Data-Driven methods for respiratory signal detection in Positron Emission Tomography

Ottavia Bertolli

A dissertation submitted in partial fulfilment
of the requirements for the degree of
Doctor of Philosophy
in Medical Physics

Institute of Nuclear Medicine
Division of Medicine
University College London

March 14, 2018

I, Ottavia Bertolli, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the work.

What would you do if you weren't afraid?

— *Sheryl Sandberg*

Abstract

Positron Emission Tomography (PET) is a nuclear medicine imaging technique which allows quantitative assessment of functional processes, by determining the distribution of radioactive tracers inside the patient body. It is mainly used in oncology. Respiration during PET data acquisition of the chest leads to blurring and other artefacts in the images, lowering their quantitative accuracy. If a respiratory signal is available, these issues can be overcome by splitting the data into different motion states. In current clinical practice this signal is obtained using external devices. However, these are expensive, require prior setup and can cause patient discomfort. This thesis develops and evaluates Data-Driven (DD) techniques based on Principal Component Analysis (PCA) to generate the signal directly from the PET data.

Firstly, the arbitrary relation between the sign of the PCA signal and the respiratory motion is addressed: a maximum in the signal could refer either to end-inspiration or end-expiration, possibly causing inaccurate motion correction. A new correction method is proposed and compared with two already existing methods. Subsequently, the methods are extended to Time-of-Flight (TOF) PET data, proposing a data processing step prior to using PCA, in order to benefit from the increased spatial information provided by TOF. The proposed methods are then extensively tested on lower lung patient data (non-TOF and TOF). The obtained respiratory signal is compared with that of an external device and with internal motion observed with Magnetic Resonance Imaging (MRI). Lastly, to investigate the performance of PCA where respiratory motion is minimal, the methods are applied to patient and simulation data of the upper lung, showing that they could potentially be utilised for detecting respiratory-induced density variations in the upper lung. This study shows that the presented methods could replace external devices for obtaining a respiratory signal, providing a simple and cost-effective tool for motion management in PET.

Impact Statement

The work presented in this thesis provides proof of the robustness of a DD method based on PCA as a tool for motion detection in PET lung imaging. Furthermore, it presents techniques for the improvement of the method. All the proposed methodologies are thoroughly evaluated and tested on patient data, yielding promising results.

The presented algorithms are easily implementable and they could be readily included within the clinical pipeline for respiratory motion detection in lung PET imaging. The availability of these tools would enable an easier detection of respiratory motion during the patient's acquisition, avoiding the cost of external devices and the additional time needed for their setup, therefore reduce the dose staff are subjected to. Moreover, it would eliminate the need for training staff in the use of the external devices.

Therefore, the use of the presented algorithms would result in a reduced cost, and in the increase of the number of patients acquired per day, because of the reduced time required for the whole acquisition.

Moreover, the use in clinical practice of the methods presented in this work would allow a wider application of respiratory motion detection, that in turn would result in a broader application of respiratory motion correction in lung imaging. This would also consequently allow to extend the knowledge of the impact that motion correction brings to PET images, both in clinical setting but also in research studies. Finally, these methods could also be extended to other types of motion and other modalities, such as Single Photon Emission Computed Tomography (SPECT) and also Cone Beam Computed Tomography (CBCT), which especially when utilised for radiotherapy planning would greatly benefit from motion compensation.

In conclusion, the methodologies presented in this thesis could be utilised by researchers for investigation purposes, and could be implemented into commercially available software and made available to clinical departments, resulting in a positive impact to public health in the field of medical imaging on a national and international level.

Acknowledgements

I believe it is never a one person's job to get to the end of a PhD, that is why I wouldn't be here if it weren't for all the great people I have in my life.

First of all I would like to thank my primary supervisor, Dr. Kris Thielemans, for his support and mentorship that allowed me to grow as a researcher, and for all that he has thought me. I would also like to thank my second supervisor, Prof. Simon Arridge, for his interest in my work, for his valuable input and for the interesting meetings we had, and Prof. Brian Hutton, for his guidance and advice.

Thank you to my colleagues from GE, Chuck Stearns and Scott Wollenweber, who have believed in my ideas and have helped and supported me from the beginning.

I would also like to thank all the people at INM, and specifically those who have expressed interest in my research and have helped me with my projects, sharing their time and knowledge: Alexandre Bousse, Elise Emond, Tom Sanderson, John Dickson, Simon Wan, Raymond Endozo and Bettina Ferre. I am very grateful to you all. And thank you to Brice Fernandez, from the Orsay group, without whom an interesting part of this work wouldn't have been possible.

A very big thank you to my girls, Debora, Alaleh, Tsai and Beverley, you know how much you did for me and I will be forever grateful. Debora, I literally wouldn't have made it here without you being my coach! And thank you Ludovica, you made me smile when I really needed to.

Thank you to my Italian friends, Caterina, Flaminia, Claudia, Alessia, Chiara, Cinzia, Antonietta, Emanuela and to the others who are far away but always close, Maria, Vesna, Afroditi.

Last but not least, the most important thank you goes to my wonderful family, whose love and support are the strength of my life: Mamma, Babbo, Martina, Erica, Edoardo, and my husband Giorgio, who has always pushed me further than I though I could ever go. To you I dedicate my thesis.

I acknowledge funding support from the Engineering and Physical Sciences Research Council Industrial CASE studentship (13220093), co-sponsored by GE Healthcare. Also, this project is supported by researchers at the National Institute for Health Research, University College London Hospitals Biomedical Research Centre.

Contents

List of Figures

List of Tables

List of Abbreviations

1D One-Dimensional

2D Two-Dimensional

3D Three-Dimensional

CBCT Cone Beam Computed Tomography

COD Centroid of Distribution

COM Centre of Mass

CT Computed Tomography

DD Data-Driven

FBP Filtered Back-Projection

FOV Field of View

FT Fourier Transform

FWHM Full Width at Half Maximum

LOR Line of Response

MATLAB Matrix Laboratory, The Mathworks Inc., Natick MA, USA

MIP Maximum Intensity Projection

ML Maximum Likelihood

MLA Maximum likely annihilation

MLEM Maximum Likelihood Expectation Maximization

MRI Magnetic Resonance Imaging

NAC Non Attenuation Corrected

OSEM Ordered Subsets Expectation Maximization

PCA Principal Component Analysis

PET Positron Emission Tomography

PSF Point Spread Function

RF radio-frequency

ROI Region of Interest

RPC Respiratory Principal Component

RPM Real-Time Position Management

SAM Spectral Analysis Method

SD Standard Deviation

SNR Signal to Noise Ratio

SPECT Single Photon Emission Computed Tomography

SPM Statistical Parametric Mapping

SSRB Single Slice Rebinning Algorithm

STIR Software for Tomographic Image Reconstruction

SUV Standard Uptake Value

TOF Time-of-Flight

VOI Volume of Interest

XCAT 4D extended cardiac-torso

Chapter 1

Introduction

In Positron Emission Tomography (PET), respiratory motion is one of the main causes of artefacts when imaging the thorax. The patient's motion during the acquisition leads to blurring in the reconstructed images, resulting in decreased lesion detectability and inaccurate Standard Uptake Value (SUV) calculation, and it can cause potential errors in the attenuation correction, therefore lowering the quantitative accuracy of this imaging modality. When PET images are utilised for treatment planning in radiation therapy, respiratory motion effects can hinder the accurate delineation of the target area, therefore causing healthy tissues to be needlessly irradiated.

In order to reduce the impact of respiratory motion effects, the acquired PET data can be grouped into approximately "motion free" bins, referred to as *gates*, through a binning process (*gating*) that depends on the respiratory cycle. Several studies have shown that this approach leads to increased SUV values, decreased lesion volumes and in an overall improvement in the quality of the images (Nehmeh et al., 2002a,b; Liu et al., 2009). Furthermore, several techniques have been proposed to perform motion correction, making use of the entirety of the PET data and a motion signal, in order to produce an image which is free of motion artefacts (Pepin et al., 2014).

The gating process, and also motion correction when applied, require a respiratory signal: in clinical practice this is obtained making use of external devices, such as a pressure belt to be placed around the patient's chest, an infra-red camera that tracks a reflective marker positioned on the abdomen, or devices that measure the air flow. However, the use of these devices incurs an added cost, additional time required for the acquisition setup (which needs to be handled by trained staff, increasing the time they are exposed to a radioactive patient) and could cause patient discomfort. For these reasons, in recent years there has been growing interest into solutions that would not require any external device and Data-Driven (DD) methods have been introduced. DD methods aim to produce a respiratory signal making

use of the PET raw data only, and could therefore constitute a useful and cost-effective tool for motion detection. Furthermore, as they originate directly from the data and not from external measurements, the DD signals could be more correlated to the internal motion than a signal based on external movement.

PCA is a "dimensionality reduction" method which is able to extract the respiratory motion by detecting the biggest changes that occur in the data, which in PET chest acquisition are generally caused by respiration. It has firstly been utilised for the purpose of respiratory motion detection in Thielemans et al. (2011), and it has provided promising results when comparing its signal to that of other DD methods, to an external device signal and also to the internal organ's motion.

The aim of this work is to thoroughly investigate the performance of the PCA method for respiratory motion detection in PET and to evaluate its performance on patient data. The method is also extended to the use of Time-of-Flight (TOF) PET data and further improvements are suggested to achieve better performance. Finally, the method is applied to PET data of the upper chest, and the ability of PCA to detect density changes in the lungs is investigated.

1.1 Overview of Thesis

Chapter 2 provides a general introduction to the Physics of PET and the theory of image reconstruction. The TOF capability in PET scanners is explained. The impact of respiratory motion in PET imaging is discussed and various approaches are described to tackle it. Respiratory motion detection methods are presented, discriminating between approaches that require external devices and DD methods. Finally, the PCA method for respiratory motion detection is described in detail.

Chapter 3 addresses an issue related to PCA and other DD methods: the sign of the generated respiratory signal is arbitrary, therefore its direction with respect to the physical motion is uncertain. Consequently, a maximum in the signal could refer either to the end-inspiration or end-expiration phase, possibly causing inaccurate motion correction and mismatch between PET and CT gates for attenuation correction (if CT gating is also applied). Therefore the sign of the signal needs to be determined. In this chapter, two sign-determination methods are presented, one of which is only applicable to correct the PCA signal and one that can be utilised with any DD method. They are sinogram-based and rely on the assumption that the major cause of change in the data is respiratory motion in the axial direction. Another already published registration-based method is also implemented, that requires image reconstruction. The performance of all methods is evaluated firstly on a simulation study, and finally on a cohort of PET patient datasets which were monitored by an external tracking device.

Chapter 4 presents the extension of PCA and the sinogram-based sign-determination methods to the use of TOF PET data. The process of unlisting data into TOF sinograms is illustrated. A masking process is presented to disregard the outer parts of the sinogram that do not include the lungs. A simple simulation study is performed, comparing the PCA results when using non-TOF and TOF sinograms, without and with the masking process.

Chapter 5 presents a thorough evaluation of the application of PCA and the sign-determination methods. Firstly, PET/MR data are utilised to compare the PCA signal extracted from the PET data with the internal motion extracted from MR images, and to the signal of an external device. Secondly, TOF PET data are utilised and the performance of the methods is investigated with respect to level of noise and duration of the data interval taken into account, and with respect to using non-TOF and TOF sinograms, without and with the masking process. Lastly, the PCA signal is compared to that of an external device for respiratory gating purposes, analysing the detected motion obtained from reconstructed gated data.

Chapter 6 presents the results obtained by applying PCA on PET data of the upper chest, where respiratory motion is small and lung density variations are expected to cause the biggest variation in the data. A simulation study is performed, simulating density-only changes and the PCA results are analysed. Finally, PCA is applied to the upper lung bed position of TOF PET patient datasets. PCA signals are compared to those of an external device. The signal given by the sum of the counts in time is also calculated and its behaviour investigated.

Finally, in Chapter 7 the main results and findings of the work presented in this thesis are summarised and the potential future work is discussed.

Chapter 2

Background

In this chapter the principles of PET imaging are described, addressing firstly the Physics at its foundation and secondly image reconstruction theory and algorithms. Subsequently, the basics of Computed Tomography (CT) and Magnetic Resonance Imaging (MRI), that are imaging modalities often used in conjunction with PET, are presented. Lastly, the issue of respiratory motion in PET is addressed. Methods used to detect respiratory motion in clinical practice are presented and DD methods are introduced.

2.1 Positron Emission Tomography

PET is a nuclear medicine imaging technique which allows non-invasive quantitative assessment of biochemical and functional processes. It achieves this by determining the distribution of radioactive tracers inside the patient body. Clinical PET imaging is a very important tool mainly used in three medical areas:

- in oncology: to differentiate malignant from benign nodules, for cancer staging, to detect recurrences, to measure the response to therapy, for radiotherapy planning;
- in cardiology: to measure myocardial perfusion and viability;
- in neurology and psychiatry: for the pre-surgical evaluation of epilepsy, for early diagnosis of types of dementia.

PET tracers are radioactive chemical compounds in which one or more atoms have been replaced by a positron emitting radioisotope. The principle at the base of the image production is the detection of the two photons generated by electron-positron annihilation, which are emitted at approximately 180° to each other and therefore allow to localize their source along a straight line of coincidence, referred to as Line of Response (LOR), without the need for physical collimation. The kind of radiotracer that is utilised depends on the target of interest. The most commonly used tracer in oncology is ^{18}F-FDG (Fluorodeoxyglucose):

it is a glucose analogue and is taken up by cells with high glucose metabolism, in particular by many types of cancer cells, and remains trapped in the cell. The positron emitting radioisotope in ^{18}F-FDG is ^{18}F, which has a half-life of 109.8 minutes (Bailey et al., 2005b).

2.1.1 PET principles

2.1.1.1 Beta decay and positron annihilation

Positrons in PET are produced by β^+ decay of the injected radioisotope, which is a particular type of radioactive decay in which a proton inside a radionuclide nucleus is converted into a neutron while releasing a positron and an electron neutrino:

$$ {}^{A}_{Z}X \rightarrow {}^{A}_{Z-1}Y + {}^{0}_{1}\beta^+ + \nu_e \tag{2.1} $$

The atom X is proton-rich and achieves stability by converting a proton to a neutron. Beta decay is a three-particle decay where the kinetic energy is mainly shared between the lighter particles, the positron and the neutrino (Knoll, 2010). The generated positron will lose its kinetic energy via interactions with the surrounding matter. Eventually the positron will lose almost all of its kinetic energy and combine with an electron, annihilating with it when both are essentially at rest and resulting in the production of two annihilation photons of 511 keV of energy, see Figure 2.1.

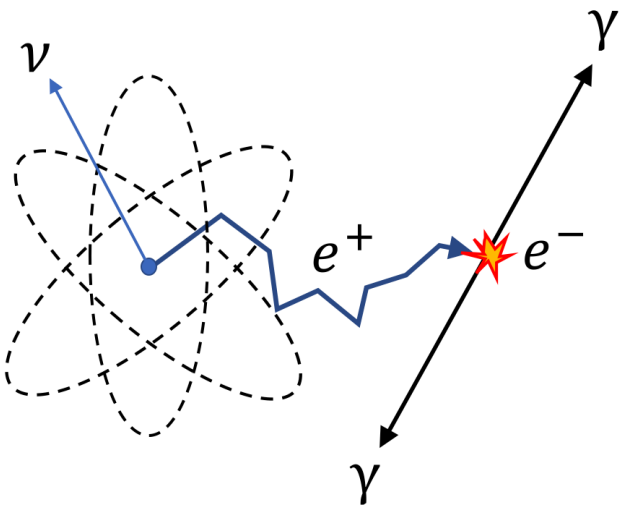

Figure 2.1: Positron-electron annihilation.

The emission of positrons, and therefore the decay of the radioactive nuclei which generate them, is a random event and it is governed by the exponential law:

$$ N = N_0 e^{-\lambda t} \tag{2.2} $$

where N_0 and N represent the number of radionuclides present at time t_0 and at time $t > t_0$

and λ is the decay constant of the radioisotope. The source *activity* A is defined as the number of disintegrations per second, that is:

$$A = \frac{\mathrm{dN}}{\mathrm{dt}} = -\lambda N \tag{2.3}$$

Since positron sources are not stable in nature, they have to be artificially produced: charged particles are accelerated by cyclotrons or linear accelerators and then irradiate specific inert matter, giving rise to nuclear reactions that eventually produce the needed radionuclides.

Radionuclide	Half-life (min)	E_{max} keV	Production reaction
^{18}F	109.8	633.5	$^{18}O(p,n)^{18}F$ $^{20}Ne(d,\alpha)^{18}F$
^{15}O	2.04	1732	$^{14}N(d,n)^{15}O$ $^{15}N(p,n)^{5}O$
^{13}N	9.97	1198.5	$^{13}C(p,n)^{13}N$ $^{16}O(p,\alpha)^{13}N$
^{11}C	20.38	960.2	$^{14}N(p,\alpha)^{11}C$
^{68}Ga	67.7	1899.1	Generator: ^{68}Ge
^{82}Rb	1.3	3378	Generator: ^{82}Sr

Table 2.1: Cyclotron-produced PET radionuclides (Del Guerra, 2004).

Rubidium-82 has a very short half-life, fortunately it has a parent, Strontium-82, from which a convenient generator system can be produced. Strontium-82 is usually produced in an accelerator by a spallation reaction using high-energy proton bombardment of Rubidium-85: $^{85}Rb(p,4n)^{82}Sr$ (Jadvar and Parker, 2006).

Gallium-68 is produced in the decay of Germanium-68 that is obtained, for example, through proton irradiation of Gallium-69: $^{69}Ga(p,2n)^{68}Ge$.

2.1.1.2 Interaction of photons with matter

High-energy photons interact with matter via various mechanisms, depending on their energy. The more pertinent interactions for photons with energy of 511 keV are the *photoelectric effect* and the *Compton scattering*.

The *photoelectric effect* is an interaction of a photon with an inner-shell orbital electron. All the photon energy is transferred to the electron, which is then able to overcome the binding energy of its orbital and is ejected from the atom. The interacting photon therefore disappears (it is *absorbed*). This effect dominates in human tissue for photons with energy below 100 keV and is therefore negligible in the patient in PET, whereas it is of primary importance in the functioning of radiation detectors.

In the *Compton scattering* the photon interacts with a loosely bound outer-shell electron. The photon transfers some of its kinetic energy to the electron that is then ejected. The

photon retains the rest of the energy and undergoes a change in direction. The energy of the photon after interaction is given by the following, where E_γ is the energy of the photon prior to the interaction and θ the scattering angle:

$$E_{\gamma'} = \frac{E_\gamma}{1 + (E_\gamma/511\text{keV})(1 - cos\theta)} \tag{2.4}$$

This effect dominates in human tissue at energies between approximatively 100 keV and ~ 2 MeV thus represent the primary mechanism of interaction for annihilation photons.

The photon cross section (i.e. probability of interaction) for these two effects is displayed for water in Figure 2.2 with respect to the energy.

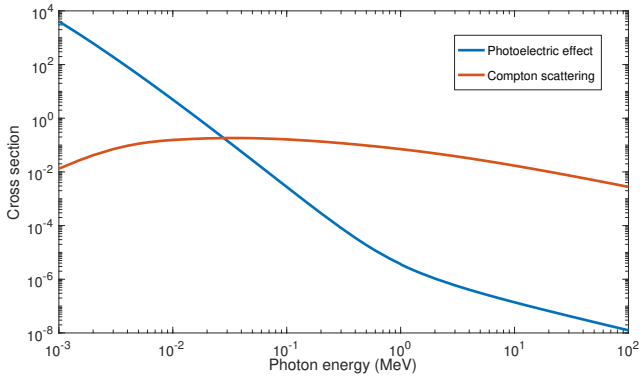

Figure 2.2: Photon cross section in water.

The loss of a photon from a photon beam due to either the photoelectric or Compton effect is defined as *attenuation*. Attenuation depends on the energy of the photon and on the characteristics of the material that the photon is passing through. The likelihood of an attenuation event taking place per *cm* of absorbing tissue is defined as the *linear attenuation coefficient* μ. Given N_0 initial photons, the number of non-attenuated photons after having crossed a path of length L follows the Beer-Lambert exponential law:

$$N = N_0 \exp\left(-\int_L \mu(x)\mathrm{d}x\right) \tag{2.5}$$

where $\mu(x)$ is the attenuation coefficient of the media crossed by the photons.

2.1.1.3 Photon detection

The detection of the annihilation photons in PET is achieved via the use of scintillation crystals coupled to photodetectors.

A scintillator is a material that exhibits *scintillation* when excited by ionizing radiation, i.e. it absorbs the incoming energy and re-emits it in the form of visible light. The main properties required for a scintillator to be used in PET imaging are: high stopping power for 511 keV photons; a short decay time, to process each pulse individually even at high

count rates; high light output and therefore high intrinsic energy resolution (Knoll, 2010). Important inorganic scintillator crystal materials used in PET detectors are: *BGO*, *LSO*, *LYSO* and *YAP:Ce*.

The weak light output generated by the scintillators is converted to an electrical signal by photodetectors. Photomultiplier tubes (PMT) and solid-state photodetectors are the two main groups of devices for detecting low light levels. Among the solid-state photodetectors, a recent type is represented by the SiPM (Silicon photomultiplier), which consist of a densely packed matrix of small APD (avalanche photodiode) cells.

2.1.2 Data Acquisition

As previously mentioned, PET does not need physical collimators as it relies on electronic collimation of the detected signals. A coincidence event is considered as valid and referred to as *prompt* if it satisfies the following characteristics:

- two photons are detected within a chosen coincidence time window, whose width 2τ depends on the timing resolution of the detectors;
- the LOR formed between the two photons is within a valid acceptance angle of the tomograph;
- the energy deposited in the crystal by each photon is within a selected energy window.

Despite these selection criteria some of the prompt events are actually unwanted, namely when they do not arise from the detection of photons arising from the same annihilation event. The various kinds of possible detected events are the following (Bailey, 2005):

- *single event*: it consists of a single photon counted by a detector;
- *true coincidence event*: it derives from the detection of the two photons originated from a single positron-electron annihilation; when both the photons reach the detectors without undergoing scattering it is an *unscattered event* whereas if scatter does occur it is referred to as a *scattered event*;
- *random coincidence*: it arises when two unrelated photons enter opposing detectors and are temporally close enough to be recorded within the coincidence timing window;
- *multiple events*: they occur when three or more events from two annihilations are detected within the coincidence timing window.

The *prompt* count rate is given by the sum of the true and the random and scattered events rates (Bailey et al., 2005b):

$$\mathrm{P} = \mathrm{T} + \mathrm{R} + \mathrm{S} \tag{2.6}$$

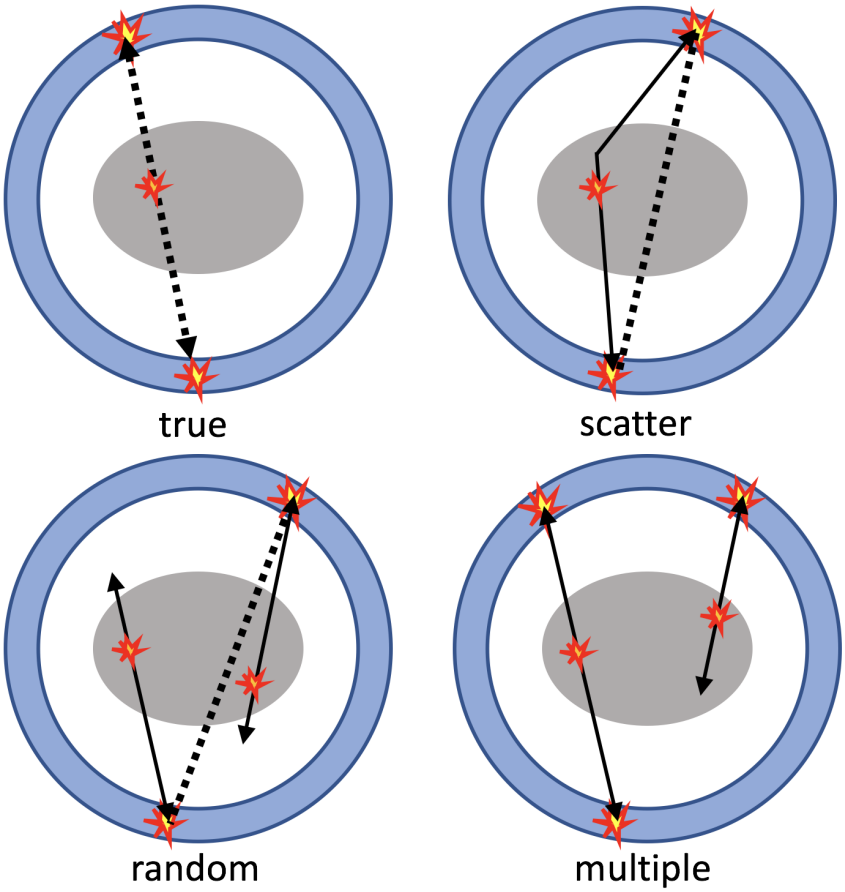

Figure 2.3: PET events.

Random and scattered events add background counts to the PET acquisition, therefore decrease the diagnostic accuracy of the exam when they are not accounted for.

The number of scattered events depends on the size of the patient, as the larger the patient the more likely it is that a photon will undergo scattering prior to entering the detector, and on the density of the tissues, as the probability for Compton scatter is approximately proportional to the atomic number of the medium (Bailey et al., 2005a).

2.1.2.1 Data organization

PET data consist in detected coincidence events that are assigned to the LOR joining the two relevant detectors. The amount of detected events related to each LOR are organized into a dataset called *sinogram*. In this thesis a *sinogram* is defined as $p(s, \phi, \theta, z)$, where s is the signed distance between the LOR and the centre of the coordinate system, ϕ is the angle measured around the ring, θ is the polar angle between rings and z is the axial coordinate (see Figure 2.4 for the scanner coordinates and a visual explanation of the sinogram). Two-Dimensional (2D) PET imaging only considers LORs lying within a specified imaging plane, i.e. with polar angle $\theta = 0$, while with Three-Dimensional (3D) PET, projections are measured also for

$\theta > 0$, to increase the sensitivity and improve the signal-to-noise ratio of the reconstructed images. In the 2D case, considering a transaxial section with $z = z_0$, the relationship between the sinogram elements and the tracer distribution $f(x, y)$ is expressed by:

$$p(s, \phi, z_0) = \int_{-\infty}^{\infty} f(x = s\cos\phi - t\sin\phi, y = s\sin\phi + t\cos\phi, z = z_0)\,\mathrm{d}t \tag{2.7}$$

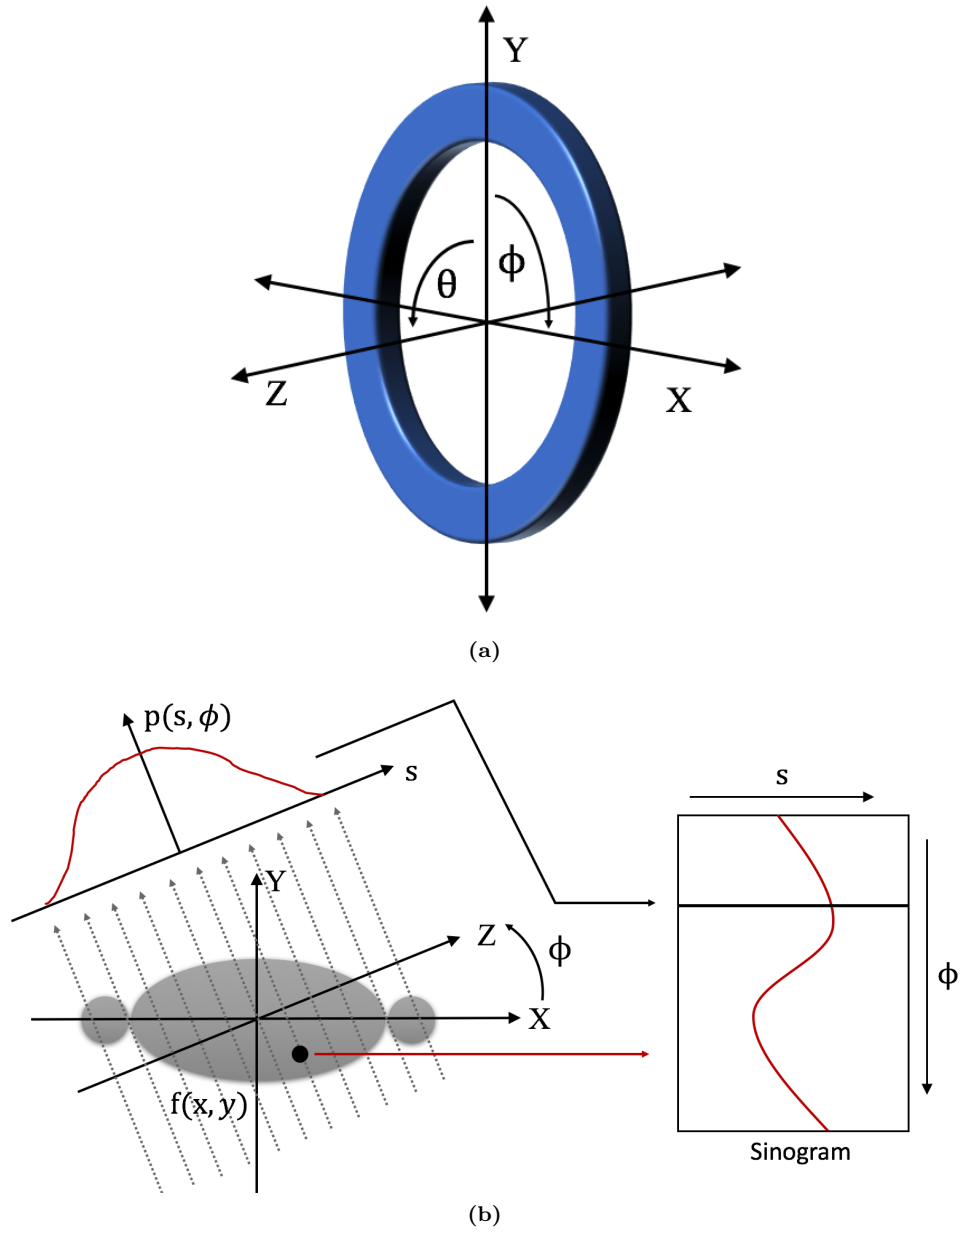

Figure 2.4: (a) Coordinate system of the scanner (adapted from Bailey et al. (2005b)); (b) two-dimensional sinogram (adapted from Alessio and Kinahan (2006)).

Since in 3D PET the sinogram datasets demand a big amount of storage space, and considering that for acquisition time frames of 5-10 seconds very few elements of the sinogram will contain an event, PET data can also be stored in the *listmode* format. In listmode

acquisition the information related to the detector pair corresponding to each individual event is explicitly stored, together with its time stamp (with resolution of ∼1 ms). Once acquired, the listmode data can be regrouped into sinograms with the required spatial and temporal sampling for reconstruction, therefore offering the capability of exploring finer temporal sampling while making a more efficient use of storage capacity (Bendriem and Townsend, 2013).

2.1.3 Time-of-Flight PET

As it has been described in the previous section, in conventional PET imaging the detection of a coincidence event is the result of the acquisition of the two annihilation photons that have interacted with the scanner detectors within a defined time window. The timing information of each of the two photons can also be further exploited to determine the position along the LOR where the annihilation event took place, as the measured difference in arrival times is directly related to the distance of the annihilation location from the midpoint of the LOR. Using Figure 2.5, the following relationships can be obtained:

$$D_1 = L + x$$
$$D_2 = L - x$$

where D_1 and D_2 are the distances travelled by the photon that interacted with detector 1 and 2 respectively. Subtracting the two equations, the result is that $D_1 - D_2 = 2x$. Given that the time for the photons to reach the detectors is given by $t_1 = D_1/c$ and $t_2 = D_2/c$ respectively, where t_1 and t_2 are the arrival times and c is the speed of light, the distance of the annihilation location from the midpoint of the LOR is given by:

$$x = c\frac{t_1 - t_2}{2} = c\frac{\Delta T}{2} \tag{2.8}$$

If the timing resolution of the PET scanner were perfect it would therefore be possible to precisely locate the annihilation event along the LOR. In reality, there are limitations caused by the uncertainty in the measured difference ΔT that is characterised by the timing resolution of the system Δt. In TOF PET the annihilation point is placed on the LOR in the point corresponding to Eq. 2.8 with a probability function whose Full Width at Half Maximum (FWHM) is equal to $c\Delta t/2$ (Lewellen, 1998; Conti, 2009). This means that, for example, a timing resolution Δt of 600 ps corresponds to a positional uncertainty of 9 cm along the LOR.

When the scanner is capable of measuring TOF information, there is therefore an additional coordinate for the PET event, given by its position along the LOR defined by the detector

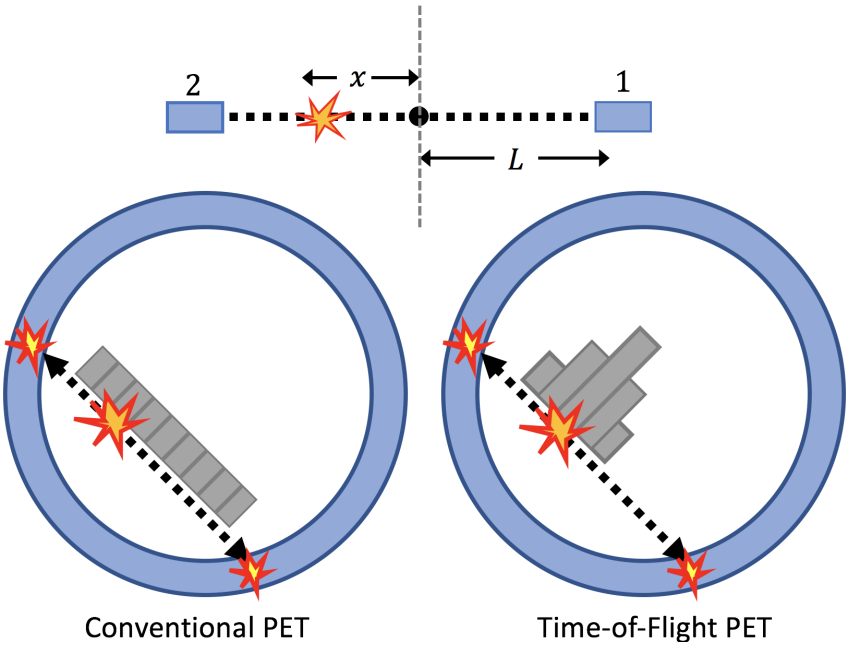

Figure 2.5: Explanation of the principles of TOF PET. In conventional PET the probability of the event taking place is the same along the whole LOR, whereas in TOF PET the probability function is centred around the distance given by Eq. 2.8 with a FWHM of $c\Delta t/2$.

pair. The TOF information is acquired with an accuracy that depends on the scanner. For example, the accuracy is equal to 25 ps for the Philips Ingenuity TF PET/CT scanner (Vandenberghe et al., 2016; Kolthammer et al., 2014) and it is 89 ps for the GE Discovery 690 PET/CT scanner (Bettinardi et al., 2011). This value is referred to as *TOF bin size* and the total number of TOF bins utilised by the scanner to store the data is given by:

$$\text{number of TOF bins} = \frac{\text{coincidence window width}}{\text{bin size}} \tag{2.9}$$

The TOF bin adds another dimension to the PET dataset, which will be referred to with the letter τ. In this work the TOF bins will be numbered assigning a value equal to 0 to the central bin, which represents the annihilation events that took place around the centre of the LOR, and then assigning positive ascending values moving further away from the centre on one side of the LOR and negative values on the opposite side. A schematic representation of the TOF bins numbering is displayed in Figure 2.6, where the values of τ range from -3 to 3, depending on the distance from the centre of the LOR.

The success of TOF PET depends primarily on the ability of accurately determining ΔT and therefore on the timing characteristic of the detectors. With the advent of newer scintillators that bring together fast timing and high light output, such as lutetium orthosilicate (LSO) and lutetium-yttrium orthosilicate (LYSO), the performance of TOF PET scanners has surpassed that of the conventional ones (Vandenberghe et al., 2016).

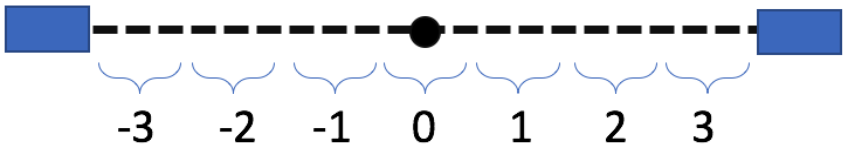

Figure 2.6: Schematic description of the TOF coordinate τ along a LOR. The greater the τ in absolute value, the more distant the event location is from the centre of the LOR.

2.1.4 Image Reconstruction

There are two main families of reconstruction algorithms for PET data: *deterministic methods* and *statistical methods*. The former rely on the assumption that the acquired data is deterministic, therefore not affected by statistical noise, and offer a direct mathematical solution to obtain the image from the known projections $\mathbf{p}$. An important category of the *deterministic methods* is that of the *analytical methods*. The *statistical methods* take into consideration the intrinsically stochastic nature of the PET data, and provide estimation processes to obtain an approximate solution, as the result of the maximization or minimization of a cost function. These methods often make use of iterative processes, and are in that case referred to as *iterative algorithms*.

In the next sections the image reconstruction algorithms will be explained for non-TOF PET data. They can be adapted to the TOF case with some modifications that will be highlighted at the end.

2.1.4.1 Analytic image reconstruction

Analytic reconstruction algorithms rely on Fourier analysis exploiting a fundamental relationship between a function and its projections. This is represented by the *Fourier central section* theorem, and one of the most commonly used analytic algorithms that is based on it is Filtered Back-Projection (FBP).

