paralyzable or non-paralyzable) or the policy for recording multi-coincidences.

Data availability

All data generated or analyzed during this study are included in this published article. Additionally, the ACCES software used for optimisation and the calibrated GATE model of the ADAC Forte have been made available through the University of Birmingham Positron Imaging Centre's GitHub Repository: https://github.com/uob-positron-imaging-centre/ACCES-CoExSiST and https://github.com/uob-positron-imaging-centre/GATE_Models.

Received: 11 July 2022; Accepted: 8 November 2022
Published online: 14 November 2022

References

1. Windows-Yule, C. R. K. *et al.* Recent advances in positron emission particle tracking: A comparative review. *Rep. Progress Phys.* **85**, 016101. https://doi.org/10.1088/1361-6633/ac3c4c (2022).
2. Grevillot, L. *et al.* Gate-RTion: A gate/geant4 release for clinical applications in scanned ion beam therapy. *Med. Phys.* **47**, 3675–3681. https://doi.org/10.1002/mp.14242 (2020).
3. Agostini, M. *et al.* The monte Carlo simulation of the borexino detector. *Astropart. Phys.* **97**, 136–159. https://doi.org/10.1016/j.astropartphys.2017.10.003 (2018).
4. Apostolakis, J. *Detector Simulation* 485–531 (Springer, Cham, 2020).
5. Jan, S. *et al.* Gate: A simulation toolkit for pet and spect. *Phys. Med. Biol.* **49**, 4543–4561. https://doi.org/10.1088/0031-9155/49/19/007 (2004).
6. Sarrut, D. *et al.* Advanced Monte Carlo simulations of emission tomography imaging systems with gate. *Phys. Med. Biol.* **66**, 10TR03. https://doi.org/10.1088/1361-6560/abf276 (2021).
7. Guez, D. *et al.* Counting rates modeling for pet scanners with gate. *IEEE Trans. Nucl. Sci.* **55**, 516–523. https://doi.org/10.1109/TNS.2007.910880 (2008).
8. Herald, M., Wheldon, T. & Windows-Yule, C. Monte carlo model validation of a detector system used for positron emission particle tracking. *Nucl. Instrum. Methods Phys. Res. Sect. A* **993**, 165073. https://doi.org/10.1016/j.nima.2021.165073 (2021).
9. Lee, S., Gregor, J. & Osborne, D. Development and validation of a complete gate model of the siemens inveon trimodal imaging platform. *Mol. Imaging* **12**, 7290.2013.00058. https://doi.org/10.2310/7290.2013.00058 (2013).
10. Salvadori, J. *et al.* Monte Carlo simulation of digital photon counting pet. *EJNMMI Phys.* **7**, 23. https://doi.org/10.1186/s40658-020-00288-w (2020).
11. NEMA. Performance measurements of positron emission tomographs (pet) 2018. Standard, National Electrical Manufacturers Association (2018).
12. Strugari, M., DeBay, D., Beyea, S. & Brewer, K. Nema nu 1-2018 performance characterization and Monte Carlo model validation of the cubresa spark sipm-based preclinical spect scanner. *EJNMMI Phys.*. https://doi.org/10.21203/rs.3.rs-1946160/v1 (2022).
13. Tiwari, A., Merrick, M., Graves, S. A. & Sunderland, J. Monte carlo evaluation of hypothetical long axial field-of-view pet scanner using ge discovery mi pet front-end architecture. *Med. Phys.* **49**, 1139–1152. https://doi.org/10.1002/mp.15422 (2022).

14. Usman, S. & Patil, A. Radiation detector deadtime and pile up: A review of the status of science. *Nucl. Eng. Technol.* **50**, 1006–1016. https://doi.org/10.1016/j.net.2018.06.014 (2018).
15. Hansen, N. *The CMA Evolution Strategy: A Comparing Review*, 75–102. Studies in Fuzziness and Soft Computing (Springer, Berlin, 2006).
16. Hansen, N. *et al. CMA-ES/pycma: r3.2.2* (Zenodo, 2022).
17. Venter, G. *Review of Optimization Techniques* (Wiley, Chichester, 2010).
18. Nicusan, A.-L., Werner, D., Sykes, J. A., Seville, J. & Windows-Yule, K. Acces: Autonomous characterisation and calibration via evolutionary simulation. https://doi.org/10.5281/ZENODO.7215377 (2022).
19. Allison, J. *et al.* Recent developments in geant4. *Nucl. Instrum. Methods Phys. Res. Sect. A* **835**, 186–225. https://doi.org/10.1016/j.nima.2016.06.125 (2016).
20. Cullen, D. E., Hubbell, J. H. & Kissel, L. Epdl97: The evaluated photo data library '97 version. *United States Department of Energy Office of Scientific and Technical Information*https://doi.org/10.2172/295438 *(1997).*
21. Herald, M., Bingham, Z., Santos, R. & Ruggles, A. Simulated time-dependent data to estimate uncertainty in fluid flow measurements. *Nucl. Eng. Des.* **337**, 221–227. https://doi.org/10.1016/j.nucengdes.2018.07.005 (2018).
22. Herald, M. T., Sykes, J. A., Werner, D., Seville, J. P. K. & Windows-Yule, C. R. K. Dem2gate: Combining discrete element method simulation with virtual positron emission particle tracking experiments. *Powder Technol.* **401**, 117302. https://doi.org/10.1016/j.powtec.2022.117302 (2022).
23. Parker, D. J., Forster, R. N., Fowles, P. & Takhar, P. S. Positron emission particle tracking using the new Birmingham positron camera. *Nucl. Instrum. Methods Phys. Res. Sect. A* **477**, 540–545. https://doi.org/10.1016/S0168-9002(01)01919-2 (2002).
24. Al-Shemmeri, M., Windows-Yule, K., Lopez-Quiroga, E. & Fryer, P. J. Coffee bean particle motion in a spouted bed measured using positron emission particle tracking (pept). *J. Food Eng.* **311**, 110709. https://doi.org/10.1016/j.jfoodeng.2021.110709 (2021).
25. Windows-Yule, C. R. K., Seville, J. P. K., Ingram, A. & Parker, D. J. Positron emission particle tracking of granular flows. *Annu. Rev. Chem. Biomol. Eng.* **11**, 367–396. https://doi.org/10.1146/annurev-chembioeng-011620-120633 (2020).
26. Ljungberg, M. *Handbook of Nuclear Medicine and Molecular Imaging for Physicists: Instrumentation and Imaging Procedures* 1st edn, Vol. I (CRC Press, New York, 2021).
27. Colsher, J. *et al.* Adjusting the low energy threshold for large bodies in pet. *IEEE Symposium Conference Record Nuclear Science 2004* **5**, 2872–2876. https://doi.org/10.1109/NSSMIC.2004.1466286 (2004).
28. Yao, S. *et al.* Timing resolution measurement for tof pet. *J. Nucl. Med.* **57**, 2726 (2016).
29. Hong, W.-J., Yang, P. & Tang, K. Evolutionary computation for large-scale multi-objective optimization: A decade of progresses. *Int. J. Autom. Comput.* **18**, 155–169. https://doi.org/10.1007/s11633-020-1253-0 (2021).
30. Rios, L. M. & Sahinidis, N. V. Derivative-free optimization: A review of algorithms and comparison of software implementations. *J. Global Optim.* **56**, 1247–1293. https://doi.org/10.1007/s10898-012-9951-y (2013).
31. Ackley, D. H. *A Connectionist Machine for Genetic Hillclimbing*, vol. 28 of *The Kluwer International Series in Engineering and Computer Science* (Springer, Boston)

Acknowledgements

The computations described in this paper were performed using the University of Birmingham's BlueBEAR service, which provides High-Performance Computing to the University's research community. See their website http://www.birmingham.ac.uk/bear for more details. This work is supported in part by a grant from the Engineering and Physical Science Research Council EP/T034327/1, Advanced Imaging and Numerical Modelling of Segregation and Transport of Plastics in Fluidised Beds: Toward a Circular Economy for Plastics.

Author contributions

M.H. conceived and conducted the experiments, M.H. and A.N. developed the methodology, A.N. created the ACCES software, M.H. designed the detector model, T.W. provided the resources, C.W. and J.S. provided supervision, M.H. analysed the results. M.H. drafted the original manuscript, all authors reviewed the manuscript.

Competing interest

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to M.H.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

© The Author(s) 2022

.4 Improving the accuracy of PEPT algorithms through dynamic parameter optimisation

Nuclear Inst. and Methods in Physics Research, A 1047 (2023) 167831

Contents lists available at ScienceDirect

Nuclear Inst. and Methods in Physics Research, A

journal homepage: www.elsevier.com/locate/nima

Improving the accuracy of PEPT algorithms through dynamic parameter optimisation

Matthew Herald [a,*], Jack Sykes [b,a], David Parker [b], Jonathan Seville [a,c], Tzany Wheldon [b,c], Christopher Windows-Yule [a,c]

[a] School of Chemical Engineering, University of Birmingham, Birmingham, United Kingdom
[b] School of Physics and Astronomy, University of Birmingham, Birmingham, United Kingdom
[c] Positron Imaging Centre, University of Birmingham, Birmingham, United Kingdom

ARTICLE INFO

Keywords:
PEPT
Compton scattering
Monte Carlo
GATE
Digital-Twin

ABSTRACT

Positron emission particle tracking (PEPT) is used to study a wide range of scientific, industrial, and biomedical systems, typically those inaccessible through conventional optical particle tracking techniques. However, in dense or thick-walled systems a fraction of the coincident gamma-rays emitted from a PEPT tracer, called Lines-of-Response (LoRs), are attenuated via Compton scattering. Additionally, at high source activity, random LoRs may be formed by two unrelated events. The incorporation of scattered or random LoRs decreases PEPT spatial accuracy and can distort the trajectory. In this work, we use validation experiments and simulations to investigate the spatial accuracy of the Birmingham Method (BM) PEPT algorithm when two key free parameters are changed: the total number of LoRs in the sample and the fraction of LoRs in the sample used to locate the tracer. Our results show that the default algorithm parameters are not suitable for all cases, however, Monte Carlo simulations of PEPT experiments can be used to estimate the optimal parameter values. Ultimately a variant of the BM, called Dynamic-BM, is demonstrated in a virtual PEPT experiment. Dynamic-BM uses the optimal parameters on a sample-by-sample basis improving PEPT accuracy in this case by 4.03% over the best constant parameters and 76.5% over the default parameters. These improvements make PEPT a more accurate and thus more useful tool.

1. Introduction

Many types of systems used in scientific, industrial, and biomedical applications pose a challenge for those who study them since the internal dynamics of fluids and solids are often difficult to observe directly. This is because the materials themselves are opaque, they are inaccessible behind the system's opaque walls or both. To better understand what is happening inside these systems, a fully three-dimensional, Lagrangian particle tracking technique called positron emission particle tracking (PEPT) was developed [1]. The PEPT technique locates a quasi-point, positron-emitting tracer by detecting coincident 511 keV gamma-rays generated from positron-annihilation [2]. Using samples on the order of 100 of these rays, which are termed lines-of-response (LoRs), the location where the LoRs converge is determined to be the tracer's position. Tracers are tracked over time using subsequent samples of LoRs to develop a trajectory. The PEPT technique is similar to positron emission tomography (PET), in so far as it uses the same types of tracers and detectors, but where PET produces images of the 3-dimensional radio-nuclide concentration throughout the imaging volume at a low temporal resolution (<1 Hz), PEPT assumes a quasi-point tracer and returns the 3-dimensional centre of the discrete tracer's radio-nuclide distribution at a comparatively high temporal resolution (>100 Hz) [3]. This allows PEPT to extract information from systems such as the velocities of individual tracers, granular temperature, and flow dynamics in engineering systems that PET cannot provide. In the decades of research since PEPT was first introduced, PEPT has been shown to capture tracer trajectories with high spatiotemporal resolution and used to study a wide range of systems [4]. For example, PEPT has been used to characterise flow regimes in gas-fluidised and vibro-fluidised beds, analyse the shape and behaviour of granular beds in rotating drums, and measure the Reynolds stress and turbulent kinetic energy budget in pipe-flow [5–8]. The basic principle of PEPT is shown in Fig. 1.

Ideally, all of the LoRs should lead back to a single point which is the tracer location and these are termed 'true' LoRs [9]. However, LoRs may also be 'scattered' or 'random'. Scattered LoRs are formed when one or both gamma rays in an LoR undergo scattering before reaching the detector. The most common type of scattering for 511 keV gamma-rays is Compton scattering, which is an inelastic scattering process that occurs when a photon interacts with an electron, lowering

* Corresponding author.
E-mail address: mxh1092@student.bham.ac.uk (M. Herald).

https://doi.org/10.1016/j.nima.2022.167831
Received 29 April 2022; Received in revised form 9 November 2022; Accepted 15 November 2022
Available online 22 November 2022
0168-9002/© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

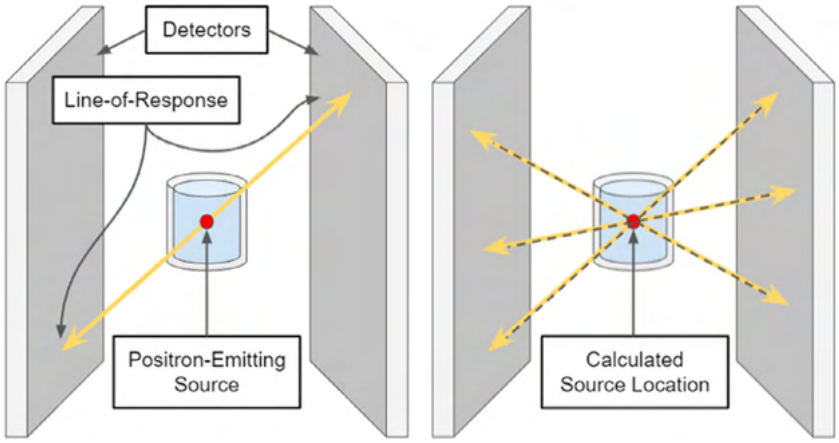

Fig. 1. 511 keV gamma-rays are collected using two opposing radiation detectors and form an LoR (left). When processed with a PEPT algorithm, the LoRs reveal the tracer location (right).