There are several limitations of analytical reconstruction methods that affect their performance. In particular, they do not allow accurate modelling either of the detector system or of the inherent statistical fluctuations in the data. Nevertheless they are important when computation time is limited and are also useful for initializing iterative algorithms, that will be discussed later in the chapter (Bailey et al., 2005b).

The Fourier Central Section Theorem and Filtered Back-Projection

The Fourier central section theorem provides a useful relationship between a two-dimensional function $f(x,y)$ and its projections $p(s,\phi)$: it states that the Fourier transform of a one-dimensional projection is equivalent to a section at the same angle through the center of the two-dimensional Fourier transform of the function. Being $\mathcal{F}_1\{p(s,\phi)\}$ and $\mathcal{F}_2\{f(x,y)\}$ respectively the Fourier transform of the projection $p(s,\phi)$ and of the two-dimensional

function $f(x,y)$:

$$\mathcal{F}_1\{p(s,\phi)\} = P(v_s,\phi) \tag{2.10}$$

$$\mathcal{F}_2\{f(x,y)\} = F(v_x,v_y) \tag{2.11}$$

the Fourier central section theorem states that:

$$P(v_s,\phi) = F(v_s\cos\phi, v_s\sin\phi) \tag{2.12}$$

Thanks to this equivalence, it is possible to recover $F(v_x,v_y)$ for all frequencies (v_x,v_y), measuring all projections for $\phi\in[0,\pi]$ and evaluating their Fourier transform. The image f may be then reconstructed by inverse 2D Fourier transform:

$$(\mathcal{F}^{-1}F)(x,y) = f(x,y) = \iint_{\mathbb{R}^2} F(v_x,v_y)\mathrm{e}^{2\pi i(xv_x+yv_y)}\,\mathrm{d}v_x\mathrm{d}v_y \tag{2.13}$$

For PET image reconstruction, it is convenient to change to polar coordinates, where $s = x\cos\phi + y\sin\phi$:

$$f(x,y) = \int_0^\pi p^F(s,\phi)\,\mathrm{d}\phi \tag{2.14}$$

$$p^F(s,\phi) = \mathcal{F}_1^{-1}\{|v_s|\mathcal{F}_1\{p(s,\phi)\}\}. \tag{2.15}$$

where $p^F(s,\phi)$ are referred to as the filtered projections and $|v_s|$ as the *ramp filter*. This kind of data processing is accomplished by the FBP. The ramp filter compensates for sampling density differences in the frequency space. This filter accentuates edges (high-frequencies) while decreasing blurring (low-frequencies). Consequently, the ramp filter will also amplify the noise, as it dominates the Fourier transform of the measured data at increasing frequencies, therefore some form of regularization with a low-pass filter is needed. Equation 2.14 becomes:

$$f(x,y) \approx \widetilde{f}(x,y) = \int_0^\pi \mathcal{F}_1^{-1}\{W(v_s)|v_s|\mathcal{F}_1\{p(s,\phi)\}\}\,\mathrm{d}\phi \tag{2.16}$$

where $W(v_s)$ is a low-pass apodizing window, such as the Hamming, Hanning or Butterworth window.

2.1.4.2 Three-dimensional analytic reconstruction

The methods dealt with so far are relative to the 2D case and are not directly applicable to the fully 3D PET imaging. The LORs measured by a volume PET scanner can be grouped into 2D sets of parallel projections, although some projections are truncated due to the

finite axial extent of the scanner. The central section theorem can be generalized to 3D and, together with the assumption of non-truncated 2D projections, leads to a method analogous to the two-dimensional FBP, that is 3D FBP. Since the available projections are actually truncated, algorithms have been proposed to overcome this issue. One common method is the *Three-dimensional reprojection algorithm* (3DRP) (Kinahan and Rogers, 1989) that estimates the unmeasured regions forward-projecting an initial image obtained applying the 2D FBP algorithm only to non-truncated projections. Then the true and estimated data are reconstructed using the 3D FBP.

Another method to obtain transaxial sections from the 3D ones is that presented by *rebinning algorithms*.

Rebinning algorithms These algorithms estimate the ordinary sinogram of transaxial sections by means of some form of signal averaging, therefore permitting to reduce the size of the data and to use analytic or iterative 2D reconstruction methods. Two of the best known approximate algorithms are the Single Slice Rebinning Algorithm (SSRB) (Daube-Witherspoon and Muehllehner, 1987), and the Fourier Rebinning Algorithm (FORE) (Defrise, 1995). The SSRB generates the rebinned sinograms through an averaging of all the oblique sinograms that intersect the direct plane at the centre of the transaxial FOV: a coincidence event is considered to have originated from a parallel plane (i.e. perpendicular to the z-axis), which is located in the middle between the axial coordinates of the points of interaction of the two gamma rays in the detectors, as displayed in Figure 2.7. This strategy results in the creation of a series of 2D projections that are assumed to be parallel. This algorithm is going to be extensively used for the data processing in this thesis. The FORE algorithm exploits an equivalence between the Fourier transforms of oblique and transverse sinograms to obtain direct sinograms (Bailey et al., 2005b).

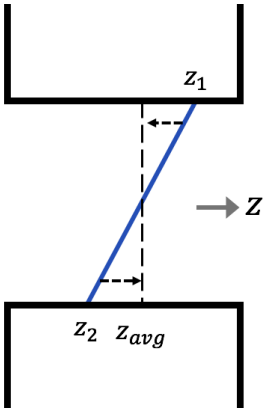

Figure 2.7: Illustration of the "compression" of the axial information when applying SSRB: the LOR shown in the figure is considered to belong to the direct plane depicted with the dashed line, corresponding to the axial coordinated given by z_{avg}, the average of axial coordinates z_1 and z_2 that define the LOR.

2.1.4.3 Iterative image reconstruction

Iterative algorithms allow more accurate modelling of the data acquisition, therefore offering a higher image quality with respect to the analytical methods. Their model includes the detection process, the photon transport in tissues and the statistical distribution of the acquired data, taking into account the noise structure in the observations. Iterative reconstruction methods are conceptually independent of whether the data is 2D or 3D.

Image and data model The image is represented using a finite set of basis functions, most commonly chosen as square pixels (2D) or cubic voxels (3D), which are indicator functions on a square or cubic region centred at one of the image sampling points j ($j = 1 \dots N$) in a regular 2D or 3D lattice. For the 2D case, the value of the tracer concentration in (x, y) is defined by:

$$f(x,y) \simeq \sum_{j=1}^{N} f_j b_j(x,y) \tag{2.17}$$

where the square pixels functions are defined as follows, with the centre of the pixel j being (x_j, y_j):

$$b_j(x,y) = \begin{cases} 1 & |x - x_j| < \Delta x/2 \quad \text{and} \quad |y - y_j| < \Delta y/2 \\ 0 & |x - x_j| \geqslant \Delta x/2 \quad \text{or} \quad |y - y_j| \geqslant \Delta y/2 \end{cases} \tag{2.18}$$

The value f_j relative to voxel j is proportional to the total number of positron-emitting nuclei contained in the volume spanned by the voxel, and the sinogram data $\boldsymbol{p}$ is related to the image values through the projection operation (see 2.1.4.3). The index j will be used to enumerate the ordered elements of the image ($j = 1 \dots N$), whereas i ($i = 1 \dots M$) will be used to enumerate the LORs.

The statistical distribution of each projection data element around its mean value can also be modelled. The number of positron emissions of a large number of radioactive nuclei can be described, in good approximation, as following a Poisson distribution (Barrett and Myers, 2013) and, supposing that each LOR measurement made by the system is independent, then the sinogram data $\boldsymbol{p}$ are a collection of Poisson random variables. The probability that the random vector of Poisson distributed photon counts (with mean value $\langle p_i \rangle$ for each LOR), equals the collected data $\boldsymbol{p}$, is given by:

$$\mathcal{L}(\boldsymbol{p}) = \prod_{i=1}^{M} \frac{\langle p_i \rangle^{p_i} \mathrm{e}^{-\langle p_i \rangle}}{p_i!} \tag{2.19}$$

where the relationship between $\boldsymbol{p}$ and the image f is defined in the following section.

System model The *system model* relates the image to the data. Since the data are discrete and the detection process is approximately linear, the relationship between the source image and the expected value of the true coincidence data can be represented by a forward projection matrix, $\boldsymbol{H} \in \mathbb{R}^{M \times N}$. $\boldsymbol{H}$ represents the imaging system and its elements, h_{ij} , contain the probabilities of detecting an emission from voxel site j in projection i, so that the expected value for the projection data p_i is given by:

$$\langle p_i \rangle = \sum_{j=1}^{N} h_{ij} f_j + s_i \tag{2.20}$$

where s_i is the expected number of background events, due to scatter and random coincidences. The system matrix should in principle take into account all of the physics processes that are at the base of data production. A useful approach is to factor $\boldsymbol{H}$ as a product of matrices, each modelling a specific aspect of the data acquisition process. The system model can be obtained via measurements or via Monte Carlo simulations, while other approaches make use of multi-ray methods, incorporating accurate analytical (geometric) considerations as well as crystal depth and crystal scatter effects (Moehrs et al., 2008).

Cost function - MLEM A fundamental ingredient of iterative algorithms is the *cost function* that depends on the unknown image values and on the measured data, that is to be minimized or maximized in order to obtain the best image estimate. The most common principle for iterative reconstruction is the *Maximum Likelihood* (ML) approach. It searches for the image that maximizes the likelihood of obtaining the measured data $\boldsymbol{p}$ that, for independent Poisson variables, is given by (2.19). Then substituting $\langle p \rangle$ with the expected value defined in (2.20), taking the logarithm and disregarding the terms that do not depend on the unknown f, the following cost function is obtained:

$$\mathcal{Q}(\vec{f},\vec{p}) = \sum_{i=1}^{M} \left[p_i \log\big(\sum_{j=1}^{N} h_{ij} f_j + s_i\big) - \big(\sum_{j=1}^{N} h_{ij} f_j + s_i\big) \right] \tag{2.21}$$

The solution of the maximization of $\mathcal{Q}$ defines the desired image estimate. One of the most commonly used algorithms for optimizing the cost function is Maximum Likelihood Expectation Maximization (MLEM), which was introduced by Dempster in 1977 (Dempster et al., 1977) and first applied to PET by Shepp & Vardi in 1982 (Shepp and Vardi, 1982) and by Lange and Carson in 1984 (Lange et al., 1984). The MLEM iterative equation is given by the following:

$$f_j^{n+1} = \frac{f_j^n}{\sum\limits_{i'}^{M} h_{i'j}} \sum_{i}^{M} h_{ij} \frac{p_i}{\sum\limits_{k=1}^{N} h_{ik} f_k^n + s_i} \tag{2.22}$$

where f_j^{n+1} is the next estimate of the image, obtained from an update of the current estimate f_j^{n+1} making use of the acquired data p_i.

Ordered Subsets Expectation Maximization (OSEM) is an accelerated version of MLEM proposed in 1994 by Hudson and Larkin (Hudson and Larkin, 1994), with the aim of reducing the reconstruction time. In the OSEM algorithm, the projection views are grouped in S different subsets, then the MLEM algorithm is applied incorporating the data from one subset only and goes through the subsets in a specified order. The image is updated after each subset is considered.

2.1.4.4 Time-of-Flight image reconstruction

When TOF PET was introduced analytic image reconstruction algorithms were the only available methods. A variation of FBP was therefore proposed, where a one-dimensional weight is applied on the LOR, centered around the TOF detected location (Tomitani, 1981). With the advent of iterative reconstruction algorithms, such as MLEM, an extension of these methods was implemented to exploit the TOF added information in Snyder and Politte (1983), where a point spread function whose width depends on the time resolution of the scanner is applied in the projector. When they were first introduced, iterative algorithms for TOF data required lengthy computational times but with the current scanners reconstruction times have decreased considerably. At present, iterative reconstruction of TOF data is the standard choice in clinical practice (Conti, 2009).

2.1.5 PET data corrections

In order to obtain quantitative measures of the the tracer concentration from PET data, corrections need to be applied. The main effects that are required to be modelled, and corrected for, are the following: the attenuation of the annihilation photons in the body, which lowers the resulting number of detected events, and the presence of scattered and random photons, which identify inaccurate LORs.

2.1.5.1 Attenuation correction

Photons emitted by the radiotracer have to pass through the patient body prior to reaching the detectors, and during their path they can be absorbed or scattered away from their initial track, as described in Paragraph 2.1.1.2. Some of the emitted photons will therefore not reach any detector and the corresponding emission event will not result in a detected coincidence, causing the acquired data to be less representative of the actual activity present in the field of view. The probability of detecting an annihilation event is given by the combined probability of detecting both the annihilation photons. Defining the thickness of the body with D and one photon's path by L_1, the other photon's path will be given by $D - L_1$, and using Eq. 2.5

the combined probability is therefore given by the following:

$$P = \exp\left(-\int_{L_1} \mu(x)\mathrm{d}x\right) \times \exp\left(-\int_{D-L_1} \mu(x)\mathrm{d}x\right) = \exp\left(-\int_{D} \mu(x)\mathrm{d}x\right)$$

This probability corresponds to the attenuation correction factor for the LOR taken into account.

Attenuation factors can be obtained by comparing the detected counts obtained from an external source when the patient is in the scanner to the counts obtained when the scanner is empty. This can be achieved by using a rotating positron emitting source, and then detecting either the singles or the coincidence events with and without the patient (these are called PET transmission scans) (Bailey et al., 2005b). Alternatively, the attenuation correction factors can be directly obtained from a CT image, exploiting the higher spatial resolution and much faster acquisition of CT compared to the PET transmission acquisitions (Kinahan et al., 1998). The ensemble of the attenuation factors are referred to as *attenuation map*.

In the context of iterative reconstruction algorithms, the attenuation correction is applied in sinogram space, multiplying the forward projection of the image estimate (at each iteration) with the attenuation factors.

2.1.5.2 Random coincidence estimation

In order to restore the quantitative information in the PET data, the contribution of the random coincidence events needs to be subtracted from the acquired data. One of the available methods to estimate the rate of random events, related to a particular LOR, exploits the single rates of the detectors, with the following:

$$R_{ij} = 2tr_i r_j \tag{2.23}$$

where $2t$ is the coincidence timing window and r_i and r_j the single rates of the detectors defining the LOR (Meikle and Badawi, 2005). The estimated random coincidences can then be subtracted from the measured data. A narrow time window is desirable to minimize the random collection.

2.1.5.3 Scatter fraction estimation

The contribution to the total detected events in 3D PET given by scattered photons can be as large as 40%, and it is referred to as *scattered fraction* (Meikle and Badawi, 2005). One of the most commonly utilised algorithms to estimate the scatter fraction is the Single Scatter Simulation (SSS) method (Ollinger, 1996; Watson, 2000). This method utilises the emission data, the attenuation map, a model of the scanner geometry and detector system in order to provide an estimate of the percentage of scatter photons that fall on each detector. It makes use of the Klein-Nishina formula, which returns the differential cross section of

photons scattered from a single free electron, as a function of the photon scattering angle and initial energy.

2.2 Other imaging modalities

PET imaging displays the physiology and metabolism of the radioactive tracers but lacks anatomical information, which is needed both for identifying the locations of increased activity (and therefore tracer uptake) and to perform attenuation correction. In the past, radiologists used to look at PET images side by side with anatomic images, such as Computed Tomography (CT) or Magnetic Resonance Imaging (MRI) acquired in separate exams, gathering information from the two modalities. Another approach is that of registering and fusing images acquired on different systems, but this process can be non-trivial and time consuming. The ideal solution is therefore to have two modalities combined into one machine, hence PET/CT scanners (and more recently also PET/MR ones) have been developed and manufactured. The basic principles of CT and MR will be addressed in the following sections.

2.2.1 Computed Tomography

CT is a non invasive technique that makes use of ionizing radiation to gain information about the structure of the body. It discriminates between tissues depending on their attenuation properties.

2.2.1.1 Basic principles

The basic concept of CT imaging is that of using X-rays to create an attenuation map of the imaged object, by detecting from many angles the number of X-rays that pass through the object unabsorbed. X-rays are created by an X-ray tube, where electrons are generated in the cathode by thermoionic emission and then accelerated by a high potential difference (50 to 150 kV) towards a tungsten anode. The accelerated electrons interact with the anode, and produce X-rays as a result of ionization and Bremsstrahlung interactions (i.e. when the electron loses energy because it is decelerated when passing close to the tungsten nuclei). A filter of low-atomic number absorbs the unwanted X-rays of low energy, as they would only contribute to the dose to the patient without adding diagnostic information, and a collimator defines the beam shape. The energy spectrum of the resulting X-rays can be tuned by changing the potential difference that accelerates the electrons in the X-ray tube, and the current in the cathode filament. The former determines the maximum energy of the generated X-rays and the latter determines their total number.

The X-ray detection is achieved via the use of scintillating detectors (usually ceramic material or cadmium tungstate), coupled with photodiodes. The detectors (with a size of the order of ~ 0.5 to 1 mm on a side) are grouped together in large arrays, positioned in order to create a circumference, and several rows of these (64 for example) compose the CT

tomograph. The most recent CT scanners have fixed detectors and a rotating X-ray beam, and the data are acquired in *helical mode*, that consists in the table and gantry moving at the same time (Hu, 1999). Helical acquisition allows a scan to be performed in only a few seconds, e.g. 3s to cover the lung region.

2.2.1.2 Image generation

The different absorption of X-rays by the various types of tissues (bone, lung, air, soft tissue) represents the means to extrapolate information about the morphology of the imaged body. The CT image reconstruction is based on the attenuation equation Eq. 2.5, that relates the intensity of the detected attenuated beam, here referred to as $I(x)$, with that of the unattenuated beam and the attenuation coefficient and thickness of the tissues crossed by the photons. In CT, the X-rays tube rotates around the patient therefore the value of $I(x)$ is known for many angles. The acquired data is then reconstructed with algorithms similar to those described for PET imaging (see paragraph 2.1.4) to obtain a 3D matrix (the image) of attenuation coefficient values. The values of each image voxel is expressed in *Hounsfield Units* (HU), that are defined by the following:

$$HU_{ij} = 1000 \frac{\mu_{ij} - \mu_{\text{water}}}{\mu_{\text{water}}} \tag{2.24}$$

where μ_{ij} is the linear attenuation of voxel ij and μ_{water} is the linear attenuation of water.

2.2.1.3 Data acquisition in PET/CT

In PET/CT scanners the two tomographs are combined on a single gantry, one next to the other, and the table can move axially to cover the desired area of the patient to be imaged with either of the modalities (Beyer et al., 2000). The acquisition of CT and PET data is therefore sequential and not simultaneous. A typical PET/CT exam starts with the acquisitions of a scout scan, that is an X-ray overview scan carried out by translating the patient through the CT gantry with a fixed X-ray tube position. This scan is used to determine the axial range of the examination and the number and location of the bed positions. After this, the transmission CT scan is performed for each bed position and finally the PET acquisition is performed. By the time the first emission scan (1^{st} bed) is completed, the reconstruction of the CT images has already been completed and they can be used for attenuation correction. Soon after the completion of the last bed position PET scan, all PET images are reconstructed and attenuation correction can be applied (Beyer et al., 2004).

2.2.1.4 Four-dimensional Computed Tomography

Conventional CT scans are not time-resolved, i.e. they do not capture several images in time but only deliver one image, which is generated by the data acquired throughout the duration of the acquisition. If motion occurs during the scan, this is going to be averaged out in one

single image.

For the purpose of decreasing motion artefacts in CT imaging, the time-resolved CT scheme has been introduced, often referred to as *cine-CT* (Lipton, 1987). A cine-CT scan acquires the projection data continuously over many gantry rotation cycles when the imaging table is stationary. The acquisition time is therefore longer than a normal CT scan. For example the acquisition of 100 slices (2.5mm thickness), with a respiratory period of 4s, requires about 2.5 minutes on a 4-slice scanner, as reported in Rietzel et al. (2005). It therefore also implies that the patient is exposed to a higher dose than in a normal CT scan. The cine-CT dataset can subsequently be sorted into respiratory gates, when a respiratory signal is available, similarly to what is described for PET data in section 2.3.1.1.

Additionally, time-resolved CT can also be obtained in helical mode with very low pitch (i.e. the motion of the bed is very slow compared to the rotation of the gantry), even though not all scanners commonly support the required low pitch values. Furthermore it has been demonstrated that cine-CT outperforms such methodology, as it can scan faster and it is more dose efficient (Pan, 2005).

2.2.2 Magnetic Resonance Imaging

MRI is a non-invasive technique that utilises the magnetic properties of certain nuclei, mostly hydrogen, in order to gain information on the internal structure of the body. MRI makes use of electromagnetic radiation that is not ionizing, therefore this technique is not harmful and can be repeated many times without effects on the patient, as long as certain safety limits related to heating are observed.

2.2.2.1 Basic principles

The fundamental physical effect exploited by MRI is that some nuclei exhibit magnetic properties. Among these, hydrogen is the most abundant in the human body therefore it is the most used in this technique. Protons exhibits a *magnetic moment* when placed in an external magnetic field, therefore when a hydrogen atom is in the presence of an external magnetic field, its proton's *magnetic moment* interacts with it, either aligning itself with or against the external magnetic field direction. The first situation is a low-energy state, whereas the second is a high-energy one. The net magnetic moment M_z is defined as the vector addiction of the magnetic moment of all the hydrogen spins.

At the same time as aligning with (or against) the external magnetic field, the spins precess around the magnetic field direction, and they do so with an angular frequency given by the *Larmor equation* (Larmor, 1897):

$$\omega_0 = \gamma B_0 \tag{2.25}$$

where B_0 is the external magnetic field and γ is a factor called *gyromagnetic ratio*. The value of γ is different for spins of different atoms. For the hydrogen γ is equal to 42.5 MHz/T. This equation constitutes the foundation of MR imaging, as it provides a feature (the precession frequency) to discriminate between different tissues at different locations in the body.

2.2.2.2 MR signal

When the net magnetization is aligned with the external magnetic field B_0 there is no detectable signal: the spins are precessing but they are out of phase, therefore there is no magnetization in the X-Y plane. In order to gain information about the tissue, the net magnetization needs to be perturbed (excited) so that magnetization in the transverse plane is generated. This is achieved by applying magnetic radiation perpendicularly to B_0, with a magnetic field called B_1 that has a frequency that matches the resonance frequency of the spins, which is given by w_0. This allows all the spins to precess in phase with one another, causing the net magnetization to precess in the X-Y plane. The net magnetic field will therefore exhibit a component in the transverse plane, M_{XY}. The disturbance radiation is called *Radio-frequency pulse*, or more often *RF pulse*.

Once the *RF pulse* stops being applied, the net magnetization gradually re-aligns with B_0, therefore losing its transverse component M_{XY}. This loss of transverse magnetization is called *free induction decay* (FID) and generates the measurable signal in MR imaging. It is detected with an antenna, called receiver coil. The speed at which the FID decays depends on the type of excited tissue.

A series of *RF pulses* is called a *pulse sequence*, and different pulse sequences are designed in order to generate images which provide various kinds of information. The signal that is detected after each sequence is called *echo*. The time from the excitation pulse to the acquisition of the echo is called *echo time* (TE), and the time between one excitation pulse and the following is called *repetition time* (TR) (Prekeges et al., 2012).

2.2.2.3 Image generation

To spatially encode the information in MR, magnetic field gradients are utilised. The variations in the applied magnetic field cause the precession frequency of the net magnetization to change. Depending on what the value of the magnetic field is in a particular area of the body, the spins will precess differently and the net magnetization will change accordingly.

Three different gradients are applied in other to spatially localize the echo in the three directions: the Z gradient (or *slice selection* gradient), the X gradient (or *phase encoding* gradient) and the Y gradient (or *frequency encoding* gradient). The data is stored in the so called *k-space* matrix and its dimensions are given by the number of frequency and phase encodings of the acquisition. There is one *k-space* matrix per slice. The final image is then generated through a process involving *Fourier transformation* (Prekeges et al., 2012).

In 2007 the first human tomograph capable of simultaneous PET and MR imaging was presented by Siemens (Schmand et al., 2007). In 2010 the first clinical commercial PET/MR scanner was launched by Siemens, with the Biograph mMR scanner (Delso et al., 2011). In 2014 General Electric introduced their PET/MR scanner with also TOF capability, the GE SIGNA scanner (Levin et al., 2016).

2.3 Respiratory motion in PET

For a diagnostic PET scan a typical acquisition duration is 2 minutes per bed position, therefore making the patient respiratory motion unavoidable. The use of the "breath-hold" technique, often applied to CT acquisitions (Fin et al., 2008), has been attempted in Nehmeh et al. (2007), but presents some difficulties, mainly the need for coaching the patient to reproduce the same breathing amplitude, and the potential inability of the patient to perform the required task. The acquired PET data is therefore averaged over several breathing cycles, during which the diaphragm can move as much as 20mm in the cranio-caudal direction (Wade, 1954), and if respiratory motion is not accounted for, it will result in a blurring of the reconstructed images and in an overall degradation of the contrast. When a lesion is moving throughout the acquisition, the number of counts related to it are spread over the range of motion. As a result the activity concentration distribution appears larger and its maximum decreases, therefore the lesion volume is overestimated and there is a reduction in the observed SUV[1] (Nehmeh et al., 2002b; Liu et al., 2009).

A measured SUV greater than 2.5 g/ml in FDG-PET (with FBP reconstruction) has been identified as being an indicator of higher possibility of malignancy for certain tumour types (Lowe et al., 1998), therefore underestimating the SUV values could potentially lead to misclassification of lesions, consequently resulting in incorrect diagnosis or, when considering follow-up studies, to inaccurate staging of tumors (Nehmeh and Erdi, 2008). This effect has been found to be more pronounced in small lesions, when their size is comparable to the motion amplitude, in which case the blurring can cause the lesions to be undetected (Daouk et al., 2011).

In combined PET/CT imaging accurate spatial registration of the images produced by the two modalities is crucial, as the latter is used for better localization of the lesions and attenuation correction purposes. In this instance respiratory motion may lead to mismatch between CT and PET datasets, which can result in artefacts in the reconstructed images generated by inaccurate attenuation correction, therefore affecting SUV quantitation. Furthermore, CT and PET acquisition generally have different acquisition durations: the data of the former is acquired much faster (with a standard helical CT for attenuation

[1]The Standard Uptake Value represents the estimation of the concentration of activity in a region of interest, normalized to the patient weight and injected activity. SUV = Activity concentration(kBq/ml)/injection dose(MBq)/patient's weight(kg). Unit of measure = g/ml.

correction taking about 3-4s to cover the whole lung) and in some cases while the patient is in breath-hold state, whereas that of the latter is time-averaged over many breathing cycles. This represents another source of misalignment between the two modalities (Erdi et al., 2004; Nehmeh and Erdi, 2008).

When using PET/CT in radiotherapy treatment planning, overestimating the size of the lesion results in larger target volumes, thus causing healthy tissues to be needlessly irradiated and increasing the probability of treatment-related complications (Keall et al., 2006). The next section presents a review of some of the methods that have been implemented to correct for respiratory motion in PET.

2.3.1 Respiratory motion correction

The first step to generate approximately "motion free" PET data is that of subdividing the detected coincidence events with respect to the breathing state they were acquired in. Consequently, motion correction algorithm can be applied.

2.3.1.1 Respiratory gating

In order to reduce the artefacts induced by the patient breathing, PET data can be acquired in discrete bins in correlation with the respiratory cycle, therefore minimizing the lesion motion within a single bin. This process is referred to as respiratory gating. As a result, multiple PET images are reconstructed, each related to a section of the respiratory cycle, and motion artefacts and image blurring are expected to decrease compared to the non-gated image.

The acquired data can be binned following various criteria, and the main difference between the available methods is whether they are based on time information (*phase gating*) or on amplitude information (*displacement gating*). In both cases the respiratory cycle is divided in a defined number of gates, each corresponding either to a specific phase of the cycle (for phase gating) or to a specific interval in the value of the signal (displacement gating) (Dawood et al., 2007). In Figure 2.8a an example of phase gating is shown: each respiratory cycle is divided in 5 gates and the data corresponding to the same gate will be averaged and is expected to be less affected by respiratory motion; Figure 2.8b displays an example of displacement gating.

Studies have shown, both in phantom experiments and with lung cancer patient data, that the application of respiratory gating reduces the degrading effect of breathing motion in PET images. In Nehmeh et al. (2002b), the PET data of 5 lung cancer patients were acquired into discrete bins in synchronization with the respiratory signal, and the lesion SUV_{max}[2] and volume in the first bin were compared to those of the PET averaged over the breathing cycles. A reduction between 13% and 34.6% in lesion volume and an increase in SUV_{max}

[2]The maximum value of SUV within a lesion.

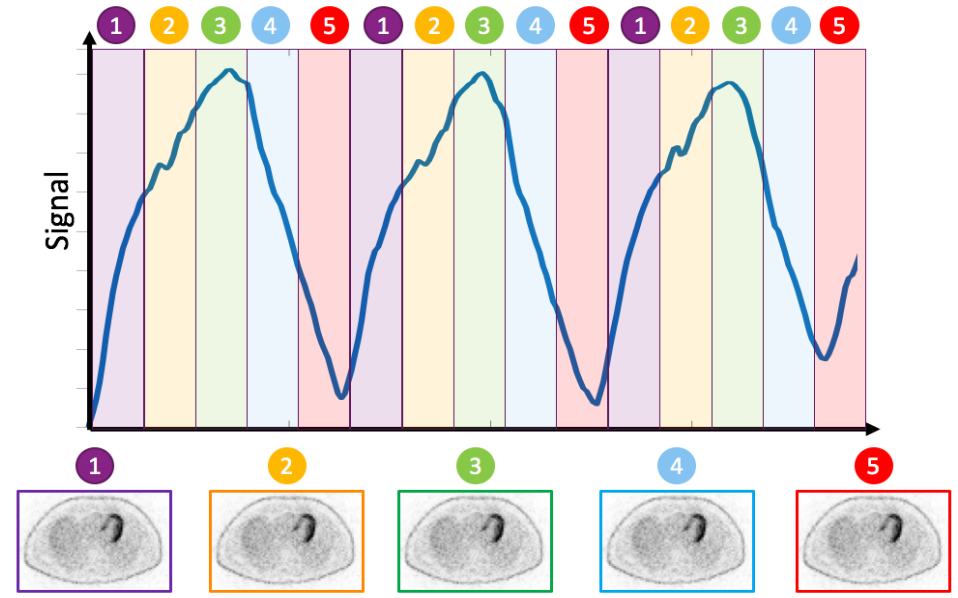

(a)

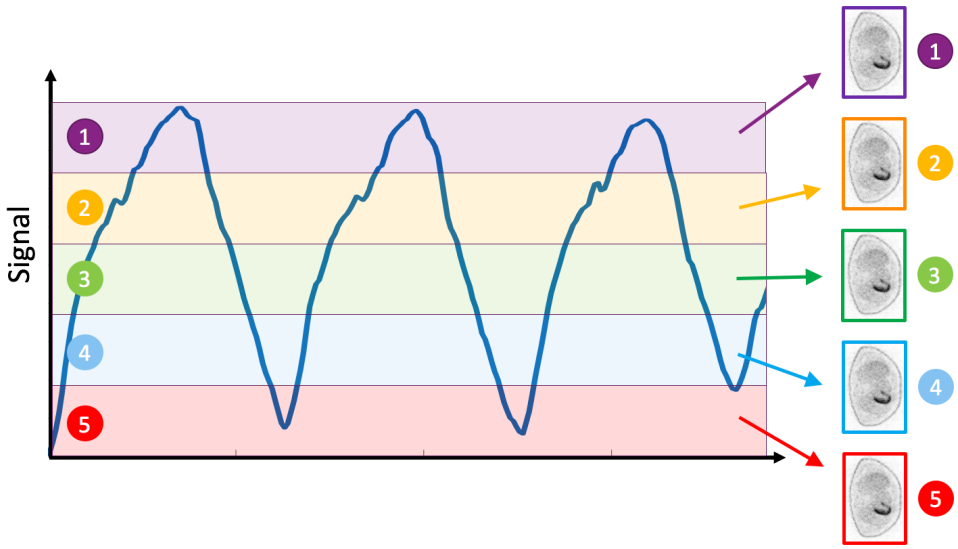

(b)

Figure 2.8: Example of phase (a) and displacement (b) gating.

ranging between 7.46% and 156.16% was achieved when binning the data. In Liu et al. (2009) 1295 patient respiratory traces and the 4D extended cardiac-torso (XCAT) phantom were used to simulate a whole-body PET/CT acquisition. They investigated the effect of motion on lesions of different size and located in the lungs and in the liver, and concluded that respiratory motion can cause on average 28% underestimation of lesion SUV_{max} and 130% overestimation in lesion size.

2.3.1.2 Motion corrected PET image reconstruction

A disadvantage related to the binning of the data is that the reconstructed images are affected by a higher level of noise, as a consequence of the decreased amount of counts utilised for the reconstruction of each gate, whereas techniques that make use of all the available data should result in images of higher quality. There are two main approaches for achieving motion corrected PET images exploiting the whole acquired dataset: one is referred to as RTA (*reconstruct-transform-average*) (Picard and Thompson, 1997; Dawood et al., 2006; Thorndyke et al., 2006) and requires the gating of the data, the individual reconstruction of each gate and then the average of the gates after having registered them to a reference respiratory phase (see next section on registration); the other approach is referred to as MCIR (*motion-compensated-image-reconstruction*) and it requires an estimate of the motion between respiratory gates and includes the motion in the forward/backward transformation operation used in the iterative reconstruction (Qiao et al., 2006; Manjeshwar et al., 2006; Lamare et al., 2006; Li et al., 2006). In both cases the statistics is greatly improved. A further development of the MCIR is represented by the joint estimation of both the deformation parameters and the image parameters from the acquired data within the reconstruction process, as in Jacobson and Fessler (2003); Blume et al. (2010). In Bousse et al. (2016a,b) the probabilistic model of the joint estimation also accounts for the warping of the attenuation data (obtained from CT).

2.3.1.3 Image registration

With image registration one intends the task of matching two or more pictures of the same object, which have been taken for example at different times, with multiple cameras or from different viewpoints. This is often a problem that needs solving in medical imaging, for instance when studies of the same patients acquired on different scanners need to be aligned with respect to each other (i.e. PET and MR exams), or when motion occurred during the data acquisition (i.e. because of patient's respiration). The purpose of image registration is generally to find the optimal spatial transformation (where what is "optimal" often depends on the specific application) between two images, that maps the locations of points in an image to new locations of points in another image. Image registration constitutes a very broad field of study, and reviews of the proposed approaches and available algorithms can be found in Brown (1992) and Zitova and Flusser (2003). For the purposes of the research presented in the next chapters of this thesis, image registration is going to be utilised to find the spatial transformation that aligns respiratory gates to a reference gate.

2.3.2 Respiratory signal detection

In order to perform the gating of the acquired data, a respiratory signal is required. There are several ways of detecting the patient's respiratory motion and they are generally divided

between the methods requiring external devices and those that rely on the PET data only.

2.3.2.1 External devices

In clinical practice, respiratory gating techniques rely on external devices to measure the respiratory motion. Several solutions have been implemented on commercial imaging systems and studies have been carried out in order to assess their performance (Rahmim et al., 2007; Visvikis et al., 2006; Pepin et al., 2014). Common utilised devices are the Anzai AZ-733V system (Anzai Medical Corp., Tokyo, Japan), which is a pressure belt that measures the expansion of the patient's chest (Riedel et al., 2006) and the Real-Time Position Management (RPM) System (Varian Medical Systems, Palo Alto, California, USA), which uses a video camera to track the motion of an infrared marker placed on the patient's abdomen (Nehmeh et al., 2002b; Otani et al., 2010). Alternatively other methods have been proposed that measure the temperature (Boucher et al., 2004) or flow of the air (Guivarc'h et al., 2004) during patient respiration, via a thermistor and a spirometer respectively, then relating those changes to respiratory motion. More recent approaches propose the use of the Microsoft Kinect camera (3D camera system that provides color and depth images simultaneously, via the use of an infrared projector and an infrared sensor) to generate a global motion field for the whole torso surface (Noonan et al., 2012; Heß et al., 2015).

The main aim of respiratory gating is to minimize the effect of motion on the reconstructed images, therefore the signal that is utilised is assumed to be a good estimate of the internal motion of the organs. Several studies have investigated the relationship between the internal motion, which can be tracked with the use of internal fiducial markers such as radiopaque tumor clips (Gierga et al., 2005) or by detecting anatomical landmarks (Fayad et al., 2011), and the external motion that can be detected with devices such as the Anzai belt (Kanoulas et al., 2007), external radiopaque markers (Gierga et al., 2005), the RPM (Vedam et al., 2003) or also from the external surface motion, obtained by thresholding 4D CT images (Fayad et al., 2011). The general conclusion of these studies is that external and internal motion are strongly correlated, even though variability in the correlation might occur depending on the location of the external markers and on the internal region taken into account (Ozhasoglu and Murphy, 2002).

2.3.2.2 Data-Driven respiratory detection methods

There are several reasons why external devices are not optimal: the associated cost, the additional time required for setting up the device (that also generates added dose for the radiographer when dealing with the patient) and the discomfort that the patient might experience. As a consequence, device-less respiratory motion detection methods have been suggested. These methods extract the motion directly from the data, they are referred to as Data-Driven (DD) methods and an extensive review is given in Kesner et al. (2014). Here a

few methods are highlighted to illustrate the range of methodologies that have been proposed in recent years.

The following are methods that require image reconstruction, and rely on either the presence of an external marker placed on the patient or on manual input to select regions of interest: in Nehmeh et al. (2003) an ^{18}F-FDG point source was set on the patient's abdomen and its position used to track respiratory motion through the consecutive reconstructed dynamic frames; in Visvikis et al. (2003) time activity curves are obtained from a Region of Interest (ROI) selected on the series of dynamic images and a Fourier transform is performed in order to estimate the frequency of the motion, assuming periodicity; in Bundschuh et al. (2007) a Volume of Interest (VOI) is manually defined in a summed image around the lesion, the Centre of Mass (COM) of the activity distribution inside the VOI is determined and the z-coordinate (see 2.4a for spatial coordinate definition) of the COM as a function of time delivers the respiratory signal.