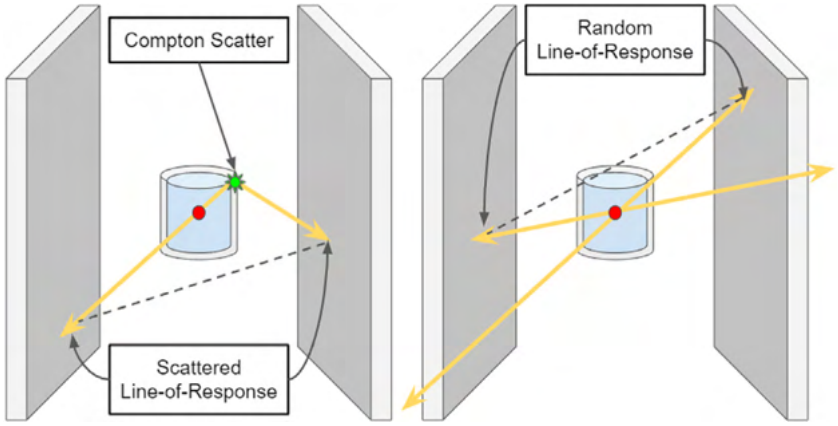

Fig. 2. A scattered LoR is formed when one or both gamma rays undergo Compton scatter (left). A random LoR is formed when two unrelated gamma rays are detected within the coincidence window (right).

the energy of the photon and changing its direction [10,11]. A scattered LoR, shown in Fig. 2, does not intersect the source location and is therefore corrupted for PEPT. While 511 keV gamma-rays are much more penetrating than visible light and X-rays, like all forms of electromagnetic radiation, they exponentially attenuate. For a beam of 511 keV gamma-rays with initial intensity I_0, the amount that penetrates without attenuating is dependent only on the material thickness, Δx, and the linear attenuation coefficient, μ, as shown in Eq. (1) [12].

$$I = I_0 \exp(-\mu \Delta x) \tag{1}$$

A 'random' LoR may also be formed between two unrelated events within the energy window. Since the two gamma rays in a random LoR originate from separate positron annihilations, the LoR is corrupt and will not intersect the source location. The rate of random LoRs is predictable and determined by the count-rate of individual 'singles' events between two detectors, S, and the coincidence window, τ, shown in Eq. (2) [13]. Since the rate of random LoRs squares with the singles rate, at high count-rates random LoRs may form a large fraction of the LoRs in a sample. A diagram of a scattered and random LoR is shown in Fig. 2.

$$R = 2\tau S_1 S_2 \tag{2}$$

The Birmingham Method (BM) PEPT algorithm was designed with the fact in mind that many LoRs in a sample may be corrupted, thus only a fraction of the LoRs in a sample should be used to find the tracer position [2]. The BM works by minimising the sum of the distances of an estimated tracer position to each LoR in the sample, described by Eq. (3), where $D(N)$ is the sum of all the distances in the sample of LoRs, N, and $\delta_{LoR}(m)$ being the three-dimensional distance of an individual LoR to the estimated tracer location, m [8]. Once the tracer position is estimated and the distances of LoRs to the position known, the LoR furthest from this point is removed and the remaining LoRs are recycled to update the estimated position. This iterates until only a user-specified fraction of the LoRs remain, f. This process is shown in Fig. 3. The BM has two free parameters: the fraction of the LoRs remaining in the initial sample, f, and the total number of LoRs in a sample, N_{LoRs}. The default parameters of the BM are 0.05 and 250 for f and N_{LoRs}, respectively. These parameters are conservative so that they can be applied to many different systems and produce reasonable trajectories.

$$D(N) = \sum_{N} \delta_{LoR}(m) \tag{3}$$

It is known that thicker and denser systems will cause more scattered LoRs and that more active tracers will cause more random LoRs. However, little has been done to develop ways to understand how this affects PEPT measurements and, further, predict the values for f and N_{LoRs} which will maximise the spatial accuracy of the BM. In recent work, a method to find the optimal N_{LoRs} has been proposed [14]. This method relies on using a large enough sample to give adequate statistics but is limited such that the tracer does not move significantly compared to the measurement uncertainty during the time used to locate the tracer. This can be summarised in Eq. (4), where w is a detector-specific parameter for spatial resolution, R is the detection rate of LoRs, and v is

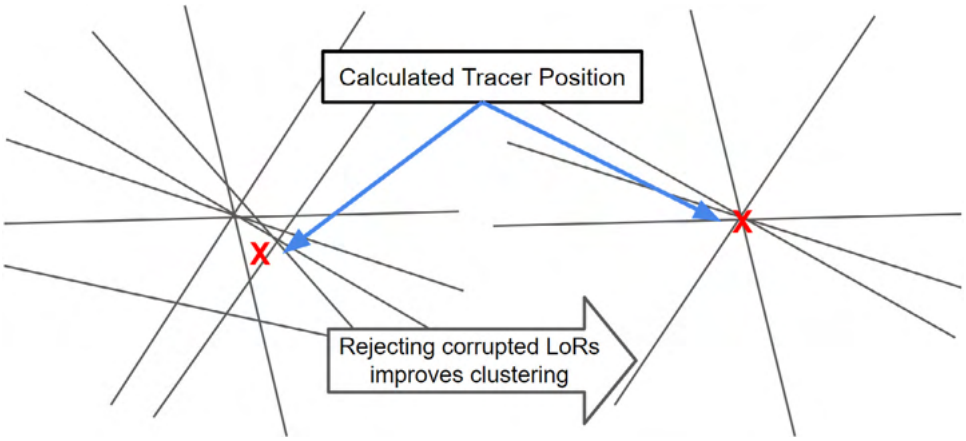

Fig. 3. A sample of LoRs and the initial guess for the tracer position (left). After a fraction of the LoRs furthest away from the initial guess are removed, the remaining LoRs are re-clustered for a more accurate position (right).

the tracer velocity. However, this approach can only be used accurately if the tracer velocity and f are known, which is often not the case.

$$N = f^{\frac{-1}{3}} \left[\frac{Rw}{v}\right]^{\frac{2}{3}} \tag{4}$$

In this work, we use PEPT measurements and simulations to investigate how the spatial accuracy of the BM is affected by the parameters f and N_{LoRs} under a range of experimental conditions, including both stationary and moving tracers. In addition to real experiments, a Geant4 Application for Tomographic Emission (GATE) v 9.1 model of the ADAC Forte, validated to within 10% of experimental measurements, is used to recreate the experiments [15,16]. These serve as further validation for the model and are later expanded to test a range of conditions not explicitly considered during the experiment. The values for f and N_{LoRs} which maximise the spatial accuracy for each PEPT trajectory are extracted as a function of the tracer position. To utilise this information, a new version of the BM, called the Dynamic-BM (DBM), is developed which dynamically changes the fraction of LoRs remaining in the sample and the sample size of LoRs based on an estimate of the tracer position found using non-optimal parameters. This new method is expected to provide both higher spatial and temporal resolution than the original BM algorithm since it can use the optimal parameter combination on a sample-by-sample basis. Moreover, this removes the necessity for trial and error optimisation as well as the cost this engenders and provides a justifiable reason why certain parameters were chosen.

2. Methods

Two experiments are created to investigate how N_{LoRs} and f vary under different conditions. The first experiment consists of low-activity, static tracers placed in the centre of the ADAC Forte dual-headed positron camera. This system is depicted in Fig. 4. The tracer is surrounded by a cylinder of material to induce Compton scattering. The experiment is repeated with six materials detailed in Table 1; this shows how the fraction of corrupted LoRs affects the optimal f parameter. The second experiment is a moderately high-activity tracer fixed to the end of an impeller rotating at 100 RPM in an attritor mill. The BM is used with a range of f and N_{LoRs} to track the tracer, and then the PEPT trajectories are compared to the predicted tracer position based on the known rotation rate and initial position to calculate the mean spatial error. The combination of parameters that maximises the mean spatial accuracy of the trajectory is found. This shows how the optimal N_{LoRs} is affected by tracer motion.

Both the static tracer experiments and the attritor mill experiment are recreated in GATE to serve as validation of the model. Subsequently, the GATE model of the attritor mill is then used again to place the tracer in a range of possible initial positions which will result in different fractions of corrupted LoRs and a range of tracer velocities. Each trajectory is analysed using the BM and the optimal parameter combination is extracted as a function of the tracer position. Ultimately, these virtual PEPT experiments are used to observe how f and N_{LoRs} change throughout the attritor mill system and are used to develop the DBM algorithm. The performance of the DBM will be compared to the best constant parameters and also the default algorithm parameters.

2.1. Static tracer experiment

A small tracer, on the order of the range of positrons in the material and PEPT uncertainty, should be used to ensure the detected LoRs form a tight cluster near the tracer. The tracer chosen for this experiment is a 1 mm diameter sphere of anion exchange resin labelled with fluorine-18 (F-18). F-18 is an ideal positron-emitting isotope for PEPT because it has one of the lowest energy spectra for positrons and thus a low range [17]. The anion exchange resin adsorbs F-18 ions from a solution of radioactive water produced on-site at the Positron Imaging Centre [18]. On the day of the experiment, the tracer was activated with an initial activity of 2.8 MBq and placed inside a 0.5 ml plastic vial for handling. According to a recent characterisation of the ADAC Forte, a tracer of 2.8 MBq will produce less than 5% random LoRs, meaning nearly all LoRs not intersecting the tracer location will most likely have undergone Compton scattering before being detected [15].

To attenuate the 511 keV gamma-rays, the vial is placed in the centre of an 800 ml cylindrical glass beaker filled to 500 ml with bulk density attenuating material. The inner diameter of the beaker is 100 mm and filled to the height of 65 mm, confirming that a volume of material of approximately 500 ml is used. The wall thickness of the beaker is 5 mm. Once filled, the beaker's mass was measured. The initial mass of the beaker was subtracted to find the total mass of the attenuating material. The density of the materials, ρ, is calculated by dividing the mass by the volume. The materials and their properties are listed in Table 1. Materials like air and high-density polyethene (HDPE) have low linear attenuation coefficients, μ, while steel and copper have larger coefficients, meaning they will attenuate a larger fraction of the gamma rays [11]. A small sample of some of these materials and the filled beaker is shown in Fig. 4. The fact that materials are bulk and porous is not expected to have an effect in this experiment since the diameters of the materials are much smaller than the diameter of the beaker. The attenuation of gamma rays in this experiment will approximate those of a perfectly continuous attenuation medium.

During the experiment, the beaker is placed in the centre of the field of view (FOV) of the ADAC Forte. The Forte is a dual-headed positron camera used for PEPT at the Positron Imaging Centre and is the most extensively used detector system for performing PEPT experiments [8,

Fig. 4. A small sample of four of the materials used to induce attenuation (left). From left to right: steel, copper, glass, and MCC. The beaker is filled with steel balls and the source is placed in the centre of the field-of-view (FOV) of the Forte (right).

Table 1
List of the bulk materials and their attenuating properties [11].

Material	ρ_{bulk} (g/cm^3)	μ/ρ_{bulk} (cm^2/g)	μ (cm^{-1})	Attenuation (%)
Air	0.00129	0.0806	0.000111	19.4
HDPE	0.890	0.0931	0.0828	64.8
MCC	1.421	0.0915	0.130	78.0
Glass	1.661	0.0858	0.143	80.7
Steel	4.425	0.0832	0.368	98.0
Copper	5.025	0.0827	0.415	98.7

19,20]. It is comprised of two opposing large-area sodium iodide crystals, 16 mm thick, each with an active area of 380 × 510 mm [15]. The intrinsic efficiency of the Forte is reported to be approximately 23% and it has an energy resolution of 14% [15]. Ideally, an energy window should be set as narrowly as possible around 511 keV, to exclude photons that have Compton scattered [21], but broad enough not to exclude valid annihilation pairs. The energy window is set to 50% to capture all true LoRs, yet inevitably recording some scattered LoRs which can later be discarded by the Birmingham algorithm. The two heads of the Forte have an adjustable separation and for this experiment are set to their maximum separation of 800 mm to achieve the most uniform illumination. Each material is imaged until more than 1,000,000 LoRs have been collected, which is enough to locate the tracer several times using the BM across a range of N_{LoRs} in the sample size. Since the position of the tracer is not known exactly, the standard deviation of the PEPT detected position, the PEPT precision, σ, is used to quantify the performance of different f and N_{LoRs} parameter combinations. This is calculated using Eqs. (5) and (6), which is the three-dimensional standard deviation of the detected positions.

$$\sigma_x = \sqrt{\frac{\Sigma(x_i - \bar{x})^2}{N}} \tag{5}$$

$$\sigma = \sqrt{\sigma_x^2 + \sigma_y^2 + \sigma_z^2} \tag{6}$$

2.2. Moving tracer experiment

In previous work, the spatial accuracy for *static* tracers has been shown to improve when more LoRs are used per sample [22]. However, in real PEPT experiments, the tracer typically moves throughout the system. This movement limits the N_{LoRs} per sample if the tracer moves more than a few millimetres between the detection of the first and last LoRs in the sample. If the tracer moves more than this, particularly if the tracer is changing direction or accelerating, the PEPT-detected position will incur higher spatial errors [2,23]. Thus there should always exist an optimum N_{LoRs} per sample for a given system at a given point in time.

To investigate this in a system representative of a PEPT experiment, an attritor mill is placed near the centre of FOV and a tracer is fixed to the end of the impeller as it rotates in the mill at 100 RPM. Mills of this type are used across a wide range of industrial and pharmaceutical applications and have been studied in the past using PEPT [24,25]. Due to the thick steel walls and predictable circular rotation of a particle fixed to the impeller, this system is an ideal candidate to investigate the effect of N_{LoRs} on the spatial accuracy of the BM. In other work, similar rotating systems have been used to better understand how changing the parameters f and N_{LoRs} affects PEPT measurements [2,26]. The attritor mill and a schematic of the mill dimensions are shown in Fig. 5.

The tracer used for this experiment is a 1.2 mm diameter MCC bead activated with a solution of F-18 and water to an initial activity of 22 MBq. A tracer of this activity is ideal for PEPT experiments in the Forte since this is approximately the activity which will produce the highest true LoR count-rate before dead-time and random LoRs degrade the measurement [15]. This tracer is taped to the end of the upper impeller in the attritor mill at a radius of 63 mm. At 100 RPM the tracer will rotate at a constant velocity of 660 mm/s. The mill is then placed near the centre of the FOV of ADAC Forte at a head separation of 510 mm. The mill is imaged over approximately 1 min (i.e. 100 rotations) to be able to develop good statistics of the tracer locations as it rotates.

To assess the performance of different f and N_{LoRs} combinations using the BM, first, each directional component of the PEPT trajectory is fitted to a sinusoidal equation as a function of time, t, as shown in Eq. (7). The amplitude, A, in the x and z components should be approximately 63 mm and in the y component, 0 if the mill is perfectly oriented with respect to the detector axes. Since the mill is likely not perfectly level a sinusoidal equation is fit to the y-component as well. The rotation rate, ω, is approximately 100 RPM, and ϕ is the phase shift which depends on the initial position of the tracer. The PEPT deviation from this motion is calculated by comparing each PEPT detected position to the predicted position using the fitted equations. In this way, the mean spatial error of the trajectory is calculated using Eqs. (8) and (9).

$$x_{Fit}(t) = A\sin(\omega t + \phi) + c \tag{7}$$

$$\varepsilon_{x_i} = |x_{PEPT_i} - x_{Fit_i}| \tag{8}$$

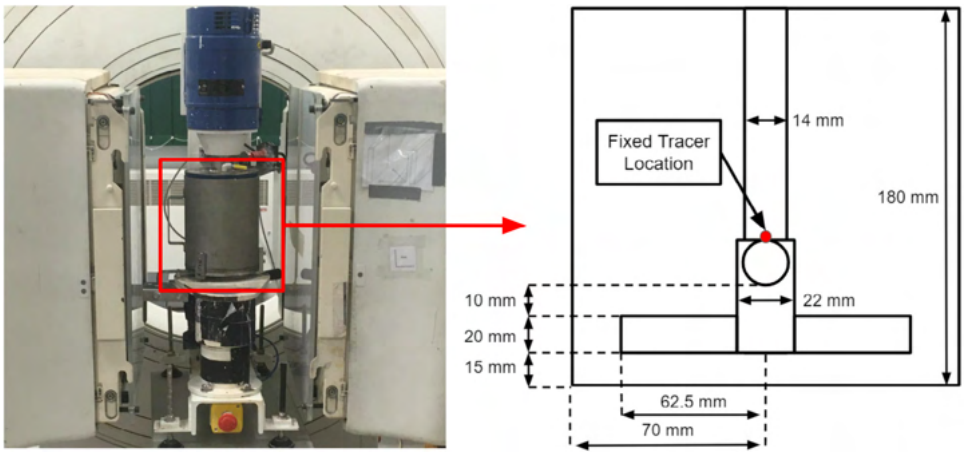

Fig. 5. The attritor mill near the centre of the FOV of ADAC Forte (left). A schematic of the dimensions of the mill and the tracer fixed to the impeller (right).

$$\varepsilon = \frac{1}{N}\sum_{i}^{N}\sqrt{\varepsilon_{x_i}^2 + \varepsilon_{y_i}^2 + \varepsilon_{z_i}^2} \tag{9}$$

2.3. Monte Carlo model

GATE v 9.1 is a powerful tool used to simulate radioactive sources, detectors, and geometries commonly found in medical imaging and radiotherapy applications [27]. PEPT relies on the same equipment as in medical imaging, thus GATE can reliably be used to create virtual PEPT experiments without having to extend the existing toolkit [22,28]. The GATE model of the ADAC Forte used in this work has been validated using a characterisation of the detector when it was installed at the Positron Imaging Centre and calibration experiments following an industry-standard protocol, showing agreement between simulation and experiment to within 10% across all tested metrics [15]. The model includes all of the major structural components of the ADAC Forte such as the sodium-iodide scintillation crystal, back compartment, and cover around the detector heads. The model also includes the 'digitizer' which is based on the pulse-processing stages of the detector, crucial for replicating the spatial resolution, sensitivity, and count-rate response of the detector [29].