Other methods do not require image reconstruction or manual selection of regions in the data, and only rely on the acquired raw data. For example, in Schleyer et al. (2009) the variation of the counts within automatically selected regions subject to respiratory motion is used to estimate the respiratory signal in the Spectral Analysis Method (SAM); in Thielemans et al. (2011) Principal Component Analysis (PCA) is applied to the listmode PET data (unlisted into dynamic sinograms), and the respiratory signal is obtained as the principal component weight factor whose frequency spectrum has the highest peak in the respiratory band; in Wachinger et al. (2012) an alternative dimensionality reduction technique is utilised, *i.e.* Laplacian Eigenmaps (LE) decomposition, in order to obtain the respiratory signal from the raw data; in He et al. (2008) the non-uniformity of the geometric sensitivity of the scanner is exploited as a means of obtaining information on motion, since the count rate for a given organ will depend on the axial location of the organ within the scanner, the respiratory phase is determined from count rate changes in the listmode data (SENS method); in Kesner and Kuntner (2010) a respiratory trace is obtained from the rebinned sinogram by combining together information from the time activity curves extracted from all the voxels of the sinogram (Sinogram Region Fluctuation method, SRF); in Jin (2013) the Centroid of Distribution (COD) for all events is computed with fine temporal resolution from list-mode data and its motion in time provides the respiratory signal. In Büther et al. (2010), the SENS and COM method were also further refined with the combined use of a segmentation process in sinogram space, to only utilise coincidences originating from a ROI delineated around the tumour region, showing improved results.

The use of DD methods has been extensively analysed and their performance compared to that of external devices. In Büther et al. (2016) the impact of using DD methods (namely SENS, COM and their version with the ROI in sinogram space) for respiratory motion

detection in clinical practice is investigated, comparing it to the use of the respiratory belt. Their results show that there are considerable improvements in diagnostic image quality when applying respiratory gating, without there being significant differences between the use of the external device or DD signal.

In Thielemans et al. (2013) the PCA, SAM, SENS and LE methods are applied on patient data and compared to the RPM signal. The SENS method is found to be the worst performing one, the other three provide comparable results under normal conditions and finally PCA provides the best performance when using noisier data. Regarding the PCA method, in Manber et al. (2015) a comparison is shown between its signal and the respiratory signal obtained from an MR-navigator, that represents the internal motion, with a mean correlation of 0.89 over 9 patients.

The PCA method, as implemented in Thielemans et al. (2011), has provided encouraging results when compared to other DD methods, to an external device signal (RPM) and additionally to MR-derived motion signal. For these reasons its thorough evaluation and further development constitutes the body of this work, while also presenting methodologies that could be applicable to other DD methods.

2.3.3 Principal Component Analysis for Data-Driven respiratory gating

PCA is a "dimensionality reduction" technique that attempts to find a linear transformation between the original data and a space of lower dimension, where the structure of the original data can be better observed (Pearson, 1901). PCA represents the data as a combination of orthogonal basis vectors that are ordered with respect to descending variance, such that the first vector has the largest variance and thus represents the maximum observed variation. If applied to PET lung acquisition data, for which the biggest variation in time is generally caused by the breathing of the patient, one of the first Principal Components is expected to represent respiratory motion.

To describe the application of PCA to PET data, the following coordinate system will be used (Bailey, 2005): the radial position r, which is the radial displacement of the Line of Response (LOR) in the transaxial plane; the axial coordinate z; the azimuthal angle ϕ, which is the angle formed between the LOR, projected onto the transaxial plane, and the y-axis of the scanner; the time index t. All coordinates are discretized. The method is applied on dynamic sinograms (i.e. one sinogram per time frame) that have undergone SSRB (Daube-Witherspoon and Muehllehner, 1987). The sinograms are described by the matrix $d(r, z, \phi, t)$, whose dimensions are equal to $N_r \times N_z \times N_\phi \times N_t$, where N_r, N_z, N_ϕ, N_t are the number of radial positions, transaxial planes, angular views and time frames. For the purpose of applying PCA, the matrix is to be considered as a series of N_t realisations of the

acquisition of a sinograms of dimension $N_r \times N_z \times N_\phi$. The empirical temporal mean on the total number of time frames N_t is given by:

$$\bar{d}(r,z,\phi) = \frac{1}{N_t}\sum_{t=1}^{N_t} d(r,z,\phi,t). \tag{2.26}$$

By applying PCA to the data $d(r,z,\phi,t)$, the method provides the following linear expansion:

$$d(r,z,\phi,t) \approx \bar{d}(r,z,\phi) + \sum_{k=1}^{K} p_k(r,z,\phi) w_k(t) \tag{2.27a}$$

$$w_k(t) = \sum_{r,z,\phi} p_k(r,z,\phi)\left(d(r,z,\phi,t) - \bar{d}(r,z,\phi)\right) \tag{2.27b}$$

where $p_k(r,z,\phi)$ are the principal components[3] and $w_k(t)$ are the weight factors, that are the coordinates of the sinogram data described in the new basis, that has K vectors. The relation in Eq. 2.27a would be an equality if the number K of principal components utilised in the expansion were equivalent to the dimensionality of the sinograms space, given by the total number of radial, axial and angular bins. Nonetheless the first 3 components are usually enough to describe the most important features in the data (that in chest PET data is expected to be motion) and their behaviour in time, whereas in practice the subsequent components represent predominantly noise. Therefore in our implementation $K = 3$. For all the patients analysed in the rest of this work the respiratory PC has been found to be among the first 3, so to the best of the author's knowledge there is no evidence that increasing K would have any impact on the method's performance, when applied for respiratory motion detection purposes. The equality in Eq. 2.27b derives from the orthonormality of the Principal Components, which are basis vectors.

PCA can be directly applied to PET dynamic sinograms as they are extracted from scanners however there are some related issues: if the full dimensionality of the sinograms is utilised, the noise in the data hinders the detection of the motion that occurs, and the high spatial resolution implies the need for a considerable large amount of storage and computing time. Furthermore, the full sinograms provide more information than it is required by the method to detect the biggest variations in time, and making use of them would waste computation time needlessly. In our implementation the PET listmode data is therefore unlisted into dynamic sinograms with reduced spatial resolution. The time frame duration is chosen to be 500ms, in order for the cardiac motion to be averaged in each frame, while

[3] In the previous paragraphs the letter p was utilised to refer to the PET projections, whereas from here onward it will be used to describe the Principal Components.

still keeping the respiratory changes between frames. Furthermore, these are processed using the Freeman-Tukey variance stabilisation technique (Freeman and Tukey, 1950) to obtain approximately normally distributed samples, as described in (Thielemans et al., 2011):

$$d(r,z,\phi,t)_{FT} = \sqrt{d(r,z,\phi,t)} + \sqrt{d(r,z,\phi,t)+1} \tag{2.28}$$

then spatial filtering is performed on $d(r,z,\phi,t)_{FT}$ and finally PCA is applied. Note that this variance stabilization process effectively pre-whitens the sinogram data before the application of PCA.

The most respiratory-like signal is obtained by selecting the weight factor $w_k(t)$ whose main frequency most closely corresponds to respiration according to the following procedure (see Figure 2.9): the power spectrum $\mathcal{S}_k(f) = |\hat{w}_k(f)|$ (where $\hat{w}$ is the FFT of the signal at frequency f) is computed for the three weight factors, the ratio between the peak of each $\mathcal{S}_k(f)$ in the respiratory frequency band ($[0.1, 0.4]$ Hz) and the mean power above 0.4 Hz is evaluated and defined as R_k :

$$R_k = \frac{\text{maximum power in [0.1,0.4]Hz}}{\text{mean above 0.4 Hz}} \tag{2.29}$$

The $w_k(t)$ that provides the highest R_k is selected and used as respiratory signal. The PC corresponding to the selected weight factor is defined as the respiratory PC (RPC) and contains information on the respiratory motion.

Since this method is applied to PET acquisitions of the chest, cardiac motion will also be present in the data. DD methods have proven to be capable of detecting cardiac motion (Thielemans et al., 2014; Büther et al., 2009; Feng et al., 2017b), however the temporal resolution of the unlisted sinograms has to be high enough in order to be able to detect changes with cardiac frequencies, that are generally in the range of 1-2 Hz. In this work, the temporal resolution of 500ms is designed to reduce the sensitivity of PCA to the cardiac contraction.

The further development of the PCA methodology for respiratory motion extraction from PET data, and the evaluation of its performance on patient data, constitute the object of this thesis.

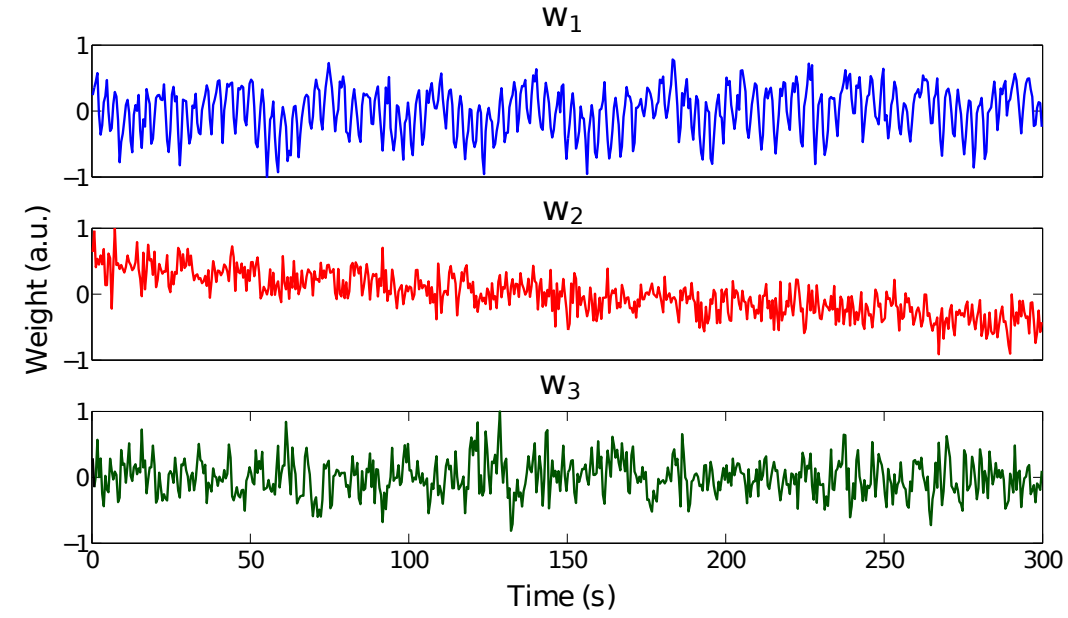

(a)

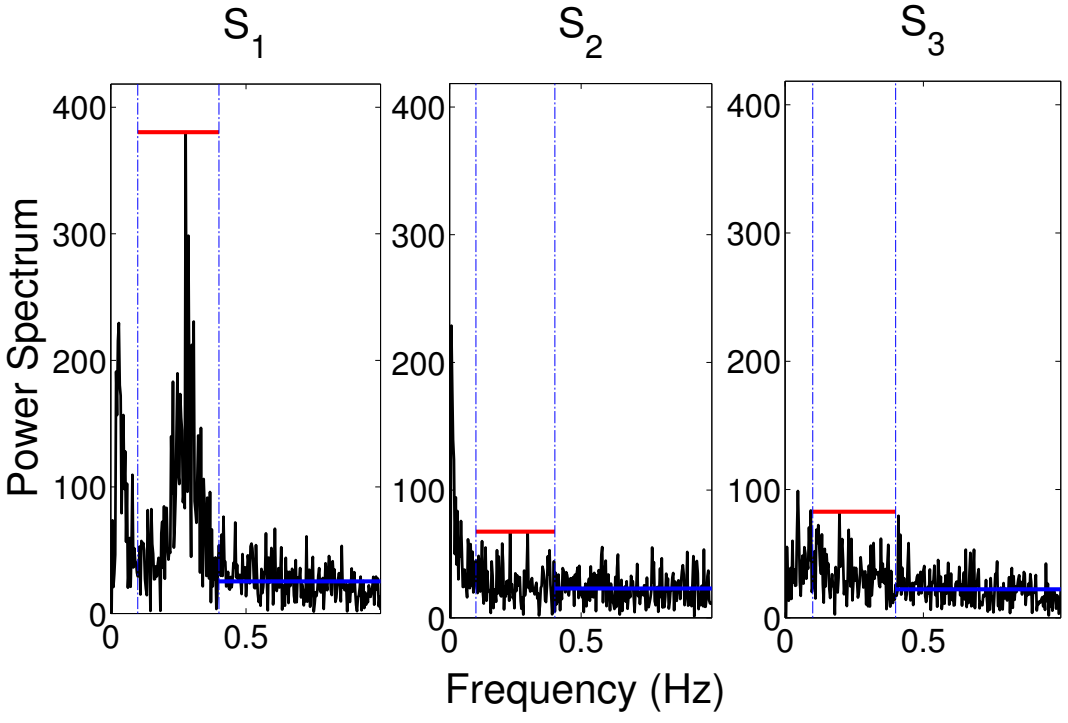

(b)

Figure 2.9: (a) The weights w_1, w_2 and w_3 related to the first three PCs, obtained from one patient dataset. (b) Power spectrum of the 3 weights. The peak within the respiratory frequency band is highlighted in red, the mean power above 0.4 Hz is displayed in blue and the value of their ratio for each weight is: $R_1 = 15.0$, $R_2 = 2.9$ and $R_3 = 3.7$. In this case, w_1 would be chosen as the respiratory signal.

Chapter 3

Sign Determination Methods for the Respiratory Signal in Data-Driven PET Gating

In this chapter, an issue related to the relationship between the physical direction of the patient's respiratory motion and the signals generated by some DD methods is addressed. Methods to correct for it are proposed, focussing on the use of the PCA respiratory signal.

3.1 Introduction

To reduce respiratory motion degrading effects in PET images, data can be gated into (nearly) "motion-free" gates prior to reconstruction by selecting the events with respect to the breathing phase. On current scanners, the respiratory signal needed for this gating procedure is obtained from an external device (such as the RPM and the Anzai belt). As discussed in the previous chapter (section 2.3.2.2) several authors have proposed Data-Driven methods (DD) that can produce a respiratory signal directly from the raw PET data. DD methods thus eliminate the use of external equipment, which is often expensive, needs prior setup and can cause patient discomfort, and they could also potentially provide increased fidelity to the internal movement.

However many of these methods, and in particular SAM (Schleyer et al., 2009) and PCA (Thielemans et al., 2011), provide signals whose direction with respect to the physical motion is uncertain (i.e. their sign is arbitrary), therefore a maximum in the signal could refer either to the end-inspiration or end-expiration phase of the breathing cycle. The possibility of obtaining a signal that represents motion in the opposite direction compared to the truth could lead to inaccurate motion correction, inconsistencies between adjacent beds and to an incorrect matching between PET and CT gates when performing attenuation correction, e.g.

when using an end-expiration CT. In order to correctly identify the gates, the correspondence between the sign of the DD signal and the direction of motion has to be determined and fixed. Figure 3.1 shows an exemplification of the relationship between the DD signal and the direction of respiratory motion, in the case of displacement gating.

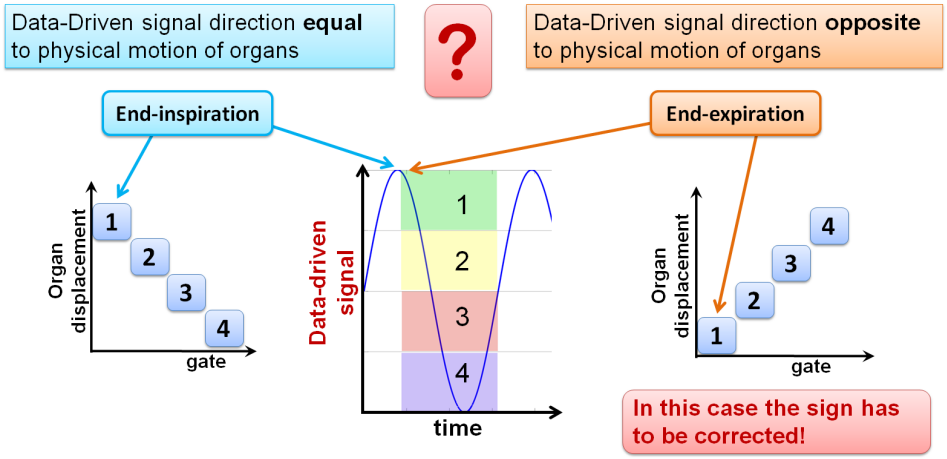

Figure 3.1: Illustration of the issue related to the arbitrary nature of the sign of the DD respiratory signal. Using a signal whose direction is opposite compared to the true respiratory motion would result in an inaccurate relation between respiratory gates and organ position.

In this work the DD method of choice is PCA therefore the specifics of its algorithm are presented here more in detail. As described in Paragraph 2.3.3, the PCA signals are obtained as a result of Eq. 2.27, which is reported here for convenience:

$$d(r,z,\phi,t) \approx \bar{d}(r,z,\phi) + \sum_{k=1}^{K} p_k(r,z,\phi) w_k(t) \tag{3.1a}$$

$$w_k(t) = \sum_{r,z,\phi} p_k(r,z,\phi) \left(d(r,z,\phi,t) - \bar{d}(r,z,\phi) \right) \tag{3.1b}$$

It can be noted that the result of the product $p_k(r,z,\phi)w_k(t)$ will not change if, at the same time, $w_k(t) \to -w_k(t)$ and $p_k(r,z,\phi) \to -p_k(r,z,\phi)$, therefore making the sign of the signal returned by PCA arbitrary, i.e. dependent on noise and implementation details. Since the RPC weight factor (hereinafter referred to as Respiratory Principal Component (RPC) signal) needs to represent the physical respiratory motion, where maxima and minima should correspond to the end inspiration and end expiration phase respectively, the correct sign of $w_k(t)$ has to be determined prior to using it as a gating signal.

In the next sections two methods are proposed to determine the correct direction of the motion represented by the DD respiratory signal, as presented in Bertolli et al. (2017), one of which is only applicable to the use of PCA, whereas the other one can be applied to any

DD method. Their performance is evaluated on patient data and compared to two versions of a published registration-based method, which require image reconstruction.

3.2 Sign determination methods

The purpose of DD methods it that of giving a reliable and accurate respiratory signal, that could potentially replace the use of external motion detection devices. Consequently, the correct determination of the motion direction equally needs to be derived by making use of the data only.

This issue has been previously addressed in Schleyer et al. (2011), where a procedure is proposed to be applied when using the Spectral Analysis Method (SAM, see paragraph 2.3.2.2). In that work, it was assumed that respiration causes a cranio-caudal motion of the internal organs, and the direction of the SAM respiratory signal is determined via registration of gated reconstructed images. In this section, two new methods are proposed, based on the same assumption that axial motion is the major cause of changes in the acquired data. However, these methods only require the PET sinograms and neither gating nor reconstruction are needed.

3.2.1 Registration-based method

The method proposed in Schleyer et al. (2011) performs registration via 1D rigid translation on gated images reconstructed without attenuation correction (to avoid issues with misalignment with the CT). The acquired data are gated with displacement gating (see section 2.3.1.1 for gating methods) based on the DD signal, with equal amplitude gates. The gates are then normalized for detection efficiencies and reconstructed with FBP (with ramp filter). The registration process is performed on the sagittal and coronal Maximum Intensity Projection (MIP) images of the gates, in order to reduce noise. For each gate the displacement in axial direction (z_{coronal} or z_{sagittal}) with respect to a reference gate (chosen as the one with longest acquisition time) is estimated, shifting the former over a range of z-locations from -5 to +5 pixels and minimising the ℓ^2 norm of the difference. The final registration displacement is given as the average of the two. See Figure 3.2 for a depiction of the registration process. Finally, the direction of the motion of the DD signal is determined by looking at:

$$\text{dir} = \text{sign}(\text{z}_{\text{coronal}} + \text{z}_{\text{sagittal}}) \tag{3.2}$$

when this value is negative the DD signal needs to be reversed (simply multiplied by -1), as the motion detected from one gate to the other goes in the opposite direction compared to the amplitude of the signal.

In this work the method was implemented with minor modifications. Each of the obtained gates is in turn taken as the reference for the registration and an average of the

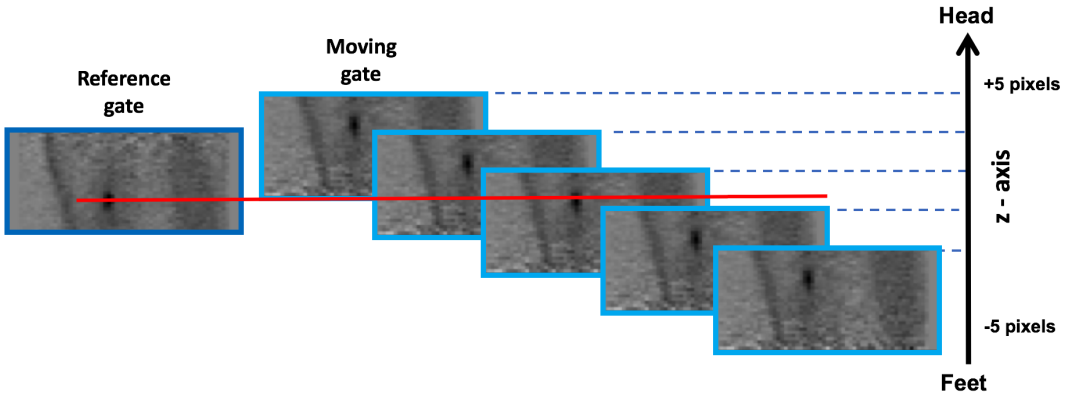

Figure 3.2: The gate to be registered is moved along the z direction across the reference gate. The displacement between the two gates that provides the smallest difference between the images (in terms of ℓ^2 norm) in the sagittal and coronal case, is chosen as the $z_{coronal}$ and $z_{sagittal}$ respectively.

obtained displacements is obtained, rather than choosing as single reference the gate with most counts. The reason for this is that the chosen displacement gating was set to distribute the counts among the gates as evenly as possible to avoid large differences in noise between the gates. In addition, the ℓ^1 norm of the sum of the absolute differences in pixel values was minimised, rather than the ℓ^2 norm, to increase robustness. This process is performed both with sagittal and coronal MIPs (obtaining $z_{coronal,avg}$ and $z_{sagittal,avg}$) and the direction of motion is defined by looking at $dir = sign(z_{coronal,avg} + z_{sagittal,avg})$. If $sign > 0$ the sign of the PC respiratory signal is kept as it is, if $sign < 0$ then the respiratory signal is reversed. A variation of this method was also implemented by registering, as opposed to the MIPs, the images obtained by collapsing the 3D images in the sagittal or coronal plane (equivalent to summing together all the 2D sagittal or coronal planes). The two implementations are referred to as *MIP* and *SUM*. All was implemented in the software Matrix Laboratory, The Mathworks Inc., Natick MA, USA (MATLAB).

3.2.2 Sinogram-based methods

Under the assumption that motion occurs primarily in the z-direction, an approximation for the dynamic sinograms was derived. This approximation then forms the foundation for two sinogram-based sign-determination methods.

The sinograms are considered to be obtained via unlisting the listmode data into dynamic sinograms $d(r,z,\phi,t)$, with a temporal resolution of 500ms, followed by the application of SSRB and radial and angular mashing (i.e. the summation of adjacent sinogram elements with respect to r and ϕ to reduce the dimension).

If it is assumed that motion occurs only as a shift in the the z-direction, therefore the change in the dynamic sinogram can be represented by $d(z,t)$, omitting for simplicity the dependencies on r and ϕ in the following steps. Since the sensitivity of the scanner is constant

in time and depends on z, and the sinograms need to be corrected for it, the change over time of the sinograms normalised for the axial pattern sensitivity is considered:

$$d(z,t) = s(z)n(z,t) \qquad \Longrightarrow n(z,t) = \frac{d(z,t)}{s(z)} \tag{3.3}$$

where $s(z)$ is the sensitivity of the scanner determined from the number of rings and axial acceptance of the scanner) and $n(z,t)$ the normalized sinograms, that are going to be used in the following steps (to be noted that both $n(z,t)$ and $s(z)$ are sinograms with the r,ϕ coordinates as well). The normalization of the sinograms is performed after SSRB. Assuming small translations $\delta z(t)$ with respect to the normalized sinogram in $t=0$, the normalized dynamic sinogram can be described as follows:

$$n(z,t) = n_0(z+\delta z(t)) \quad \text{where} \quad \delta z(0)=0,\ n_0(z) = n(z,0) \tag{3.4}$$

and writing the Taylor expansion up to the first order provides the following:

$$n(z,t) \sim n_0(z) + \delta z(t) n_0'(z) \tag{3.5}$$

with $n'(z)$ the gradient of the sinograms along the z-direction. Subtracting the temporal mean the result is given by:

$$n(z,t) - \bar{n}(z) \sim (\delta z(t) - \overline{\delta z}) n_0'(z) \tag{3.6}$$

where $\bar{n}(z)$ and $\overline{\delta z}$ are the temporal mean of $n(z,t)$ and $\delta z(t)$. The left-hand side term of 3.6 can be obtained directly from the dynamic sinogram, whereas if $n_0'(z)$ were computed from the initial temporal frame ($t=0$) it would be too noisy. Therefore, in the gradient computation, $n_0(z)$ has been approximated with the temporal mean of the entire sinograms $\bar{n}(z)$, consequently decreasing the noise. The gradient along z can be approximated using finite differences and is defined by *GradSino*(z):

$$\mathit{GradSino}(z) = \frac{\partial}{\partial z}\bar{n}(z) \sim \frac{\bar{n}(z+1) - \bar{n}(z-1)}{2}. \tag{3.7}$$

As a result the "component" $n_0'(z)$ in 3.6, that is multiplied by the function describing the axial motion, is approximated by *GradSino*:

$$n(z,t) - \bar{n}(z) \sim \left(\delta z(t) - \overline{\delta z}\right) \mathit{GradSino}(z). \tag{3.8}$$

In order to reproduce the process performed on the sinograms that is ultimately utilised for PCA (see Paragraph 2.3.3), the Freeman-Tukey transformation on the original dynamic sinogram $d(z,t)$ is approximated and a Taylor expansion is applied, given the assumption that

the motion in the z-direction is small. The Freeman-Tukey transformation is approximated as follows:

$$d(z,t)_{FT} \simeq 2\sqrt{d(z,t)} = 2\sqrt{n(z,t)s(z)} \tag{3.9}$$

and $\sqrt{n(z,t)}$ can be written using 3.8:

$$\sqrt{n(z,t)} \simeq \sqrt{\bar{n}(z) + \left(\delta z(t) - \overline{\delta z}\right) GradSino(z)}. \tag{3.10}$$

Assuming that the motion is small, the variation over time in the dynamic sinogram will be small as well and consequently the Taylor expansion of $\sqrt{x+\epsilon}$ can be utilised, yielding:

$$\sqrt{n(z,t)} \simeq \sqrt{\bar{n}(z)} + \frac{1}{2}\frac{\left(\delta z(t) - \overline{\delta z}\right) GradSino(z)}{\sqrt{\bar{n}(z)}}. \tag{3.11}$$

Ultimately the relevant result is that of $d(z,t)_{FT} - \bar{d}(z)_{FT}$, and using the equivalence $d(z,t)_{FT} \simeq 2\sqrt{n(z,t)s(z)}$ and applying the time average to the latter, the outcome is given by the following:

$$d(z,t)_{FT} - \bar{d}(z)_{FT} \simeq \left(\delta z(t) - \overline{\delta z}\right)\frac{GradSino(z)\sqrt{s(z)}}{\sqrt{\bar{n}(z)}}. \tag{3.12}$$

Bringing back the full description of the sinogram for completeness, with the r and ϕ coordinates as well, the final result is that the dynamic sinogram that is given to PCA can be represented by the following (the temporal mean is subtracted):

$$d(r,z,\phi,t)_{FT} - \bar{d}(r,z,\phi)_{FT} \simeq \left(\delta z(t) - \overline{\delta z}\right) WeightedGradSino(r,z,\phi). \tag{3.13}$$

where:

$$WeightedGradSino(r,z,\phi) = \frac{GradSino(r,z,\phi)\sqrt{s(r,z,\phi)}}{\sqrt{\bar{n}(r,z,\phi)}}. \tag{3.14}$$

This description of how the sinogram changes in time compared to its mean (described by $d(r,z,\phi,t)_{FT} - \bar{d}(r,z,\phi)_{FT}$) resembles the result expected from PCA: it represents the dynamic sinogram with a product of a temporal one-dimensional signal with a component in the sinogram space. This similarity can be exploited to compare the assumed motion along the axial direction to the PCA output. The respiratory signal of interest is represented by the term $\delta z(t) - \overline{\delta z}$, see Eq. 3.13, and can be described as the "weight" of the component represented by *WeightedGradSino*.

An example of the result of this procedure on patient data is displayed in Figure 3.3: Figure 3.3a shows the RPC obtained from the application of PCA on a time interval (i.e.

temporal portion) of the acquisition of a patient, Figure 3.3b shows the temporal mean of the sinograms and Figure 3.3c shows the resulting *WeightedGradSino* (where it can be seen that 4 planes are set to zero at the top and at the bottom of the projection, in order to avoid issues with boundary conditions when applying Eq. 3.7). All the sinograms shown here have undergone sensitivity normalisation.

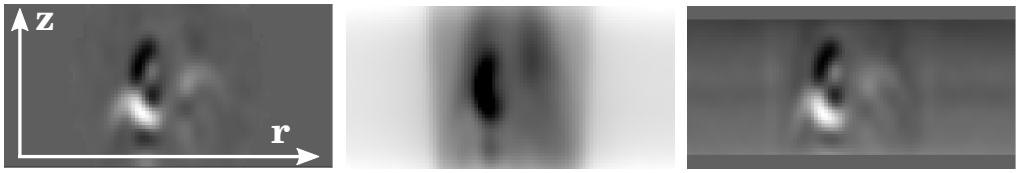

(a) (b) (c)

Figure 3.3: (a) example of RPC for a time interval of the dynamic sinogram; (b) temporal mean of the dynamic sinogram; (c) *WeightedGradSino* obtained from the mean. Grey scale range for (a) is $[-0.05, 0.03]$ and for (c) is $[-0.65, 0.40]$ and for (b) is $[0.07, 0.86]$. RPC, mean and *WeightedGradSino* are 3D sinograms and what is shown here are the projections in the anterior-posterior direction. The horizontal and vertical coordinates in these images refer to the r and z of the sinograms. Truncation in the *WeightedGradSino* is performed to avoid issues with border conditions.

The above approximations motivate two new methods that only require the PET raw data (listmode data that are consequently unlisted into sinograms) to determine the sign of the RPC signal, and these are presented in the following sections.

CorrWeights method The correspondence between the *WeightedGradSino* and the direction of motion is known as it derives from the application of the gradient in the z-direction, see Eq. 3.7, therefore its weight is related to the correct direction of the motion. For a "head-first" patient position, an increase in the signal would correspond to motion towards the head ("up") while a decrease corresponds to motion towards the feet ("down"), thus maxima and minima would correspond to the end-expiration and end-inspiration phase respectively.

As a consequence of 3.13, the *WeightedGradSino* weight can be obtained making use of 2.27b, therefore evaluating the inner-product between the *WeightedGradSino* and the dynamic sinogram on the left-hand side of 3.13:

$$u(t) = \sum_{r,z,\phi} \mathit{WeightedGradSino}(r,z,\phi)\big(d(r,z,\phi,t)_{FT} - \bar{d}(r,z,\phi)_{FT}\big) \tag{3.15}$$

The sums are over all sinogram elements. An example of an RPC signal w and the *WeightedGradSino* weight u corresponding to the same time interval of PET data is shown in Figure 3.4, where high similarity between the two signals can be observed.

In practice, the signal u cannot usually be used as a good respiratory signal since Eq. 3.15 is based on various approximations (*viz.* that the changes in the dynamic sinogram are small and only in the axial direction, see Eq. 3.4 and 3.5). Its importance lies in its

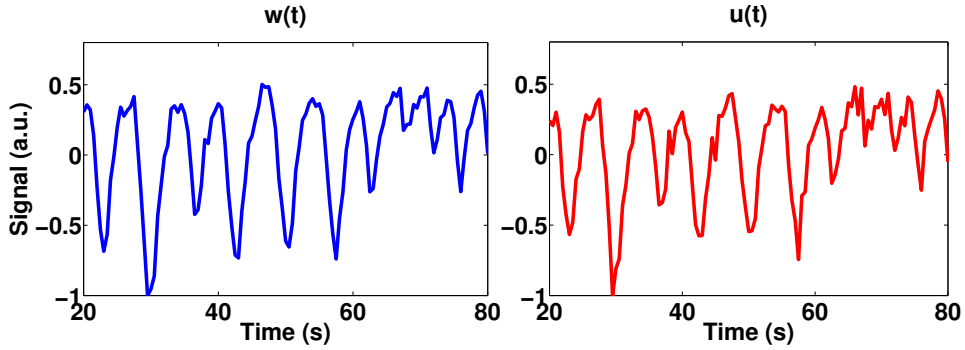

Figure 3.4: On the left the RPC signal, $w(t)$; on the right the *WeightedGradSino* signal $u(t)$, corresponding to the same time interval of the dynamic sinogram.

known relationship with the direction of the axial motion, that makes it a comparison tool to determine whether w correctly represents the direction of the physical motion of the patient.

These arguments lead us to propose the following method to determine the direction of the DD respiratory signal: if the correlation between the two signals w and u is calculated, the sign of the correlation gives an indication about whether the DD respiratory signal and *WeightedGradSino* signal represent motion in the same direction. The Pearson correlation between the two signals is given by:

$$\mathrm{corr}(w,u) = \frac{\sum_t (w(t)-\overline{w})(u(t)-\overline{u})}{\sqrt{\sum_t (w(t)-\overline{w})^2 (u(t)-\overline{u})^2}}. \tag{3.16}$$

If the correlation is negative it is an indicator that w represents a motion in the opposite direction compared to the motion related to u, thus the sign of w (i.e. the respiratory signal used for gating) has to be reversed. For the example shown in Figure 3.4 (where only a 60s interval of the total 300s signals is displayed) the correlation is 0.98, while in other cases it can be lower.

This method is not specifically related to PCA, therefore could be utilised to determine the correct direction of the motion of respiratory signals obtained with any other DD method.

CorrSino method Considering the similarity between Eq. 2.27b and Eq. 3.15, where the latter is expected to hold when the assumptions of small motion along the z-direction are met, it could be argued that *WeightedGradSino* should be similar to the RPC to a certain extent (see Figure 3.3). *WeightedGradSino* could therefore be expected to be representative of the biggest changes that occur in the dynamic sinogram because of respiration, while additionally having a known relationship with the direction of axial motion. For this reason, to check if the sign of the RPC signal corresponds to axial motion in the desired direction, the RPC itself can be directly compared to the *WeightedGradSino*, evaluating the correlation between the two sinograms:

$$\text{corr}(\text{RPC}, \textit{WeightedGradSino}) = \frac{\sum_{r,z,\phi} \text{RPC}(r,z,\phi)\text{WeightedGradSino}(r,z,\phi)}{\sqrt{\sum_{r,z,\phi} \text{RPC}(r,z,\phi)^2 \text{WeightedGradSino}(r,z,\phi)^2}} \quad (3.17)$$

The RPC and *WeightedGradSino* undergo a masking process prior to the evaluation of the correlation, that consists in selecting only the non-zero bins. The absolute value of the correlation relates to the degree of agreement between the RPC and the *WeightedGradSino* sinogram, taking higher values in the case the gradient sinogram adequately represents the motion detected by the respiratory PC. If the correlation is negative the sign of the RPC signal has to be reversed, as in the *CorrWeights* method.

In Figure 3.5 an example is shown of the appearance of the RPC when it has the wrong sign: the features and intensity patterns look similar in *WeightedGradSino* (3.5b) and in RPC (3.5a) and their correlation, from Equation 3.17, is equal to 0.57, whereas 3.5c is the negative version of 3.5a (−RPC) and would be the result of the PCA related to the "wrong" direction of motion. The correlation between 3.5b and 3.5c would result in −0.57 and the method would in this case reverse the corresponding respiratory signal w.

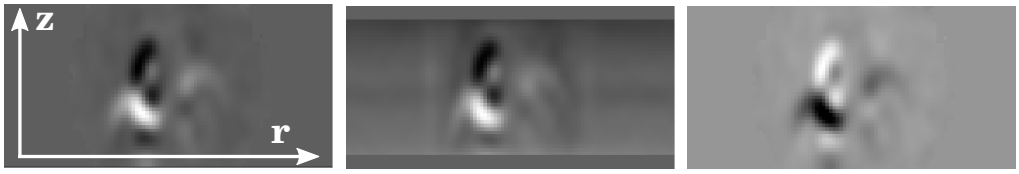

Figure 3.5: (a) and (b) are the RPC and *WeightedGradSino* as in Figure 3.3; (c) is the opposite of (a) and its grey scale is $[-0.05, 0.03]$. It can be seen that (a) is much more similar to (b) (correlation equal to +0.57) than (c) is (correlation equal to −0.57).

This method is specific to PCA, as it relies on the comparison of the RPC to the *WeightedGradSino*.

3.2.3 Data

3.2.3.1 Phantom simulations

As an initial test of the performance of the proposed methods, they were applied on simulated data produced with the XCAT phantom (Segars et al., 2010). In order to make the simulation realistic, the RPM respiratory trace of a normal-breathing patient was used as a template for the phantom motion. Diaphragm maximum displacement was set to 2cm and anterior-posterior maximum displacement to 1.2cm (default in the XCAT parameters). The duration was set to 50s and the temporal sampling to 500ms. Activity and attenuation images were produced. The activity images were projected through the use of the utilities provided by Software for Tomographic Image Reconstruction (STIR) (Thielemans et al., 2012), simulating

3D projections in a GE Discovery STE PET/CT scanner (Teräs et al., 2007), and their projections were attenuated using attenuation coefficients obtained from the attenuation images. Poisson noise was added to the projections, and no random nor scattered events were simulated (total number of counts per frame $\sim 1.2 \times 10^5$, which is similar to half the counts in clinical data, considering there are no random nor scattered events).

3.2.3.2 Patient data

To test the performance of the four sign-determination methods, they were applied on FDG data of oncology patients. Patient datasets were acquired in 3D listmode on GE Discovery STE (21 datasets) and GE Discovery 690 (Bettinardi et al., 2011) (16 datasets) PET/CT scanners using activity levels for routine clinical protocols. The acquisition times ranged from 240s to 720s per bed position and only the lower chest bed positions were selected for this study. The acquisitions were monitored by the RPM device, whose signal was used as the comparative standard for the evaluation process (to be noted that the peaks in the RPM signal correspond to the situation in which the patient's chest is expanded, and therefore the reflective marker positioned on it is at its maximum height). For each patient the selected bed acquisition was also subdivided to independent shorter time intervals of 50, 100, 200 and 300s in order to assess the performance of the methods on smaller amounts of data (i.e. for a 360s acquisition, 7 intervals of 50s, 3 of 100s, 1 of 200s and 1 of 300s were selected).

3.2.3.3 Data processing

The 3D sinograms provided by the scanner (and those obtained from the XCAT simulation) consist of 553 2D sinograms of 329×280 in r, ϕ for Discovery STE and 553 2D sinograms of 381×288 in r, ϕ for Discovery 690, the axial field-of-view is for both scanners 157mm and the listmode data were unlisted into sinograms with time frames of 500ms. Before the application of PCA and consequently the sinogram-based sign determination methods, the sinograms underwent SSRB and radial and angular mashing: the number of summed angles was 40 for the Discovery STE and 32 for Discovery 690 and the number of summed radial bins was 4 in both cases, so that the final data format was $82 \times 7 \times 47$ (radial positions $\times$ angles $\times$ transaxial planes) for Discovery STE datasets and $95 \times 9 \times 47$ for Discovery 690 datasets.

For the registration-based method, the procedure of section 3.2.1 was used and images were reconstructed with voxel dimensions of $5.99 \times 5.99 \times 3.27\text{mm}^3$ for Discovery STE datasets and $5.33 \times 5.33 \times 3.27\text{mm}^3$ for Discovery 690 datasets.

3.2.4 Evaluation

The evaluation followed the following steps:

1. for the selected interval of the patient listmode file and for the projections of the

simulation data, the PCA method described in Thielemans et al. (2011) was used to produce the RPC and its respiratory signal, with a time frame duration of 500ms;

2. the four methods (*CorrSino*, *CorrWeights*, *MIP* and *SUM*) were then applied on the selected intervals and the obtained RPC signal was corrected according to the method's response;

3. the Pearson correlation was evaluated between each of the obtained respiratory signals and the corresponding RPM signal (downsampled to the same time resolution as the RPC signal). When the correlation is positive, the respiratory signal is correct, when the correlation is negative the respiratory signal is opposite to the RPM, therefore the method failed in correcting the sign.

3.3 Results

The application of the presented sign determination methods on XCAT data resulted in each case in the correct determination of the signal, and the correlation values obtained from the application of *CorrWeights* and *CorrSino* were 0.81 and 0.3 respectively (see 3.16 and 3.17).

For the patient data, examples of the appearance of the RPC and *WeightedGradSino* sinograms, and of the RPC signal and $u(t)$ are shown in Figures 3.3 and 3.4. The application of the methods resulted in the failure rates reported in Table 3.1, which includes all the acquired patients and all the different time duration intervals.

Table 3.1: Percentage of failures that were obtained for each group of intervals (that vary in time duration). The numbers in brackets refer to the amount of evaluated intervals for each group.

	Duration (n.intervals)				
Method	50 s (242)	100 s (113)	200 s (44)	300 s (32)	All (431)
CorrSino	5.8	4.4	6.8	9.4	5.8
CorrWeights	4.9	4.4	6.8	9.4	5.3
MIP	21.9	15.0	9.0	12.5	18.0
SUM	14.9	12.4	6.8	6.3	12.8

Table 3.1 shows that methods *CorrSino* and *CorrWeights* have a more stable behavior with respect to the duration of the time intervals, compared to the registration-based methods *MIP* and *SUM* that show much higher failure rates in the 50s and 100s cases. The reduced amount of data that is utilised in such cases affects the quality of the reconstructed gated images, therefore the registration process is more likely to fail.

The DD output depends on the actual presence of respiratory motion and on the assumption that it represents one of the main causes of variation in the acquired data. The quality of the provided respiratory trace, considered as degree of agreement with the RPM

signal, is therefore not predictable and can vary between different patient acquisitions and also between different time intervals of the same study. The same issues exist for the sign determination methods, and their outcome might vary with respect to different acquisitions and different respiratory signals.

In order to better understand the trend of the results with respect to the "respiratory-likeness" of the PCA signal $w(t)$, the failure rates have been analysed with respect to the correlation between $w(t)$ and the RPM signal. The majority of the obtained signals $w(t)$ (401 cases, equal to 93%) are highly similar to the RPM trace, showing correlations ranging from 0.75 to a maximum of 0.99, while in the remaining 7% of the intervals (30 cases) the quality of the PCA signal is sub-optimal, with 21 intervals with correlation included in $[0.5, 0.75]$, 3 cases in $[0.25, 0.5]$ and 6 cases in $[0, 0.25]$. The failure rates corresponding to the 401 cases with DD-RPM correlation values above 0.75 and the 30 cases with DD-RPM correlation values below 0.75 are shown in Table 3.2.

Table 3.2: Failure rates (%) over all intervals with respect to the correlation with the RPM signal.

	Correlation range (n.intervals)	
Method	$[0, 0.75]$ (30)	$[0.75, 1.00]$ (401)
CorrSino	46.0	2.7
CorrWeights	36.6	2.9
MIP	36.7	16.6
SUM	46.7	10.1

To compare in more detail the performance and the reliability of *CorrSino* and *CorrWeights*, the numerical results for each patient were examined. In Figure 3.6 the results obtained by the application of both methods are displayed per patient (for the Discovery STE datasets): each point in the plot corresponds to one evaluated interval and its absolute value corresponds to the absolute value of the methods' outcome (see 3.16 for *CorrWeights* and 3.17 for *CorrSino*). Each point is assigned a positive or negative sign, depending on whether the method succeeded or failed in detecting the correct direction of motion in that specific interval. For each patient, the displayed values correspond to the evaluated intervals ordered from the shortest to the longest in duration, and intervals of the same duration are ordered chronologically (i.e. $[250, 300]$, $[300, 350]$, $[350, 400]$,..).

Figure 3.6 shows that the sinogram-based methods *CorrWeights* and *CorrSino* fail in the same cases (apart from two 50s intervals where *CorrSino* fails and *CorrWeights* succeeds). This could be expected considering that both methods rely on similar assumptions and therefore are potentially vulnerable to cases where the assumptions are not met, e.g. when the biggest changes in the acquired data are caused by bulk motion. Nevertheless it is clearly visible that *CorrWeights* values are considerably higher than those provided by *CorrSino*,

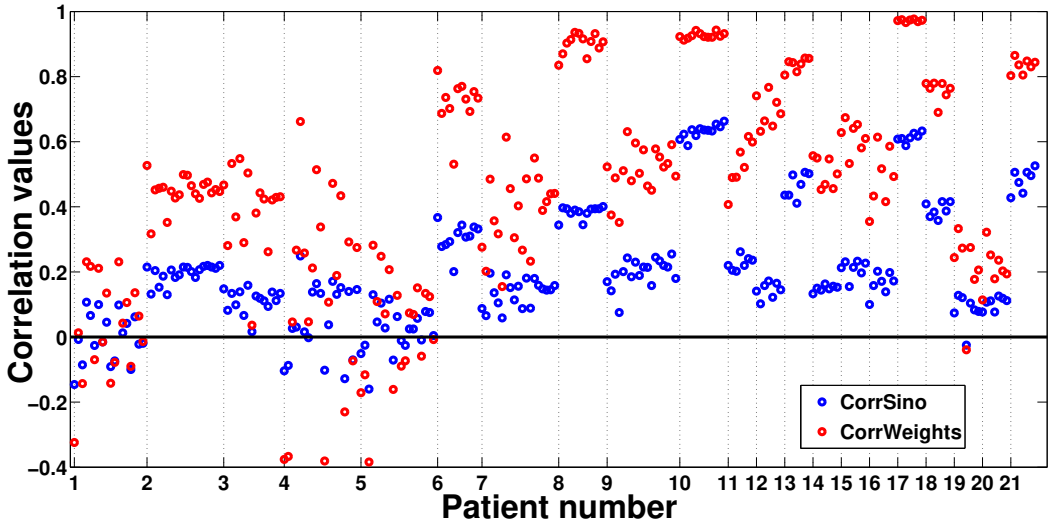

Figure 3.6: Correlation values for *CorrWeights*, see 3.16, and *CorrSino*, see 3.17, for all interval durations, for Discovery STE datasets. When values are negative, the method failed in detecting the correct direction of the RPC signal.

therefore suggesting *CorrWeights* might be considered more reliable. For Patient 1, Patient 4 and Patient 5 both methods failed in several time intervals, possibly due to the patient moving during the acquisition. The plot also suggests that the results can vary considerably from patient to patient. This can be expected considering that the breathing patterns and the activity distribution can be noticeably different between patients and also that interference due to other types of motion can be quite variable.

In order to investigate whether there is a difference between the two methods with regards to variation in values within each patient, the coefficient of variation (defined as the ratio of the standard deviation to the mean) of the output values of both methods for each patient was calculated for the patients displayed in Figure 3.6, disregarding Patient 1, 4 and 5. The mean value of the coefficient of variation has been found to be 0.15 and 0.10 for *CorrSino* and *CorrWeights* respectively. However, the results have been found to be highly patient dependent, with some patients showing higher coefficient of variation with *CorrSino* and some others with *CorrWeights*, and no clear pattern can be extracted from the limited number of analysed patients.

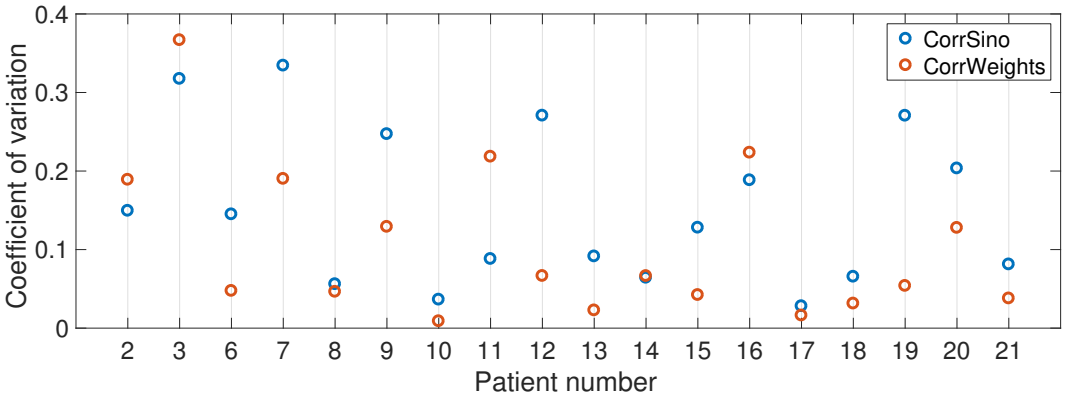

Figure 3.7: Coefficient of variation of the output values of *CorrSino* and *CorrWeights*, which are shown in Figure 3.6, for the patients with no failures (so as to not evaluate standard deviation and mean with negative values). For Patient 19 the only failure was disregarded.

3.4 Discussion

When using the simulation data, all the presented sign-determination methods correctly detected the direction of axial motion, even in the presence of anterior-posterior motion. This suggests that the rationale for the two sinogram-based methods is reasonable. Moreover, the motion in the XCAT simulation data is non-rigid and relatively large (2cm diaphragm movement and 1.2cm chest expansion), therefore it has been shown that the methods, that suppose motion is in axial direction, small and rigid, perform well even in this instance. Nevertheless, the structure and motion of the XCAT phantom is relatively simple and the evaluation of the methods on the patient data is therefore essential.

When applying the methods on the available patients datasets, some failures occurred. *CorrSino* and *CorrWeights* proved to be generally more accurate than *SUM* and *MIP* in detecting the correct sign of the respiratory signal, providing an overall failure rate of 5.8% and 5.3% on the 431 evaluated intervals, compared to 18.0% and 12.8%. The former two methods also have the advantage of avoiding the reconstruction step making use of the raw PET data only. Table 3.1 also shows that the sinogram-based methods are more stable compared to the registration-based ones with respect to the amount of data that is used.

In the majority of the cases in our cohort, which includes the lower chest bed position of FDG scans of oncology patients, the RPC signal was highly correlated with the external RPM signal (93% had correlation higher than 0.75).

The failure rates displayed in Table 3.2 show that all sign-determination methods are more likely to fail as the correlation between RPC signal and RPM decreases, i.e. when the RPC signal is not adequately representing the external motion of the chest, although there is insufficient data (only 30 of the total 431 evaluated intervals) to draw significant conclusions about which method is best in these circumstances. For the cases with correlation

higher than 0.75, the sinogram-based methods are clearly performing better compared to the registration-based methods, both failing in less than 3% of the cases.

The more detailed analysis carried out on *CorrWeights* and *CorrSino* shows that the correlation between the *WeightedGradSino* weight and the RPC signal is generally higher than the correlation between their corresponding sinograms (see Figure 3.6). This indicates that *CorrWeights* could be a more reliable technique than *CorrSino*. Furthermore, the majority of the failures of *CorrWeights* arise when its result has an absolute value below 0.2, suggesting that a threshold could be defined to only use the method when it produces high values and is therefore highly likely to succeed.

In addition, it is worth noting that the *CorrWeights* method can be applied to fix the sign of the signal obtained from any DD method as it does not rely specifically on the use of PCA.

In general, failures in determining the sign of the DD respiratory motion with the presented techniques could be due to several reasons. Low contrast and noise in the data hinder the extraction of a good DD respiratory signal. Minimal motion along the z-direction contradicts the assumption at the core of the investigated methods. Also other types of patient movements during the acquisition might contaminate the RPC signal. Figure 3.6 clearly shows that for the sinogram-based methods the strongest variability in the outcome is found between different patients and also that the duration of the interval of data taken into account does not have a noticeable impact.

In our implementation, all of the data in the sinograms are used both for the extraction of the respiratory signal and for the determination of its sign. Selecting regions of interest (ROI) known to include respiratory motion might potentially improve both the quality of the signal and the performance of the methods. However, defining ROIs would usually require reconstruction of the data, as anatomical features are difficult to detect in sinogram space, and would therefore require additional processing time and external input from trained personnel. Alternatively, image reconstruction could be avoided by making assumptions about the part of the sinogram which is likely to exhibit respiratory changes, as shown in Schleyer et al. (2009), where the utilised part of the data are selected through an iterative process aimed at filling a predetermined (i.e. user defined) proportion of the whole dataset. As the goal of our work is to provide a methodology which is exclusively dependent on the acquired PET raw data, this has not been put into practice. Developments regarding other methodologies for the pre-selection of parts of the data are addressed in the next chapter.

The evaluation on patient data of the proposed methods has been performed via comparison with the RPM signal. This has been shown to potentially exhibit a time lag when compared to the internal organ motion (Ionascu et al., 2007), while DD methods are expected to provide signals that are closely representative of the latter. However, as this time lag is

expected to be small (below 0.2s in Ionascu et al. (2007)), it is not expected to undermine the conclusions of this work, as it would only decrease the absolute value of the evaluated correlation between RPM and RPC signal but not its sign. Nevertheless a limitation of this study is that it relies on the RPM signal correctly representing the respiratory motion of the patient. A poor correlation between RPC signal and RPM could therefore be due to a deficient RPM rather than inaccurate RPC signal. However, the RPM signal is currently utilised in clinical practice and is usually considered to be a reliable method.

As for potential irregularity of the breathing in time, the methods are not expected to be affected by it, as long as the respiratory motion is similar throughout the acquisition. What could potentially change is the value of the ratio shown in Figure 1 for an example of PCs: if the breathing is irregular the peak in the frequency band would be expected to decrease, and so would the resulting ratio value.

3.5 Conclusion

In this Chapter two new methods were presented, *CorrWeights* and *CorrSino*, for the determination of the direction of the DD respiratory signal obtained with PCA. The new methods provided lower failure rates compared to the previously published registration-based methods. When the extracted respiratory signal is highly correlated with the RPM they fail in less than 3% of the cases, compared to more than 10%. Among the two proposed methods, *CorrWeights* is not specifically related to PCA, therefore it could be utilised to determine the correct direction of the motion of respiratory signals obtained with any other DD method. Considering its simple implementation, it could easily be applied together with any DD method in clinical practice.

In the next Chapter further refinements of the presented methods, together with their extension to TOF data, will be presented.

Chapter 4

Time-of-Flight PCA and sign-determination methods

In the previous chapter the issue of the sign-determination for some DD methods has been addressed and methods to correct for it have been presented. In this chapter, the PCA methodology and the proposed sign-determination methods are extended to TOF data.

4.1 Introduction

With the advent of TOF PET scanners more spatial information has become directly available from the raw data. By exploiting the TOF value it is possible to improve the localization of the PET annihilation event along the LOR (see section 2.1.3), and this increased ability is expected to improve the performance of Data-Driven methods in motion detection. Some methods that do utilise the TOF information are presented here.

In Xu and Tsui (2015) the DD method of choice is COM in sinogram space, and they extend to 3D the methodology presented in Kao (2008), where the TOF information is utilised to perform ROI image reconstruction of 2D PET data. In Xu and Tsui (2015) the TOF information is utilised to generate the non-TOF PET data of comprised within a user-defined VOI, that is weighted to decrease the stationary background and increase the amount of detected motion. Their method is effective in decreasing influence of the background but has not been adequately tested, as the results are shown for one patient only. In Feng et al. (2017a) firstly Maximum likely annihilation (MLA) is used to map each listmode event to 3D volume space, where the location of the annihilation event are estimated from the coordinates stored in the listmode file, exploiting the TOF information to locate the event along the LOR. Secondly, the best VOI for the COM respiratory signal extraction is selected through an iterative process, that maximises the SNR or the generated respiratory signal. The use of the selected VOI with COM improves the signal compared to when no

VOI is applied and provides signals highly similar to the output of external motion trackers (mean correlation value over 8 patients equal to 0.85). In Wang et al. (2015) 3D TOF PET data is rebinned into 2D non-TOF PET with the FORET[1] algorithm (Cho et al., 2009) and COM and PCA are compared on one patient dataset. The performance of both methods improve when compared to the use of non-TOF data. In Ren et al. (2017) the COD method presented in Jin (2013) (mentioned in section 2.3.2.2) is updated to the use of TOF PET data, improving the determination of the 3D coordinates of the annihilation events. The method is applied on 10 abdominal patient studies, and the generated respiratory signals provide a mean correlation with the Anzai belt signal of 0.79.

In the next sections the extension of PCA the and sign-determination methods to TOF PET data is presented, together with the illustration of further data processing steps to improve their performance. Lastly, an illustrative example of the benefit of using TOF information is provided, using a simulation study. A thorough evaluation on patient data is the topic of the next chapter.

4.2 Methods

4.2.1 TOF unlisting

In order to apply PCA as described in 2.3.3 the listmode files need to be unlisted into sinograms. This process bins together all the events that were acquired within a determined time interval into the corresponding time frame. With TOF data, for each event additional information is available: the difference in arrival time of the 2 annihilation photons. This added information enables the location of the annihilation event along the LOR, placing it in the appropriate TOF bin (see Figure 2.6 in section 2.1.3. Different events can therefore be distinguished from one another by their TOF value.

In our implementation of the unlisting process, the events are split into the 500ms time frames, and subsequently into a number of TOF bins which is related to the TOF resolution of the scanner. The end result is a dataset given by a 3D sinogram for each TOF bin for each time frame, see Figure 4.1. In the rest of this thesis, the dataset composed of a 3D sinogram per TOF bin will be referred to as *TOF sinogram*. When considering the series of TOF sinograms in time, the temporal coordinate t will be added and it is referred to as *dynamic TOF sinogram*.

As previously mentioned, PCA does not require the full-size sinograms to successfully detect respiratory motion. For each TOF bin and time frame, the same process as in the non-TOF case is performed: the 3D sinograms undergo non-TOF SSRB and consequently radial and view mashing. Furthermore, mashing in the TOF dimension is also applied, to

[1] FOurier REbinning of Time-of-flight data to 3D non-time-of-flight

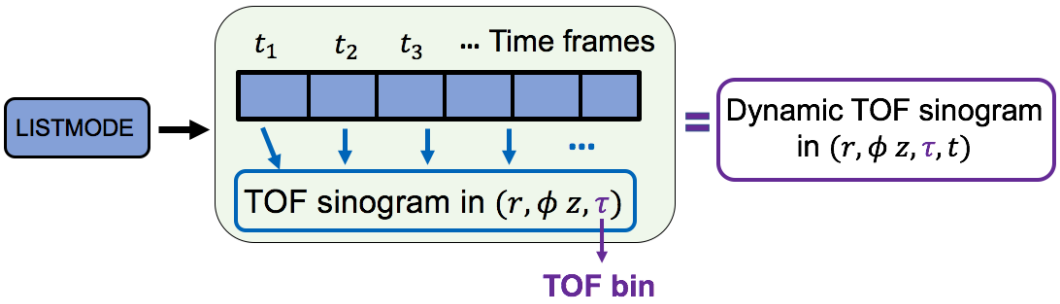

Figure 4.1: Illustration of the TOF unlisting process. One *TOF sinogram* is generated per time frame. When a dynamic series of *TOF sinograms* is considered together, it is referred to as *dynamic TOF sinogram*.

reduce computational demand, by summing together a number of adjacent TOF bins (more details will be given in the patient data description). Finally PCA is applied:

$$d(r,z,\phi,\tau,t) \approx \bar{d}(r,z,\phi,\tau) + \sum_{k=1}^{K} p_k(r,z,\phi,\tau) w_k(t) \tag{4.1a}$$

$$w_k(t) = \sum_{r,z,\phi,\tau} p_k(r,z,\phi,\tau) \left(d(r,z,\phi,\tau,t) - \bar{d}(r,z,\phi,\tau) \right) \tag{4.1b}$$

4.2.2 Sign-determination methods

The methods described in the previous chapter (and published in Bertolli et al. (2017)) can be directly extended to TOF sinograms as the TOF information is hardly affected by axial motion. The steps detailed in 3.2.2 are performed on the TOF sinograms, after SSRB and spatial downsampling, and result in a factor $WeightedGradSino(r,z,\phi,\tau)$ that has an extra dimension given by the TOF coordinate. Its signal in time is therefore given by the following TOF-version of Eq. 3.15:

$$u(t) = \sum_{r,z,\phi,\tau} WeightedGradSino(r,z,\phi,\tau) d(r,z,\phi,\tau,t) \tag{4.2}$$

The methods *CorrWeights* and *CorrSino* are applied as previously described.

4.2.3 Cylindrical masking

Since PCA detects the biggest changes that occur in the data, selecting the parts of the sinogram that more likely include respiratory motion has the potential to decrease the amount of detected non-respiratory change (e.g. bulk motion, arm motion). Therefore, by more specifically targeting the desired area of the body, both the quality of PCA signal and the performance of the sign-determination methods are expected to improve. Given the added spatial information available with TOF data, the selection process is foreseen to

be particularly effective, and to achieve even greater localisation of the targeted areas. In previous work, some authors have shown the positive effect of selecting an ROI in image space for DD methods, both on reconstructed images (Büther et al., 2010) and in the space of the back-projected events (Feng et al., 2017a). Additionally, specific areas in the PET data, such as the whole lung region or a specific tumour, can also be selected with more refined segmentation algorithms, that exploit the morphological information of the CT image alone or together with the reconstructed PET image (Han et al., 2011; Mokri et al., 2012). Here we suggest a simpler approach that does not require either image reconstruction or back-projection, and aims to automatically disregard the outer parts of the sinograms before applying PCA.

The selection of the data to be given as input to PCA (for the respiratory signal extraction) and for the creation of the *WeightedGradSino* and its signal $u(t)$, is performed as follows:

- A cylinder is created in image space with value 1 and 0 outside, placed in the centre of the FOV and with diameter of 20cm.
- Non-TOF and TOF forward-projection are applied to the cylinder image to obtain the areas in the sinogram that are comprised within its volume. A binary mask is then created assigning a value of 1 to all the non-zero elements in the sinogram and 0 elsewhere.
- The non-TOF and TOF sinograms data are masked with the corresponding cylinder sinogram prior to using PCA.

A cylindrical volume of 20cm diameter is supposed to include the central part of a normal-sized patient and effectively removes the arm area, when the patient is acquired with arms-down. This process can be easily and automatically applied on any patient dataset, without requiring any manual input or additional data processing steps. An example of the result of this masking step is shown in Fig.4.2: for a patient sinogram that was unlisted into 5 TOF bins the central and one of the outer TOF bins are displayed without and with the application of the mask.

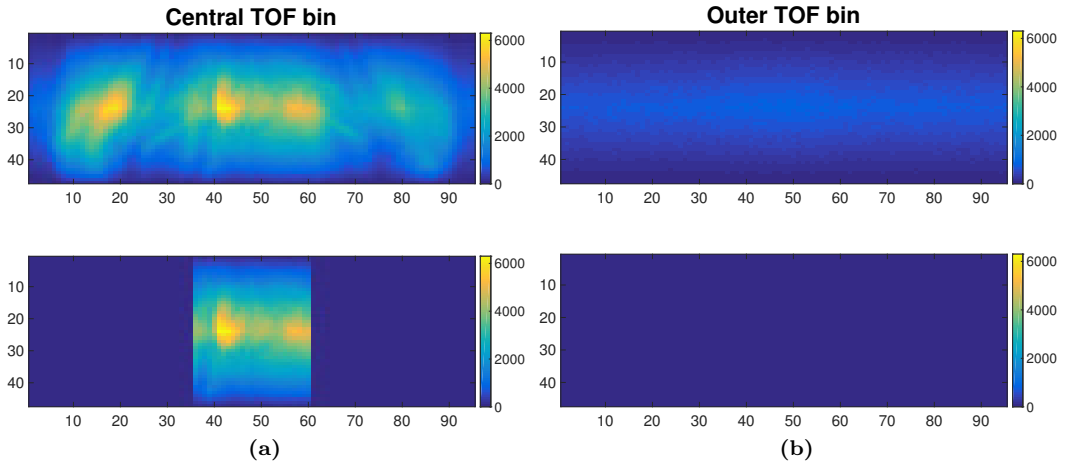

Figure 4.2: (a) Central TOF bin without and with the masking on the top and bottom row respectively. It can be clearly seen that the PET counts in the areas at the borders that correspond to the patient arms are removed when using the mask. (b) In the outer bin the masking process removes additional counts that are not related to the central area of the patient.

4.3 Explanatory example with simulation data: two cylinders with different motion

A simulation study has been performed in order to show the performance of PCA on TOF data. The aim is to demonstrate that, with the added spatial information provided by the TOF capability, it is possible to discriminate between motions that occurred in different parts of the FOV to an extent that is not attainable with non-TOF data.

4.3.1 Simulation design

The simulation includes two separate objects of considerably different sizes, both moving in the z-direction following a sinusoidal function, with same frequency but different phases. They are located far apart from each other and the challenge is that of retrieving the motion of the smaller object.

The simulation is designed as follows:

- Objects simulated with the `generate_image` utility provided by STIR:
  - one cylinder (5mm radius and 5mm width) is placed in the centre of the FOV (axially and transversally);
  - one bigger cylinder (35mm radius and 35mm width) is placed in the centre of the FOV axially, and 25cm off-centre to the right transversally, see Figure 4.3.

  Both cylinders had their axis perpendicular to the z-axis of the scanner. The generated images had $233 \times 233 \times 47$ voxels, with $3 \times 3 \times 3.27$ mm^3 voxel size.

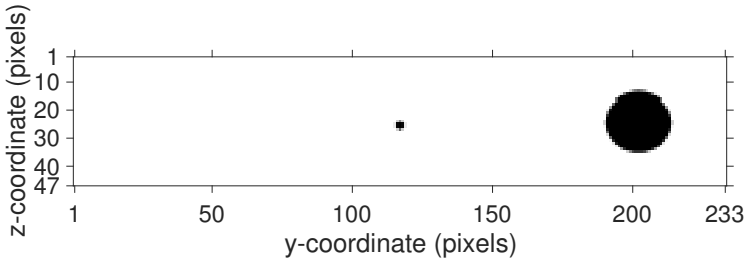

Figure 4.3: Frontal view of the simulated objects: the smallest cylinder is placed in the centre of the FOV and the biggest one off-centre to the right. They move in the z-direction only, with motion with identical amplitude.

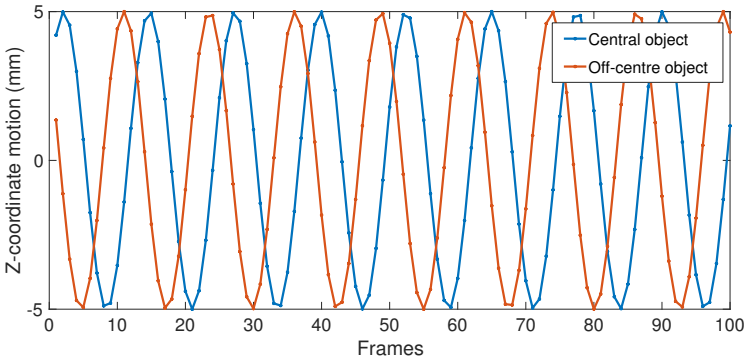

Figure 4.4: Motion of the z-coordinate of the two objects.

- Simulated motion: both cylinders move only in the axial direction, following a sinusoidal path (ranging from -5 to +5 mm from the centre) with same frequency and different phase, see Figure 4.4. The motion is simulated for 100 frames.

- Data projection simulation:

  - the `create_projdata_template` utility from STIR is utilised to create a template for the sinogram as it would be created by the Discovery 690 PET scanner, for 2D projections (i.e. the maximum ring difference taken into account is equal to 1). The TOF resolution of the scanner is also simulated, which is equal to 550ps (Bettinardi et al., 2011). The TOF bin width is 89ps therefore the total number of TOF bins is 55 (see Eq. 2.9).

  - the `forward_project` utility from STIR is utilised to simulate the projection data, both with the non-TOF and TOF projector.

4.3.2 Application of PCA

For each simulated frame, the generated sinograms underwent spatial downsampling (radial and view mashing), and the TOF sinograms also had TOF bin mashing (summing 5 adjacent bins together) and their final dimensions are: $95 \times 16 \times 47$ (radial positions $\times$ angles $\times$ transaxial plane) and $95 \times 16 \times 47 \times 11$ for non-TOF and TOF sinograms respectively, where

the latter ones have the TOF bin dimension (with TOF value ranging from -6 to +6, where bin 0 is the one in the centre of the LOR).

PCA was then applied on the non-TOF and TOF sinograms, as described in Eq. 2.27, with $K = 5$. The biggest change in the data in this simulation is evidently generated by the motion of the bigger cylinder, as its motion is the same as that of the other object, but its volume much bigger. Therefore, to evaluate whether it is possible to avoid the detection of the big cylinder's motion and only retrieve that of the small one, a selection in sinogram space is applied before using PCA:

- radial selection: out of the 95 radial positions in which the sinogram is subdivided, only those between the 37^{th} and the 57^{th} are utilised;
- TOF bin selection: in the TOF sinograms, additionally to the radial selection, only the central TOF bin (number 0) is selected.

The aim of the selection is to disregard the contribution of the big cylinder to the projection data. Figure 4.5 shows the effect of the radial and TOF bin selection in the case of the simulated objects. Of the 16 angular views the sinograms are composed of (after mashing), three are shown as example. It is possible to see that when only radial selection is applied, the off-centre object is still captured by LORs at certain angles, whereas when simultaneously selecting the central TOF bin the goal is more effectively achieved, eliminating most of the off-centre object contribution.
After the application of PCA, the obtained signals (the weights of the Principal Components) were compared with the ground-truth motion signals of the two moving objects (see Figure 4.4).

4.3.3 Results

The weights of the 5 PCs obtained with non-TOF and TOF sinograms, and without and with the application of the selection process, were compared with the ground-truth motion signals of the small and big cylinders. It is worth remembering that the Principal Components are ordered with respect to descending variance, meaning that the first will be the one representing the biggest variation in the data, and the following ones will represent effects with gradually decreasing relevance.

In Tables 4.1 and 4.2 the correlation values are shown for the non-TOF and TOF data case respectively. When the correlations values were below 0.01 they were reported as ~ 0. Table 4.1 shows that when no selection is applied, the first PC represents the motion of the big cylinder, as the correlation value of its weight with the sinusoidal motion utilised in the simulation for that object is equal to 0.99. The motion of the small cylinder is instead captured predominantly by the 5^{th} PC, showing a correlation of 0.93. With the application

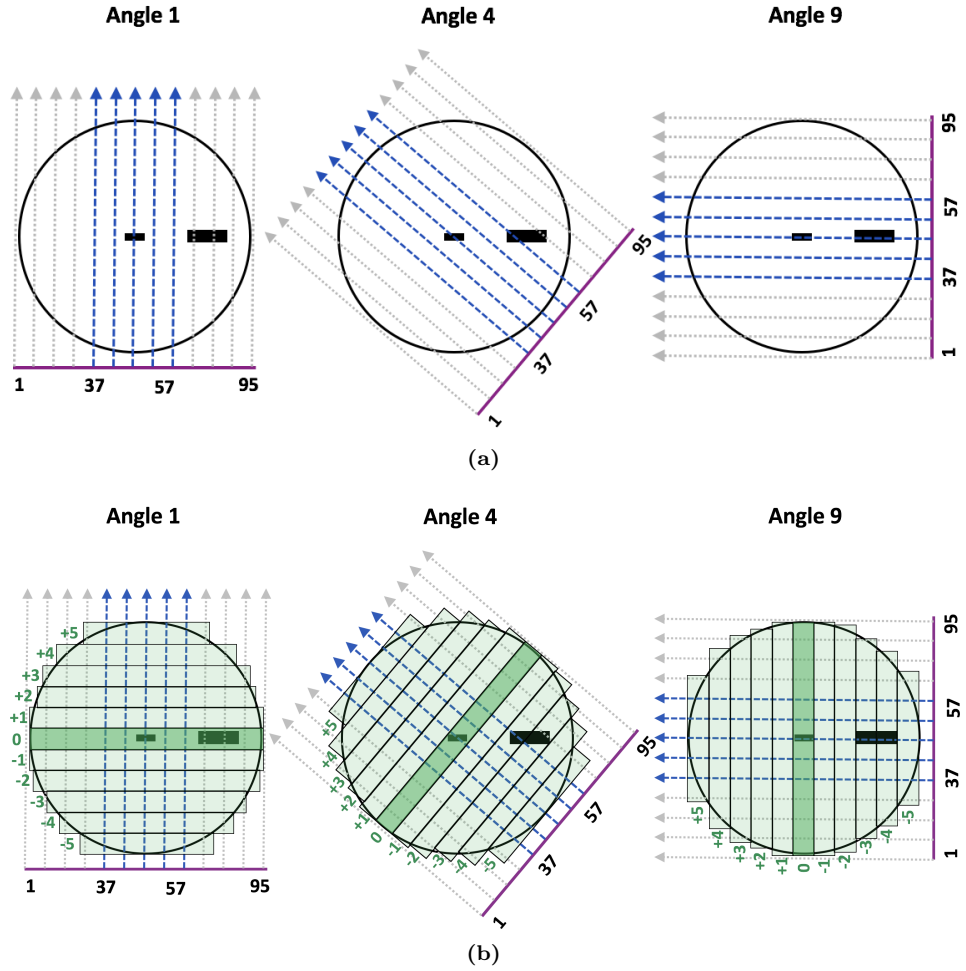

Figure 4.5: (a) Illustration of the effect of the radial selection on the projections, where the blue lines are the LORs included in the selection and the grey ones are disregarded. With respect to some angles, the radial selection successfully disregards the LORs that cross the big cylinder, nevertheless for some other angles it is not sufficient to completely disregard its contribution to the projection data. (b) Effect of the radial selection together with the selection of the central TOF bin only. To be noted that for the example at *Angle 9* the radial selection does not succeed in disregarding the outer object, whereas with the simultaneous selection of the TOF bin, only the central object is is within the selected area. The depiction of the TOF bin is here only explanatory, for the purpose of showing the effect of TOF bin selection, and it is not spatially accurate.

of the selection process, the big cylinder's motion is still represented by the 1st PC, therefore still appearing as the biggest change in the dataset, whereas the signal representing the motion of the small cylinder shows the maximum correlation when compared to the weight of the 3rd PC.

Table 4.2 shows that with no sinogram selection the result is almost identical with the use of TOF or non-TOF data. However, if the radial and TOF bin selection are applied, the motion of the central cylinder, even though much smaller than the off-centre one, is perfectly represented by the first PC, whereas the motion of the big cylinder is depicted by the last

PC.

See Figures 4.6 and 4.7 for the visual comparison of the true motion of the objects and the PC weights.