The ADAC Forte's digitizer is responsible for converting the interactions of the simulated gamma-rays with the scintillation crystals, termed 'hits', into a 'pulse' which is analogous to what would be produced by the real detector by passing the time, position, and energy of the hit through blurring filters. This is needed because the simulation produces exact values, but in reality, some characteristic imprecision is present. A flow diagram of the ADAC Forte's digitizer is shown in Fig. 6. Once the hit is registered, the information from the interaction such as the time, position, and energy absorbed by the crystal is recorded as a GATE pulse. Next, a series of blurring filters are added to match the time resolution, spatial resolution, and energy resolution of the detector. In the simulation, the values for these pulse properties are known absolutely and must be blurred to mimic the imprecision of real detectors. Next, the pulses which fall near to each other in a short time window of 400 ns are allowed to pile-up on one another forming a combined signal. After this step, energy thresholds for recording the pulse are added which ensures that only pulses falling within the threshold will trigger a detector response, excluding pulses below 250 keV and 950 keV. Pulses falling within this range trigger the detector to record them and this creates a period where no pulses can be recorded called dead-time. The dead-time model used for the single pulses is a paralysable model with a dead-time of approximately 1.2 μs which can be restarted by another pulse [30]. Of the recorded singles pulses, only those falling within a 50% energy window of the 511 keV photo-peak, which is 360 keV to 640 keV, are considered for forming a coincidence

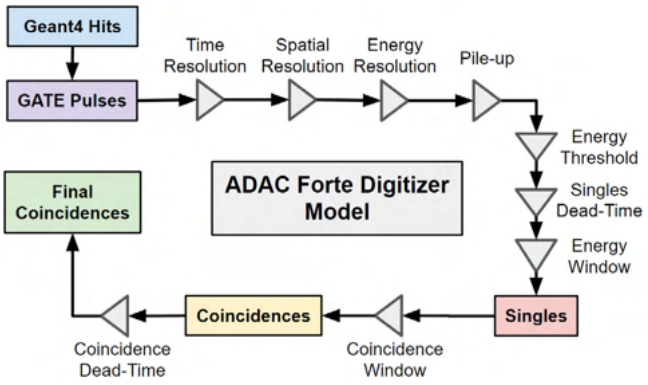

Fig. 6. The ADAC Forte's digitizer is represented as a flow diagram of the pulse-processing stages.

if another pulse on the opposite head of the detector records a single within the coincidence window of 15 ns. Finally, from the coincidences that are formed, an additional non-paralysable coincidence dead-time of 1.2 μs is used to filter the data following the data write speed limitation of the detector, which forms the final coincidences. These final coincidences are what are ultimately considered the LoRs which are used for PEPT [15].

This GATE model of the ADAC Forte is first used to repeat the experiments described in Sections 2.1 and 2.2 for comparison and further validation. After this, the GATE model and recreated attritor mill are used to expand the tested conditions to observe how the optimal combination of f and N_{LoRs} changes throughout the system. These optimal parameters are used in Section 2.4 to develop the DBM algorithm.

2.3.1. Static tracer simulation

Following the previously described experiments, the tracer, geometry, and detector for the static tracer experiments are replicated in GATE. This starts by creating a radioactive tracer defined as a 1 mm diameter resin sphere, emitting positrons with an F-18 energy spectrum. The tracer is placed inside a geometric model of the glass beaker which has been filled with 500 ml of bulk-density material. A cross-section of the beaker and source geometry is shown in Fig. 7. To replicate the attenuation of each material, new material definitions are added to the GATE material database which is described in Table 1. The tracer is prescribed an activity equal to that of the tracer's activity at the beginning of each data acquisition. Initially, the tracer was approximately 2.8 MBq at 11:48 am, but by the end of all experiments

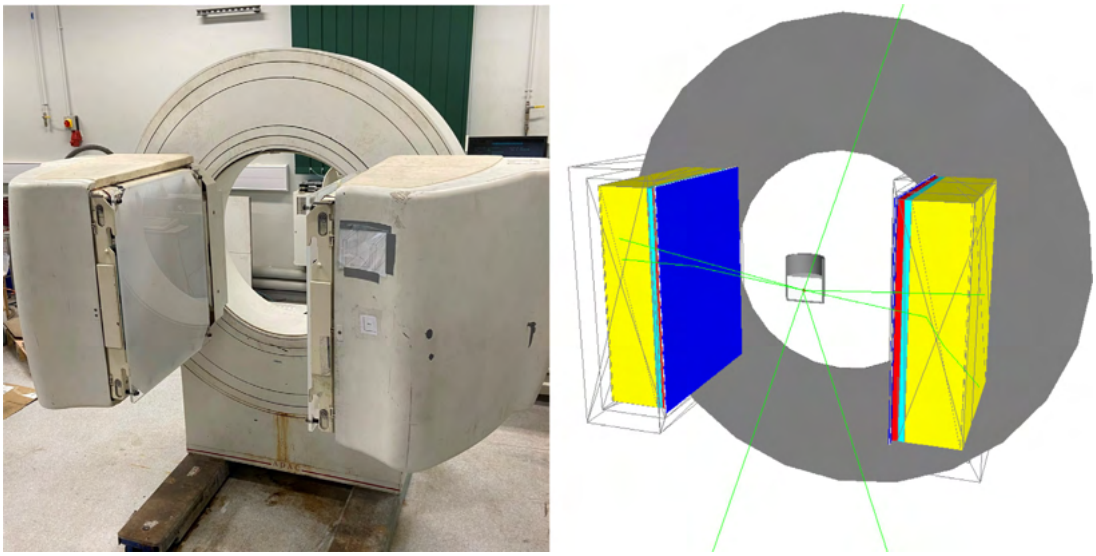

Fig. 7. The ADAC Forte dual-headed positron camera at the Positron Imaging Centre (left). GATE model of the Forte with a cutaway of the experimental geometry and source in the centre of the FOV (right).

decayed to approximately 1.3 MBq. For each simulation, the activity is adjusted accordingly to compensate for decay. Finally, the tracer and beaker are placed in the centre of the FOV of the GATE model of the ADAC Forte. The detector separation is set to 800 mm. While the exact positions of the tracers simulated in GATE are known, Eq. (6) is still used to calculate the PEPT precision to compare the real experiment and simulations directly.

2.3.2. Moving tracer simulation

In a similar manner as before, the moving tracer experiment is recreated in GATE. The model of the ADAC Forte is the same as the model used for the static tracer experiment and the head separation is set to 510 mm, which is as close to the system as the detectors could be set. A lower head separation results in higher geometric efficiency, meaning LoRs can be collected at a higher rate. At the time of the experiment, the tracer was measured to approximately 20 MBq. This is modelled in GATE as a 1.2 mm diameter sphere emitting positrons with an F-18 energy spectrum.

The mill dimensions can be found on the right-hand side of Fig. 5 and the GATE model of the mill is shown in Fig. 8. It consists primarily of a 70 mm diameter vessel lined with polyethylene and a stainless-steel impeller with 4 pins. The polyethylene liner is 10 mm thick and the outer stainless-steel wall is 5 mm thick. The inside of the vessel is filled with air to match the experimental conditions. In a real experiment, the mill will be filled with grinding material. To recreate this in GATE, the model could be filled with a volume of the bulk density material, or a Discrete Element Method simulation can be used to calculate the three-dimensional density distribution and this can be imported into the simulation using a voxel array [23] (see Fig. 8).

The whole mill is rotated about the y-axis at 100 RPM to induce particle motion. This is achieved through discrete rotations of every simulation time-step of 0.0005 s. For the tracer fixed at a 63 mm radius, this results in a change in position of approximately 0.33 mm per time-step. This is smaller than the tracer diameter and can thus be safely used to mimic continuous tracer motion. The GATE simulation produces LoRs which are processed in the same way as real PEPT data. A range of different f and N_{LoRs} are used to locate the tracer. The PEPT trajectory is then compared back to the GATE-prescribed tracer positions using Eqs. (7)–(9). The GATE simulated tracer's position is known exactly and there is no variability in its rotation rate or system vibration present. As such, the spatial accuracy of the GATE simulated tracer is expected to be somewhat higher than in the real experiment.

2.4. Dynamic Birmingham Method

In PEPT experiments, the scattering environment and detector sensitivity change as a function of the tracer position. The amount of corrupted LoRs in a sample affects the optimal f while the sensitivity affects the optimal N_{LoRs}. However, the BM uses constant parameter values. This means users must choose a conservative parameter combination that will work over the whole data set. Inevitably, this will return trajectories with lower spatial accuracy than is theoretically possible. To solve this problem, a variant of BM is developed called Dynamic-BM (DBM), 'Dynamic' because it can dynamically change f and N_{LoRs} to the optimal values determined by a Monte Carlo simulation of the experiment as the tracer moves through the system.

Initially, the DBM uses constant values of f and N_{LoRs} to estimate the position of the tracer. Then, for each initially-estimated tracer position, the optimal parameter values are looked up from a table produced by Monte Carlo simulation. To use this table, the predicted 3-dimensional location of the tracer and other optional information is input, then the closest simulated position is found and the optimal parameters are output. This approach relies on the user simulating a sufficiently high number of possible tracer positions within the system such that the change in parameter values between adjacent simulated positions is smooth. The Monte Carlo simulations must be analysed before using the DBM to generate the optimal parameter look-up table.

To test the DBM, the attritor mill is simulated using GATE with the tracer placed in a range of initial positions. This is done to investigate the whole system's behaviour rather than only a single region. To match the experimental conditions, the mill is rotated at 100 RPM and the tracer activity is set to 22 MBq. Each tracer position is a new GATE simulation. The initial tracer positions are seeded in one quadrant of the mill to take advantage of the symmetry of the system. The positions are created in regular intervals in the x and z direction from 0–60 mm in 20 mm steps and the y-direction from −50–50 mm in 12.5 mm steps. The positions falling outside of the system or intersecting the impeller are removed, leaving a total of 84 positions. Since the impeller is rotating the position of the tracer is important for the optimal f and N_{LoRs} values since it will cause different amounts of scattering depending on its rotated angle. In the simulation, the rotation of the impeller is prescribed so it is known, but during a PEPT experiment, this can be more difficult to ascertain. However, the ADAC Forte can record readings from an optical switch directly into the data file, allowing the impeller rotation angle to be recorded throughout an experiment [31].

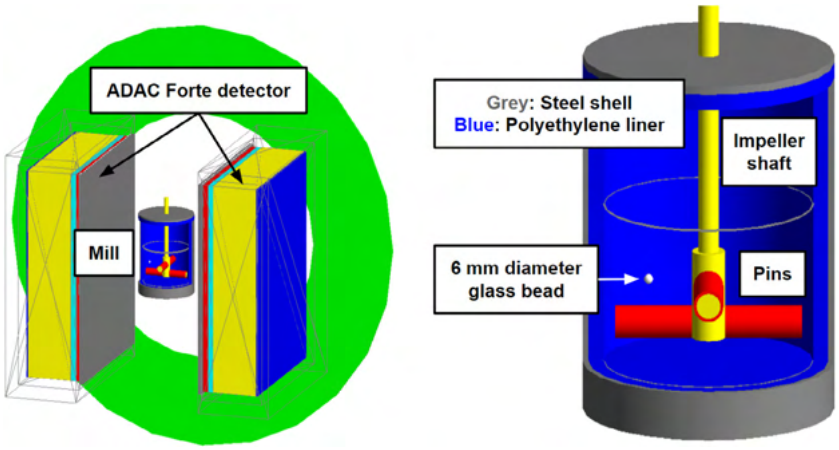

Fig. 8. The GATE model of the ADAC Forte dual-headed positron camera (left) and a cut-away of the attritor mill (right). In the cut-away, the major components are labelled.

The 360° is divided into 30° increments and the optimal f and N_{LoRs} are calculated over this range.

To use the look-up tables, first, the amount the mill is rotated must be determined. This is done by multiplying the time by 360° and dividing by the rotation period, 0.6 s. Since all rotations beyond 360° are duplicates of previous rotations, a modulo operation is performed. Then all the simulated positions within the nearest rotation are compiled and the values of f or N_{LoRs} for the position closest to the PEPT estimated position are extracted. Using this method, the optimal f and N_{LoRs} of any point within the system can be determined. Due to the change in N_{LoRs} per sample, the number of detected positions in a trajectory processed using the BM and the DBM may vary. To ensure the changing parameters are used at the appropriate position within the data set, the f and N_{LoRs} are linearly interpolated as a function of time and the average time of the LoRs in each sample is used to compute the exact parameter values.

To assess the performance of the DBM, it is compared to both the default BM parameter values of 0.05 f and 250 N_{LoRs} as well as the best constant parameter values. This represents the default and the best possible BM performance, showing how dynamically changing the parameter values improves PEPT algorithms. The method of comparison between the PEPT-detected trajectories and the GATE-prescribed trajectory is the same as in Section 2.3.2.

3. Results and discussion

3.1. Model validation results

Each of the static tracer experiments described in Sections 2.1 and 2.2 produce a unique amount of attenuation. Some of the coincident 511 keV gamma-rays attenuated via Compton scattering are inadvertently passed to the PEPT algorithm as corrupted LoRs. When the attenuating medium is air, the amount of corrupted LoRs is relatively low, but as the medium becomes more attenuating, a larger fraction of LoRs are corrupted. This is clearly shown in Fig. 9 where the LoRs from the air, glass, and copper attenuation experiments are plotted.

Since these materials produce different fractions of corrupted LoRs, the optimal f for each experiment should be a unique value. The LoRs from the real experiment and simulations were both processed using the BM under a range of f and N_{LoRs} and the variation of the standard deviation in position for different combinations of the two parameters is plotted as colour variation in Fig. 10. These plots show that for static tracers the optimal value of f is decreased when more corrupted LoRs are present and the standard deviation in the position decreases as more N_{LoRs} per sample are used. The experiment and GATE simulations closely agree across all the parameter combinations, both in the optimal values for f and in the standard deviation in the tracer position. For a static tracer, the optimal value for f remains constant and the standard deviation in position will always decrease with greater N_{LoRs}. The optimal values for f across all the materials tested for both the experiment and the simulation are shown in Fig. 11. Additionally, the values for f which minimise the uncertainty are plotted against the fraction of true LoRs in Fig. 12 demonstrating that f must lower when more attenuation occurs. Moreover, this shows that the values for f which minimise position uncertainty are approximately equivalent to the fraction of true LoRs in the sample. This, until now, has been an assumption of the BM, but this provides the first direct evidence that this assumption is true. The experiment and simulation provide similar values in both the overall PEPT precision across all the parameter combinations tested and also the values of f which minimise the uncertainty for a given N_{LoRs}.

Table 2
Comparisons of the experiment and simulation in the optimal parameters for the moving tracer experiment.