Table 4.1: Values of the Pearson correlation evaluated between the weights of the PCs and the ground-truth motion of the small (central) and big (off-centre) object, when using non-TOF data without and with the radial selection.

	non-TOF without selection					non-TOF with selection				
	PC 1	PC 2	PC 3	PC 4	PC 5	PC 1	PC 2	PC 3	PC 4	PC 5
Small	0.30	∼0	0.02	0.20	0.93	0.30	∼0	0.87	0.37	0.05
Big	0.99	∼0	0.04	0.01	∼0	0.99	∼0	∼0	0.03	0.01

Table 4.2: Values of the Pearson correlation evaluated between the weights of the PCs and the ground-truth motion of the small (central) and big (off-centre) object, when using TOF data without and with the radial and TOF bin selection.

	TOF without selection					TOF with selection				
	PC 1	PC 2	PC 3	PC 4	PC 5	PC 1	PC 2	PC 3	PC 4	PC 5
Small	0.30	∼0	0.02	0.25	0.92	0.99	∼0	0.12	0.04	∼0
Big	0.99	∼0	0.04	0.13	∼0	0.30	∼0	0.03	0.12	0.95

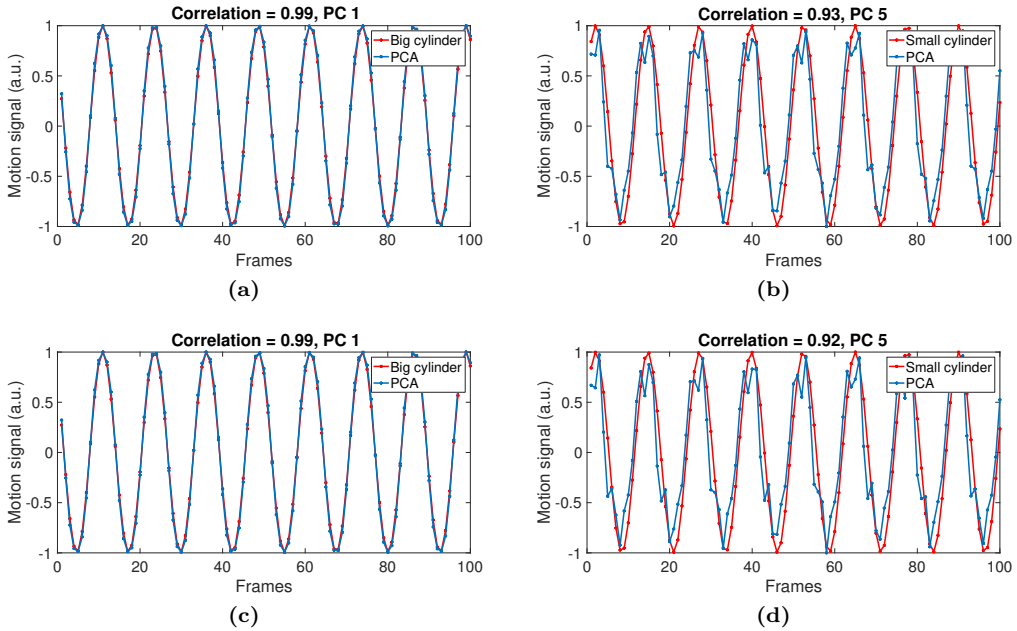

Figure 4.6: In (a) and (b) the comparison between the ground-truth motion signals of the big and small cylinder with the PC weights that represent them the better, obtained from non-TOF data without the radial selection. In (c) and (d) the TOF case, without radial and TOF bin selection.

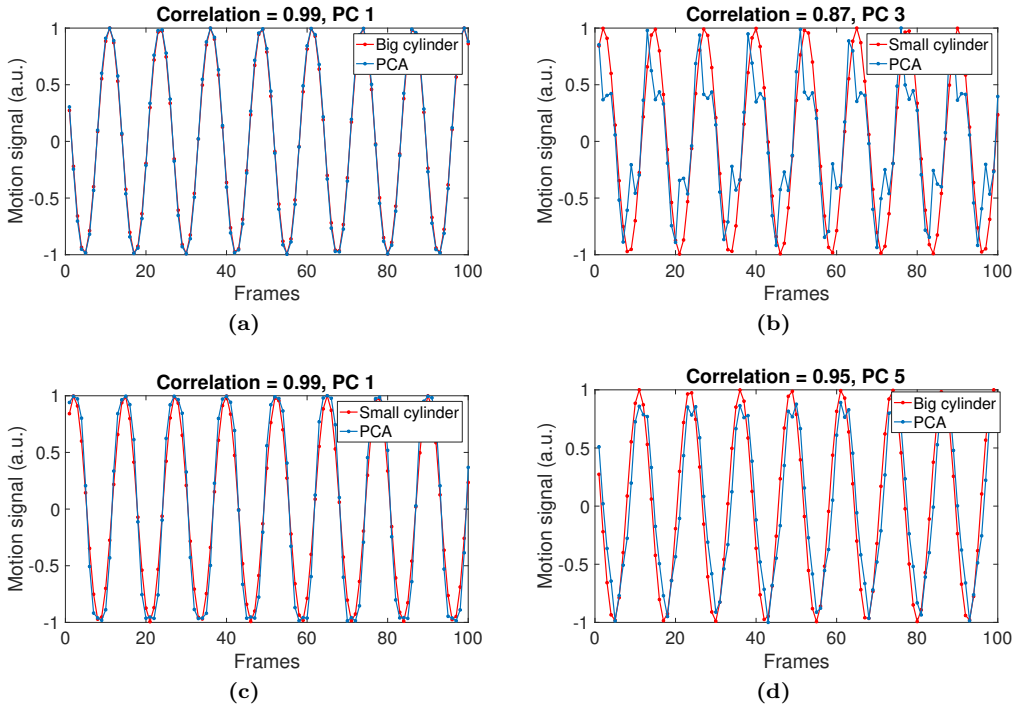

Figure 4.7: In (a) and (b) the comparison between the ground-truth motion signals of the big and small cylinder with the PC weights that represent them the better, obtained from non-TOF data with the radial selection in sinogram space. In (c) and (d) the TOF case, with the radial and TOF bin selection. To be noted that in the TOF case, when applying the selection, the 1^{st} PC represents the motion of the small cylinder, whereas the motion of the off-centre big cylinder is described by the 5^{th} PC.

4.3.4 Discussion

The correlation values in Tables 4.1 show that in the non-TOF case the motion of the big cylinder is always the biggest change in the projection data, even when the selection in sinogram space is applied. In the latter case, the only notable effect is that the change in the data generated by the small cylinder "gains" in importance, as it is better described by the 3^{rd} PC as opposed to the 5^{th}. Nevertheless, the greater cause of change in the data is still represented by the off-centre object.

When using TOF data without radial and TOF bin selection the results are the same as in the non-TOF case. When the radial and TOF bin selection are applied, the 1^{st} PC describes the motion of the small cylinder, whereas the motion of the bigger and off-centre object is represented by the last PC. This proves that applying the selection effectively boosts the contribution to the change in the data given by the motion of the object in the area of interest, even when its effect is in reality smaller than that of other objects.

The motion signals of the two objects were retrieved both with non-TOF and TOF data, and also without the selection process. What is relevant in the TOF case is that, when the radial and TOF bin selection are applied, the motion of the object of interest becomes the

most importance source of change in the data, whereas the change happening in the outer areas contributes to a smaller degree. When only using radial selection with non-TOF data, the first PC describing the biggest variance in the data has been found to always be the one representing the big cylinder's motion, therefore the selection process failed in achieving the desired goal. Furthermore, this example clearly showed that when no masking is applied, as there are two sources of considerable change in the data, it might be challenging to correctly retrieve the correct shape of the motion of the smaller object. This can be seen from the signal in Figures 4.6b and 4.6d where, even though the correlations are high, the true shape of the motion signal is not entirely recovered. The extent of this effect is going to depend on the sizes of the moving objects, on the range of motion and on the contrast.

To be noted that the effect of the use of TOF sinograms and masking process on the sign-determination methods has not been evaluated here, and also that no noise was added to the simulated data therefore its impact on the results has not been analysed. Both aspects are going to be thoroughly investigated on patient data in the next chapter.

This illustrative example shows an example of what can be achieved with the additional use of the TOF information, and the obtained results constitute the reason why the cylindrical masking described in 4.2.3 is suggested for respiratory motion detection in PCA.

4.4 Conclusions

In this chapter the use of TOF information in respiratory DD methods has been discussed, and the extension of PCA and the sign-determination methods to TOF sinograms has been presented.

Furthermore, a simulation study has been performed to illustrate the benefits of exploiting TOF information, when aiming to retrieve the motion that occurs in a specific part of the FOV. The results show that the application of a radial and TOF bin selection, prior to the application of PCA, successfully brings the motion of the desired object in *first place*, i.e. it is described by the 1^{st} PC. This implies that the effect of the other changes that are not targeted, even though they are in fact bigger, has been effectively disregarded. This outcome suggest that this selection process could be of considerable benefit to the respiratory detection with PCA, since the sought-after motion is expected to occur in the central part of the patient, whereas the outer areas might be affected by bulk or arm motion.

In the next chapter these methods will be put into practice with a large cohort of clinical patient datasets, and a thorough evaluation will be presented.

Chapter 5

Evaluation on patient data

In this chapter the performance of PCA and the sign-determination methods is thoroughly evaluated on patient data. As a first step, PCA respiratory signals are compared to the patient internal motion extracted from MR dynamic acquisitions acquired on a PET/MR system. Subsequently, the methods are applied on a large cohort of PET/CT patient data, and the obtained respiratory signals are compared with the RPM signal for different acquisition durations and different total number of utilised counts. Finally, the respiratory signals are used for gating the PET data and the resulting reconstructed images are compared.

5.1 Comparison with internal motion extracted from MR data

Several studies have investigated whether the signals provided by external devices (such as the Anzai belt, RPM, etc) correctly represent the internal respiratory motion of the organs, drawing a generally positive conclusion (see Fayad et al. (2011) and others in Section 2.3.2). Furthermore, in Manber et al. (2015), the PCA signal was compared to the respiratory signal extracted from an MR-navigator for 9 patients, showing a mean correlation between the two signals of 0.89.

Nevertheless, a combined analysis of the results of DD methods (PCA in this case), external device signal and internal motion has never been performed and it is the object of the work presented in this section.

5.1.1 PET/MR patient data

Thirteen patients were acquired on a GE SIGNA PET/MR system (Levin et al., 2016), at the Service Hospitalier Frédéric Joliot (SHFJ) in Orsay, France. They were cancer patients (not all with lung cancer) imaged for cancer staging, in an exam that includes the acquisition of the thorax, the abdomen and the pelvis. All datasets were acquired with approval of the local Institutional Review Board and patients consented to the use of their data for research

purposes[1]. The details of the thorax acquisitions are the following, see Figure 5.1 for a diagram:

- The patients were injected with FDG using activity levels according to routine clinical protocols.
- A pneumatic respiratory belt was positioned around the patient chest to acquire the respiratory signal: the belt measures the expansion and contraction of the chest and is acquired throughout the whole data acquisition, with temporal resolution of 40ms.
- Multiple MR acquisitions (3 for 9 patients and 2 for 4 patients) were performed at different stages, simultaneously to the PET scan. The MR acquisition consisted in a single-slice fast acquisition performed on the right liver-lung interface. The pulse sequence was a balanced SSFP (Steady State Free Precession) which on the GE scanner is called FIESTA (Fast Imaging Employing Steady-state Acquisition), with TE/TR=2.176/5.365ms, BW/2=125kHz, slice thickness=10 mm, flip angle=35 degrees, FOV=256 x 256mm, matrix 256 x 64 with acceleration factor of 3 in phase encoding directions (Anterior/Posterior). It was continuously repeated 512 times, with a temporal resolution of 140ms per frame, for a total duration of approximately 71.7s. Each image has a spatial resolution of 1mm in the read-out direction.
- The PET data were acquired for 300s.

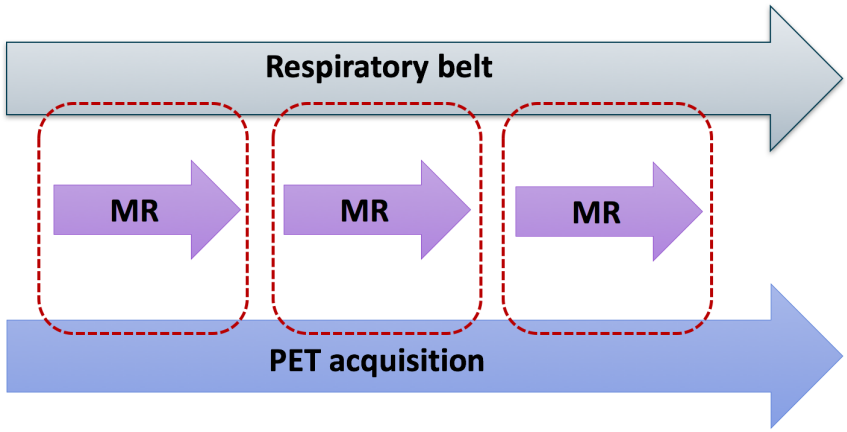

Figure 5.1: PET/MR acquisition setup. There are 3 (or 2) intervals of approximately 72 seconds duration during which PET, MR and the pressure belt signal are acquired simultaneously.

5.1.2 Extraction of internal motion from MR data

The fast MR acquisitions were utilised to extract a signal representing the motion of the diaphragm, with high temporal resolution (140ms). An example of one MR frame is shown in

[1]Thanks to Dr. Brice Fernandez, Dr. Michael Soussan, Dr. Claude Comtat and Dr. Irène Buvat

Figure 5.2 and the diaphragm can be clearly identified. The motion was obtained by applying an edge detection process along a line (depicted here in red and superimposed to the MR data) crossing the diaphragm , as the MR values drop considerably between the diaphragm and the lung area. Figure 5.3a shows the MR image intensity values along the line for three different points in time. The edge corresponds to the border between the diaphragm and the lungs and it moves in time. The displacement of the edge in time constitutes the respiratory signal representing the internal motion of the organs, see Figure 5.3b. This signal will be referred to as MR-derived.

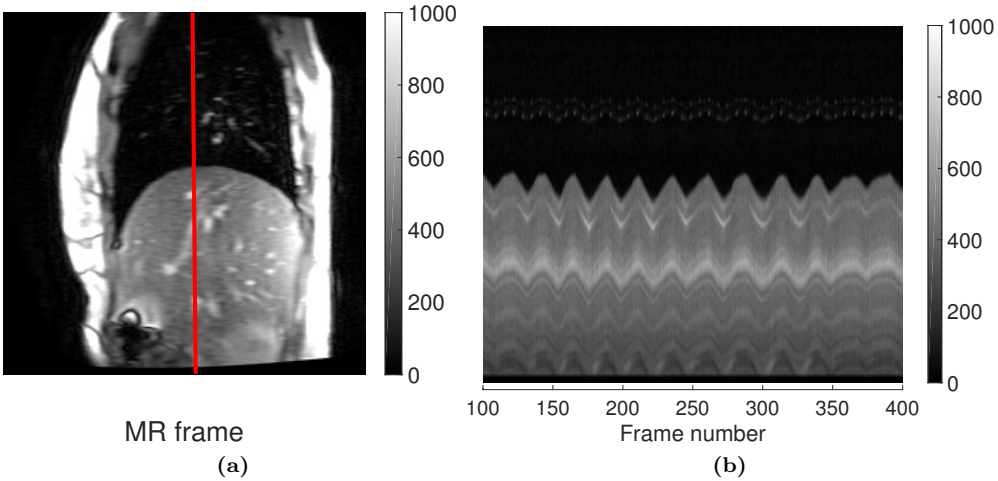

Figure 5.2: (a) Example frame of the MR acquisition with line along which the edge detection process is applied. (b) For each of the 512 frames the data along that line are selected and placed sequentially to illustrate how the edge moves in time.

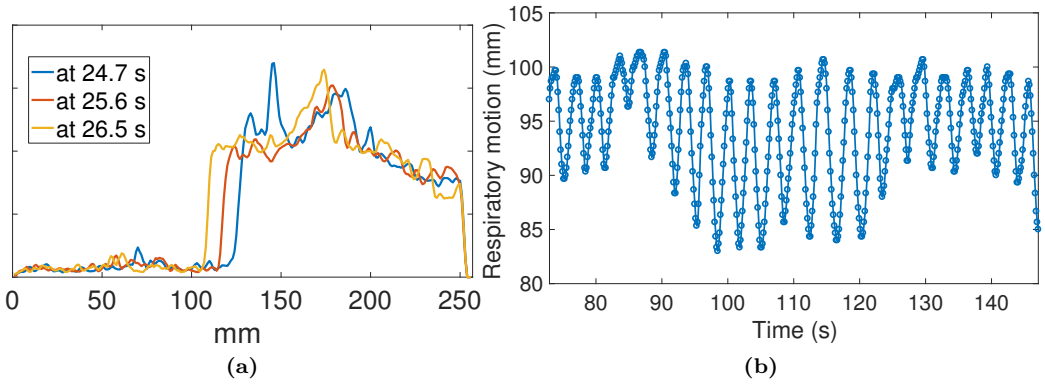

Figure 5.3: (b) Profiles along the red line at three different points in time. The changing position of the edge in the profile represents the motion. Tracking that position in time provides a respiratory motion signal, like the one shown in (b).

5.1.3 Extraction of respiratory motion from PET data with PCA

The PET data generated by the GE SIGNA PET/MR system is organised into 351 TOF bins of 13ps width. For the application of PCA the whole dataset duration of 300s was

unlisted into dynamic sinograms with 27 TOF bins (with TOF bin mashing factor of 13), with a 500ms time frame. No cylindrical masking was applied to the sinograms. PCA was then applied as described in 2.3.3, and the sign of the signal was correctly identified by the *CorrWeights* method in all cases. The temporal sampling of the signal was upsampled from 500ms (the duration of each time frame) to 250ms, with linear interpolation. The portions of the PCA respiratory signal related to the MR acquisitions were subsequently selected from the total signal of duration 300s.

5.1.4 Comparison of respiratory signals

For each portion of the acquisition during which the PET, the MR and the respiratory belt signal are acquired simultaneously (see Figure 5.1), three respiratory signals can be extracted: the PET-derived obtained with the application of PCA, the MR-derived obtained as described in the previous section, and the pressure belt output. The PCA signal has a sampling rate of 250ms, whereas for the belt and MR-derived signals the rate is equal to 40ms and 140ms respectively. The latter were temporally downsampled to the PCA signal sampling rate of 250ms. The three signals were then synchronised to all be consistent with the time given by the PET listmode. The MR derived signal was synchronised exploiting the temporal information found in the MR dicom files. The synchronisation of the belt proved to be the most challenging process, as the only available temporal information is given by the end time (with a precision of only a second), the duration of the belt acquisition and its sampling rate. This information was sufficient to approximately place the start of the belt signal along the PET acquisition, but manual adjustment was required for a more accurate synchronisation. This was done by visually comparing the belt and PCA respiratory signals, and shifting the belt signal with steps of 0.1 to reach the best match. Finally, the three synchronised signals were compared among each other, with the evaluation of the Pearson correlation.

5.1.5 Results

For 9 out of the total 13 patients, the MR FIESTA acquisition was performed 3 times during the PET acquisition, whereas for the remaining 4 only 2 were within the PET duration. The total number of intervals with PET, MR and belt simultaneously is therefore 35. Figure 5.4 shows the correlation values with respect to each interval for all patients, and Table 5.1 reports the mean of all the correlation values. Figure 5.6 shows the whole duration of the respiratory signal obtained from PCA together with the signal extracted from the three MR acquisitions, for 4 patients. The belt signal is not superimposed in order to show the PCA signal more clearly. Figure 5.5 shows an example of the respiratory signals given by PCA, the belt and one of the MR acquisitions, together with a closer selection of a 20s interval. Figure 5.7 shows the full duration of the PCA and the belt signal for the same patients as in

Figure 5.6.

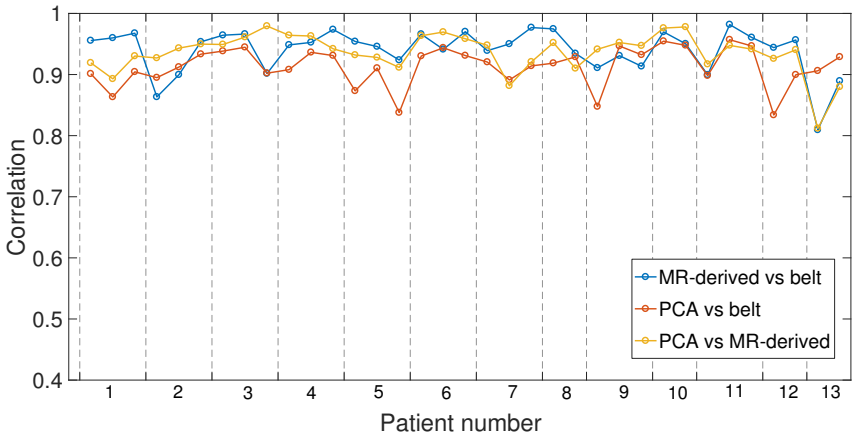

Figure 5.4: Value of the Pearson correlation between the signals for all the 13 patients (35 intervals in total).

Table 5.1: Mean of the Pearson correlation values over all the 13 patients (35 intervals in total).

	Mean correlation
MR-derived vs belt	0.94 ± 0.04
PCA vs belt	0.91 ± 0.03
PCA vs MR-derived	0.94 ± 0.03

5.1.6 Discussion

These results show that the PCA signal, the internal motion extracted from the MR data and the pneumatic belt signal are very similar to each other, as showed in Figure 5.5 for an example interval. The correlation values are in all cases extremely high, as seen from Figure 5.4 and Table 5.1, with mean values above 0.9 for the three comparisons: PCA with the MR-derived signal, PCA with the belt and MR-derived signal with the belt.

In order to assess the general trend of the PCA signal, comparing it to the internal motion extracted from intervals far apart during the PET acquisition, and with respect to the signal given by the pressure belt, some patients have been further investigated. The presence of a drift in the PCA signal with respect to the other signals has been examined. From Figure 5.6 the whole duration of the PCA signal can be appreciated for 4 example patients, together with the three MR-derived signals. Patients 1 and 7 show an irregular breathing pattern, nevertheless the PCA method produces signals that are highly representative of the internal motion during the MR acquisitions (as also proved by the correlation values in Figure 5.5). Furthermore, for these two patients it appears as if some unexpected change in the data was detected by the PCA signal in the periods between the MR acquisitions, as a sudden change in the signal amplitude can be observed. The PCA signal for patients 2 and

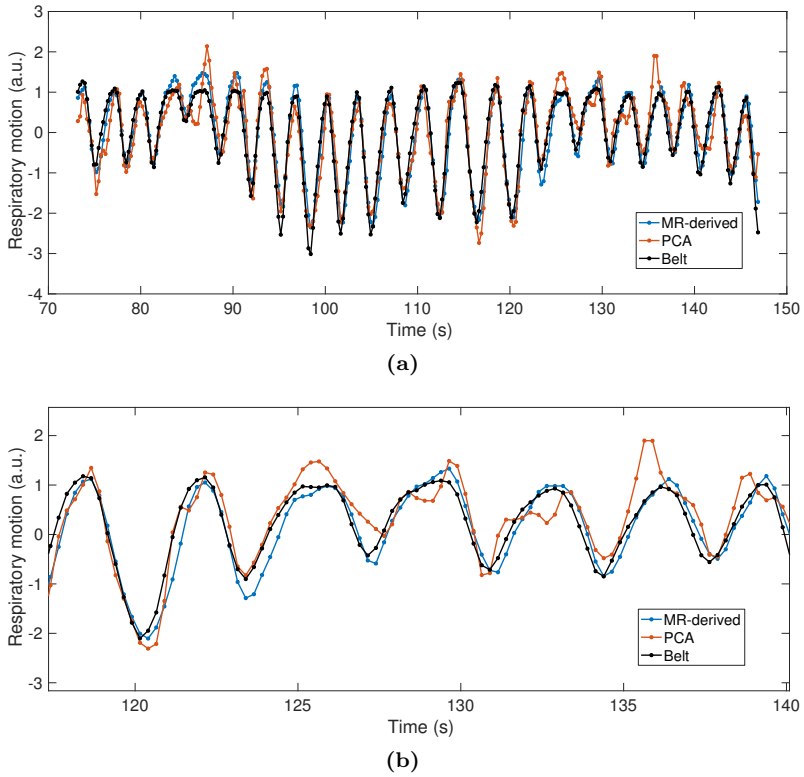

Figure 5.5: (a) The three respiratory signals for patient number 1: MR-derived, PCA and belt. Correlations between the signals: MR-derived vs belt = 0.96, PCA vs belt = 0.90, PCA vs MR-derived = 0.92. (b) Enlarged selection of a portion of (a), to better show the similarity of the traces. No smoothing is applied to any of the traces.

9 does not show any detectable drift in amplitude. Similar conclusions can be drawn from Figure 5.7, where the PCA signal is compared to the full duration of the belt signal. Patients 1 and 7 are again those presenting the most different behaviour with respect to the periods between the MR acquisitions, whereas the correlation between PCA and belt signal for the subintervals corresponding to the MR duration are high (see Figure 5.5). Patients 2 and 9 show very high similarity with the full length of the belt acquisition, and this has been found to be the case for all the other patients. Considerable drift in the PCA values has not been found in any of the analysed patients, apart from patients 1 and 7 where parts of the PCA signal seem to have been affected by other changes in the data (i.e. bulk motion). Moreover, also the motion extracted from the MR and that of the belt might be affected by drift in amplitude, the former because of patient motion that causes the detected edge to shift, and the latter because of possible mechanical issues. Therefore, the "best" respiratory signal cannot be determined by only looking at the trace itself, as no ground truth is available. An analysis on gated reconstructions is needed and it will be addressed in the last section of this chapter.

In summary, this analysis proves that PCA produces an accurate and reliable surrogate

respiratory signal, while also confirming that the external motion of the chest and the internal motion of the organs are both good measures of the breathing pattern of the patient. These promising results constitute the starting point from which to proceed towards a more thorough analysis of PCA, taking into account various aspects of its performance. This will be addressed in the next section, where the PCA signal will be evaluated on a larger cohort of patients and compared to the RPM signal (Real-Time Position Management device, see Paragraph 2.3.2).

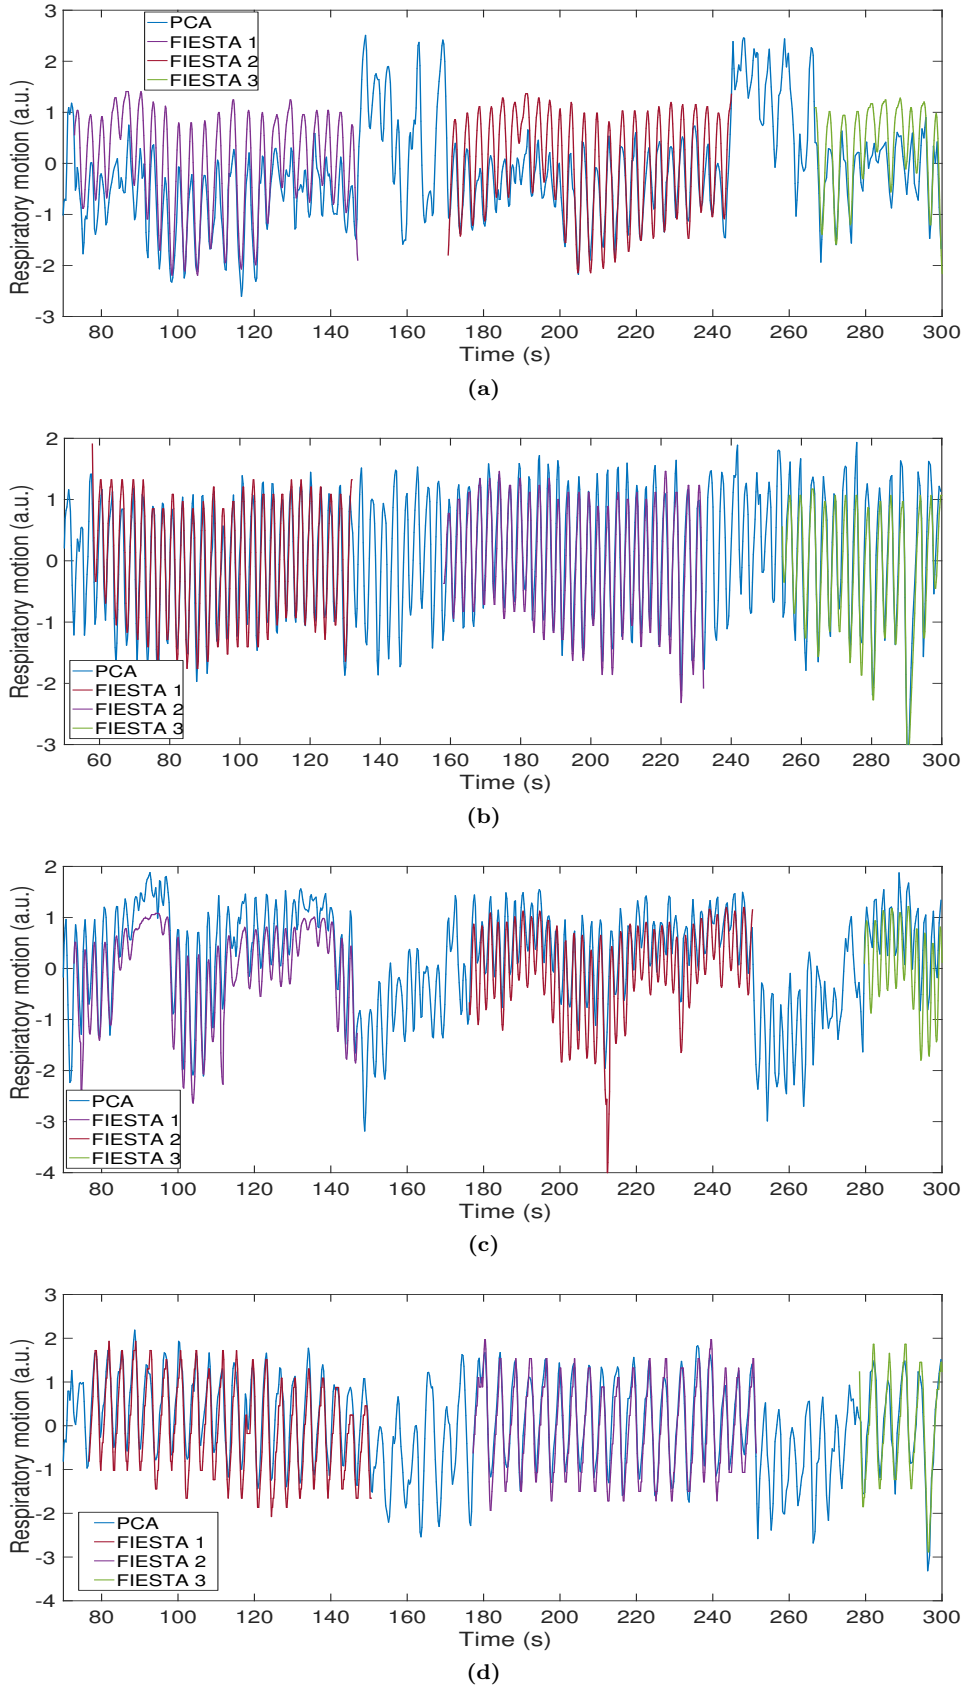

Figure 5.6: Respiratory signal obtained from PCA (in blue), together with the signal obtained processing the three FIESTA acquisitions, for patients number (a) 1, (b) 2, (c) 7 and (d) 9.

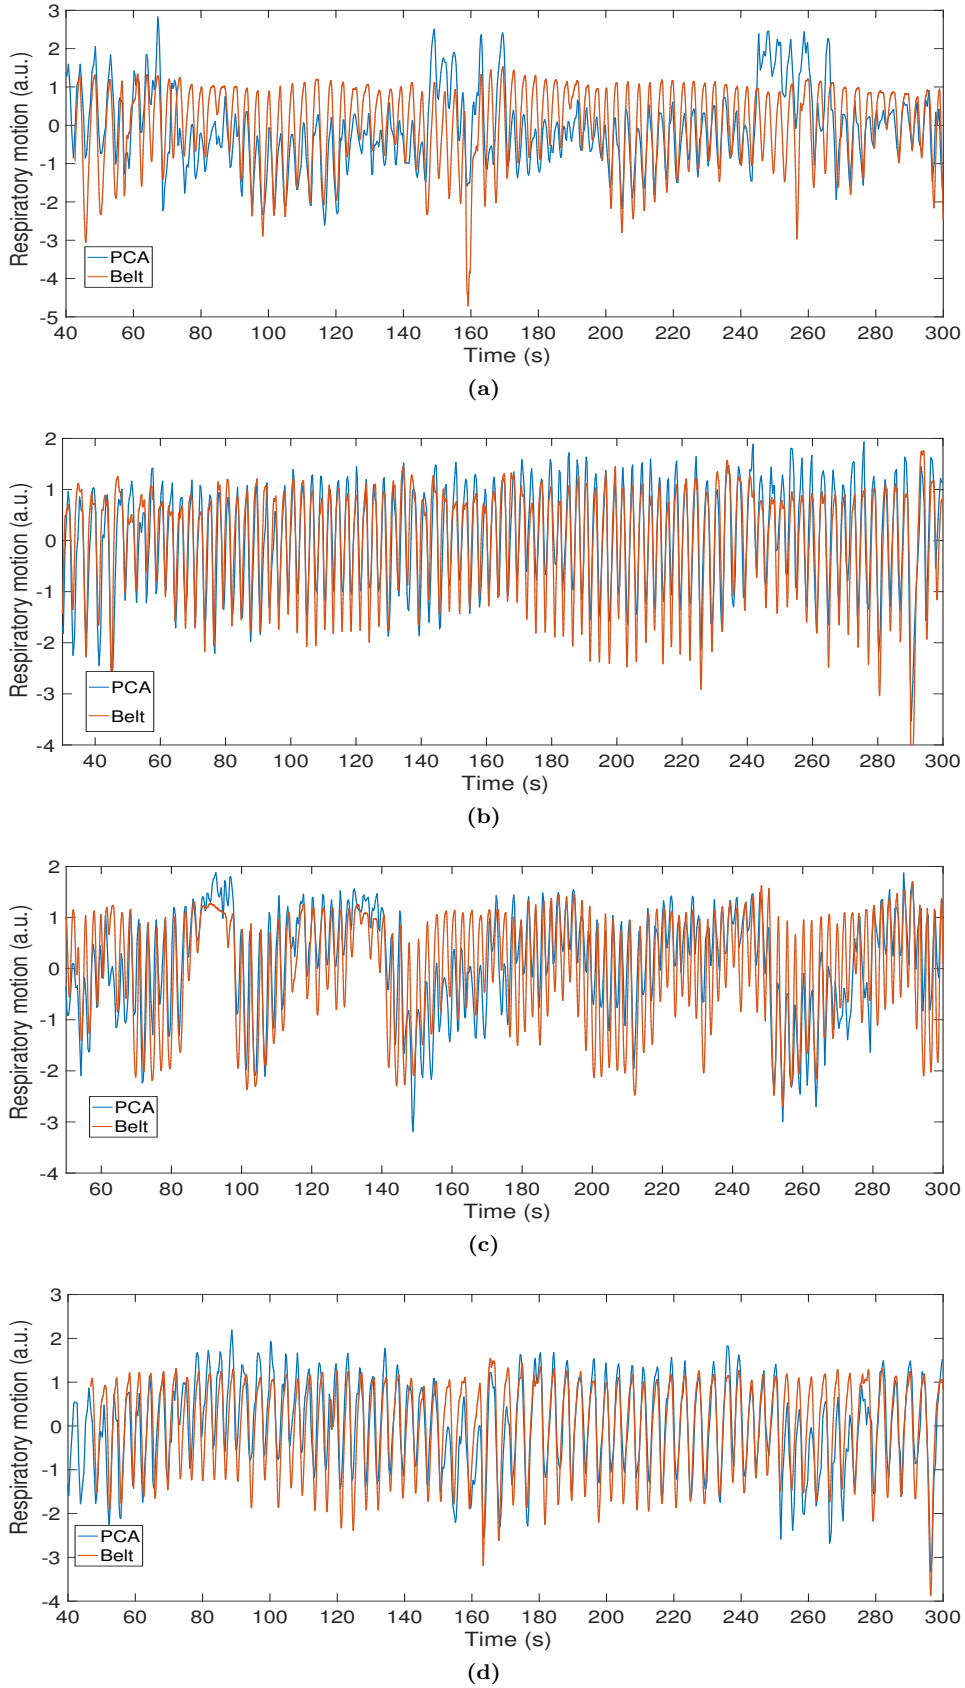

Figure 5.7: Respiratory signal obtained from PCA (in blue), together with the belt signal for the same patients as in Figure 5.6, number (a) 1, (b) 2, (c) 7 and (d) 9.

5.2 Evaluation of PCA signals and sign-determination methods on PET data with RPM

In this section the performance of PCA and the sign-determination methods is thoroughly investigated. In the first part, a comparison is performed between the DD and the RPM signals when varying the duration of the data interval taken into account; in the second part, a similar analysis is performed with respect to different amount of utilised counts. In both cases the impact of using non-TOF and TOF data, with and without the use of the masking described in 4.2.3, is also investigated.

5.2.1 Patient data

A cohort of 28 FDG oncology patients PET/CT datasets has been used: 16 patients were acquired on a Discovery 690 PET/CT scanner[2] and 12 patients acquired on a Discovery 710 PET/CT[3]. All datasets were acquired with approval of the local Institutional Review Board and patients consented to the use of their data for research purposes. The acquisitions were monitored by the RPM device. All studies included multiple bed positions, and for each study only the one related to the lower lung area was chosen. Both scanners have the same hardware and data format. The TOF resolution is 550ps (Bettinardi et al., 2011) and the data are acquired in 55 TOF bins of 89ps width.