Method	Optimal f	Optimal N_{LoRs}
Experiment	0.275	1400
Simulation	0.25	1300
Percent error (%)	−9.09	−7.14

A similar analysis was conducted for the moving tracer experiment. A 22 MBq tracer was fixed to the impeller of an attritor mill and rotated at 100 RPM. In this case, a position error was calculated using Eq. (7) for the experiment and simulation. A range of constant f and N_{LoRs} values are used with the BM to assess the spatial errors produced under different parameter combinations. Trajectories of the experiment and simulation are presented in Fig. 13 showing that the rates of detection are approximately the same. In Fig. 14, the parameter values for f and N_{LoRs} are varied to assess their effect on the spatial error of the reconstructed PEPT trajectory and to find the optimal combination of parameters. The experimental plot has a generally higher error because the error was calculated using fitted functions, assuming perfect circular motion, whereas the GATE simulations benefit from having analytical functions describing the tracer motion. From Fig. 14, the optimal values for f and N_{LoRs} for the real experiment are determined to be 0.275 and 1400, respectively. Similarly, the optimal parameters for the simulation are determined to be 0.25 and 1300, respectively. The optimal parameters and their percent errors are shown in Table 2.

3.2. Applying the Dynamic Birmingham Method to an attritor mill

The GATE simulations of the attritor mill are expanded to test a range of initial particle positions. In a real PEPT experiment on this mill, tracers near the centre of the mill will move slower than tracers

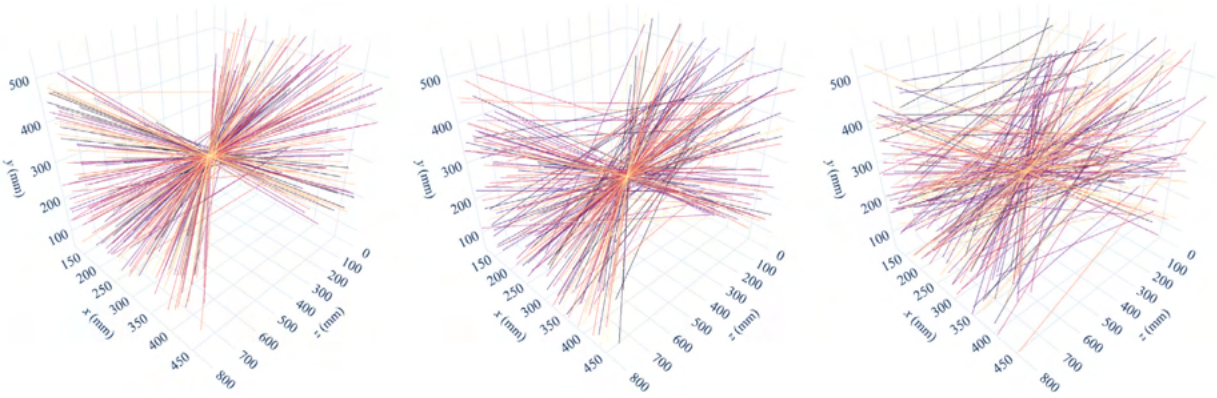

Fig. 9. Three samples of 200 experimentally acquired LoRs from the static tracer experiment for air (left), glass (middle), and copper (right) show that the amount of corrupted LoRs in the sample increases with more attenuation.

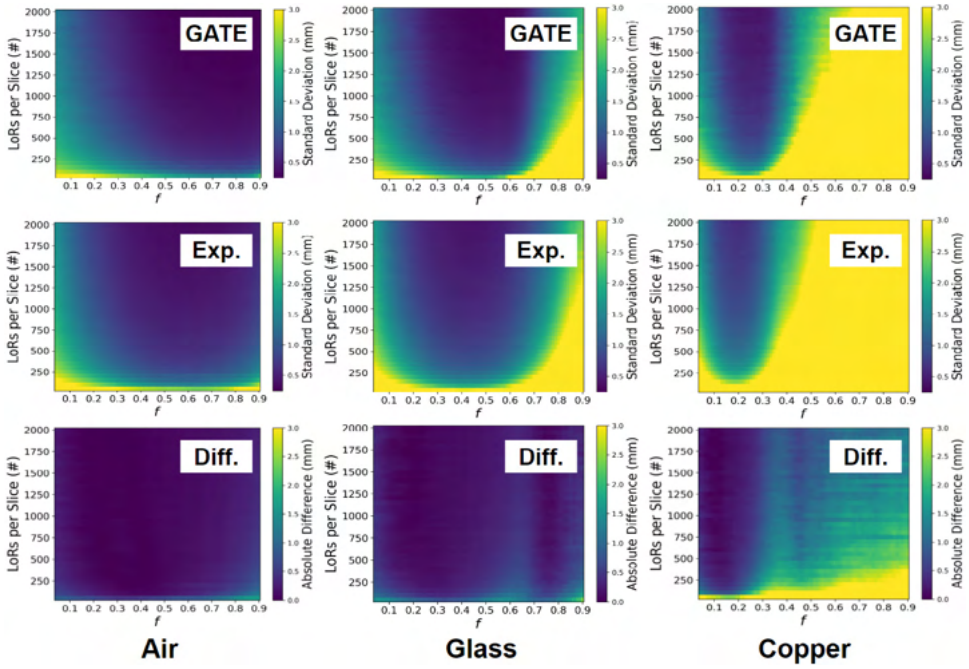

Fig. 10. Position uncertainty in the static tracer experiments and simulations for air (left), glass (middle), and copper (right), and the absolute difference between the simulation and experiment (bottom).

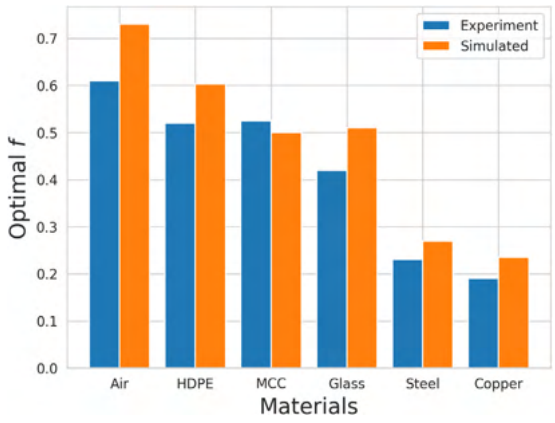

Fig. 11. A comparison of the optimal values for f across all the materials tested.

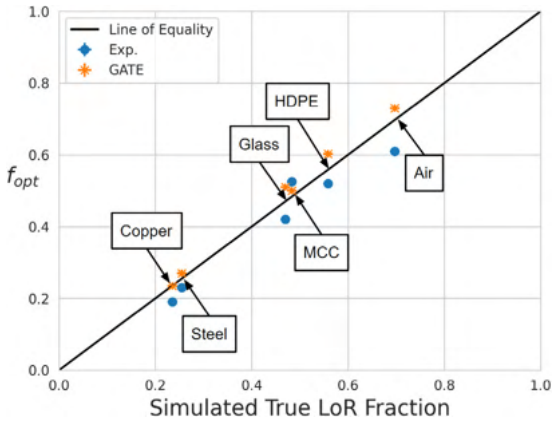

Fig. 12. A comparison of the optimal values for f across all the materials as a function of the fraction of true LoRs in the sample.

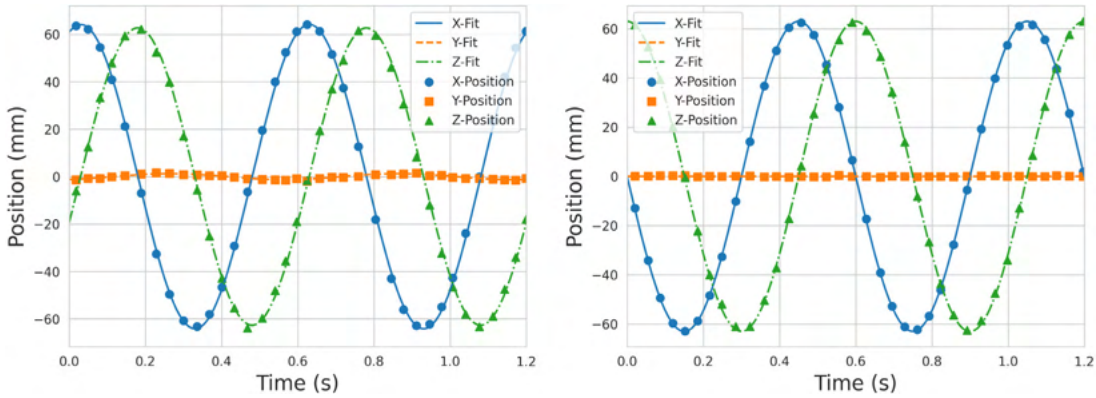

Fig. 13. Trajectories for the experimental tracer in the attritor mill under a constant f of 0.25 and N_{LoRs} of 1500 (left) compared to the simulated tracer trajectory under the same parameters (right).

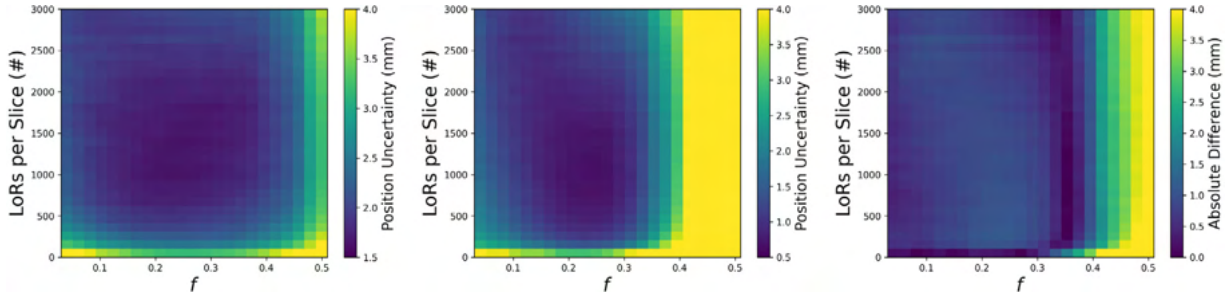

Fig. 14. The parameter values for f and N_{LoRs} are varied for the real experiment (left) and GATE simulation (middle) of the moving source, affecting the spatial error of the reconstructed PEPT trajectory, and the difference between the two is shown (right).

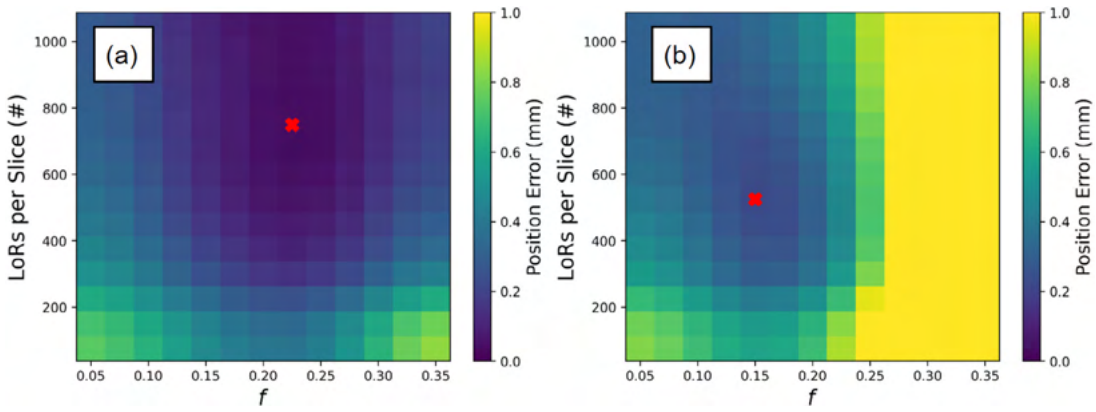

Fig. 15. Panel (a) shows the results of a parameter sweep when the tracer is beside the impeller shaft where the least amount of LoRs have been attenuated while panel (b) shows the effects of the tracer passing behind the impeller shaft where a larger fraction of LoRs are attenuated.

near the tip of the impeller blade, which is where the maximum velocities are recorded [25]. Therefore, the GATE simulated movement of the tracers in this work is similar to that of a real experiment. For each of the tested positions, the optimal values for f and N_{LoRs} were recorded as the tracer moves through the system. These optimal parameters were calculated for every 30° rotation to update their values continuously. Fig. 15 shows the change of optimal parameters at two different degrees of rotation, one where the least amount of attenuation occurs and another where the tracer is behind the impeller shaft, with respect to the detectors, where the most attenuation occurs. As the mill rotates, the optimal parameters fluctuate, becoming more or less conservative when more corrupted events are recorded or the count-rate decreases, such as when the tracer passes in front or behind the impeller shaft.

After the optimal parameters have been extracted, these are used to inform the DBM algorithm. Each simulated trajectory is reprocessed with an estimate of the optimal parameters and the mean spatial error is calculated over the trajectory. These are compared to the trajectories extracted using the default and the best constant parameters. The results of this comparison are in Table 3. The mean spatial error of the default BM parameters, best constant parameters, and the DBM are 2.20 mm, 0.544 mm, and 0.517 mm, respectively. While the default parameters can reconstruct the trajectory, when tailored parameters extracted from Monte Carlo simulation are used, in this case at least, the errors decrease to nearly a quarter of their original values. When dynamic parameters are used, the errors decrease by 76.5% over the default parameters and by 4.03% over the best constant parameters. A histogram of the percent changes in spatial error between the trajectories produced with the best constant parameter and the DBM is shown in Fig. 16. This plot shows the DBM increases the accuracy of nearly all trajectories average and also that some individual trajectories are

Table 3
Mean trajectory comparisons using different algorithm parameter methods.

Parameters	Spatial error (mm)	STD (mm)	Locations (N)
Default	2.20	1.48	1337
Best constant	0.544	0.372	223
Dynamic	0.517	0.3481	236

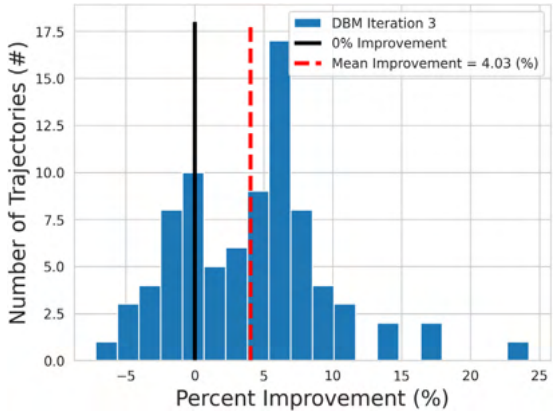

Fig. 16. The change in spatial accuracy for each trajectory. The mean improvement is 4.03% over the best constant parameters.

improved by over 10%. The trajectories which are the most improved are from areas of the system around the impeller blades where the local optimal parameters deviate the most from the best constant parameters.

This work is significant because it demonstrates that the optimal parameter values for the BM can be predicted and that using these improves spatial accuracy over default parameter values. Moreover, not only can the best constant parameter values be estimated, but also the BM can be extended to use the local optimal parameters based on an estimate of the tracer's position, producing the best possible trajectory. In the future, using the methods presented here, Monte Carlo simulations conducted in conjunction with PEPT experiments can be used to remove the guesswork of choosing f and N_{LoRs}. This means users of PEPT algorithms do not need to be experts to produce good trajectories, making the Birmingham Method PEPT algorithm more rigorous and more accessible. While producing simulations of the systems being imaged with PEPT increases the amount of work being done, this work shows that it is worth doing this additional step because of the improvement in trajectories. Additionally, while this work was conducted using the BM PEPT algorithm and the Forte detector system, a similar workflow could be applied to other PEPT algorithms which have free parameters, such as the Feature Point Identification PEPT algorithm, and other detector systems for which a validated Monte Carlo model exists, such as the Phillips Vereos digital photon counting system [32,33].

4. Conclusions

This work shows that the Monte Carlo simulation of PEPT experiments can be used to predict the parameter values of the Birmingham Method PEPT algorithm which provide the best trajectory spatial accuracy. This is demonstrated through experiments and simulations which show that the values for the two key free parameters, f and N_{LoRs}, can be found by comparing the known positions to the reconstructed position and identifying the parameter combination which minimises the spatial error. Further, the values for these parameters are found to be sensitive to the amount of corrupted LoRs in a sample of LoRs and also the velocity and activity of the tracer. For static tracers, the optimal value for f is a balance between removing as many corrupted LoRs as possible while still preserving as many true LoRs and for N_{LoRs} larger values always decrease the spatial error. If too many LoRs are removed (i.e. setting f too low) the spatial error increases due to the statistical uncertainty. Other work has shown uncertainty is proportional to the inverse square root of the number of LoRs remaining in the sample. However, if the tracer is moving, a global optimal combination of f and N_{LoRs} exists. The value for these two parameters is a complex relationship between the amount of corrupted LoRs, tracer velocity, and tracer activity. However, the optimum combination can be found through Monte Carlo simulation of the system and trying a range of parameter value combinations then selecting the one which minimises the spatial error.