Prior to the application of PCA, the listmode data underwent unlisting both into non-TOF and TOF sinograms. In the TOF unlisting, the 55 total TOF bins are mashed (downsampled) into either 5 or 11 TOF bins, therefore applying a mashing factor equal to 11 and 5 respectively. The final sinogram dimensions after SSRB (see 4.2.1) are therefore $95 \times 9 \times 47$ (radial positions $\times$ angles $\times$ transaxial planes) for the non-TOF case, and $95 \times 9 \times 47 \times 5$ and $95 \times 9 \times 47 \times 11$ for the two TOF configurations, with the final dimension being the TOF bin number. The sinograms with 5 and 11 TOF bins will be hereinafter referred to as 5TOF and 11TOF, in Tables and Figures.

Acquisition durations range from 180s to 360s, and for each patient the RPM motion tracking device was utilised during the whole duration of the PET acquisition.

5.2.2 Analysis with respect to duration

Considering that PCA extracts motion information directly from the patient data, it is important to assess its performance with respect to varying conditions, such as when the available dataset has a shorter or longer duration. To do this, similarly to what has been presented in 3.2.3.2, the available patient sinograms (non-TOF and TOF) have been

[2]Anonymised data from the BCCA, Vancouver, CA (thanks to Dr. François Benard).

[3]Five anonymised datasets Oxford University Hospitals NHS Foundation Trust, (thanks to Dr. Daniel McGowan), and 7 patients from the Institute of Nuclear Medicine, UCLH, London (thanks to Dr. Simon Wan, Dr. Abdulrhman Alnaim, Tom Sanderson, Raymond Endozo and Bettina Ferre).

subdivided in independent shorter time intervals of 25s, 50s, 100s, 200s and 300s (these last two only for the longest acquisitions). The PCA method and the sign-determination methods were then applied on all the sinograms.

5.2.2.1 Correlation with RPM

As in Chapter 3, the quality of the respiratory signal produced by PCA is assessed via comparison with the RPM signal. The metric of the evaluation is the Pearson correlation coefficient.

In Figure 5.8 the mean values of the correlation of the PCA respiratory signal with the RPM are shown, averaged over all the available interval durations and patients. Figure 5.9 shows the boxplots[4] of the correlation values for the 25s and 300s intervals, that are the shortest and longest. These plots show that there is a general and notable increase in the correlation with the RPM with respect to increasing duration. This is to be expected: the more data is available to the method and the better it will detect the changes that occur. With shorter time intervals there is also an increased probability that other types of motion (e.g. bulk motion) would contribute to the overall change in the data in addition to respiration. In this case the method would generate Principal Components that represent a mixture of effects, without succeeding in extrapolating respiration only.

The correlation of the PCA signal with the RPM is positively affected by the use of TOF data, as opposed to non-TOF, as the mean values increase (see Figure 5.8). Even though the increase has not been found to be considerable, it is nevertheless informative of the fact that using TOF sinograms is beneficial and improves the results. The application of the cylindrical mask, on the other hand, does not produce noticeable changes. The values in Figure 5.8 are also reported in Table 5.2, from which it is possible to appreciate that the biggest impact on the correlation between the PCA signal and the RPM is given by the use of TOF data. Furthermore, the effect is more pronounced when using shorter time intervals.

Table 5.2: Mean and standard deviation of the Pearson correlation values over all patients for all interval durations. The number of total available intervals for each duration is reported in brackets.

Duration (s) (num. intervals)	non-TOF w/o m	non-TOF w/ m	5TOF w/o m	5TOF w/ m	11TOF w/o m	11TOF w/ m
25 (322)	0.70 ± 0.22	0.70 ± 0.21	0.74 ± 0.18	0.73 ± 0.20	0.74 ± 0.18	0.74 ± 0.20
50 (155)	0.76±0.16	0.80±0.16	0.82±0.15	0.82±0.15	0.82±0.15	0.82±0.15
100 (67)	0.80±0.16	0.79±0.15	0.82±0.15	0.82±0.15	0.82±0.15	0.82±0.15
200 (23)	0.88±0.14	0.87±0.16	0.89±0.18	0.92±0.08	0.92±0.09	0.92±0.08
300 (16)	0.94±0.05	0.93±0.05	0.95±0.05	0.95±0.05	0.95±0.05	0.95±0.05

[4] The central mark indicates the median, and the bottom and top edges of the box indicate the 25th and 75th percentiles, respectively. The whiskers extend to the most extreme data points not considered outliers, and the outliers are plotted individually using the '+' symbol. The same applies to all the following boxplots.

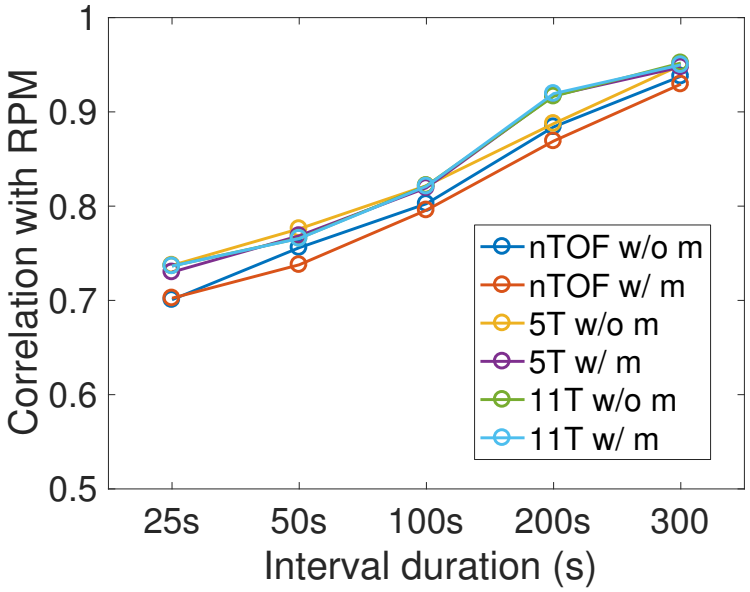

Figure 5.8: The mean values for each interval duration are obtained averaging the correlation with the RPM of the signals obtained for all patients. This process is applied for all the various data configurations: non -TOF sinograms without and with masking (respectively *nTOF w/o m* and *nTOF w/ m*); TOF sinograms with 5 and 11 TOF bins without and with masking (respectively *5TOF w/o m* and *w/ m* and *11TOF w/o m* and *w/ m*).

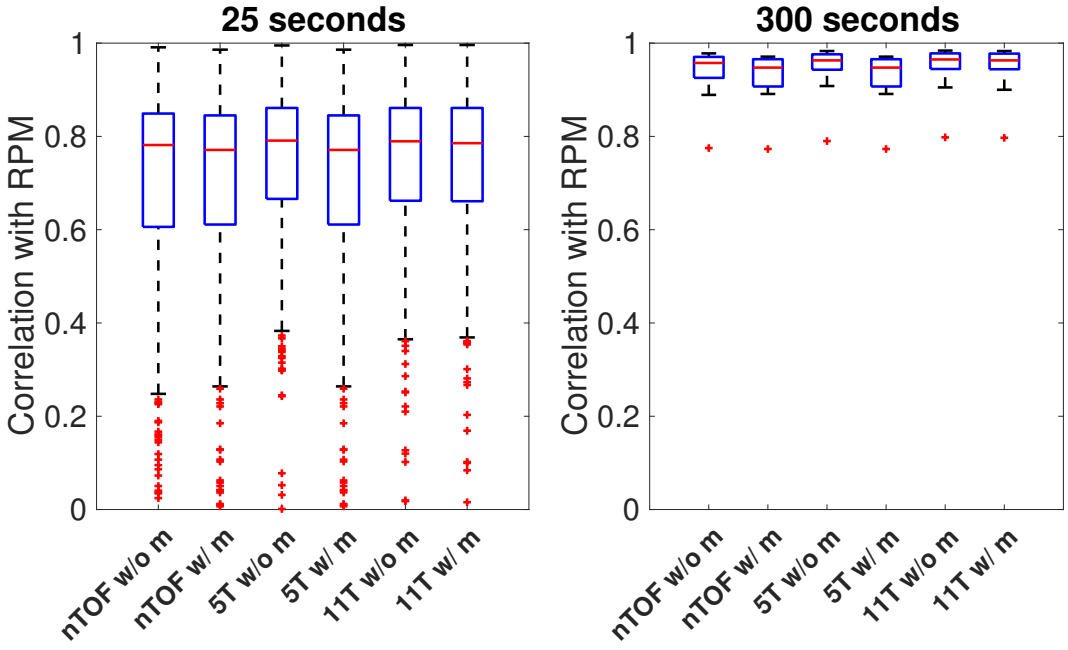

Figure 5.9: Box-plot of the correlation with the RPM for the two most extreme cases: the shortest intervals (25s, 322 total intervals) and the longest ones (300s, 16 total intervals).

5.2.2.2 Sign-determination methods

The PCA respiratory signal needs to have high correlation with the RPM but, as discussed in Chapter 3, it also needs to represent the correct direction of motion. The sinogram-based

sign-determination methods presented in Chapter 3 were therefore assessed via checking the sign of the signal with the external device. The methods were also compared between each other in respect of their failure rates and their output value, that is corr(w,u) for the method *CorrSino* (where u is the *WeightedGradSino* signal and w is the respiratory PC signal), and corr(RPC, *WeightedGradSino*) for *CorrWeights*. To be noted that in this analysis the output of the sign-determination methods are analysed using their absolute values, as their sign is taken into consideration when evaluating the failure rates. In 3.3 it has already been shown on non-TOF data that the output value generated by the *CorrWeights* method is notably higher than that produced by *CorrSino*, therefore suggesting that the former could be a more reliable tool for sign-determination, even though the observed failure rates were similar (on the 37 analysed patients). In this section the comparison is therefore expanded to TOF data and more patients.

In Figure 5.10 the mean of the output value of both methods is shown, averaged over all patients for each of the interval durations taken into account. The closer the value is to 1 and the better the method is expected to perform. The values of the *CorrWeights* method are always considerably higher than those of *CorrSino*, in accordance to the results presented in Chapter 3.3. Furthermore, the use of the TOF sinograms and the application of the cylindrical masking notably increases the output values, and this is for both *CorrWeights* and *CorrSino*. Tables 5.3 and 5.4 show the failure rates over all patients with respect to the interval durations (number of failures divided by the total number of intervals of that length). The total failure rates per method and per type of sinograms used (i.e. not distinguishing between the durations, but evaluating the total number of failures on all the 583 total utilised intervals) is shown in Table 5.5. The values are the same for both methods, except for the 11 TOF bins case with the use of the masking, where *CorrWeights* fails less frequently. Given these results, and the fact that the *CorrWeights* is applicable to any DD method while *CorrSino* is specific to PCA, the *CorrWeights* method is designated as the method of choice. The further analysis is therefore restricted to the *CorrWeights* method.

Table 5.3: Failure rate (in percentage) of *CorrSino* in determining the correct direction of motion, with respect to the duration of the interval taken into account.

Duration (s) (num. intervals)	non-TOF w/o m	non-TOF w/ m	5TOF w/o m	5TOF w/ m	11TOF w/o m	11TOF w/ m
25 (322)	3.7	1.6	0.9	0.9	1.2	1.5
50 (155)	1.3	0	0.6	0	0.6	0.6
100 (67)	3.0	1.5	0	0	0	0
200 (23)	0	0	0	0	0	0
300 (16)	0	0	0	0	0	0

To further analyse and better appreciate the effect of TOF sinograms and masking

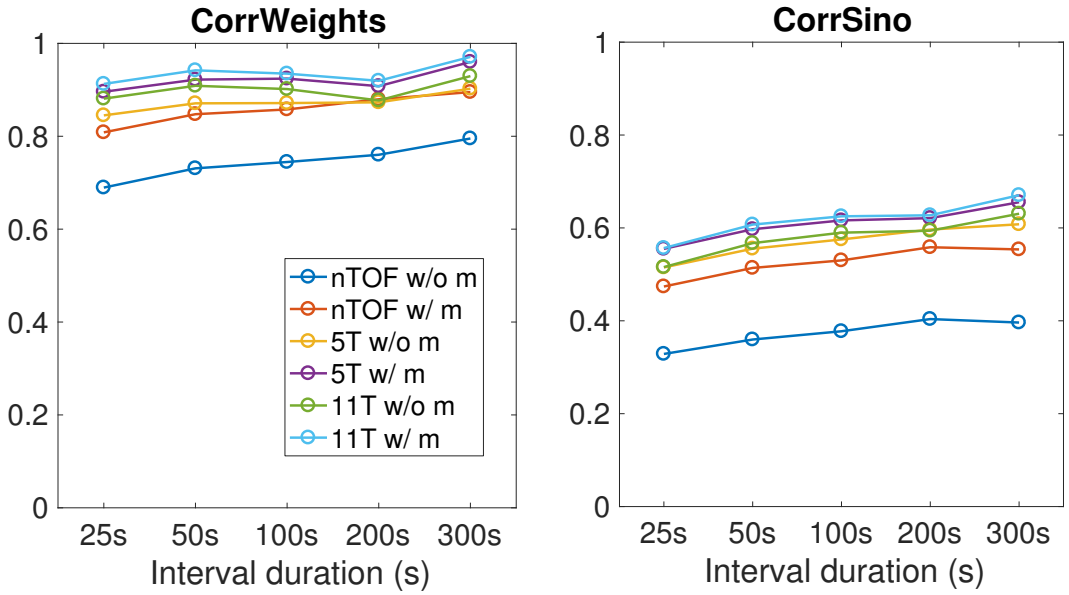

Figure 5.10: Mean of the output values of the sign-determination methods over all patients for each interval duration. The legend applies to both plots.

Table 5.4: Failure rate (in percentage) of *CorrWeights* in determining the correct direction of motion, with respect to the duration of the interval taken into account.

Duration (s) (num. intervals)	non-TOF w/o m	non-TOF w/ m	5TOF w/o m	5TOF w/ m	11TOF w/o m	11TOF w/ m
25 (322)	3.7	1.6	0.6	0.6	1.2	1.2
50 (155)	1.3	0	0.6	0	0.6	0.6
100 (67)	3.0	1.5	1.5	1.5	0	0
200 (23)	0	0	0	0	0	0
300 (16)	0	0	0	0	0	0

Table 5.5: Failure rate (in percentage) of the methods in determining the correct direction of motion, on all intervals taken into account (with duration equal to 25s, 50s, 100s, 200s and 300s, for a total of 583 intervals).

Method	non-TOF w/o m	non-TOF w/ m	5TOF w/o m	5TOF w/ m	11TOF w/o m	11TOF w/ m
CorrWeights	2.7	1.0	0.7	0.5	0.9	0.9
CorrSino	2.7	1.0	0.7	0.5	0.9	1.0

on the *CorrWeights* method's performance, the output values are displayed in boxplots in Figure 5.11, with respect to the interval duration. The plots show a clear increasing trend in the values when using TOF sinograms as opposed to non-TOF, and a decreased spread. Furthermore, for all data configurations (TOF and non-TOF) and interval durations, the application of the mask further increases the values and decreases their spread, therefore suggesting a more stable behaviour of the method. When comparing the use of 5 or 11 TOF bins, the latter case produces the highest values, as it can also be seen from Table 5.6.

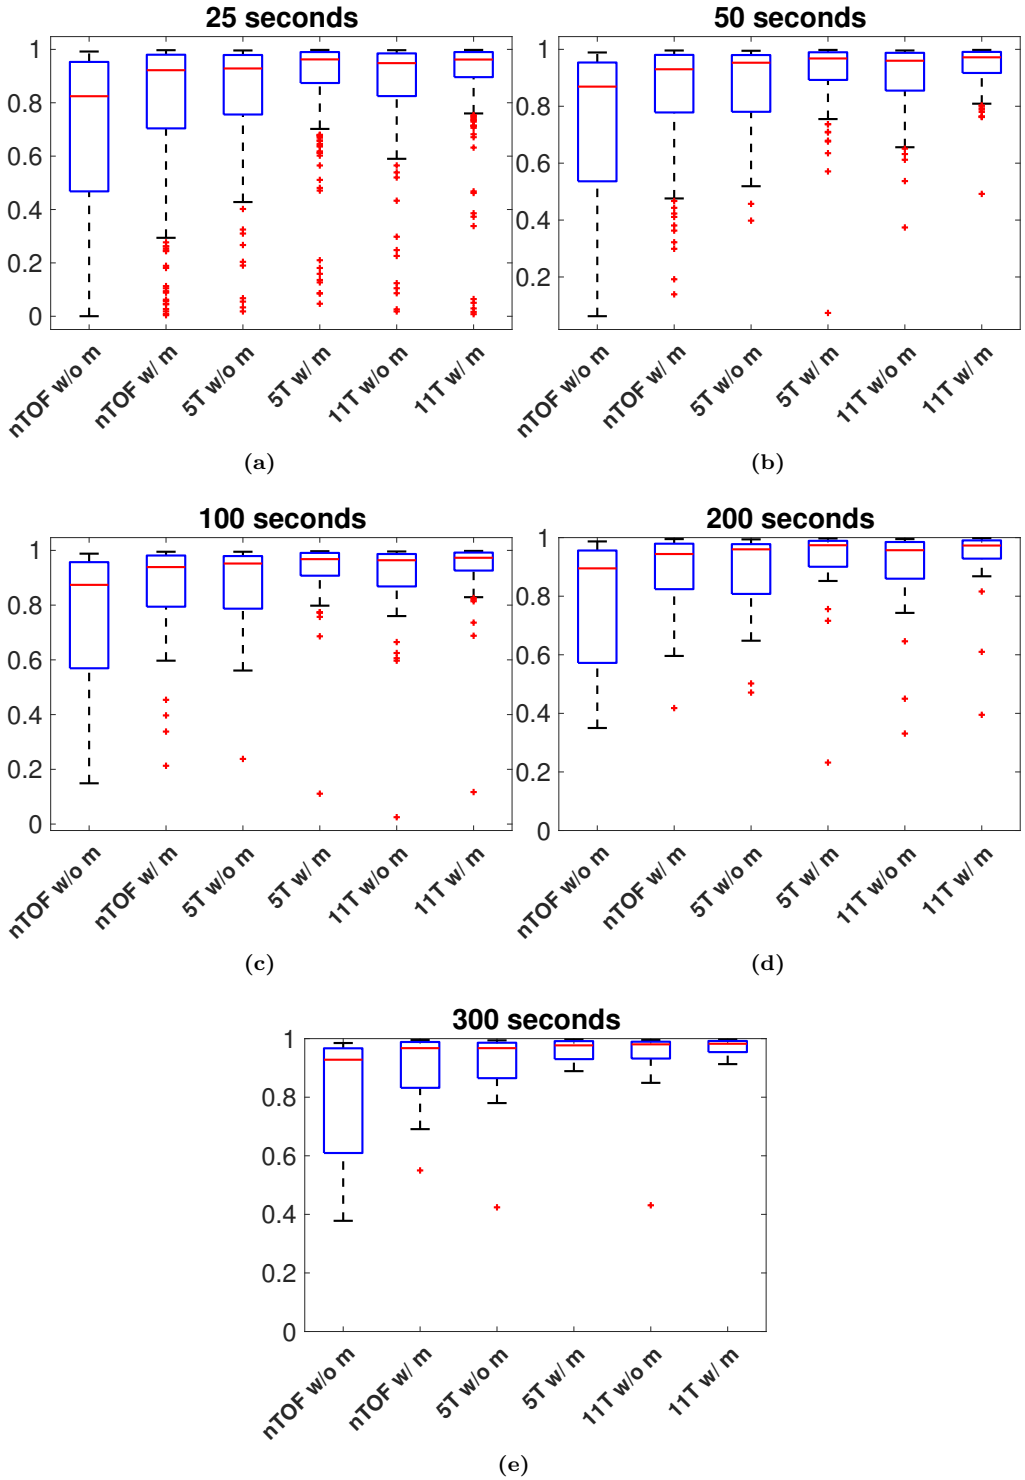

Figure 5.11: Boxplots of the *CorrWeights* values with respect to the interval duration for all the various data configurations: non -TOF sinograms without and with masking (respectively *nTOF w/o m* and *nTOF w/ m*); TOF sinograms with 5 and 11 TOF bins without and with masking (respectively *5TOF w/o m* and *w/ m* and *11TOF w/o m* and *w/ m*).

Table 5.6: Mean values and standard deviation of the output value of *CorrWeights* over all intervals (583, as previously mentioned) for all data configurations.

non-TOF w/o m	non-TOF w/ m	5TOF w/o m	5TOF w/ m	11TOF w/o m	11TOF w/ m
0.71 ± 0.27	0.83 ± 0.21	0.86 ± 0.16	0.91 ± 0.14	0.89 ± 0.15	0.93 ± 0.13

5.2.3 Analysis with respect to fewer counts

The respiratory signal detection and sign-determination are required to be stable processes with respect to count levels, which depend on the injected activity. The performance of PCA and the method of choice *CorrWeights* were therefore assessed on sinograms with various total counts, to simulate their behaviour when different activity levels are utilised. This was done via unlisting the available patient datasets into sinograms with random rejection, keeping a percentage of the total counts equal to 3%, 6%, 12.5%, 25% and 50%. This process was performed using the utility `lm_to_projdata` of the STIR library (Thielemans et al., 2012). To make use of all patient datasets and only investigate the effect of different total counts and not that of different interval durations, only 180s of each of the 28 patients' dataset were utilised (that corresponds to the acquisition duration of the shortest studies).

5.2.3.1 Correlation with RPM

The mean values of the correlation of the PCA signals with the RPM, for all the different percentage of counts utilised and averaged over all 28 patients, are reported in Table 5.7 and shown in Figure 5.12. In Figure 5.13 the correlations are displayed in boxplots, showing the difference in the spread of their value when using the various data configurations (non-TOF and TOF sinograms, and without and with the application of the cylindrical masking).

Table 5.7: Mean and standard deviation of the Pearson correlation values over all patients for all count levels.

%	non-TOF w/o m	non-TOF w/ m	5TOF w/o m	5TOF w/ m	11TOF w/o m	11TOF w/ m
3	0.35 ± 0.30	0.32 ± 0.30	0.39 ± 0.33	0.38 ± 0.32	0.40 ± 0.32	0.43 ± 0.29
6	0.47±0.30	0.40±0.33	0.52±0.32	0.51±0.31	0.54±0.32	0.55±0.30
12.5	0.61±0.27	0.58±0.28	0.69±0.26	0.68±0.26	0.69±0.27	0.72±0.22
25	0.74±0.20	0.74±0.20	0.80±0.17	0.80±0.17	0.80±0.22	0.81±017
50	0.78±0.21	0.78±0.21	0.85±0.17	0.84±0.17	0.84±0.17	0.85±0.17
100	0.85±0.17	0.84±0.18	0.87±0.19	0.89±0.13	0.86±0.20	0.89±0.13

Both Figures clearly show the increase in the correlation with the RPM when TOF sinograms are utilised. This applies for all percentages of utilised counts. The application of the cylindrical mask, however, does not contribute noticeably to the improvement of the

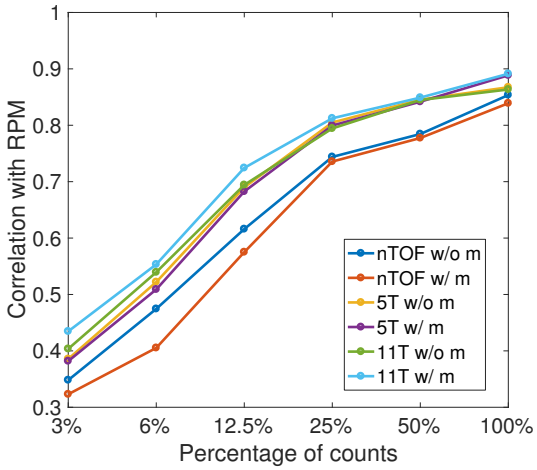

Figure 5.12: Mean value of the correlation of the PCA signal with the RPM, averaged over all patients, for all count levels.

respiratory signal.

5.2.3.2 Sign-determination method - CorrWeights

The chosen sign-determination method *CorrWeights* was applied on the sinograms with the various percentage of counts and the direction of motion compared to that of the RPM. The mean value of *CorrWeights* over all patients for each case is shown in Figure 5.14. Furthermore, similarly to the analysis of the quality of the PCA trace, the method's output values are displayed in boxplots in Figure 5.15. The use of TOF sinograms provides an improvement in the performance of the method. The only instance where this does not apply is with 3% of the total counts, in which case however also the correlation of the respiratory signal to the RPM is very low (see Figure 5.12). This suggest that the application of PCA for respiratory detection is not robust with such a low percentage of counts.

From both Figures it is possible to appreciate that the application of the cylindrical masking increases considerably the *CorrWeights* values. This occurs with all data configurations: non-TOF and 5 and 11 TOF bins sinograms. Furthermore, the results show that when using TOF sinograms, the 11 TOF case provides better results than the 5 TOF.

Table 5.8 shows the failure rate of the sign-determination methods, averaged over the 28 patient, with respect to the utilised counts. The method is more robust when using TOF sinograms, and its performance is further improved when applying the cylindrical mask. When applied to 11 TOF bins sinograms together with the cylindrical masking, the method correctly determines the direction of the PCA respiratory signal in all cases. For the 5 TOF bins sinograms, the only failures are found when using very few counts (3 and 6%), in correspondence to the decreased overall quality of the PCA signal, see Figure 5.12.

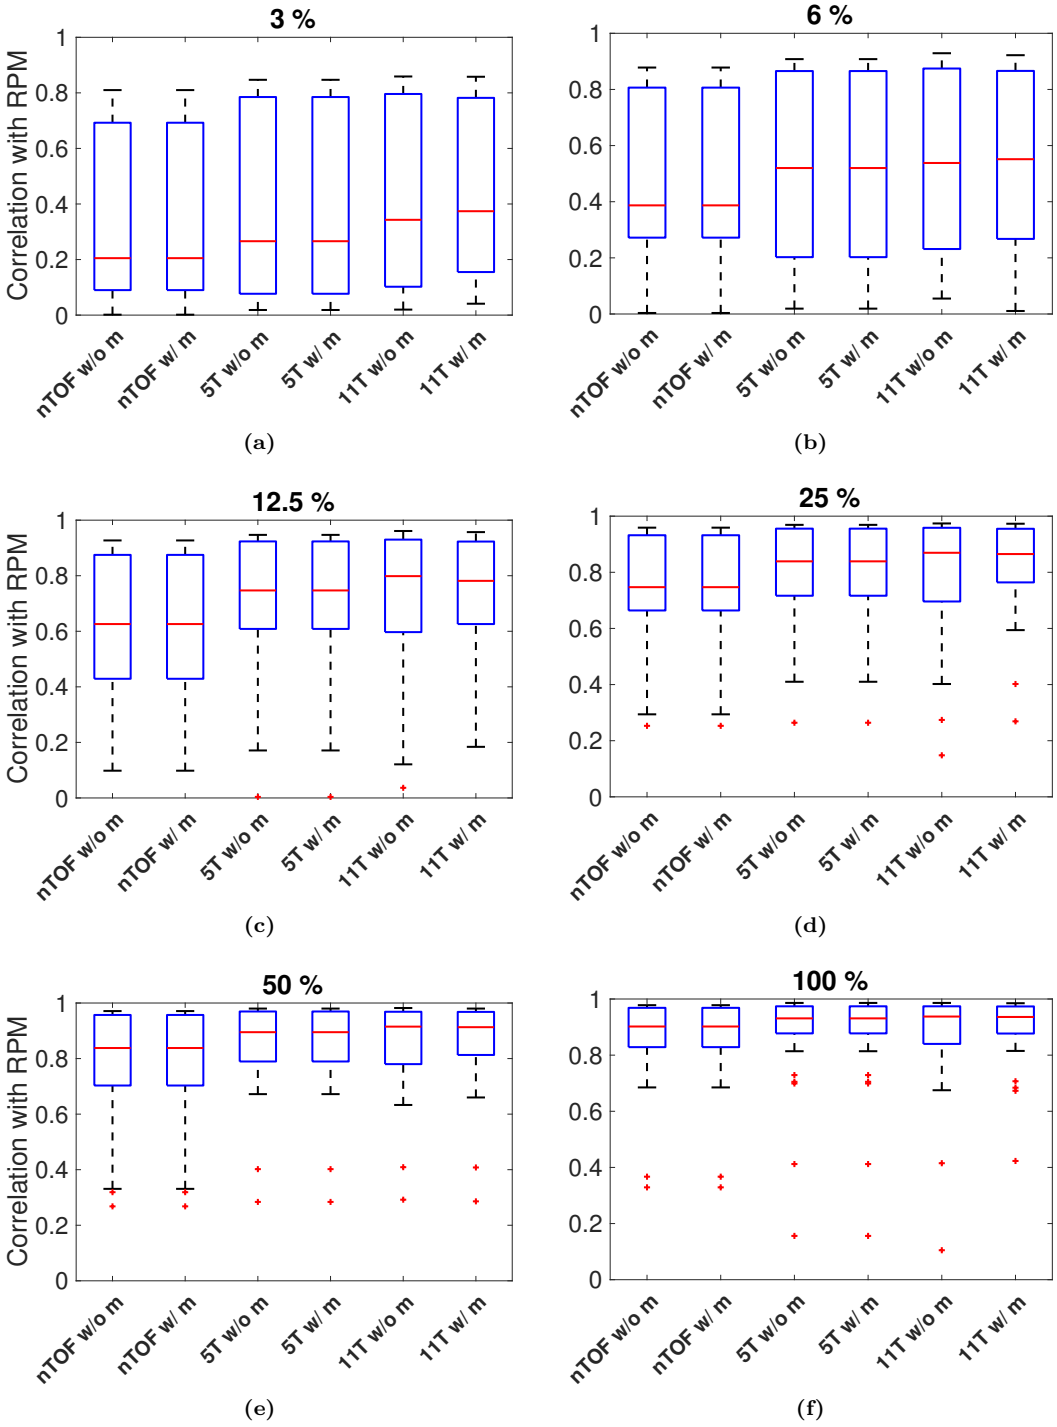

Figure 5.13: Boxplots of the correlation of the PCA signal with the RPM for all patients, with respect to the percentage of utilised counts. It can be seen that at lower count levels, using the TOF cases, without and with the cylindrical masking, have a beneficial effect on the quality of the resulting respiratory trace.

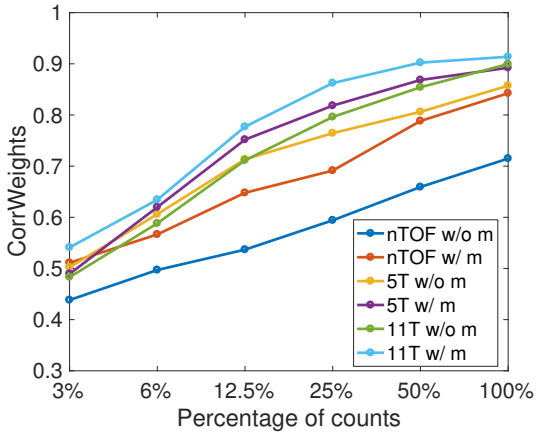

Figure 5.14: Mean of the output values of *CorrWeights*, averaged over all patients, for all count levels.

Table 5.8: Failure rate (in percentage) of *CorrWeights* in determining the correct direction of motion, on all the 28 patients, for the various percentages of utilised counts.

%	non-TOF w/o m	non-TOF w/ m	5TOF w/o m	5TOF w/ m	11TOF w/o m	11TOF w/ m
3	14	10	7	4	4	0
6	4	20	7	4	4	0
12.5	18	7	4	0	0	0
25	21	7	7	0	0	0
50	7	7	4	0	0	0
100	0	0	0	0	0	0

5.2.4 Discussion

In this section the PCA respiratory signals and the performance of the sign-determination methods have been analysed with respect to the duration of the data interval taken into account, and the amount of utilised counts. The results suggest that using TOF sinograms slightly improves the quality of the respiratory signal, evaluated via Person correlation with the RPM signal, when intervals of different durations are utilised, whereas the application of the cylindrical masking did not provide appreciable differences. The performance of the sign-determination method of choice (*CorrWeights*), with data intervals of different duration, improves when TOF data is used, furthermore also the application of the cylindrical masking proved to be beneficial. When using datasets with decreased number of counts, the TOF information and the sinogram masking increase both the correlation of the respiratory signal with the RPM and the *CorrWeights* values. The sign-determination method, if used with the 5 or 11 TOF bin sinograms with the masking, provided a zero failure rate for all the count levels apart from the most extreme cases (3% and 6% total counts). Generally, the impact of using TOF sinograms, when applying PCA and the proposed sign-determination method, has

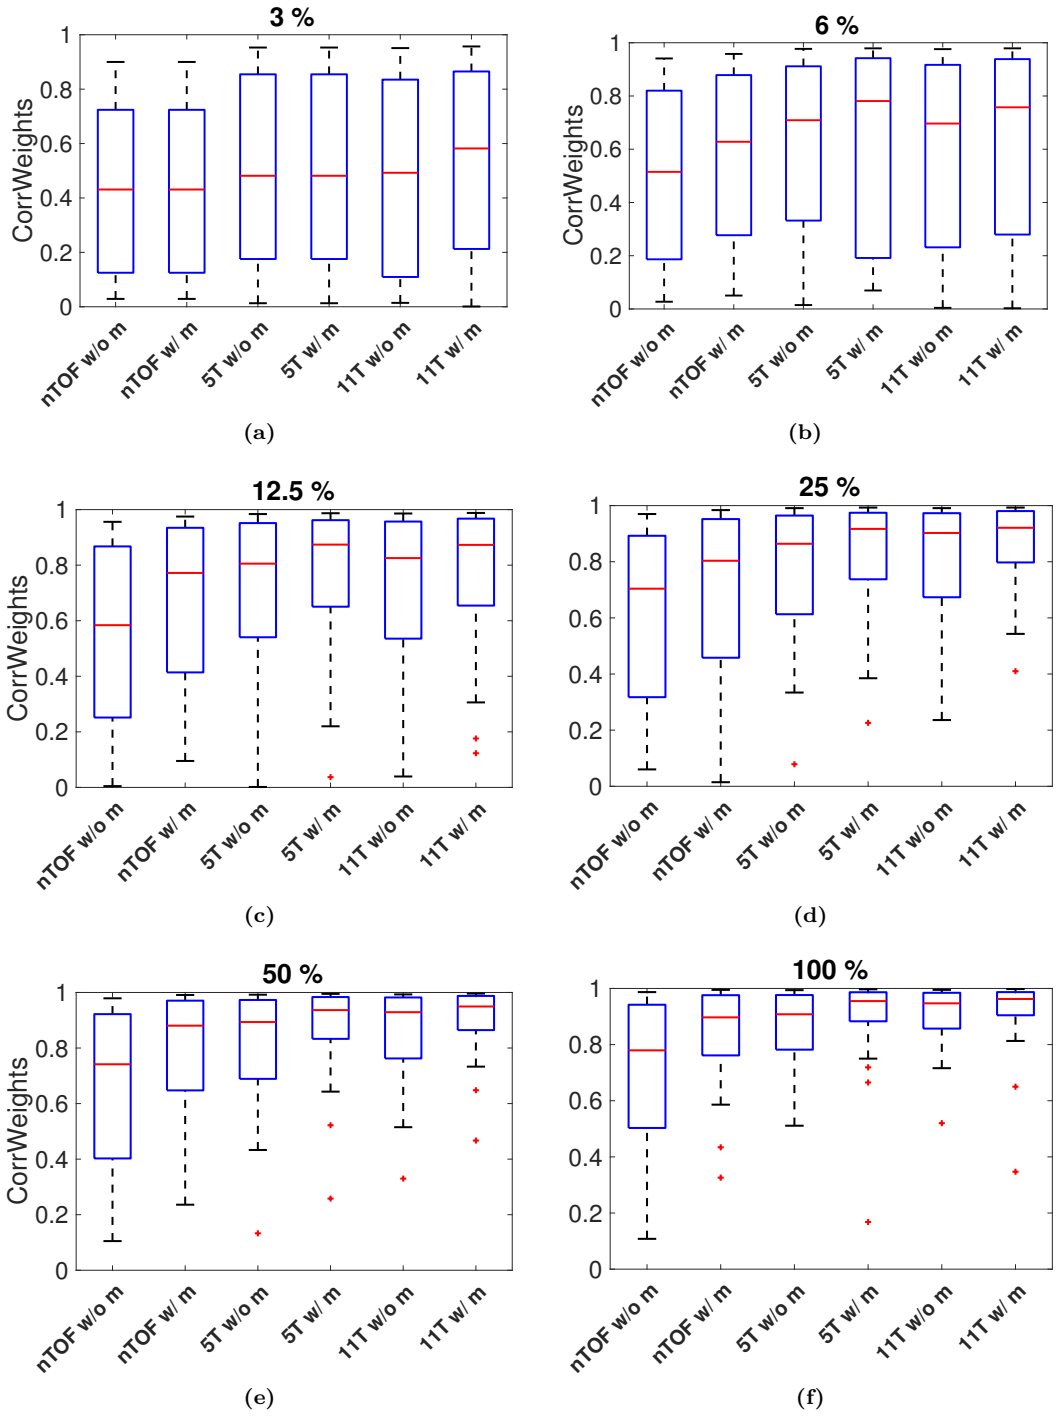

Figure 5.15: Boxplots of the *CorrWeights* values with respect to the percentage of counts for all the various data configurations.

been found to be more significant when using lower counts, as opposed to shorter intervals durations. The reason for this is expected to be related to the well known gain in Signal to Noise Ratio (SNR) when using TOF data as opposed to non-TOF, as discussed in Conti (2011), where the TOF SNR gain is described as a sensitivity amplifier therefore enabling

better results with lower counts.

5.3 Respiratory gated reconstructions

In this part of the study, the quality of the PCA signal for gating is evaluated with respect to the detected displacement in gated reconstructed images, compared to the results obtained when using the RPM signal.

When respiratory motion correction is applied in PET image reconstruction, the respiratory signal is utilised to gate the data into bins, with respect to either the phase or the amplitude of the breathing cycle, see 2.3.1.1. Therefore it needs to be verified that the final result of this process is comparable in the two cases (using RPM or PCA signal). All of the 28 patients datasets have been utilised for this analysis.

5.3.1 Respiratory gating and reconstruction: RPM vs PCA

For each patient, the gating was applied to the whole duration of the PET study. This was done in order to achieve the best possible quality in the reconstructed images for each patient. The data were gated both with the RPM signal and the PCA respiratory signal.

As the difference between using non-TOF and TOF sinograms was investigated in the previous analysis, the chosen data configuration for the PCA application was that with 5 TOF sinograms with the use of the cylindrical masking. In these conditions the *CorrWeights* method provides a null failure rate (see Table 5.8 for the 180s case, and all the other total durations were individually inspected), therefore the direction of the generated signal is consistent with that of the RPM. The 5 TOF bins configuration was chosen over the 11 TOF bins also because it require less memory allocation. The specifics of the gating are the following:

- The full amplitude range of the respiratory trace (either the RPM or PCA signal) was utilised.
- The data were gated with *displacement gating* into 5 gates.
- The gate boundaries were selected such that each gate has the same temporal duration ("equal time" mode).

The choice of the "equal time" mode guarantees that each gate has approximately the same number of counts, therefore similar noise characteristics. This is desirable for the current evaluation, as a registration process will be applied on the reconstructed gated data. Furthermore, as no threshold is applied on the amplitudes of the signals prior to gating, this setting ensures the gating process is less sensitive to outliers in the signal amplitude.

The gating generates one TOF sinogram per gate. Each sinogram was fully spatially sampled (no mashing nor SSRB), therefore the dimensions are $553 \times 381 \times 288 \times 11$ (planes $\times$

$r \times \phi \times$ TOF bins). The number of TOF bins in the sinograms for the reconstruction was chosen to be 11 to achieve better quality images.

To be noted is that the temporal resolution of the RPM and PCA signals are not the same. The RPM has much higher temporal resolution, equal to 33.4 ± 0.5 ms, whereas the PCA signal, which is obtained from sinograms unlisted to 500ms frames, is then upsampled with linear interpolation to 250ms. In this evaluation, the RPM signal was not downsampled to the lower PCA signal resolution, as the aim was to make use of the two signals at their full potential.

Each sinogram was reconstructed individually, using an OSEM algorithm, with 28 subsets and 2 iterations. As both attenuation and scatter correction require the use of the CT, and misalignment between the CT and the PET gate can cause artefacts, the reconstructions were performed without these corrections. Only random correction and sensitivity normalisation were applied. No postfiltering was applied to the reconstructed images. The TOF OSEM reconstruction was performed with proprietary software provided by General Electric, implemented in MATLAB.

5.3.2 Displacement detection

The chosen metric to compare the reconstructed gates produced with RPM-gating or PCA-gating is the *detected displacement* in the z-direction. This is obtained through a registration process, performed as follows, see also Figure 5.16:

- the 1^{st} gate (that correspond to end-inhalation) is chosen as the reference gate and the other 4 gates are registered to it;
- a point in the reconstructed image of the 1^{st} gate is manually selected, placed onto a feature that is evidently moving throughout the gates (close to the heart, the liver, the diaphragm or a visible lesion, an example is shown in Figure 5.16) and a cubic volume of $10 \times 10 \times 10$ cm^3 is selected around it;
- the volume within the cube is registered between the gates, making use of the `spm_realign` function of the Statistical Parametric Mapping (SPM) software[5] (Friston et al., 1995) implemented in MATLAB.

The `spm_realign` SPM algorithm realigns a series of images of the same subject (in this case a series of gates of the same volume within a patient), using a least squares approach and a rigid body spatial transformation, producing as an output the three angles defining the rotation and the three parameters of the translation. The algorithm default settings were utilised in this analysis, which include the choice of a Gaussian smoothing kernel (with a Full Width at Half Maximum (FWHM) of 5 voxels) to be applied to the images before

[5] http://www.fil.ion.ucl.ac.uk/spm

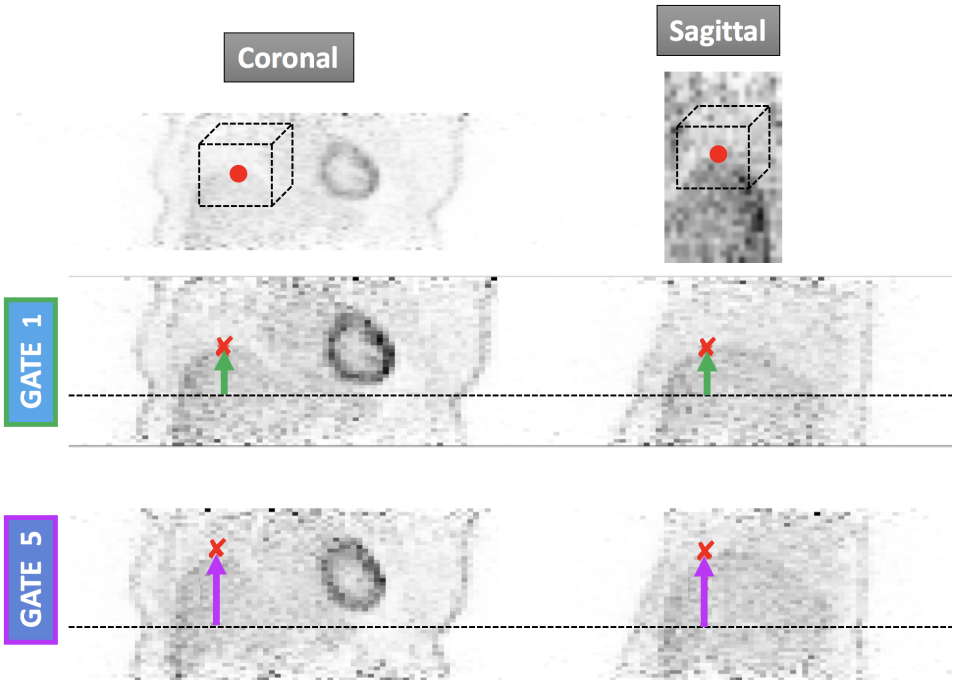

Figure 5.16: Example of the 1st and 5th gate of a patient. On the top row the representation of the process explained above: a point is chosen on the top of the liver, a cubic volume is selected and the registration operates on the voxels comprised within. The second and third row show the motion that occurred between the gates.

estimating the realignment parameters. For each of the 4 registrations the transformation matrix is then applied to the selected point, to obtain the motion that it undergoes between gates. The transformation matrix from the 1st to the nth gate is given by $M_{1,n}$:

$$M_{1,n} = \begin{bmatrix} A_{1,n} & \vec{b}_{1,n} \\ 0 & 1 \end{bmatrix} \quad ; \quad \begin{bmatrix} \vec{x}_{1\to n} \\ 1 \end{bmatrix} = M_{1,n} \begin{bmatrix} \vec{x}_1 \\ 1 \end{bmatrix}$$

where $A_{1,n}$ is a linear transformation, described by the rotation parameters, and $\vec{b}_{1,n}$ is a translation, $\vec{x}_1$ is a point in the 1st gate and $\vec{x}_{1\to n}$ is the new location of the point in gate nth. The displacement in the z-direction for each gate, compared to the reference one, is obtained as the z-component of the displacement vector between the gates:

$$(\vec{x}_{\text{disp}})_z = (\vec{x}_{1\to n} - \vec{x}_1)_z \tag{5.1}$$

5.3.3 Analysis

The displacement in the z-direction detected from the data gated with RPM and PCA was compared. The comparison was made using two different metrics, the maximum detected displacement and the displacement per gate, where the latter is the motion between each gate and the 1st gate.

5.3.3.1 Maximum detected displacement

This value only depends on the two gates between which the maximum displacement has been found (that are expected to be the 1st and 5th but this is not assumed) and provides information on the maximum range of internal motion captured by the utilised respiratory signal. This metric has also been used in the analysis presented in Thielemans et al. (2013), where different DD approaches were assessed.

5.3.3.2 Displacement per gate and linear fit

In order to make use of the result of the registration process related to all gates, the detected displacement of each gate (i.e. the motion that occurred between the 1st gate and the other 4 gates) is analysed, with respect to the mean value of the respiratory motion signal within each gate. This provides additional insight into the performance of the gating process applied with the two respiratory traces, and it does not reply solely on the motion detected between two gates.

As the boundaries of the gates are chosen such that each of them have approximately the same amount of counts, the motion within each gate is likely to be different among the gates. In Figure 5.17 a schematic representation of displacement gating in shown. It can be seen that the width of the 1st and 5th gates is larger than that of the central gates. This is because these gates comprise the highest and lowest peaks of the signal, which are not always reached during each cycle, therefore in order to contain the same amount of counts as the other gates they need to include a wider range of motion. In the figure there are 65 points and the gates' width is chosen such that there are 13 points in each of them.

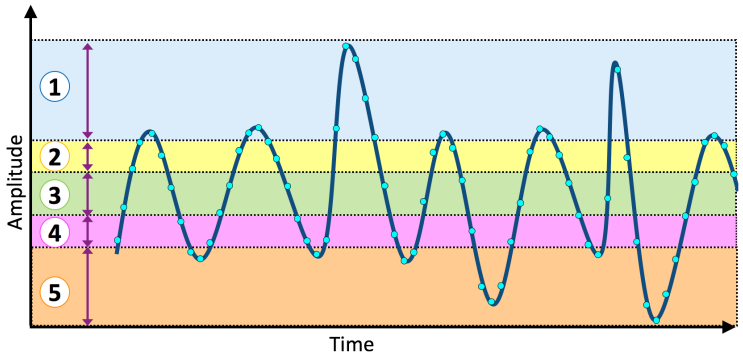

Figure 5.17: Depiction of a respiratory signal with amplitude gates and their boundaries when using *equal-time* gating. In this illustration, the 1st and 5th gates are wider than the others, as they include the highest and lowest peaks, respectively, of the respiratory signal, which are not reached in every cycle. In order to contain the same number of counts (i.e. same number of points of the respiratory signal, like the dots shown on the trace) they need to include a larger range of the signal.

Even when the respiratory signals from the RPM and PCA are very highly correlated, the width of the resulting gates might differ, and consequently also the mean value of the

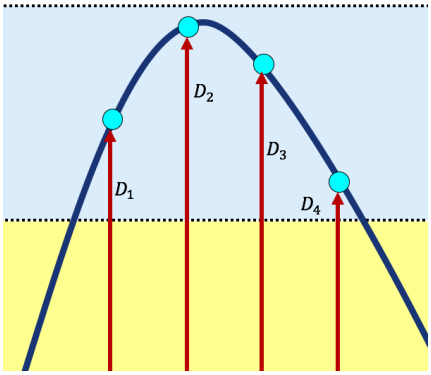

Figure 5.18: Value of the respiratory signal within one gate for 4 points in a respiratory cycle.

motion within them. For example, since the whole range of amplitude values of the signals is utilised for the gating (see 5.3.1), if there are spikes or outliers due to noise or to the poor performance of either of the methods, the borders of the gates will be affected by this. In order to perform an analysis that is not sensitive to this difference, as the aim of this investigation is to compare the amount of detected motion independent of the width of the gates, the displacement per gate was analysed with respect to the *mean value of the respiratory motion signal per gate* (abbreviated with *mean motion per gate*), which is defined as follows:

$$\langle \mathrm{D} \rangle_i = \frac{\sum_{k \in G_i} \mathrm{D}_k}{K_i} \tag{5.2}$$

where G_i contains all the samples of the respiratory signal that fall in gate i, K is the total number of time points in that gate and D_k is the value of the trace at the point k, see Figure 5.18.

The values of the signal produced by the PCA method do not have a known correspondence to the physical motion, as they are in arbitrary units. They are also dependent on the contrast in the data, as shown in Eq. 3.13, where the output of PCA is approximated with the multiplication of temporal signal with a factor that depends on the average of the data. Therefore, in order to make the comparison between the RPM and PCA case consistent, the *mean motion per gate* was evaluated after the standardization of the signals, where the mean value (over time) is subtracted and the result is divided by its standard deviation:

$$Z = \frac{X - \langle X \rangle}{\sigma(X)} \tag{5.3}$$

Finally, to utilise the information provided by the 5 gates all together, a linear fit was applied to the detected displacement in the z-direction obtained from Eq. 5.1 with respect to the *mean motion per gate*. The motivation for this analysis is that it is expected that both the RPM and the PCA signal are approximately proportional to motion. This has also been

shown in the first part of this chapter, where the internal motion extracted from the MR has been found to be highly correlated to the PCA signal, except for an arbitrary scale factor. The slope of the fit was then chosen as the parameter to be examined, as it provides general information on the amount of detected motion: the higher its value, the wider the motion that occurred between the gates. An example can be seen in Figure 5.23.

In summary, the analysis was performed with respect to the following 2 metrics:

1. comparison between the maximum detected displacement in the RPM-gated and PCA-gated images;

2. comparison of the slope of the linear fit of the detected displacement per gate with respect to the mean value of the respiratory motion within each gate (see Eq. 5.2).

5.3.4 Results

The values of the correlation between the PCA signal and the RPM are shown in Figure 5.19 for all patients, as a reference for the next sections.

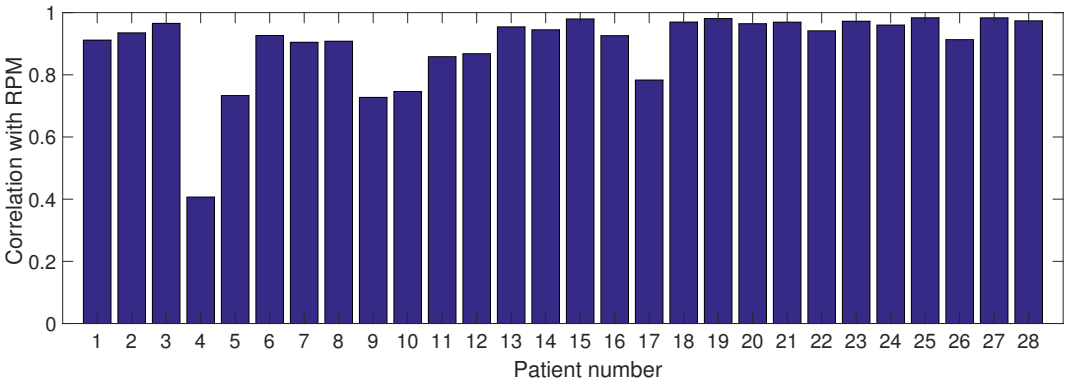

Figure 5.19: Correlation of the PCA signal with the RPM for all the 28 patients.

5.3.4.1 Maximum displacement

Figure 5.20 shows an example of the result of the gating and reconstruction process, when using the PCA respiratory signal (similar looking images are obtained with the RPM signal). This patient presented a very well defined lesion that showed considerable motion between the gates, therefore the point required for the registration and displacement detection was located at its centre. Lesion and heart motion is noticeable from gate 1 to gate 5.

In Figure 5.21 the maximum detected displacement for each patient is shown, for the PCA-gated and RPM-gated images. The wide range of values is to be expected, as the amplitude of the breathing motion is very patient-dependent. Furthermore, the results are also dependent on the point that is utilised for the registration and subsequently for the determination of the displacement within gates (see Eq. 5.1), which is not placed in the same

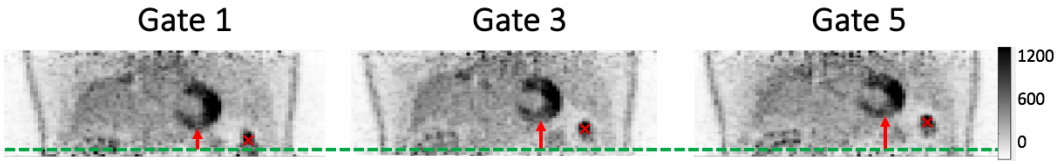

Figure 5.20: Example of three of the five reconstructed gates, using the PCA respiratory signal for the gating process. The red cross shows the selected point for registration, which for this patient corresponded to a lesion. The superimposed dotted line is for reference.

location for each patient (but the same point is used for each patient for both the RPM and PCA-gated images).

The observed displacements when using RPM or PCA gating are similar, except for a few cases that present notable differences: patients 10, 19 and 20, for which the use of the RPM signal produced a maximum displacement that is 3.7mm, 5.4mm and 8.5mm lower than in the PCA case. To be noted are the values for patient number 1, where no motion was detected with either of the methods. This is most likely due to the very low quality of the reconstructed image, where no distinct features were identifiable, whereas the correlation between the RPM and the PCA signal was equal to 0.91.

In 15 out of the 28 analysed patients, the use of the PCA respiratory signal provided a higher detected maximum displacement compared to the RPM case, and the cases where the difference between the two values is more significant (patients 6, 10, 19 and 20), the PCA gating produced the highest detected displacement. A Z-test was performed in order to evaluate whether the results from the PCA-gated images are statistically greater than those from the RPM-gated images. More specifically, the random variable given by the mean of difference between the maximum detected displacement in the PCA case and those in the RPM case was tested. The hypothesis of this variable been centred around a mean value of zero or lower is rejected with 99.14% confidence. The PCA gating therefore provided values that are statistically higher than those of the RPM.

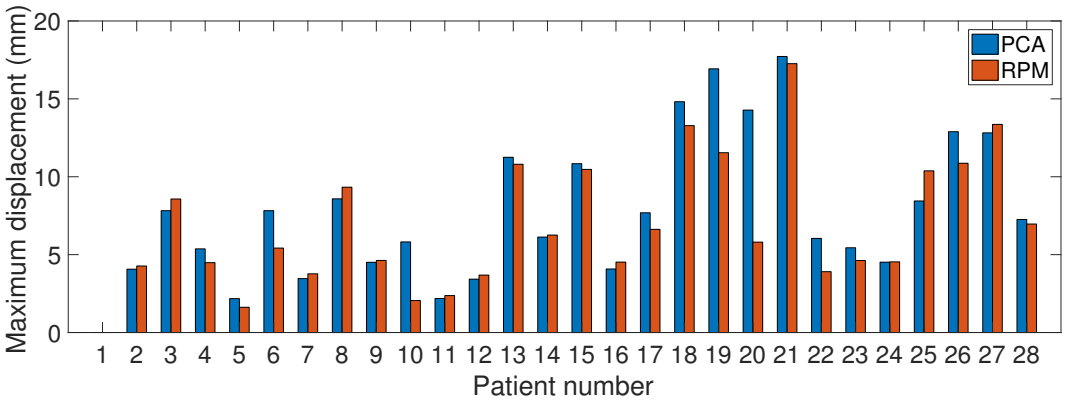

Figure 5.21: Maximum detected displacement in millimetres for each patient, from the gated images generated using the PCA or the RPM signal.

5.3.4.2 Displacement per gate and linear fit

Figure 5.22 shows the slope of the linear fit for all patients, in the PCA and RPM case. The trend of the values is very similar to that of the maximum detected displacement seen in Figure 5.21. This was to be expected as a large maximum detected displacement is presumed to be related to a steep linear fit, therefore resulting in higher slope values, conversely a low displacement is going to result in lower slope values. An example of the displacement per gate, with respect to the mean motion within each gate, can be seen in Figure 5.23. The plots show the results for three patients for which the maximum displacement was found to be the same in both cases (PCA or RPM gating), therefore comparable results were expected. The trend is in fact very similar, implying that the motion captured by the PCA signal was consistent with that of the RPM throughout the whole study. The correlation between the PCA and RPM signal for these patients is 0.98, 0.97 and 0.91 respectively, see Figure 5.19.

Figures 5.24 and 5.25 show in detail the results obtained for patients 4 and 17, for which there is a notable difference in slope value, see Figure 5.22, but the maximum detected displacement is instead more similar (within 1mm) as can be seen from Figure 5.21. Figure 5.26 shows the results for patient 10, for which both slope values and maximum detected displacement have been found to be very different between the PCA-gated and RPM-gated images.

A similar Z-test as in the previous section was performed in order to evaluate whether the results from the PCA-gated images are statistically greater than those from the RPM-gated images. The hypothesis of the variable given by the mean of difference of the two values (PCA and RPM related) been centred around a mean value of zero or lower is rejected with 99.99% confidence. The PCA gating therefore provided slope values that are statistically higher than those of the RPM.

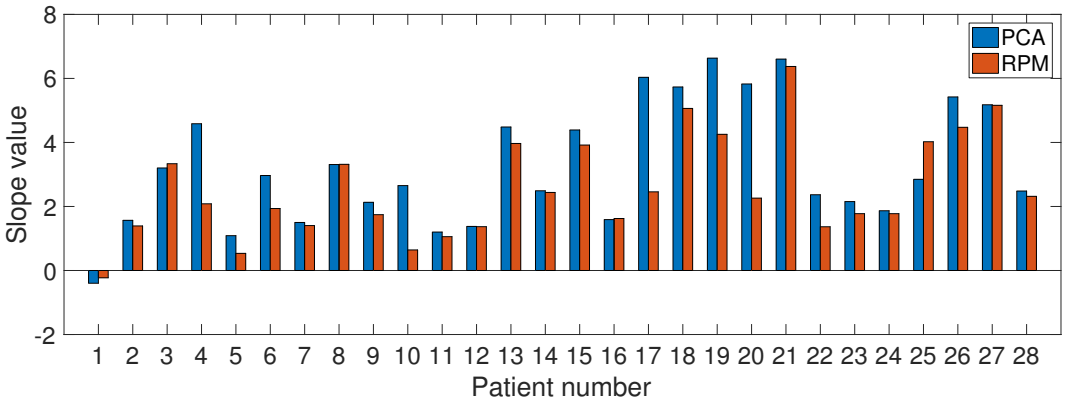

Figure 5.22: Slope of the linear fit applied to the detected displacement with respect to the mean motion per gate. Values for all patients for the gated images obtained with PCA and RPM signal.

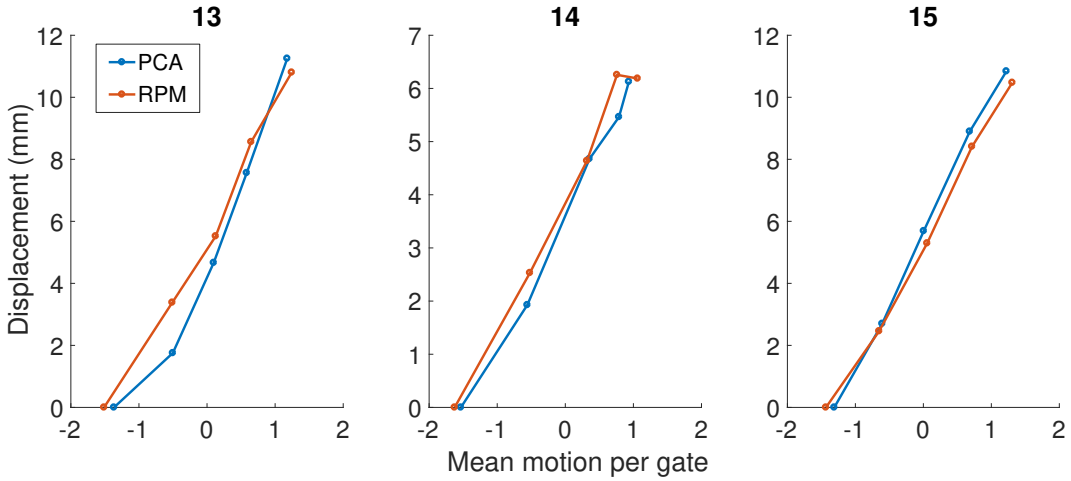

Figure 5.23: Plot of the displacement per gate with respect to the mean motion per gate for patients number 13, 14 and 15 (from Figure 5.21). The central gate, which corresponds to the 3rd gate, has *mean displacement* equal to zero. This is a consequence of the application of standardization of the respiratory signals, see Eq. 5.3.

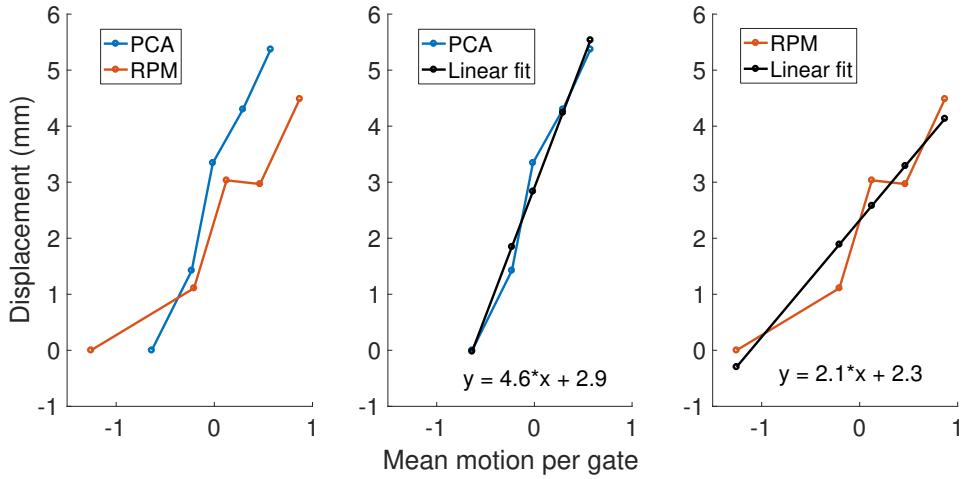

Figure 5.24: Detected displacement per gate for patient number 4 (see Figure 5.22), for which the slope value obtained from the linear fit was very different between the PCA and RPM gates: 4.6 and 2.1 respectively. Maximum displacement values: 5.4mm (PCA) and 4.5mm (RPM).

5.3.4.3 Detailed analysis of patients with notably different PCA and RPM results

In order to further investigate the patients for which the slope resulted in noticeably different values in the two cases, the assignment of the data to the respiratory gates was examined. The gate assignment consists in allotting each temporal frame of the data to a specific gate, depending on the value of the respiratory signal at that time.

In Figure 5.27 the respiratory signals, and the gate assignment they generated, are shown for patient 4. The correlation between the signals is 0.40. The difference in density of points between the PCA and RPM signal is given by the different temporal sampling,

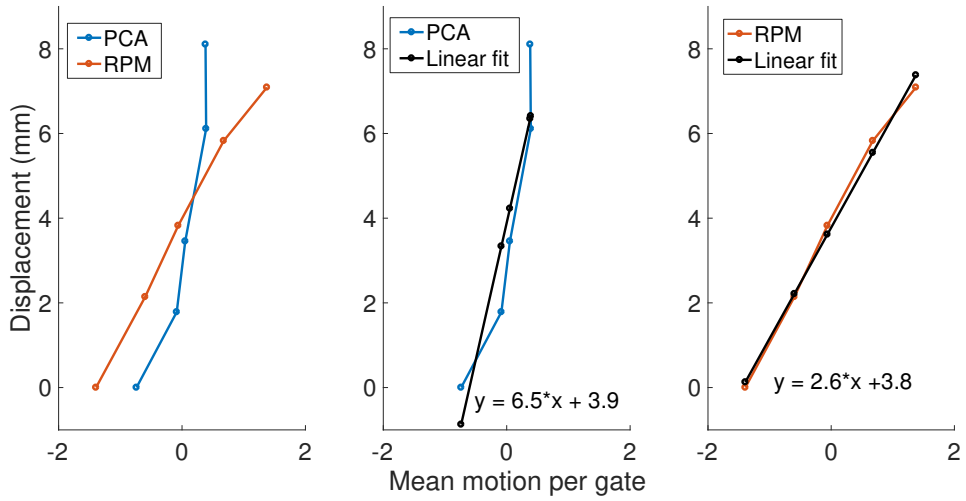

Figure 5.25: Detected displacement per gate for patient number 17 (see Figure 5.22), for which the slope value obtained from the linear fit was very different between the PCA and RPM gates: 6.5 and 2.6 respectively. Maximum displacement: 6.6mm (PCA) and 7.7mm (RPM).

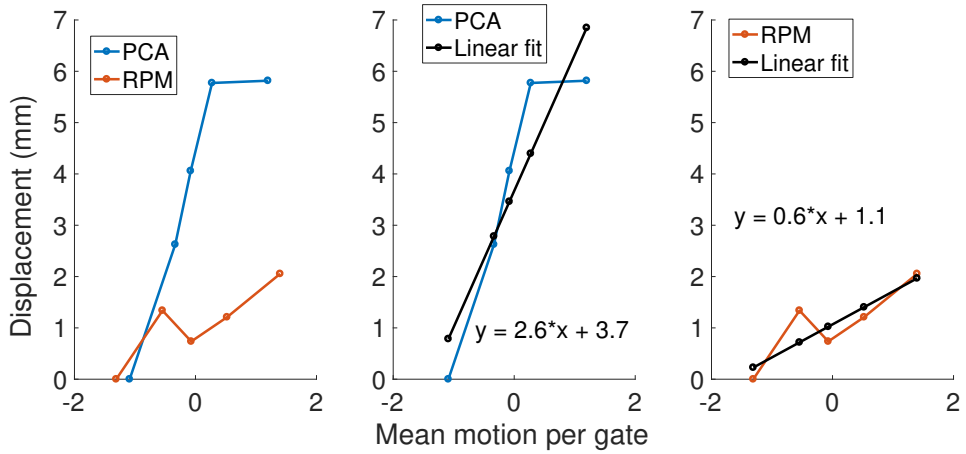

Figure 5.26: Detected displacement per gate for patient number 10, for which both the slope value and the maximum detected displacement were found to be very different (see Figures 5.22 and 5.21). Maximum displacement values: 5.8mm (PCA) and 2.0mm (RPM).

that is much higher in the latter case (∼33ms compared to 250ms), and this applies to all the following images. In the first 40s of acquisition a considerable difference can be noted between the PCA and RPM signal gate assignment, as the 1st gate is filled when using the PCA signal whereas it is not when using the RPM signal, see Figures 5.27a and 5.27b. This is due to the way the gate boundaries have been determined from the gating process, in order to spread the counts equally among them. The detected motion per gate is also affected, and results in the very different linear behaviour seen in Figure 5.24 (4.6 and 2.1 slope value for the PCA and RPM gated case respectively). The maximum detected displacement is instead almost the same in both cases (5.4mm and 4.5mm), because later on in the acquisition the 1st gate is eventually filled also when using the RPM signal. The second half of the acquisition

presents a much similar gate assignment between the two cases.

The correlation between the two signals, as previously mentioned, is 0.40. This value would have normally indicated a PCA signal of poor quality, as the RPM is considered the gold-standard signal, whereas it is possible to see that in this instance the dissimilarity between the two traces is mainly due to a short interval centred around the time equal to 80s, where the RPM signal presents a downward spike, apart from which the two signals are comparable.

Another example is that of patient 10, for which both the maximum detected displacement and the slope of the linear fit are very different in the PCA and RPM case (slope values 2.6 and 0.6 and maximum displacement 5.8mm and 2.0mm respectively). In Figure 5.28 the respiratory signals and the gate assignment are shown. This patient presented a particularly irregular breathing pattern. Nevertheless, the correlation between the RPM and PCA signal is 0.75. A significantly different gate assignment can be noted in the interval ranging from 25s to 75s of the acquisition, where in the RPM case only the 4^{th} and 5^{th} gate are filled, whereas in the PCA case the signal amplitude spans from the 2^{nd} to the 5^{th}. This seems to be caused by the RPM signal being shifted downwards during that interval, as opposed to the rest of the acquisition. Even though the signal presents a respiratory-like shape, its amplitude is only subdivided in two gates (mostly assigned to the 5^{th} gate) therefore the motion that occurred is expected to be blurred within those two gates. In this case the differences between the gate assignments with RPM or PCA signal are so significant that, as opposed to the case of patient 4, not only the slope but also the maximum detected displacement is noticeably different: 5.8mm with the PCA signal and 2.0mm with the RPM.

The results for patient 17 are shown in Figure 5.29. In this case the correlation between the signals is equal to 0.78. For this patient the slope and maximum detected displacement are 6.5 and 2.5, and 6.6mm and 7.7mm for PCA and RPM respectively. The RPM respiratory signal presents some notable shifts in value: over the breathing cycles from 150s to 250s the maxima and minima of the signal are considerably lower than those in the preceding and consecutive part of the acquisition. In correspondence to that interval, the 1^{st} gate is not filled and the motion is mostly subdivided among three gates. This effect is not replicated by the PCA signal, which presents a much more stable behaviour throughout the whole duration of the study and does not present any clear shift, and the values of the signal are assigned to the 5 gates more evenly when compared to the RPM case.

One last example is shown in Figure 5.30, with the results for patient 20. In this case the correlation between the RPM and PCA signal is extremely high, equal to 0.96. Nevertheless the amplitude of the 5^{th} gate is much larger in the RPM case, due to the the fact that the minima of the signal reach much lower values compared to the PCA signal (the signals are shown after standardization). As a consequence, this gate is likely to include a wider range

of motion when compared to the PCA equivalent gate, therefore causing it to be averaged and hindering the detection of displacement between it and the 1st gate. This has an impact on the detected maximum displacement and slope values, which have been found to be significantly different: 14.3mm and 5.8mm, 5.8 and 2.3 for PCA and RPM respectively.

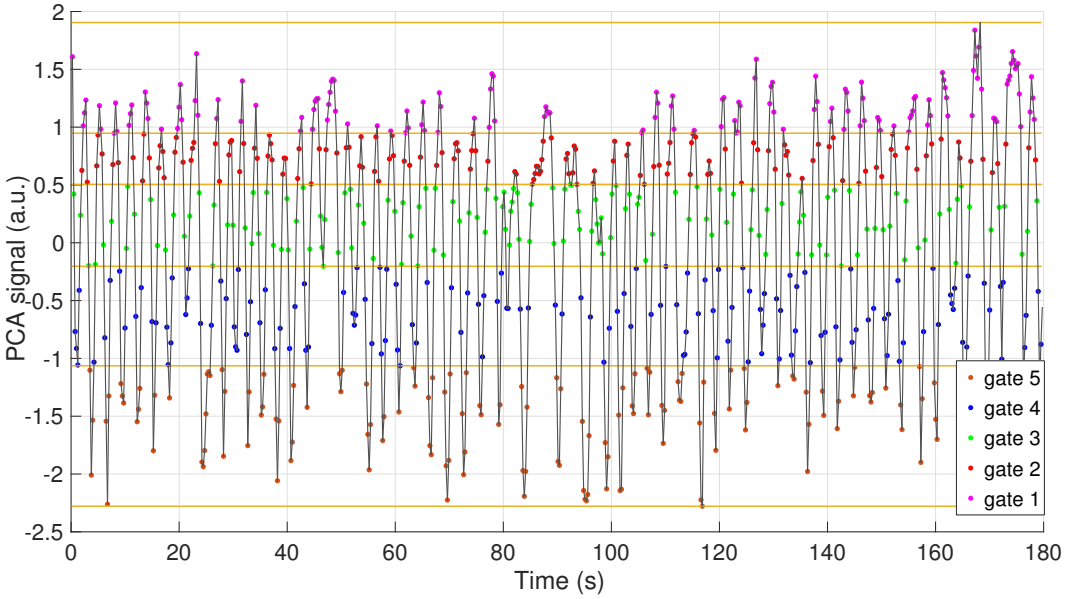

(a)

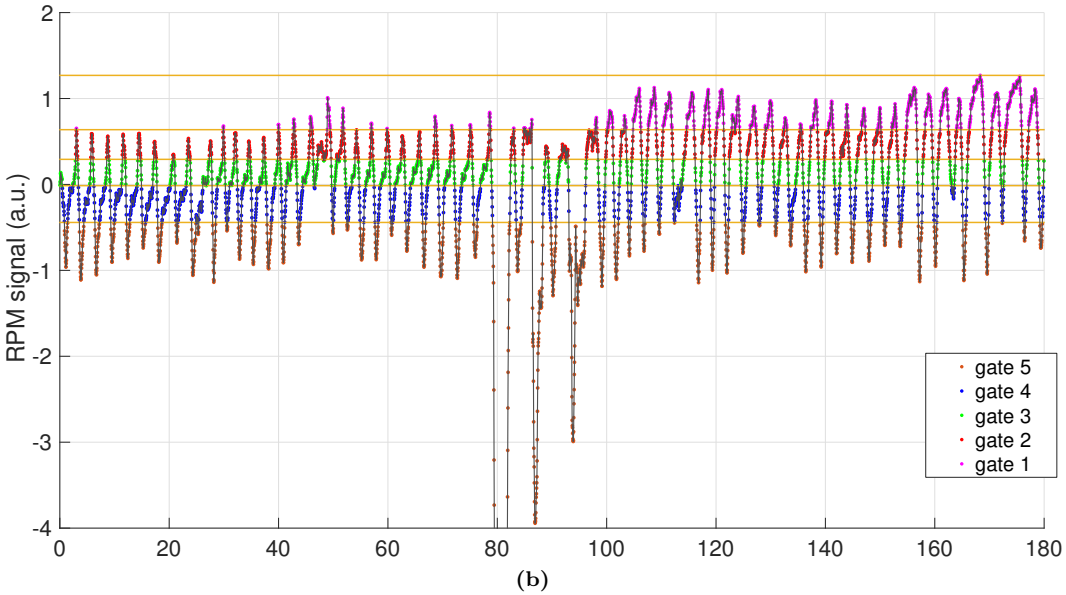

(b)

Figure 5.27: Results for patient 4. (a) PCA signal with the gate boundaries superimposed. It can be seen that each point of the signal, depending on its value, is assigned to the gate whose boundaries include that point. The temporal frame of the dataset related to that point is going to contribute to that gate. (b) Same information displayed for the RPM signal case. The downward peak at 80s is cut (it reaches value -8) in order to better show the shape of the rest of the signal.