Ultimately, the Birmingham Method is extended to update the algorithm parameters on a sample-by-sample basis using the local optimal parameter values. This is demonstrated using a simulation of an attritor mill, representative of a typical PEPT experiment. PEPT trajectories are reconstructed using the Birmingham Method with dynamically updated parameters and compared to trajectories produced with the default and also the best constant parameter combination. In this system, the PEPT trajectory spatial errors are decreased by 76.5% compared to the default parameter and 4.03% compared to the best constant parameters. In light of these results, we suggest that Monte Carlo simulations be used in conjunction with PEPT experiments to determine the algorithm parameters for the Birmingham Method. Doing this removes the guesswork for PEPT users, making the technique more rigorous.

CRediT authorship contribution statement

Matthew Herald: Conceptualization, Methodology, Software, Visualization, Writing – original draft. **Jack Sykes:** Writing – review & editing. **David Parker:** Software. **Jonathan Seville:** Supervision, Writing – review & editing. **Tzany Wheldon:** Resources. **Christopher Windows-Yule:** Supervision, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgements

The authors would like to recognise the valuable contributions made by Andrei Leonard Nicuşand for the development of the pept Python package on GitHub and to Dan Rhymer and Mondelez for the use of equipment. The computations described in this paper were performed using the University of Birmingham's BlueBEAR HPC service, which provides a High-Performance Computing service to the University's research community. See theirwebsite for more details. This work was supported by the Engineering and Physical Sciences Research Council, United Kingdom (grant number EP/T034327/1); and the Engineering and Physical Sciences Research Council, United Kingdom Centre for Doctoral Training in Topological Design (grant number EP/S02297X/1). Work at the Positron Imaging Centre is supported in part by a grant from the Engineering and Physical Science Research Council, United Kingdom EP/R045046/1, Probing Multiscale Complex Multiphase Flows with Positrons for Engineering and Biomedical Applications.

References

[1] M. Hawkesworth, D. Parker, P. Fowles, J. Crilly, N. Jefferies, G. Jonkers, Nonmedical applications of a positron camera, Nucl. Instrum. Methods Phys. Res. A 310 (1–2) (1991) 423–434, http://dx.doi.org/10.1016/0168-9002(91)91073-5, URL: https://linkinghub.elsevier.com/retrieve/pii/0168900291910735.

[2] D. Parker, C. Broadbent, P. Fowles, M. Hawkesworth, P. McNeil, Positron emission particle tracking - A technique for studying flow within engineering equipment, Nucl. Instrum. Methods Phys. Res. A 326 (3) (1993) 592–607, http://dx.doi.org/10.1016/0168-9002(93)90864-E, URL: https://linkinghub.elsevier.com/retrieve/pii/016890029390864E.

[3] n. Guobao Wang, High temporal-resolution dynamic PET image reconstruction using a new spatiotemporal kernel method, IEEE Trans. Med. Imaging 38 (3) (2019) 664–674, http://dx.doi.org/10.1109/TMI.2018.2869868.

[4] C. Windows-Yule, J. Seville, A. Ingram, D. Parker, Positron emission particle tracking of granular flows, Annu. Rev. Chem. Biomol. Eng. 11 (1) (2020) annurev–chembioeng–011620–120633, http://dx.doi.org/10.1146/annurev-chembioeng-011620-120633, URL: https://www.annualreviews.org/doi/10.1146/annurev-chembioeng-011620-120633.

[5] C.R.K. Windows-Yule, N. Rivas, D.J. Parker, A.R. Thornton, Low-frequency oscillations and convective phenomena in a density-inverted vibrofluidized granular system, Phys. Rev. E 90 (6) (2014) 062205, http://dx.doi.org/10.1103/PhysRevE.90.062205.

[6] A. Morrison, I. Govender, A. Mainza, D. Parker, The shape and behaviour of a granular bed in a rotating drum using Eulerian flow fields obtained from PEPT, Chem. Eng. Sci. 152 (2016) 186–198, http://dx.doi.org/10.1016/j.ces.2016.06.022.

[7] C. Wiggins, N. Patel, Z. Bingham, A. Ruggles, Qualification of multiple-particle positron emission particle tracking (M-PEPT) technique for measurements in turbulent wall-bounded flow, Chem. Eng. Sci. 204 (2019) 246–256, http://dx.doi.org/10.1016/j.ces.2019.04.030.

[8] D.J. Parker, Positron emission particle tracking and its application to granular media, Rev. Sci. Instrum. 88 (5) (2017) 051803, http://dx.doi.org/10.1063/1.4983046.

[9] W.W. Moses, Fundamental limits of spatial resolution in PET, Nucl. Instrum. Methods Phys. Res. A 648 (2011) S236–S240, http://dx.doi.org/10.1016/j.nima.2010.11.092.

[10] A.H. Compton, A quantum theory of the scattering of X-rays by light elements, Phys. Rev. 21 (5) (1923) 483–502, http://dx.doi.org/10.1103/PhysRev.21.483.

[11] M. Berger, J. Hubbell, S. Seltzer, Z. Chang, J. Coursey, R. Sukumar, D. Zucker, K. Olsen, XCOM: Photon cross sections database, 2009, NIST, URL: https://www.nist.gov/pml/xcom-photon-cross-sections-database.

[12] K. Singh, H. Singh, V. Sharma, R. Nathuram, A. Khanna, R. Kumar, S. Singh Bhatti, H. Singh Sahota, Gamma-ray attenuation coefficients in bismuth borate glasses, Nucl. Instrum. Methods Phys. Res. B 194 (1) (2002) 1–6, http://dx.doi.org/10.1016/S0168-583X(02)00498-6, URL: https://www.sciencedirect.com/science/article/pii/S0168583X02004986.

[13] C. Stearns, D. McDaniel, S. Kohlmyer, P. Arul, B. Geiser, V. Shanmugam, Random coincidence estimation from single event rates on the Discovery ST PET/CT scanner, in: 2003 IEEE Nuclear Science Symposium. Conference Record (IEEE Cat. No.03CH37515), IEEE, Portland, OR, USA, 2004, pp. 3067–3069, http://dx.doi.org/10.1109/NSSMIC.2003.1352545, URL: http://ieeexplore.ieee.org/document/1352545/.

[14] D.J. Parker, D.M. Hampel, T. Kokalova Wheldon, Performance evaluation of the current Birmingham PEPT cameras, Appl. Sci. 12 (14) (2022) 6833, http://dx.doi.org/10.3390/app12146833, URL: https://www.mdpi.com/2076-3417/12/14/6833.

[15] M. Herald, T. Wheldon, C. Windows-Yule, Monte Carlo model validation of a detector system used for positron emission particle tracking, Nucl. Instrum. Methods Phys. Res. A 993 (2021) 165073, http://dx.doi.org/10.1016/j.nima.2021.165073.

[16] S. Jan, G. Santin, D. Strul, S. Staelens, K. Assié, D. Autret, S. Avner, R. Barbier, M. Bardiès, P.M. Bloomfield, D. Brasse, V. Breton, P. Bruyndonckx, I. Buvat, A.F. Chatziioannou, Y. Choi, Y.H. Chung, C. Comtat, D. Donnarieix, L. Ferrer, S.J. Glick, C.J. Groiselle, D. Guez, P.-F. Honore, S. Kerhoas-Cavata, A.S. Kirov, V. Kohli, M. Koole, M. Krieguer, D.J. van der Laan, F. Lamare, G. Largeron, C. Lartizien, D. Lazaro, M.C. Maas, L. Maigne, F. Mayet, F. Melot, C. Merheb, E. Pennacchio, J. Perez, U. Pietrzyk, F.R. Rannou, M. Rey, D.R. Schaart, C.R. Schmidtlein, L. Simon, T.Y. Song, J.-M. Vieira, D. Visvikis, R. Van de Walle, E. Wieërs, C. Morel, GATE - Geant4 Application for Tomographic Emission: A simulation toolkit for PET and SPECT, Phys. Med. Biol. 49 (19) (2004) 4543–4561, URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3267383/.

[17] S. Jan, C. Comtat, D. Strul, G. Santin, R. Trebossen, Monte Carlo simulation for the ECAT EXACT HR+ system using GATE, in: 2003 IEEE Nuclear Science Symposium. Conference Record (IEEE Cat. No.03CH37515), IEEE, 2004, pp. 2545–2548, http://dx.doi.org/10.1109/NSSMIC.2003.1352409, URL: http://ieeexplore.ieee.org/document/1352409/.

[18] D.J. Parker, X. Fan, Positron emission particle tracking—Application and labelling techniques, Particuology 6 (1) (2008) 16–23, http://dx.doi.org/10.1016/j.cpart.2007.10.004.

[19] D. Parker, R. Forster, P. Fowles, P. Takhar, Positron emission particle tracking using the new Birmingham positron camera, Nucl. Instrum. Methods Phys. Res. A 477 (1–3) (2002) 540–545, http://dx.doi.org/10.1016/S0168-9002(01)01919-2, URL: https://linkinghub.elsevier.com/retrieve/pii/S0168900201019192.

[20] T.W. Leadbeater, D.J. Parker, J. Gargiuli, Positron imaging systems for studying particulate, granular and multiphase flows, Particuology 10 (2) (2012) 146–153, http://dx.doi.org/10.1016/j.partic.2011.09.006.

[21] R. Yao, J. Cadorette, J.-F. Beaudoin, R. Lecomte, Energy window optimization of PET detectors for SPECT imaging, in: 2013 IEEE Nuclear Science Symposium and Medical Imaging Conference, 2013 NSS/MIC, 2013, pp. 1–4, http://dx.doi.org/10.1109/NSSMIC.2013.6829104.

[22] C.R.K. Windows-Yule, M. Herald, L. Nicusan, C. Wiggins, G. Pratx, S. Manger, E.A. Odo, T. Leadbeater, J. Pellico, R. de Rosales, A. Renaud, I. Govender, L. Carasik, A. Ruggles, T. Kokalova-Wheldon, J. Seville, D.J. Parker, Recent advances in positron emission particle tracking: A comparative review, Rep. Progr. Phys. (2021) http://dx.doi.org/10.1088/1361-6633/ac3c4c, URL: https://iopscience.iop.org/article/10.1088/1361-6633/ac3c4c.

[23] M. Herald, J. Sykes, D. Werner, J. Seville, C. Windows-Yule, DEM2gate: Combining discrete element method simulation with virtual positron emission particle tracking experiments, Powder Technol. (2022) 117302, http://dx.doi.org/10.1016/j.powtec.2022.117302, URL: https://linkinghub.elsevier.com/retrieve/pii/S0032591022001966.

[24] Y. Yang, N.A. Rowson, R. Tamblyn, A. Ingram, Effect of operating parameters on fine particle grinding in a vertically stirred media mill, Sep. Sci. Technol. 52 (6) (2017) 1143–1152, http://dx.doi.org/10.1080/01496395.2016.1276931.

[25] D. Daraio, J. Villoria, A. Ingram, A. Alexiadis, E. Hugh Stitt, M. Marigo, Validation of a Discrete Element Method (DEM) model of the grinding media dynamics within an attritor mill using Positron Emission Particle Tracking (PEPT) measurements, Appl. Sci. 9 (22) (2019) 4816, http://dx.doi.org/10.3390/app9224816, URL: https://www.mdpi.com/2076-3417/9/22/4816.

[26] K. Cole, D.J. Barker, P.R. Brito-Parada, A. Buffler, K. Hadler, I. Mackay, D. Mesa, A.J. Morrison, S. Neethling, A. Norori-McCormac, B. Shean, J. Cilliers, Standard method for performing positron emission particle tracking (PEPT) measurements of froth flotation at PEPT Cape Town, MethodsX 9 (2022) 101680, http://dx.doi.org/10.1016/j.mex.2022.101680, URL: https://www.sciencedirect.com/science/article/pii/S2215016122000644.

[27] S. Jan, G. Santin, D. Strul, S. Staelens, K. Assié, D. Autret, S. Avner, R. Barbier, M. Bardiès, P.M. Bloomfield, et al., GATE: A simulation toolkit for PET and SPECT, Phys. Med. Biol. 49 (19) (2004) 4543–4561, http://dx.doi.org/10.1088/0031-9155/49/19/007.

[28] M. Herald, Z. Bingham, R. Santos, A. Ruggles, Simulated time-dependent data to estimate uncertainty in fluid flow measurements, Nucl. Eng. Des. 337 (2018) 221–227, http://dx.doi.org/10.1016/j.nucengdes.2018.07.005, URL: https://linkinghub.elsevier.com/retrieve/pii/S0029549318303844.

[29] S. Kerhoas-Cavata, D. Guez, Modeling electronic processing in GATE, Nucl. Instrum. Methods Phys. Res. A 569 (2) (2006) 330–334, http://dx.doi.org/10.1016/j.nima.2006.08.107.

[30] S. Usman, A. Patil, Radiation detector deadtime and pile up: A review of the status of science, Nucl. Eng. Technol. 50 (7) (2018) 1006–1016, http://dx.doi.org/10.1016/j.net.2018.06.014, URL: https://linkinghub.elsevier.com/retrieve/pii/S1738573318302596.

[31] D. Parker, D. Allen, D. Benton, P. Fowles, P. McNeil, M. Tan, T. Beynon, Developments in particle tracking using the Birmingham Positron Camera, Nucl. Instrum. Methods Phys. Res. A 392 (1–3) (1997) 421–426, http://dx.doi.org/10.1016/S0168-9002(97)00301-X, URL: https://linkinghub.elsevier.com/retrieve/pii/S016890029700301X.

[32] C. Wiggins, R. Santos, A. Ruggles, A feature point identification method for positron emission particle tracking with multiple tracers, Nucl. Instrum. Methods Phys. Res. A 843 (2017) 22–28, http://dx.doi.org/10.1016/j.nima.2016.10.057, URL: https://www.sciencedirect.com/science/article/pii/S0168900216311184.

[33] J. Salvadori, J. Labour, F. Odille, P.-Y. Marie, J.-N. Badel, L. Imbert, D. Sarrut, Monte Carlo simulation of digital photon counting PET, EJNMMI Phys. 7 (1) (2020) 23, http://dx.doi.org/10.1186/s40658-020-00288-w.

.5 Monte Carlo model of the Large Modular Array for positron emission particle tracking

IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

Received 8 December 2022, accepted 4 March 2023, date of publication 10 March 2023, date of current version 17 March 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3255505

APPLIED RESEARCH

Monte Carlo Model of the Large Modular Array for Positron Emission Particle Tracking

MATTHEW HERALD[1], DAWID HAMPEL[2], TZANY KOKALOVA WHELDON[2,3], JONATHAN SEVILLE[1,3], AND CHRISTOPHER WINDOWS-YULE[1,3]

[1]School of Chemical Engineering, University of Birmingham, B15 2TT Birmingham, U.K.
[2]School of Physics and Astronomy, University of Birmingham, B15 2TT Birmingham, U.K.
[3]Positron Imaging Centre, University of Birmingham, B15 2TT Birmingham, U.K.

Corresponding author: Matthew Herald (mxh1092@bham.ac.uk)

This work was supported in part by the Engineering and Physical Science Research Council under Grant EP/T034327/1, and in part by the Advanced Imaging and Numerical Modelling of Segregation and Transport of Plastics in Fluidised Beds: Toward a Circular Economy for Plastics.

ABSTRACT Positron emission particle tracking (PEPT) is a non-invasive technique used to study fluid, granular, and multi-phase systems of interest to academia and industry. PEPT employs position-sensitive radiation detectors to record gamma rays in coincidence and track the movement of discrete sources. A modular detector array, the Large Modular Array (LaMA), has been constructed at the University of Birmingham's Positron Imaging Centre (PIC) to enable custom detector geometries. To estimate the LaMA's performance characteristics prior to experimentation, assist in developing optimised camera geometries, and determine ideal PEPT tracer characteristics a Monte Carlo model of LaMA is created and subsequently validated with experimental measurements. Validation is achieved through comparisons of the spatial resolution and count-rate response following the National Electrical Manufacturers Association (NEMA) industry standard protocol. Notably, the model's pulse-processing chain is autonomously calibrated to match experimental measurements using a recently developed technique which applies an evolutionary algorithm. The results show the simulated spatial resolution of the validated model matches the experiment to within 5%. Additionally, the total, true, and corrupted count-rates are reproduced to a mean error of 3.41%. This calibrated detector model strengthens the PIC's modelling capabilities. To facilitate future research, this model has been made publicly available through the PIC's GitHub repository.

INDEX TERMS Digital twin, GATE, Monte Carlo, positron emission particle tracking.

I. INTRODUCTION

Positron emission particle tracking (PEPT) is an imaging technique used to study opaque engineering and scientific systems using flow-following tracers labelled with positron-emitting radionuclides [1]. Positrons annihilate with electrons in close proximity to the labelled tracer, producing back-to-back 511 keV gamma-rays which can be detected with position-sensitive radiation detectors such as gamma cameras or positron emission tomography (PET) scanners [2]. Reconstruction of the labelled tracer position requires detecting both annihilation photons in coincidence to form a line-of-response (LoR) and applying a PEPT algorithm on a sample of LoRs to find the most likely annihilation point [3]. By detecting the tracer successively, a time-dependent trajectory is developed which can be analysed to determine system properties such fully three-dimensional velocity fields, tracer re-circulation times, and diffusivity [1]. PEPT has been used extensively over the last 30 years to study a variety of equipment ranging from coffee roasters, washing machines, and liquid metal castings [4], [5], [6].

To perform a PEPT measurement, equipment must first be moved to a lab and placed in the field-of-view (FOV) of a position-sensitive detector. At the University of Birmingham's Positron Imaging Centre (PIC), detectors like the ADAC Forte and SuperPEPT have been acquired or built specifically for PEPT measurements [7]. The Forte dual-headed positron camera was acquired because the two

The associate editor coordinating the review of this manuscript and approving it for publication was Giovanni Merlino.

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

detector heads can be separated up to between 250 mm and 800 mm, which enables the accommodation of a variety of equipment while optimising detector sensitivity. This system can record LoRs up to approximately 100 kHz [8]. Additionally, SuperPEPT, which has recently been constructed using components from CTI/Siemens ECAT EXACT 31, ART, and EXACT HR+ scanners has a cylindrical geometry of about 400 mm in diameter and 544 mm in length. The diameter can be separated up to 600 mm to place an experiment in the FOV, but must be closed again to record data. SuperPEPT records LoRs up to 2500 kHz, a factor of 20 higher than the Forte, enabling improved spatiotemporal resolution of PEPT trajectories [9].

However, there are instances where experimental equipment is too large or consists of awkward geometries which cannot be easily fit into the FOV of existing systems. To address this, a modular detector array, named the Large Modular Array (LaMA), has been designed to be assembled around experimental systems in custom geometries [10]. LaMA consists of building block, called 'boxes', which contain four ECAT951 block detectors. These boxes can be placed in nearly any configuration and connected to a coincidence processor, allowing flexible geometries to be designed for imaging large-scale industrial equipment [11]. Since the LaMA is reconfigured in a different geometry for each experiment, the performance characteristics of the camera are difficult to predict. As such, estimating the spatiotemporal resolution of the expected trajectories, designing an optimised geometry, and selecting an appropriate tracer activity for a given experiment can be challenging.

In this work we describe and validate a Monte Carlo model of the LaMA which will be used to help optimise experiments in the future. This model is created using the Geant4 Application for Tomographic Emission (GATE), a Monte Carlo radiation transport software which emulates the detector geometry, radioactive sources, and electronic pulse-processing of particles interacting with the detector [12], [13]. The performance of the camera in a simple geometry is characterised following the industry-standard National Electrical Manufacturers Association (NEMA) protocol [14]. Notably, the pulse-processing model is calibrated by through a recently developed method which compares the simulated and real performance characteristics and tunes parameters of the pulse-processing model using an evolutionary algorithm [15], [16], [17].

A. THE LARGE MODULAR ARRAY (LaMA)

The idea for the LaMA was developed at the PIC in the early 2000s, growing out of the need to have a detector system which could image large or awkward industrial systems and potentially be able to be transported to equipment which could not be moved. One feature that was identified as being key to this future system was being modular so that custom geometries could be tailored to each experiment. To this end, initially, three CTI/Siemens ECAT ring scanners were acquired and dismantled to retrieve the block detectors which could then be reassembled into a new geometries [10].

Each block detector consists of an 8×8 bismuth germanate (BGO) crystal array, with each crystal measuring 6.25 mm in width, 6.75 mm in height, and 30 mm in thickness. The BGO crystals are optically coupled to a 5 mm thick glass light guide and four photo-multiplier tubes (PMTs), each 100 mm in length. Using the 192 block detectors extracted from the PET scanners, a modular unit consisting of four blocks was designed. Each of these units are termed a 'box' and 48 boxes were constructed in total. Each box is approximately 360 mm in width, 95 mm in height (including spacers), and 460 mm in thickness with each of the four block detectors spaced 90 mm apart from centre-to-centre. A single box is shown on the left side of Fig. 1 and a stack of four boxes is shown on the right. These modular units are the fundamental building blocks of LaMA. Currently, up to 32 boxes can be connected to a single coincidence processor unit to form a detector array. By using more than one coincidence processor and merging the data streams in post-processing, all 48 boxes can be used simultaneously, but coincidences can only be formed between boxes connected to a shared coincidence processor.

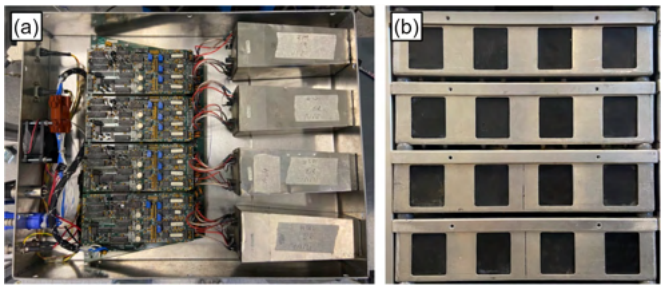

FIGURE 1. (a) A view inside a single box where four block detectors are mounted. (b) A stack of four boxes. The geometry used in this work is two stacks of four boxes separated by 500 mm.

B. Geant4 APPLICATION FOR TOMOGRAPHIC EMISSION (GATE)

GATE is a Monte Carlo radiation transport and detector simulation software designed for emulating the geometries, sources, and pulse-processing chain of imaging systems [12]. Many different types of detectors and imaging modalities have been modelled using GATE, including other detector systems used for PEPT such as the ADAC Forte and Siemens Inveon [13], [18], [19]. GATE is a useful tool for modelling PEPT experiments because it provides output comparable to real detectors and serves as 'sandbox' through which changes in the detector or source properties can be assessed, synthetic data can be processed to verify imaging techniques, or limits of techniques can be investigated without expending the considerable time and resources required for physical experimentation. Of special importance for the model presented in this work, a GATE model of a PEPT detector and experimental geometry can be used to the assess impact of changes in the geometry to the sensitivity of the system and affect this has on

tracer trajectories quantified through estimates of their spatial and temporal resolution [20], [21], [22].

One useful feature of GATE is the 'parameterisation' of the simulation scripts which allows users to quickly change aspects of the simulation through the command line. Using this, the source activity, placement of LaMA's boxes, and the values for parameters of the digitizer can be edited without having to manually change the file. In this work, we use GATE to first replicate the geometry of the LaMA and the source, then calibrate the digitizer through evolutionary an evolutionary algorithm which is able to edit the pulse-processing settings of the detector through parameterised simulation scripts. The fitness of a set of candidate solutions with tune-able digitizer parameters is then compared through the ability of the GATE model to replicate the experimentally measured performance characteristics.

C. EVOLUTIONARY ALGORITHMS

When calibrating simulated models to experimental measurements, the complex relationship between variables, noisy measurements, and the large number of solutions that need to be tried to explore the solution space leads to a difficult optimisation problem with many false local minima [23]. When there are several parameters which need to be optimised, such as encountered in a digitizer model, traditional approaches like design-of-experiments become too unwieldy and gradient-based optimisers struggle to overcome local minima in a multi-dimensional and noisy solution space. This has led most GATE models to be calibrated manually using estimates of the optimal parameters provided by manufacturers of the detectors [24], [25]. This type of manual tuning is both subjective and also not guaranteed to produce the optimal calibration.

However, a type of optimisation algorithm that has been shown to excels in these cases are evolutionary algorithms [15], [26]. Evolutionary algorithms emulate biological evolution by using a group of simulations to act as a population with varied features. In this way, the fitness of individual simulations against a selective pressure can be quantified. To improve the fitness of the next generation of simulations, the fittest simulations are allowed to reproduce which allows their features to be transferred and some random mutations added to increase diversity, potentially introducing beneficial features.

In this work, the Covariance Matrix Adaptation using Evolutionary Strategy (CMA-ES), which is a stochastic optimiser for robust non-linear non-convex numerical optimisation, is used to perform the model calibration [16]. While CMA-ES generates, assesses, and updates solutions to parameters of the digitizer, an additional software is used to couple CMA-ES to the GATE simulation. This software is the Autonomous Characterisation and Calibration via Evolutionary Simulation (ACCES) which is a Python interface to the CMA-ES algorithm specifically designed for general calibration simulations and has been previously used to calibrate a GATE model of the ADAC Forte [15], [17]. More details about our use of this software to calibrate the LaMA digitizer are provided in Section II-C.

II. METHODS

A. CHARACTERISATION EXPERIMENTS

Two sets of characterisation experiments are conducted to measure the spatial resolution and count-rate response of the LaMA in a simple geometry. These two characteristics are the most important detector characteristics in regard to PEPT experiments because spatial resolution predominately influences the ability to resolve a point-like source and the digitizer model controls the count-rate response curve. Thus, these two characteristics ultimately determine the spatial and temporal resolution of a PEPT tracer, which are the characteristics of interest for users of PEPT algorithms [7].

In order to characterise LaMA's performance characteristics, a single, simple geometry was chosen. This geometry is a dual-headed stack of four boxes on either side of a source. The two stacks are separated by 500 mm. This configuration is used because of the low number of boxes required and also because of the large solid angle it creates with the source, allowing three-dimensional tracking of a point source, which is important for PEPT algorithms. Additionally, by choosing a simple geometry and achieving a calibration with the GATE model, it is expected that when the system is scaled up for future experiments the GATE model will remain accurate. The LoRs collected from these experiments are processed according to the NEMA protocol and in all cases the centre of the FOV is defined as 0, 0, 0 mm.

1) SPATIAL RESOLUTION

The spatial resolution of the detector is defined as the full-width half-maximum (FWHM) of the point-spread-function (PSF) for a small point-like source in the detector's FOV. The source used for this experiment is a 1 mm sphere of anionic exchange resin, volumetrically activated with fluorine-18 in a solution of water produced by the University of Birmingham MC40 cyclotron [27]. For imaging, the source was placed in a small plastic sample holder and fixed to a block of polystyrene foam at six locations ranging from the centre of the FOV and locations at 1/4th of the FOV. These locations and the source activities at the time of the experiment are listed in Table 1. The source and the LaMA geometry are shown in Fig. 2.

TABLE 1. Spatial resolution test parameters.

X (mm)	Y (mm)	Z (mm)	Initial Activity (MBq)	End Activity (MBq)
0	0	0	2.26	2.07
0	77.5	0	2.06	1.91
86.25	0	0	1.88	1.69
0	0	125	1.39	1.20
0	77.5	125	1.19	1.05
86.25	0	125	1.00	0.86

To produce projection images, the LoRs collected from the experiment were first voxelised into a three-dimensional array with a 1 mm voxel size using the voxelisation method

FIGURE 2. The spatial resolution tests are conducted using a 1 mm diameter resin bead placed inside a plastic sample holder and taped to a piece of polystyrene foam.

implemented in the `pept` Python package [28]. Two-dimensional slices from this array were extracted which contain the voxel with the maximum number of LoR crossings. From these slices. A one-dimensional profile was drawn through the maximum voxel and the FWHM was extracted from each position. Following the NEMA protocol, the spatial resolution is calculated using (1-4).

$$Res_{Trans_{Center}} = (Res_{y_{x=0,y=0,z=0}} + Res_{x_{x=0,y=0,z=0}} + Res_{y_{x=1/4,y=0,z=0}} + Res_{x_{x=1/4,y=0,z=0}})/4 \quad (1)$$

$$Res_{Trans_{1/4}} = (Res_{y_{x=0,y=1/4,z=1/4}} + Res_{x_{x=0,y=1/4,z=1/4}} + Res_{y_{x=0,y=1/4,z=0}} + Res_{x_{x=0,y=1/4,z=0}})/4 \quad (2)$$

$$Res_{Axial_{Center}} = (Res_{z_{x=0,y=0,z=0}} + Res_{z_{x=1/4,y=0,z=0}})/2 \quad (3)$$

$$Res_{Axial_{1/4}} = (Res_{z_{x=0,y=0,z=1/4}} + Res_{z_{x=1/4,y=0,z=1/4}})/2 \quad (4)$$

2) COUNT-RATE RESPONSE

The count-rate experiment measures the LoR count-rate of the detector in response to a central source which is imaged over several half-lives. The total, true, and scattered plus random LoRs count-rates are extracted using the NEMA protocol and recorded at regular intervals as a function of source activity. The scattered plus random count-rate is termed the corrupted count-rate since the LoRs do not pass through the positron annihilation point due to scattering or originating from separate annihilation events and thus are treated as noise in a PEPT experiment.

The phantom is a hollow, high-density polyethylene cylinder measuring 120 mm in height and 50 mm in diameter with an inner cavity measuring 100 mm in height by 12 mm in diameter. This cavity is filled with a well-mixed solution of fluorine-18 and water. Initially, the activity of the phantom was approximately 80 MBq. This activity was chosen such that the expected count-rate will exceed the maximum rate at which LoRs can be recorded by the detector, then, as the source decays, the count-rate response curves can be developed. The phantom was imaged over several half-lives until the activity reached near that of the background. The phantom and detector geometry are shown in Fig. 3.

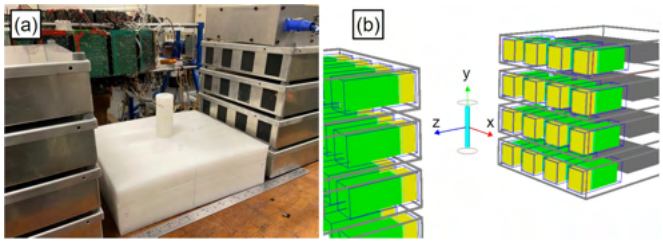

FIGURE 3. (a) The count-rate experiment is conducted with the high-density polyethylene phantom placed in the centre of FOV and imaged over several half-lives. (b) A GATE model of the same experiment is conducted.

Similarly to Section II-A1, projection images are produced by voxelising LoRs into a three-dimensional array with a 1 mm voxel size. Two-dimensional slices are extracted in the plane parallel to the detector face (XY plane) which contains the maximum voxel. The slice is then transformed into a one-dimensional profile by summing the voxels in the Y-axis which are along the cylinder's axis. The average source activity, $\bar{A}$, for each projection image is determined by (5) which calculates the average activity by using the initial activity, A_0, the initial and final time of the acquisition, t_0 and t_f, as well as the decay constant for fluorine-18, λ. The true counts are considered to be the LoRs $\pm$ 20 mm from the peak of the profile from which the background on either side of the 40 mm window is averaged and subtracted from the counts in the window. The remaining counts outside the window and including the background are considered corrupted counts.

$$\bar{A} = \frac{A_0}{t_f - t_0} \int_{t_0}^{t_f} \exp(-\lambda t)\, dt \quad (5)$$

B. GATE MODEL

The design of a GATE model for the LaMA presents several challenges since it must be easily customised to rapidly prototype new geometries, only allow specific coincidences to be formed between boxes connected to a single coincidence processor, and be able to emulate the noise, data buffer, and spatial blurring inherent in the system. Achieving these goals is accomplished using the tools available in the GATE software in addition to custom data post-processing. The geometry, digitizer, and post-processing are described in the following sections. A downloadable version of the LaMA GATE model and post-processing software is provided through a GitHub repository: LaMA Model.

1) GEOMETRY MODEL

Since any LaMA geometry is built using boxes, to build a model of the LaMA, only a single box needs to be described which can later be copied, translated, and rotated to any position and orientation using GATE's generic repeater function. The dimensions for each box are found in Section I-A and these are replicated in GATE. Material definitions for BGO, aluminium frame, and glass light guide are already included in the GATE materials database and definitions for the PMTs and electronics are added. Importantly, the four ECAT951 block detectors are included in each box and the 8 × 8 BGO crystal array in each block is defined as the 'Sensitive Detector' (SD) through which GATE records the interactions of particles. The model of a single box is shown in Fig. 4 from various viewpoints. The order in which the repeated boxes are listed determines their volume number, which will become important later when defining which pairs of boxes are valid for recording coincidences.

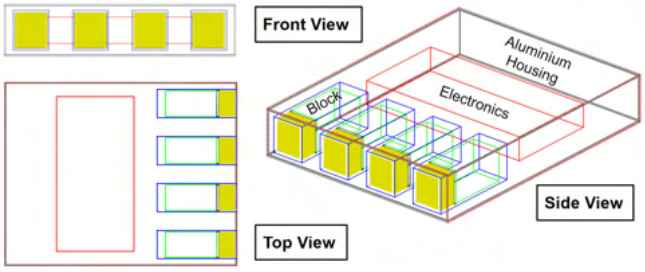

FIGURE 4. A model of a box for the LaMA consisting of four ECAT951 block detectors each with an 8 × 8 array of BGO crystals. The box is shown from various viewpoints and has major components labelled.

2) DIGITIZER MODEL

While a geometric model ensures the interactions of the source's radiation field with the detector are accurately recorded, the system's pulse-processing chain must also be properly modelled in order to emulate the detector's response. In GATE, implementing a linear pulse-processing chain to particle interactions is the role of the digitizer. The digitizer is a series of steps and filters which transform the observable information (time, energy, position) of a particle interaction with the SD into a form similar or identical to the real detector output [29]. This includes grouping interactions, flow-logic, and data-loss to mimic the real behaviour of the imaging system [30]. The digitizer for the LaMA is shown in Fig. 5.

For the LaMA, we apply a typical digitizer model for a PET system, but also implement a post-processing stage to implement aspects of the detector not directly possible through GATE [19]. This is needed to force the detector to only record coincidences between pairs of boxes which are associated with each other in the real coincidence processor, implement a random spatial-blurring of LoRs to match the experimentally observed spatial resolution, and implement a bandwidth limitation to cap the rate at which LoRs are written to file. The post-processing steps for the LaMA GATE model digitizer are the final three steps before the final simulated detector response is produced, as shown in Fig. 5.

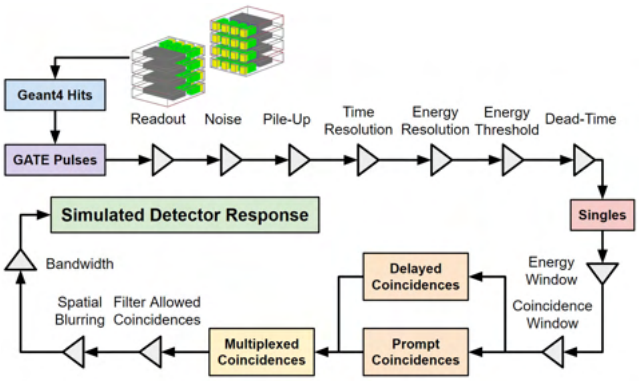

FIGURE 5. The pulse-processing digitizer of the LaMA GATE model.

C. DIGITIZER OPTIMISATION

Even if every detail of the LaMA were known, there are differences between simulation and experiment which require aspects of the digitizer to be calibrated [10], [31]. While this could be achieved through manual calibration, for this we work we use a recently developed method which applies an evolutionary algorithm to achieve a calibration to experimental performance characteristics autonomously [15].

Six stages of the digitizer are chosen to be calibrated because of the effect they have on replicating the count-rate response. These are the noise frequency, pile-up time, time resolution, lower-level energy discriminator, upper-level energy discriminator, and the non-paralysable singles dead-time. The noise frequency is the rate at which random events are generated simply by having the detectors running and is a combination of the background activity and electronic noise. The pile-up time is the time after an event is detected when other events in close succession can be added to the signal. The time resolution is the probability of precision by which two events can be distinguished in time, as defined by an FWHM. The lower and upper energy discriminators are the thresholds between which an event can trigger the dead-time and be recorded. The singles dead-time is a non-paralysable dead-time model, described by (6) which limits the rate of recording single events [32].

$$\frac{\lambda_{out}}{\lambda_{in}} = \frac{1}{(1 + \lambda_{in}\tau)} \tag{6}$$

In order to use an evolutionary algorithm to calibrate a GATE model, there must a metric through which the fitness of candidate solutions to the parameters of the digitizer can be assessed. This is achieved through a cost function, shown in (7), which is the product of the percent differences between the experiment and simulation's total, true, and corrupted count-rate response over a range of source activities, calculated using (8-10). Using this metric reduces the agreement between the experiment and simulation down to a single value

which can be optimised through minimisation of (7).

$$\varepsilon = \varepsilon_{Tot}\varepsilon_{True}\varepsilon_{Corrupt} \quad (7)$$

$$\varepsilon_{Tot} = \sum |(R_{Tot_{exp}} - R_{Tot_{sim}})/R_{Tot_{exp}}| \quad (8)$$

$$\varepsilon_{True} = \sum |(R_{True_{exp}} - R_{True_{sim}})/R_{True_{exp}}| \quad (9)$$

$$\varepsilon_{Corrupt} = \sum |(R_{Corrupt_{exp}} - R_{Corrupt_{sim}})/R_{Corrupt_{exp}}| \quad (10)$$

Parameter combinations are generated following a multivariate normal distribution with the initial uncertainty defined as the range between the upper and lower bounds of the solution space for each parameter. To stop the optimisation, the user must provide either the number of generations, called epochs, of simulations desired or provide a target uncertainty in the calibrated parameters. In this optimisation, we used 100 epochs with 100 parameter value combinations per epoch as the terminating criteria because this provides greater than 10 times the factorial of the number of free-parameters in the optimisation, sufficiently constraining the problem.

Each combination is simulated over 10 different activities, ranging from 2 MBq to 80 MBq, until 5 million LoRs are recorded at each activity. This number of LoRs provides sufficient counts such that an accurate measurement of each respective count-rate can be extracted from the projection images. The bounds of the parameters and their initial guesses are provided in Table 2. The optimisation is conducted on the University of Birmingham's high-performance computing system, BlueBEAR, on Icelake cores with 16 GB of memory each. After the optimisation is finished, the calibrated parameters are extracted and a new set of simulations are conducted with 20 activities over the same activity range until 30 million LoRs are recorded, reducing statistical error even further. These simulations are presented and compared in Section III.

TABLE 2. Digitizer parameter bounds and initial guesses for calibration.

Parameter	Lower Bound	Upper Bound	Initial Guess
Singles Dead-Time (ns)	0	10000	5000
Noise Frequency (ns)	1000	10000	5500
Pileup Time (ns)	0	1000	500
Lower Level Disc. (keV)	0	400	200
Upper Level Disc. (keV)	700	2200	1450
Time Resolution (ns)	0	30	15

III. RESULTS AND DISCUSSION

In this section, we present the results of the characterisation experiments as well as the results of the digitizer calibration. For spatial resolution, six tracer positions were imaged over several minutes then the FWHM of the 1-dimensional PSF is extracted. The FWHM at these positions are used to compute the transaxial and axial spatial resolutions in the centre of the FOV and at 1/4th of the FOV. Next, the experiments are reproduced in simulation and the crystal blurring is adjusted until the best match between the experiment and simulation was achieved. The crystal blurring that best agrees with the experiment 2 times the crystal dimensions

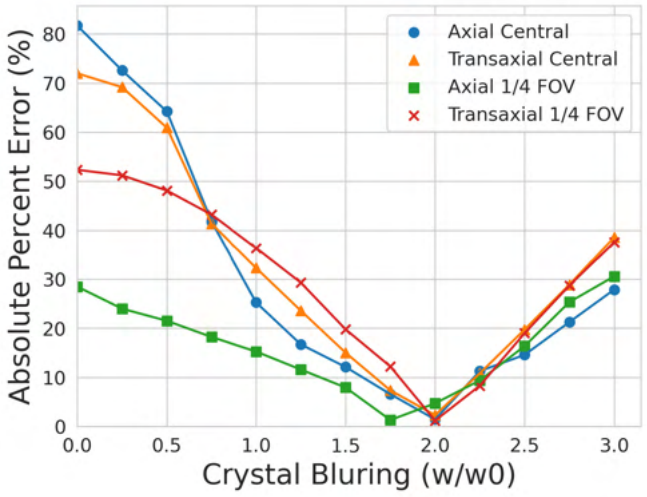

FIGURE 6. The spatial blurring is calibrated by finding the crystal blurring that minimises the absolute percent error.

TABLE 3. Results and comparisons of the spatial resolution tests for the experiment and simulation.

Spatial Resolution	Experiment (mm)	Simulation (mm)	Percent Error (%)
Central Transaxial	8.26	8.38	1.44
1/4th FOV Transaxial	8.95	9.37	4.72
Central Axial	21.83	22.34	2.32
1/4th FOV Axial	24.09	23.80	-1.20

(6.25 mm by 6.75 mm), as evidenced in Fig. 6. The experimental and simulated results are presented in Table 3 and compared through their respective percent differences.

Following the spatial resolution characterisation and crystal blurring calibration, the count-rate response experiments were analysed to be used as a comparison for the ACCES optimisation. The optimisation takes place over 100 epochs with 100 parameter value combinations tried every epoch. This results in 10,000 cost function evaluations which took approximately three days to complete running the BlueBEAR high-performance computing system. At the end of the optimisation, the final mean parameter values were extracted. The value for these parameters and their uncertainty are provided in Table 4. Additionally, the history of these parameters during the optimisation (uncertainties and mean values) is presented in Fig. 7, demonstrating that before the end of the optimisation, each parameter reaches a stable value, meaning that the parameters have been calibrated.

All parameters produced reasonable calibrations within the upper and lower bounds given to the optimiser. Interestingly, the optimised value for the time resolution falls within the 12 ns $\pm$ 2 ns measured in a previous characterisation of the LaMA [31]. This provides additional evidence that the calibrated parameters correspond to physical reality and are global solutions, not simply local solutions.

The parameter with the highest uncertainty is the upper level energy discriminator. We believe this is due to the relatively small impact of this parameter on the calibration. To illustrate this, take for example the the singles dead-time which has the lowest uncertainty. The singles dead time has

a strong pressure to be calibrated because this is applied to nearly all events which are detected. meaning small changes in the calibrated values will causes large differences in the simulated count-rate response. On the other hand, the upper energy discriminator acts on a much smaller set of events and can only be applied to events that have piled up on one another. Since the upper energy discriminator is set to 1990 keV, this means that at least four 511 keV events must be grouped together and this happens only a limited number of times in a simulation. As a result, there is not a strong pressure to calibrate this value. While this results in a higher uncertainty, because of the lower effect of this parameter on the overall response of the model this value is considered adequately calibrated.

TABLE 4. Digitizer calibration results and uncertainty.

Parameter	Calibrated Values	Uncertainty
Singles Dead-Time (ns)	6630	± 57.9
Noise Frequency (ns)	1970	± 9.39
Pileup Time (ns)	637	± 6.91
Lower Level Disc. (keV)	324	± 7.99
Upper Level Disc. (keV)	1990	± 85.5
Time Resolution (ns)	13.7	± 0.084

When the new set of simulations is conducted with the calibrated digitizer, the results match the experiment to a mean absolute difference of 3.41% over all three count-rates with the total, true, and corrupted count-rates being 2.31%, 2.18%, and 5.72%, respectively. The experimental and simulated count-rate response is shown in Fig. 8. To quantify the calibration further, it is also important to observe how the fraction of true and corrupted counts behave as a function of source activity. These results are presented in Fig. 9, showing that their behaviour is approximately the same overall activities with the true and corrupted count fractions reconstructed to 1.91% and 3.72% error, respectively.

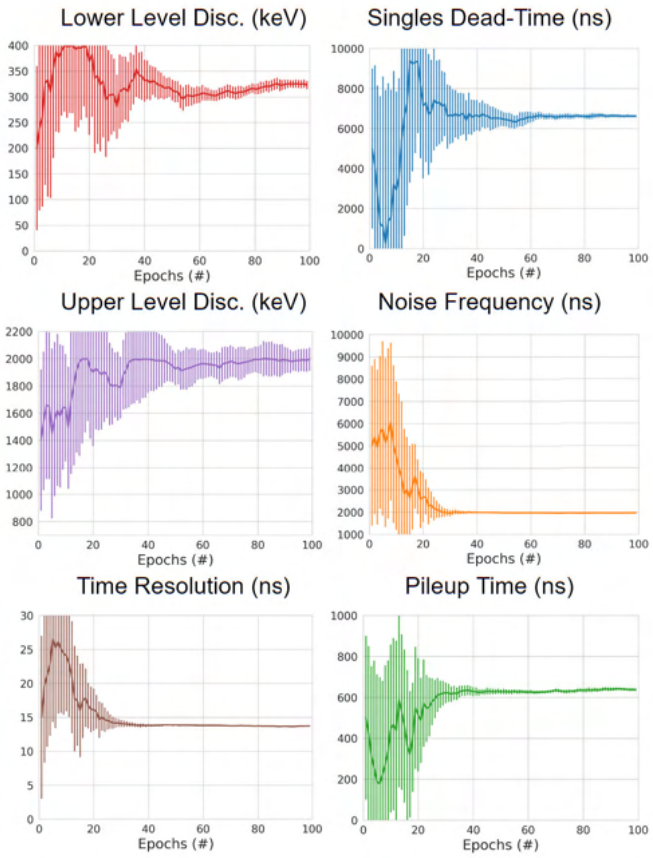

FIGURE 7. Subplots showing the mean values of calibrated parameters with the error bars as the standard deviation of solutions.

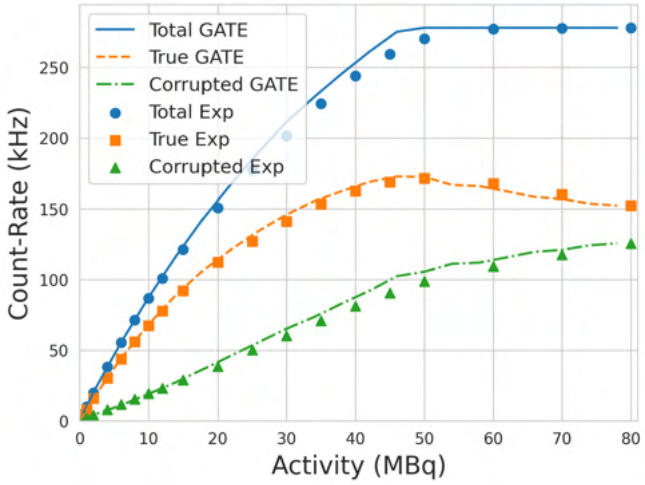

FIGURE 8. Results of the count-rate experiment and comparison of the ACCES calibrated GATE model.

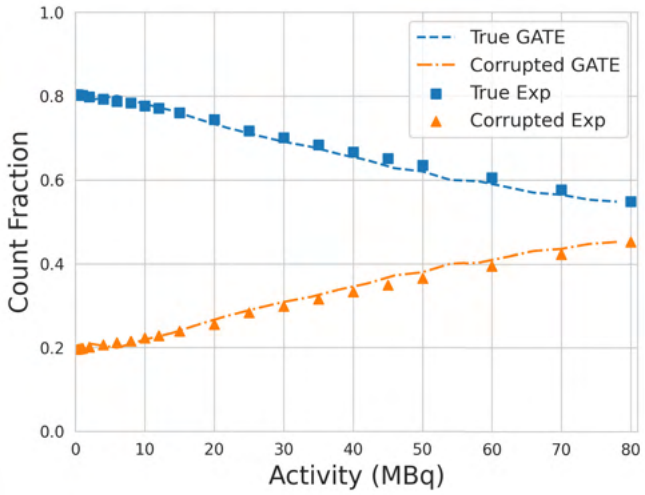

FIGURE 9. Results of the count-rate experiment in terms of the true and corrupt count fractions and comparison with the ACCES calibrated GATE model.

In summary, the GATE model of the LaMA has been characterised by experiments testing the spatial resolution and count-rate response and the digitizer pulse-pulse processing model has been calibrated using an evolutionary algorithm. In addition to calibrating the amount of crystal spatial blurring, six free-parameters of the pulse-processing digitizer model were calibrated using evolutionary simulation,

ultimately producing simulated results which match the observed real behaviour of the detector system. The ACCES software was used to perform this optimisation autonomously, needing only the number of solutions to try for each epoch, the bounds of the solution space for each parameter, and a terminating criterion.

IV. CONCLUSION

In this work, we have introduced a Monte Carlo model of the Positron Imaging Centre's LaMA and validated the model against experimental measurements. This model will bolster the modelling capabilities of the Positron Imaging Centre, complementing the existing Monte Carlo model of the ADAC Forte. A model of LaMA is particularly useful because the camera is typically configured into new geometries for every experiment in order to capture the relevant system behaviour. Prior to this model, it was not possible to quantitatively estimate the spatial resolution and count-rate response because these characteristics depend on a complex relationship between source activity, detector geometry, and gamma-ray scattering. Using this model, not only can the spatial resolution and count rate response be estimated, but moreover the tracer activity which maximises the true count-rate and the detector geometry which maximises the camera sensitivity for a given experiment can be identified.

CONFLICT OF INTEREST

The authors declare that there is no conflict of interest.

ACKNOWLEDGMENT

The authors would like to thank Andrei Nicusan for developing and maintaining the `pept` Python Library and Access Optimization Tools. The computations described in this article were performed using the University of Birmingham's BlueBEAR service, which provides High-Performance Computing to the University's research community. See their website for more details.

REFERENCES

[1] K. Windows-Yule, L. Nicusan, M. T. Herald, S. Manger, and D. Parker, *Positron Emission Particle Tracking*. Bristol, U.K.: IOP Publishing, 2022.

[2] C. R. K. Windows-Yule, J. P. K. Seville, A. Ingram, and D. J. Parker, "Positron emission particle tracking of granular flows," *Annu. Rev. Chem. Biomol. Eng.*, vol. 11, no. 1, pp. 367–396, Jun. 2020.

[3] D. J. Parker, C. J. Broadbent, P. Fowles, M. R. Hawkesworth, and P. McNeil, "Positron emission particle tracking—A technique for studying flow within engineering equipment," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 326, no. 3, pp. 592–607, Mar. 1993.

[4] M. Al-Shemmeri, K. Windows-Yule, E. Lopez-Quiroga, and P. J. Fryer, "Coffee bean particle motion in a spouted bed measured using positron emission particle tracking (PEPT)," *J. Food Eng.*, vol. 311, Dec. 2021, Art. no. 110709.

[5] C. R. Jones, A. Corona, C. Amador, and P. J. Fryer, "Dynamics of fabric and dryer sheet motion in domestic clothes dryers," *Drying Technol.*, vol. 40, no. 10, pp. 2087–2104, 2021.

[6] A. Dybalska, A. J. Caden, D. J. Parker, J. Wedderburn, and W. D. Griffiths, "Liquid metal flow studied by positron emission tracking," *Metall. Mater. Trans. B*, vol. 51, no. 5, pp. 1912–1917, Oct. 2020.

[7] D. J. Parker, D. M. Hampel, and T. K. Wheldon, "Performance evaluation of the current Birmingham PEPT cameras," *Appl. Sci.*, vol. 12, no. 14, p. 6833, Jul. 2022.

[8] D. Parker, R. Forster, P. Fowles, and P. Takhar, "Positron emission particle tracking using the new Birmingham positron camera," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 477, nos. 1–3, pp. 540–545, Jan. 2002.

[9] D. M. Hampel, S. Manger, D. J. Parker, and T. K. Wheldon, "SuperPEPT: A new tool for positron emission particle tracking; first results," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 1028, Apr. 2022, Art. no. 166254.

[10] D. J. Parker, T. W. Leadbeater, X. Fan, M. N. Hausard, A. Ingram, and Z. Yang, "Positron emission particle tracking using a modular positron camera," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 604, nos. 1–2, pp. 339–342, Jun. 2009.

[11] W. D. Griffiths, D. J. Parker, X. Fan, and M. Hausard, "Tracking inclusions in aluminium alloy castings using positron emission particle tracking (PEPT)," *Mater. Sci. Technol.*, vol. 26, no. 5, pp. 528–533, May 2010.

[12] S. Jan et al., "GATE: A simulation toolkit for PET and SPECT," *Phys. Med. Biol.*, vol. 49, no. 19, pp. 4543–4561, Oct. 2004.

[13] D. Sarrut et al., "Advanced Monte Carlo simulations of emission tomography imaging systems with gate," *Phys. Med. Biol.*, vol. 66, no. 10, May 2021, Art. no. 10TR03.

[14] *Performance Measurements of Positron Emission Tomographs (PET) 2018*, National Electrical Manufacturers Association, Rosslyn, VA, USA, 2018.

[15] M. Herald, A. Nicuşan, T. K. Wheldon, J. Seville, and C. Windows-Yule, "Autonomous digitizer calibration of a Monte Carlo detector model through evolutionary simulation," *Sci. Rep.*, vol. 12, no. 1, p. 19535, Nov. 2022.

[16] N. Hansen, K. Nozawa, L. Rolshoven, M. Chan, Y. Akimoto, and D. Brockhoff, "CMAES/pycma: R3.2.2," Zenodo, Tech. Rep., Mar. 2022.

[17] A.-L. Nicusan, D. Werner, J. A. Sykes, J. Seville, and K. Windows-Yule, "ACCES: Autonomous characterisation and calibration via evolutionary simulation," GitHub, San Francisco, CA, USA, Tech. Rep., Feb. 2022.

[18] M. Herald, T. Wheldon, and C. Windows-Yule, "Monte Carlo model validation of a detector system used for positron emission particle tracking," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 993, Mar. 2021, Art. no. 165073.

[19] S. Lee, J. Gregor, and D. Osborne, "Development and validation of a complete GATE model of the Siemens Inveon trimodal imaging platform," *Mol. Imag.*, vol. 12, no. 7, Oct. 2013, Art. no. 7290.2013.00058.

[20] M. Herald, Z. Bingham, R. Santos, and A. Ruggles, "Simulated time-dependent data to estimate uncertainty in fluid flow measurements," *Nucl. Eng. Des.*, vol. 337, pp. 221–227, Oct. 2018.

[21] C. R. K. Windows-Yule, M. T. Herald, A. L. Nicusan, C. S. Wiggins, G. Pratx, S. Manger, A. E. Odo, T. Leadbeater, J. Pellico, R. T. M. D. Rosales, A. Renaud, I. Govender, L. B. Carasik, A. E. Ruggles, T. Kokalova-Wheldon, J. P. K. Seville, and D. J. Parker, "Recent advances in positron emission particle tracking: A comparative review," *Rep. Prog. Phys.*, vol. 85, no. 1, Jan. 2022, Art. no. 016101.

[22] M. T. Herald, J. A. Sykes, D. Werner, J. P. K. Seville, and C. R. K. Windows-Yule, "DEM2GATE: Combining discrete element method simulation with virtual positron emission particle tracking experiments," *Powder Technol.*, vol. 401, Mar. 2022, Art. no. 117302.

[23] S. Shalev-Shwartz, O. Shamir, and S. Shammah, "Failures of gradient-based deep learning," in *Proc. 34th Int. Conf. Mach. Learn.*, vol. 70, D. Precup and Y. W. Teh, Eds., Aug. 2017, pp. 3067–3075.

[24] M. Strugari, D. DeBay, S. Beyea, and K. Brewer, "NEMA NU 1–2018 performance characterization and Monte Carlo model validation of the Cubresa Spark SiPM-based preclinical SPECT scanner," *EJNMMI Phys.*, Aug. 2022.

[25] A. Tiwari, M. Merrick, S. A. Graves, and J. Sunderland, "Monte Carlo evaluation of hypothetical long axial field-of-view PET scanner using GE discovery MI PET front-end architecture," *Med. Phys.*, vol. 49, no. 2, pp. 1139–1152, Feb. 2022.

[26] Y. Kwon, S. Kang, Y.-S. Choi, and I. Kim, "Evolutionary design of molecules based on deep learning and a genetic algorithm," *Sci. Rep.*, vol. 11, no. 1, p. 17304, Aug. 2021.

[27] D. J. Parker, "Positron emission particle tracking and its application to granular media," *Rev. Sci. Instrum.*, vol. 88, no. 5, May 2017, Art. no. 051803.

[28] A. L. Nicuşan and C. R. K. Windows-Yule, "Positron emission particle tracking using machine learning," *Rev. Sci. Instrum.*, vol. 91, no. 1, Jan. 2020, Art. no. 013329.
[29] S. Kerhoas-Cavata and D. Guez, "Modeling electronic processing in GATE," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 569, no. 2, pp. 330–334, Dec. 2006.
[30] D. Guez, F. Bataille, C. Comtat, P.-F. Honore, S. Jan, and S. Kerhoas, "Counting rates modeling for PET scanners with GATE," *IEEE Trans. Nucl. Sci.*, vol. 55, no. 1, pp. 516–523, Feb. 2008.
[31] T. W. Leadbeater, "The development of positron imaging systems for applications in industrial process tomography," Ph.D. thesis, Univ. Birmingham, Birmingham, U.K., Jul. 2009.
[32] S. Usman and A. Patil, "Radiation detector deadtime and pile up: A review of the status of science," *Nucl. Eng. Technol.*, vol. 50, no. 7, pp. 1006–1016, Oct. 2018.

MATTHEW HERALD was born in the USA. He received the B.S. degree in nuclear engineering from the University of Tennessee, Knoxville, in 2019. He is currently pursuing the Ph.D. degree in chemical engineering with the University of Birmingham, under the supervision of Christopher Windows-Yule.

He is a member of the Institue of Physics and currently works for Jacobs as a Physicist. His work revolves around developing Monte Carlo models of position-sensitive detectors and simulating positron emission particle tracking experiments to optimize data acquisition. He has developed a model of the ADAC Forte dual-headed positron camera and used this in an international collaboration of other institutes using positron emission particle tracking algorithms to benchmark their performance on a common, known data set. He is the coauthor on the first book written about positron emission particle tracking which was written to be used by students, users, and researchers.

DAWID HAMPEL was born in Poland, but spent most of his primary education in the USA. He received the M.Eng. degree (Hons.) in nuclear engineering from the University of Birmingham, Birmingham, U.K., in 2019, where he is currently pursuing the Ph.D. degree in cyclotron production and radiochemical purification methods of emerging radionuclides for cancer diagnosis and treatment while developing positron emission particle tracking detector arrays.

His early work includes the development of Geant4 radiation transport simulations and simulations of numerous nuclear reactions in TALYS and EMPIRE. He is the Co-Developer of UoB-TIP Toolkit used to ease the analysis of nuclear reaction data and OLab an online, interactive nuclear laboratory both used with the University of Birmingham. His current work involves the development and improvement of PEPT detector arrays in both hardware and software levels. He is working with the Positron Imaging Centre, he has built and tested SuperPEPT and redeveloped and modernized the Large Modular Array (LaMA) for PEPT with contributions to other systems used at the site.

Mr. Hampel is a member of the Institute of Physics.

TZANY KOKALOVA WHELDON received the Ph.D. degree in nuclear physics from Freie Universitaet, Berlin, Germany.

She is currently a full-time Professor of nuclear physics with the University of Birmingham and the Director of the Positron Imaging Centre. Her work spans the full breadth of nuclear physics from experimental nuclear astrophysics and machine learning, medical isotopes, applications, and industry-related nuclear decommissioning. She carries out research, teaches, and supervises eight Ph.D. students, and key roles mentoring women. She has travelled extensively, performing experiments with laboratories in Australia, the USA, and Europe. She was a recipient of the Daphne Jackson Fellowship and won a Timewise Power Returner Award celebrating, in 2020, her ability to combine career success at a senior level whilst working flexibly.

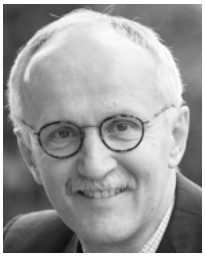

JONATHAN SEVILLE was born in Yorkshire, U.K. He received the degree in chemical engineering from the University of Cambridge and University of Surrey.

Before starting his academic career, he worked for several years for the research division of Courtaulds Ltd. Since 2017, he has been a Professor of formulation engineering with the University of Birmingham. Previously, he was the Executive Dean of Engineering and Physical Sciences with the University of Surrey (2011–2016), the Dean of Engineering with the University of Warwick (2008–2011), and the Head of the Department of Chemical Engineering, University of Birmingham (1998–2008). He is an experienced researcher in chemical engineering, particularly in processes involving particulate solids, and has pioneered the use of radioactive tracers to image flow in industrial equipment. He has around 200 publications and authored or edited five books on aspects of particle technology. Among the numerous companies which have supported his research are Unilever, Procter and Gamble, GSK, Merck Sharp and Dohme, and Pfizer. He is a fellow of the Royal Academy of Engineering and was President of the Institution of Chemical Engineers (IChemE), in 2016–2017.

CHRISTOPHER WINDOWS-YULE was born in Worcestershire. He received the B.S. and Ph.D. degrees in physics from the University of Birmingham.

Previously, he was a Postdoctoral Researcher with the University of Twente, The Netherlands. He is currently an Associate Professor with the University of Birmingham's School of Chemical Engineering. His research interests include imaging and numerical modeling of particulate and multiphase systems, employing a diverse range of techniques, notably discrete element method (DEM) modeling, computational fluid dynamics (CFD), and positron emission particle tracking (PEPT). His research aims to address significant contemporary challenges in science, medicine and industry by exploiting the synergy of experimental techniques, numerical simulation, and machine-learning methodologies. His current projects include work, funded by EPSRC, the Royal Academy of Engineering and the Royal Society, developing novel plastic recycling methods, work funded by the British Heart Foundation aiming to develop novel methods of blood-flow imaging for the diagnosis of cardiovascular disease, and diverse industry-funded projects. He is a Turing Fellow, a two-time Royal Academy of Engineering Industrial Fellow, and a Founding Member of the U.K. Young Academy.