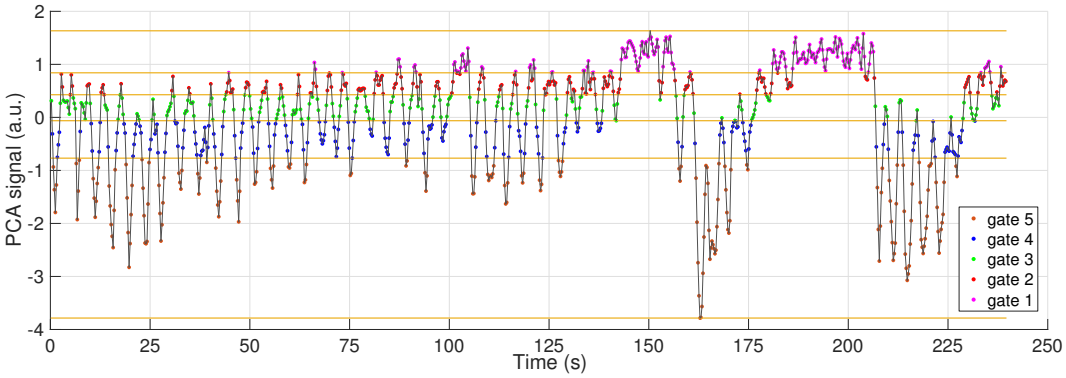

(a)

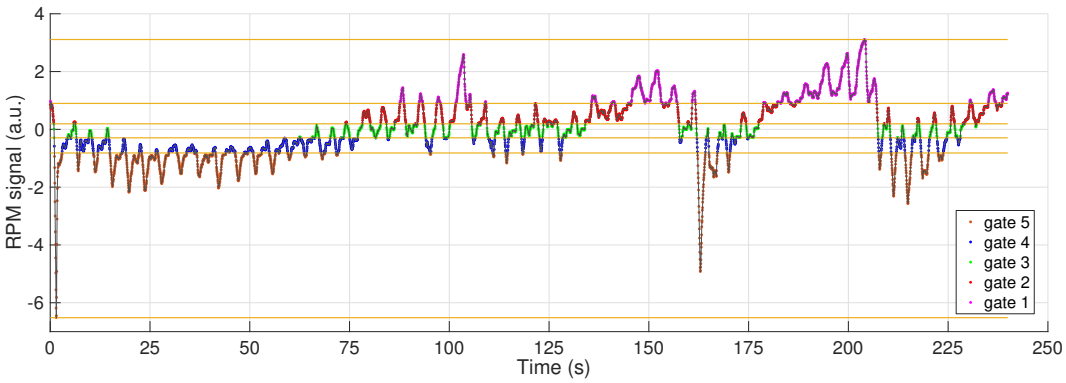

(b)

Figure 5.28: Results for patient 10. (a) PCA and (b) RPM signals and related gate boundaries.

5.3.5 Discussion

The aim of this part of the study was to investigate whether using the PCA signal for respiratory gating is able to produce comparable results to using the RPM. The same processes were applied to gate and reconstruct the data using the two signals. The maximum detected displacement, and the relationship between the detected motion per gate and the mean motion within each gate, were analysed. For the latter, a linear fit was applied and the slope value compared between the two cases.

The maximum detected displacement and the slope values proved to be similar in most of the evaluated patients datasets, with some exceptions, where the PCA gating produced values that were greater than those of the RPM gating.

Some of the patients for which the slope values were found to be very different in the two cases were analysed in more detail. For patient 4, the RPM signal presents a sharp peak at around 80s, that has no counterpart in the PCA signal. In the interval prior to that event the RPM signal never reaches the 1st gate, whereas the PCA signal assigns samples to all gates. This results in the partition of respiratory motion in effectively fewer gates, therefore in less detected motion within gates.

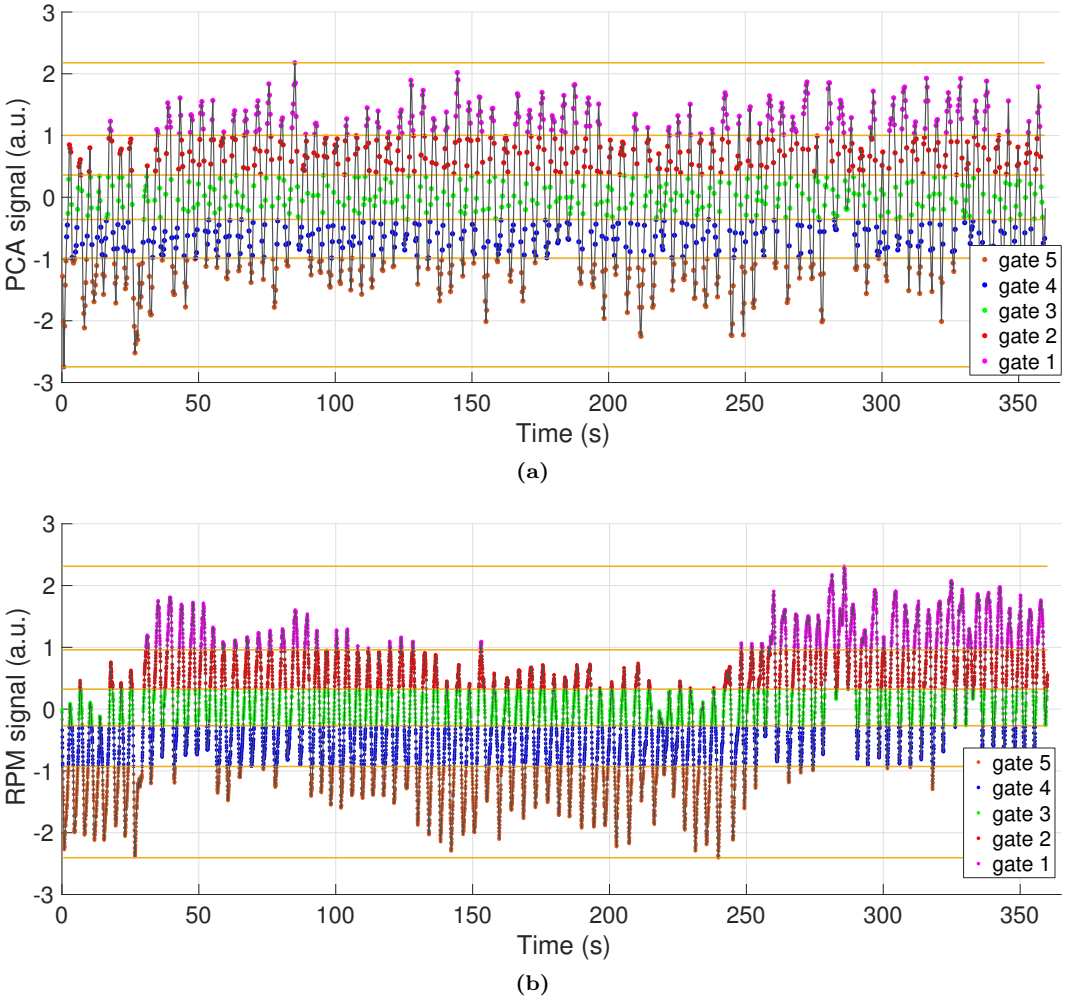

Figure 5.29: Results for patient 17. (a) PCA and (b) RPM signals and related gate boundaries.

For patients 10 and 17, the considerable difference in results between PCA and RPM case seems to have been caused by the presence of a drift in the values of the latter, which could have been due by the RPM sliding down the patient's chest. Patient 10 presents a very irregular breathing pattern, that is successfully captured by the PCA method. In the first part of the acquisition the RPM signal shows a drift that results in the assignment of most of the data frames to the 5^{th} gate, therefore hindering the differentiation of the motion among more gates. It is however impossible to determine whether there actually was a drift in the patient's motion, as there is no ground truth available to compare with.

The practice of utilising the slope value of the linear fit as a comparison metric between respiratory signals has been suggested for the first time in this work (as far as the author is aware at the time of writing). This procedure, as it has been shown, allows to compare the signals for the presence of drift in the signal, the knowledge of which is of crucial importance when constructing motion models that depend on measured respiratory signals as surrogate for the real patient motion (McClelland et al., 2013).

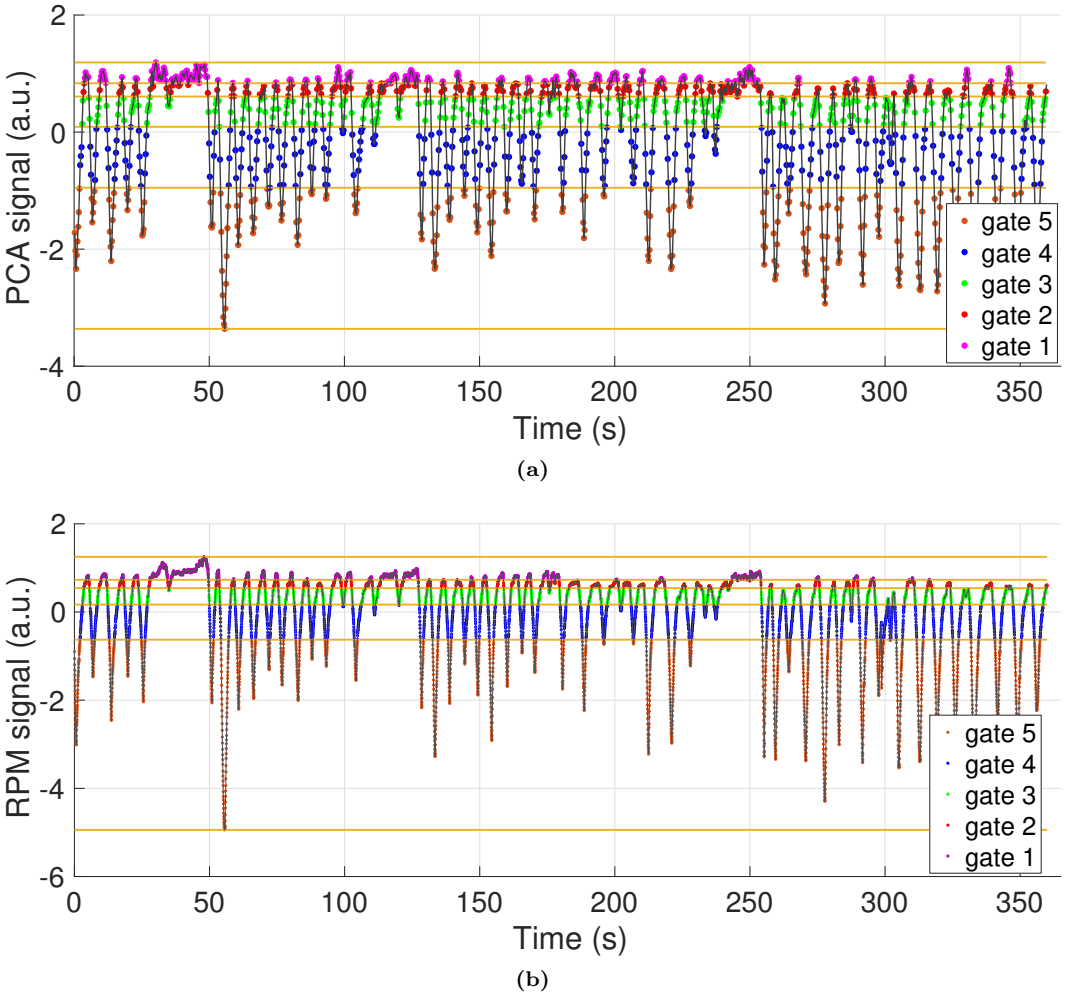

Figure 5.30: Results for patient 20. (a) PCA and (b) RPM signals and related gate boundaries.

It is to be noted that in this study the impact of choosing different gating schemes was not investigated, and the results presented here are relative to the set of parameters that has been described. Results might be sensitive to the gating scheme, and some adjustments, such as outlier rejection prior to the calculation of the gate boundaries, or the choice of phase-gating instead of displacement gating, could improve the ability to detect motion within the gates. This is outside the scope of this analysis.

5.4 Conclusions

In the first part of this chapter, the PCA respiratory signal has been compared to the internal respiratory motion extracted from MR dynamic data, and to an external device signal (pressure belt). The comparison showed that these three signals are highly correlated with each other, therefore suggesting that PCA correctly represents the internal motion that occurs during respiration. Moreover, it confirmed that the external motion of the chest determined by the belt is representative of the internal motion, therefore comparing the

PCA signal with that of an external device is a reliable method to ascertain the quality and correctness of the former. In addition, in the utilised patients' datasets (which were FDG acquisitions performed at least 60 minutes after injection), no drift has been found between the PCA signal and the internal movement.

In the second part of the chapter a thorough analysis of PCA on PET data has been performed on a large cohort of patients. The studies were monitored with the RPM device and its signal was used for comparison. The quality of the respiratory signal and the performance of the proposed sign-determination methods were evaluated. The analysis showed that the use of TOF information, together with the application of a cylindrical mask, improve the overall quality of the PCA respiratory signal and of the sign-determination methods, providing results that are more robust and reliable over shorter acquisition durations and especially for fewer counts. With regards to the two presented sign-determination methods, *CorrWeights* has been selected as the method of choice. Its output values are higher than those of *CorrSino*, suggesting a more robust performance, furthermore it can be utilised to correct the direction of motion of any DD signal, whereas the latter is only applicable to PCA.

Finally, the analysis on reconstructed data provided a thorough comparison between the use of PCA and RPM signal for gating purposes, in terms of maximum detected displacement and linear fit of the detected motion with respect to the mean motion within each gate. The results showed that the use of the PCA signal resulted in most cases in similar maximum detected motion and slope values as when using the RPM signal. For the patients where more considerable differences where found, the PCA gating provided the higher detected maximum displacement and higher slope value. The analysis of the gating assignment showed that, for those patients, the RPM signal presented drift in values that affected the gate assignment, combining temporal frames that were most likely related to different parts of the breathing cycles, therefore generating gates where motion was averaged out.

In conclusion, this study showed evidence that the PCA method, together with the sign-determination method, is able to produce a reliable respiratory signal, highly correlated to the internal motion of the organs and capable of generating respiratory gates that successfully discriminate respiratory motion. These methods therefore provide a simple and cost-effective tool for motion management in PET.

Chapter 6

Detection of Lung Density Variations With Principal Component Analysis in PET

The analysis and results shown in the previous chapters focused on the application of the method to the PET data related to the region of the thorax near the diaphragm, where organ movement is the predominant cause of changes in the data. This is generally the case when considering the lower part of the lungs, and the area of the chest that includes the dome of the liver, the diaphragm and the lower part of the heart. In these cases, the FOV comprises features that move considerably, although always to an extent that is patient dependent, therefore respiratory motion is expected to be the major cause of change in the data, and consequently it is presumed to be detected by PCA.

The subject of this chapter is instead the application of the method to the data of the upper part of the chest, where respiratory motion is expected to be small. The aim is to investigate if there are significant respiratory-induced changes in the data of the upper lungs, and whether they are detectable with the use of PCA. The method is firstly applied to simulation data and secondly to patient data. The issue of sign-determination of the PCA signal is also addressed with patient data.

6.1 Introduction

During the breathing cycle the lung density changes as a consequence of the expansion and compression of the chest volume. In PET/CT imaging, the changes in density have an impact on the attenuation effect and on the radionuclide activity concentration, therefore inducing respiratory-related changes in the acquired data of both modalities. The effect of lung expansion on CT values has been extensively investigated, for example in Verschakelen et al.

(1993) where the density of various areas in normal lung is measured using CT. This effect has also been exploited to gain additional information as in Heremans et al. (1992). In their work CT values of manually drawn regions in the lung are shown to have positive correlation with the results of pulmonary function tests, with patients with suspected pulmonary malignancy. In Chen et al. (2017) an extensive review on PET quantification approaches in lung diseases is presented, and the density changes in PET and CT are described as significant contributions to errors in lung quantification.

In Cuplov et al. (in press) the effects of respiration in PET/CT imaging are analysed in connection to the application of respiratory gating and registration. Their results show that the variations in density (measured with the CT) and activity concentration (determined from the PET) have been found to be dominated by local volume changes. Specific corrections are then suggested to account for this effect and restore quantitative accuracy in PET.

In Holman et al. (2016) the errors in PET quantification in the lung, caused by the mismatch between the PET data and the lung density attenuation maps, have been investigated. Their work specifically addresses the situation where PET attenuation correction is performed with a CT image that is not phased-matched (i.e. corresponding to the same breathing state as the PET). They utilise ungated PET data, corrected for attenuation with CT maps obtained from both respiratory-gated and ungated data, to evaluate the differences between those configurations. The results show that using an attenuation map which is not phase-matched with the PET has a considerable impact on the PET quantitative accuracy. Furthermore, they suggest that when using gated PET data, the attenuation correction obtained from a single CT (which is the case when no cine-CT is available, see section 2.2.1.4) might result in greater errors compared to using ungated PET, because of the density mismatch between the modalities.

Detection of the density variations in the lungs can therefore be beneficial when aiming at analysing the respiratory-induced effects on PET data, and furthermore when interested in performing respiratory motion correction (through gating and registration), where compensation for density variations could improve the quantitative accuracy of the resulting image. In this chapter, the use of PCA is explored as a tool to detect the respiratory-induced density variations in the upper part of the lungs, where the more significant variation is expected to be caused by lung density changes.

6.2 Methods

6.2.1 Simulations

In the previous chapters the capability of PCA to detect respiratory motion has been extensively addressed. In order to verify whether density-only changes could also be detected

by the method, simulation data were specifically generated to simulate density changes only and no motion. Although this is in reality not possible, this study provides information on the expected changes in the data and the ability of PCA to detect them.

6.2.1.1 Local volume change detection

As shown in Reinhardt et al. (2008) with CT data, and applied also in Cuplov et al. (in press), the local volume change between two different states can be extracted from the Jacobian determinant ($|J|$) of the spatial transformation between the two states. In this instance the states are given by the respiratory gates, and the spatial transformation matrix can be obtained via image registration of one gate with respect to the other. The Jacobian is given by the matrix of all the first-order partial derivatives of the transformation matrix, and for small deformations its determinant can be approximated by the following ratio:

$$|J| = \frac{\textit{deformed} \text{ voxel volume}}{\textit{reference} \text{ voxel volume}} \tag{6.1}$$

where one of the two respiratory gates of interest is taken as the *reference* and the other is considered the *deformed* version of it.

This relation was used to construct realistic simulation data, which incorporate the volume changes generated by the respiratory motion of the chest.

6.2.1.2 Data generation

The simulation was based on a patient chest FDG-PET/CT dataset, that was acquired with cine-CT on a GE Discovery STE scanner (Teräs et al., 2007). The patient presented a Solitary Pulmonary Nodule (SPN). Both the CT and the PET acquisitions were monitored with the RPM, and the PET and the CT data were gated with equal time displacement gating in 5 gates (similarly as described in section 5.3.1). The 5 CT gates were reconstructed with OSEM. The PET gate corresponding to end-inhalation (g_1) was reconstructed with OSEM with attenuation correction performed using the CT gate corresponding to the same breathing status.

To estimate the local volume changes in the lung between gates, the gated CT images were registered, as they offer a greater spatial resolution and lung contrast compared to the corresponding PET images. Non-rigid registration between the end-inhale state and the other 4 gates was applied using the NiftyReg software (Ourselin et al., 2001; Modat et al., 2010). The utilised function was the `reg_f3d` function, which uses cubic B-Splines and returns the transformation matrix (also referred to as the deformation field). The Jacobian determinant of each transformation was obtained with the `reg_jacobian` function.

In order to simulate exclusively the impact of density changes in the lungs on PET data, as opposed to the motion effects, a series of simulated PET gated data was generated

exploiting the relationship in Eq. 6.1: the reconstructed PET image corresponding to end-inhalation g_1 (smoothed with an edge-preserving filter to decrease the noise) was multiplied by the Jacobian determinants of the 4 deformation fields, producing a new series of PET gates g_2, g_3, g_4, g_5. The registration process and the creation of the PET gates are depicted in Figure 6.1.

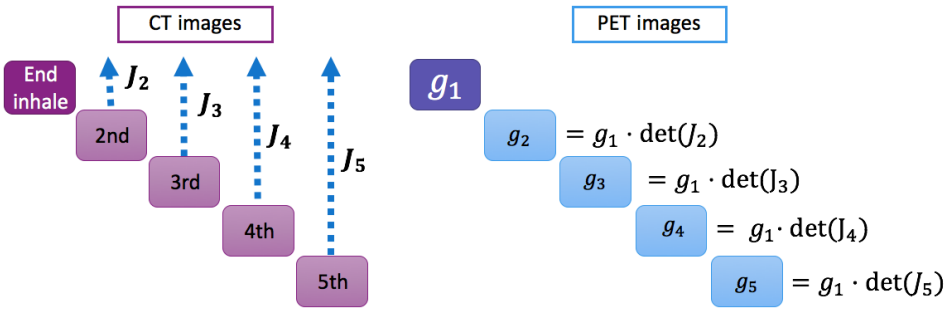

Figure 6.1: On the left: registration of the CT gates to the 1st gate, which produces the Jacobian determinants. On the right: creation of the PET gates, simulating only the volume changes and no motion effects, by application of the Jacobian determinant of the relevant transformation.

The same processes applied to g_1 were performed on its attenuation map (obtained from the matched CT gate), to generate the matching deformed attenuation maps for each gate. The simulated PET gates were then forward projected using STIR (Thielemans et al., 2012), simulating the geometry of the GE Discovery STE PET scanner. The generated sinograms were attenuated with factors obtained from the corresponding attenuation map, producing the series of sinograms s_1, s_2, s_3, s_4, s_5.

6.2.1.3 Analysis

The simulated sinograms were spatially downsampled to dimensions of $82 \times 7 \times 47$ (radial positions $\times$ angles $\times$ transaxial planes), as in Chapters 3 and 5. The area related to the upper part of the lungs was selected, by segmenting the central parts of the sinogram in the radial direction (choosing 28 of the total 82 radial positions) and selecting only LORs through the upper 15 planes of the total 47, see Figure 6.2. Therefore the final sinograms had dimensions equal to $28 \times 7 \times 15$ (they will still be referred to as s_1, s_2, s_3, s_4, s_5).

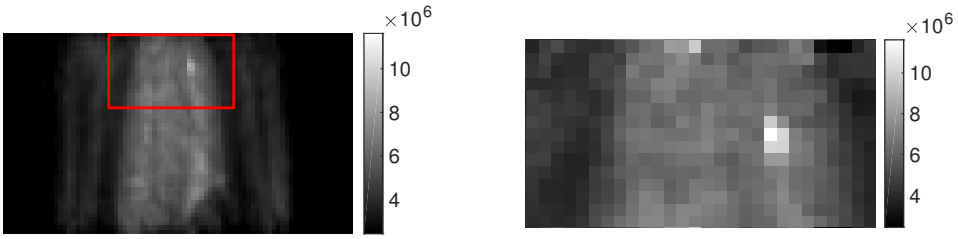

Figure 6.2: On the left: coronal projection of the entire sinogram of gate 2 with area selected for analysis. On the right: zoom in of selected area.

PCA was applied to the series of sinograms s_g, leading to the expansion:

$$s_g \approx \bar{s} + \sum_{k=1}^{K} w_{kg} p_k \tag{6.2}$$

where $\bar{s}$ is the mean of the sinograms over the gates, p_k are the Principal Components (PCs) and w_{kg} their weight factors. The first PC, that is the one corresponding to the biggest variation in the data, was examined together with its weight factor.

The total number of counts in each simulated gate (in the selected part) was evaluated and compared to the PC weight factor.

Furthermore, to discriminate further the effect of attenuation changes on the resulting projection data, the sinograms with and without the application of the attenuation were compared: the mean value of a ROI in the lung was evaluated in the end-inhalation and end-exhalation gates (g_1 and g_5) in both cases and their difference calculated. The result is expected to illustrate the difference between the impact on the projections given by change in concentration and change in attenuation, as they influence the PET data in opposite ways: when the latter increases the PET values should increase, while an increase in the former leads to more attenuated projections.

6.2.2 Patient data

PCA was applied on 15 FDG oncology patient datasets acquired in 3D listmode on a GE Discovery 690 PET/CT scanner (Bettinardi et al., 2011) of 360s duration, monitored by the RPM device. The upper chest bed position was chosen for each patient.

The method was applied as described in Chapter 5, using non-TOF and TOF sinograms, and without and with the application of the cylindrical masking. The respiratory Principal Component was selected as explained in section 2.3.3, by choosing the one corresponding to the weight factor with the highest peak in the respiratory frequency band (that is $[0.1, 0.4]$ Hz). The PCA respiratory signal was subsequently compared to the RPM via evaluation of the Pearson correlation, as in the previous chapters.

In order to investigate more thoroughly the effect of respiratory-induced lung density variations on the detected data (similarly to the ROI and total counts analysis with simulation data), the sum of the total number of counts in each temporal frame of the PET sinogram (after decay correction) was evaluated. The signal given by the variation in detected counts in time was then compared to the RPM signal. The number of detected counts is expected to be directly related to the changes in density of the lungs (and therefore to the changes in attenuation) because the analysed bed positions did not include moving organs such as the liver or the diaphragm, which can move in and out of the FOV therefore affecting the total

detected activity independently of density changes. Note that the number of detected counts was also used in the method presented in He et al. (2008). In their work a respiratory signal is obtained by exploiting the change in counts generated by the motion of the organs in the FOV, together with the fact that the scanner sensitivity is not constant along the z-direction. This produces changes in the detected counts due to respiratory motion.

6.3 Results

6.3.1 Simulations

The first PC obtained from the application of PCA on the series of simulated sinogram data is shown in Figure 6.3a, and in Figure 6.3b the PC weight factor and the total number of counts are displayed per gate. The diffused features of the PC suggest that the variations detected by PCA are related to local changes of the values in the sinogram, as opposed to moving edges. The trend of the weight factor and total counts along the gates is identical, showing that the effect detected by PCA varies correspondingly to the amount of counts during respiration. In Figure 6.4 the coronal projections with and without the application of the attenuation effect are displayed together with the ROI selected in the upper part of the lung. The relative difference between the mean in the ROI from the gates corresponding to inhalation and exhalation is -3% and +0.5% on the non attenuated and attenuated data respectively, demonstrating that the changes in density throughout the respiratory cycle have a significant effect on the detected data, and that the effect due to attenuation dominates, at least for this simulation.

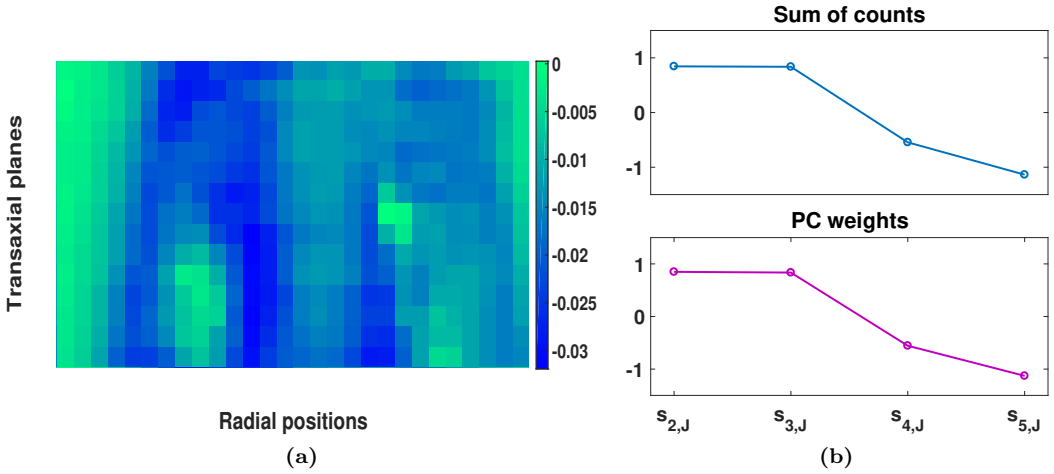

Figure 6.3: (a) Coronal view of the first PC. (b) Weight factor and change in total number of counts for each gate.

Figure 6.4: On the left the coronal view of the non attenuated projections of gate1, on the right the attenuated projections. In red the selected ROI.

6.3.2 Patient data

For each of the 15 patients datasets, the chosen bed position covered the area of the upper part of the lungs. The FOV did not include the liver or the diaphragm, see Figure 6.5, and there is little contrast in the lung regions of the sinograms. The total number of detected counts is therefore expected to be related to the changes in density of the lungs (and therefore to the changes in attenuation). The total counts per frame were then calculated and displayed with respect to time.

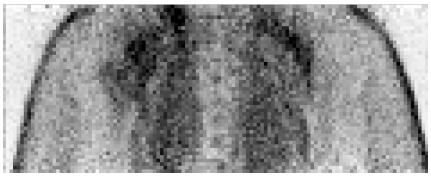

Figure 6.5: Example of Non Attenuation Corrected (NAC) PET reconstructed image of one patient (coronal projection). The FOV only includes the upper part of the lungs.

In Figure 6.6 the respiratory signal obtained with PCA is shown together with the signal given by the total number of counts in time and the external device respiratory signal, for two patients. In Figure 6.7 the value of the correlation between the RPM and the PCA signal obtained with the various data configurations (non-TOF and TOF sinograms, without and with the application of the masking) is shown per patient, together with the correlation between the RPM and the total counts signal. On the right-hand side the same values are displayed in a box-plot, to better convey the overall effect of using the different data configurations for the signal detection with PCA.

Table 6.1 shows the mean values, over all patients, of the correlation between the RPM and the PC signal, and between the RPM and the total number of counts.

6.3.2.1 Sign-determination

In this part of the study, the issue of sign-determination of the PCA signal requires careful consideration. This is because the utilised datasets are related to bed positions that cover the upper part of the lung where, as previously stated, organ motion is expected to be small. The assumption at the base of the sign-determination methods proposed in Chapter 3, specifically

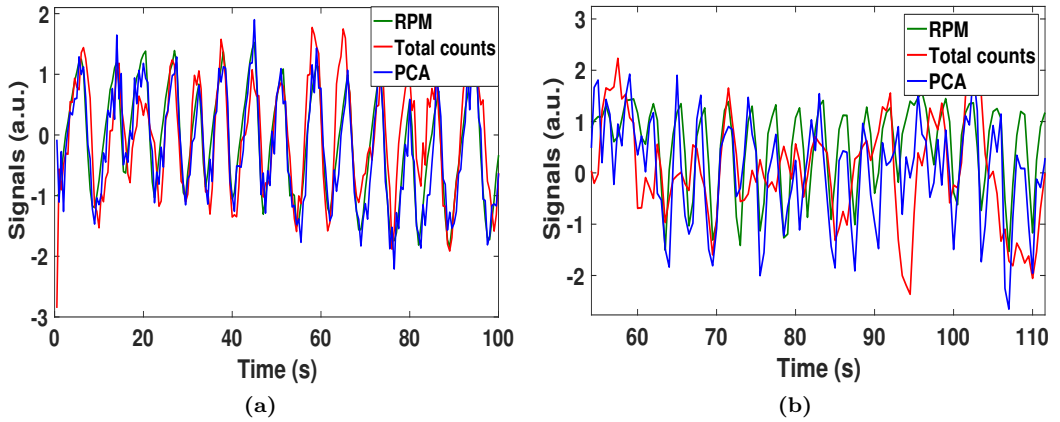

Figure 6.6: (a) Example of RPM, total counts in time and PCA signal for a patient (in this case obtained with 5 TOF bin sinograms, without mask). For this patient the correlation between RPM and PCA is 0.95 and that between the RPM and the signal given by the total counts is 0.80. This is an example of high similarity between the three signals. (b) Same signals for another patient, for which the correlation between the RPM and the PCA signal is lower than in the other case, being equal to 0.70. The correlation between RPM and total counts signal is in this case very low, 0.38.

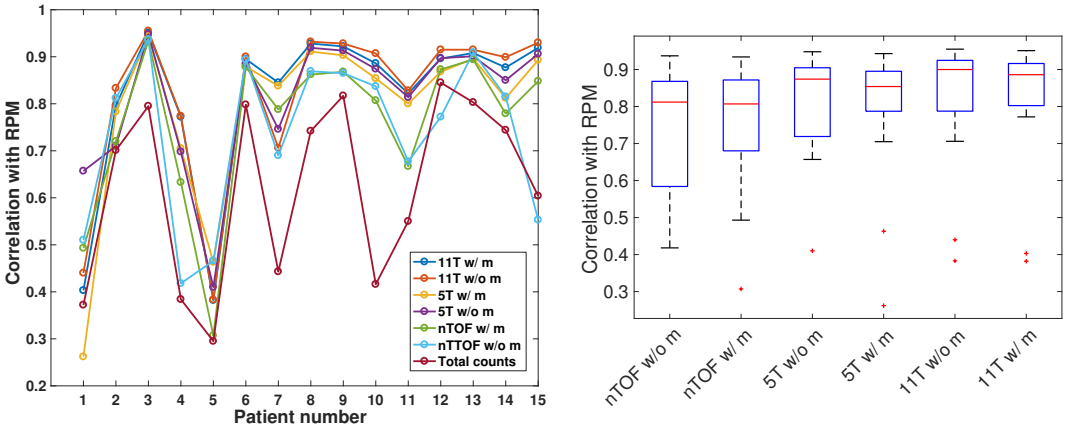

Figure 6.7: On the left-hand side, the correlation of the PCA signals with the RPM for the 15 patients with respect to the kind of utilised dataset (non-TOF or TOF sinograms, without and with the masking), and also the correlation between the total counts signal and the RPM. On the right-hand side, the correlation values between PCA signal and RPM are organised in a box-plot to better illustrate the difference between the various cases.

that the cause of the biggest changes in the data is respiratory motion in the z-direction, does not hold in this case. It is therefore presumed that the method's performance (only *CorrWeights* is considered here) will not be reliable. The analysis was performed as in the previous chapters: the correct direction of the respiratory signal is considered to be that of the RPM, therefore the method *CorrWeights* failed if the correlation between the RPM and the corrected PCA signal is negative.

In Table 6.2 the failure rates are shown with respect to the type of dataset that has been utilised.

Table 6.1: Mean and standard deviation values obtained on the 15 patients for the Pearson correlation between respiratory PC and RPM, and between total number of counts and RPM.

Data	Mean correlation
non-TOF w/o mask	0.73 ± 0.17
non-TOF w/ mask	0.76 ± 0.17
5 TOF w/o mask	0.81 ± 0.14
5 TOF w/ mask	0.79 ± 0.19
11 TOF w/o mask	0.82 ± 0.17
11 TOF w/ mask	0.81 ± 0.18
Total counts	0.62 ± 0.19

Table 6.2: *CorrWeights* failure rate in determining the correct sign of the respiratory signal produced by PCA.

Data	non-TOF w/o m	non-TOF w/ m	5TOF w/o m	5TOF w/ m	11TOF w/o m	11TOF w/ m
%	60	13	20	40	33	47

6.4 Discussion

The results of the application of PCA on the simulations proved that the algorithm is able to detect the density variations in the lung PET data.The ROI analysis provided insight into the combined effect that the change in density and in activity concentration have on the acquired data: when the projections are not attenuated, the difference in the mean ROI value between the end-inhalation and end-exhalation gate is negative, whereas it is positive when the attenuation effect is applied. This suggest that the change in attenuation, generated by volume variations (and therefore also density variations) is dominant. This is also corroborated by the behaviour of the total counts between gates, which decrease with increasing density.

It has to be noted that the only simulated variation in the data was the density change between respiratory gates, without adding the volume change given by the expansion and compression of the lungs. This variation generates from the Superior-Inferior motion of the diaphragm in the "diaphragmatic breathers", and from the Anterior-Posterior motion of the chest wall in the "chest breathers". In the former case the volume of the upper part of the chest is not presumed to exhibit significant changes, therefore the characteristics of the utilised simulation are thought to be consistent with this situation. In the "chest breathers" case, a more accurate study could include the increase and decrease of the volume of the upper chest during respiration. However, for small volume changes, the results of

the simulation study are not expected to be considerably affected: for an increase δ_l in path travelled by the photons (end-inhalation state), there is an increase in activity which is proportional to δ_l and an increase in attenuation given by $e^{-\mu\delta_l}$, leading to an overall increase in counts for increasing δ_l, therefore being consistent with the presented finding that counts are higher at end-inhalation compared to end-exhalation.

The application of PCA on patient data of the upper chest, using both non-TOF and TOF sinograms and without and with the masking, generated respiratory signals highly correlated to the RPM signal, see Figure 6.7 and Table 6.1, apart from patients 1 and 5 where all correlations are very low. The main difference is seen between the use of non-TOF and TOF sinograms, where the latter produces respiratory signals that present considerably higher correlation with the RPM, going from a mean value of 0.73 with non-TOF and no masking to a value of 0.82 when using 11 TOF bins and no masking. These results provide additional evidence that, if TOF capability is available, using TOF sinograms is beneficial for respiratory motion extraction with PCA. In the current dataset the difference between using 5 or 11 TOF bin is not significant. The application of the mask has not been found to improve the quality of the generated signals.

The signal created by summing the total counts per frame shows a similar trend to both the RPM and PCA signals, as it can be appreciated from Figure 6.6a, whereas in some case the similarity is lower, as in Figure 6.6b. The correlation with the total counts signal and the RPM over all patients has a mean value of 0.62, therefore suggesting that the changes in counts within the FOV do vary in time following a respiratory-like pattern. Since the similarity of the total counts signal to the RPM is lower than that of the PCA signals (in terms of correlation) the former cannot be considered as an adequate tool for respiratory motion representation.

The positive correlation between the total counts and the RPM signal indicates that the counts increase towards the end-inhalation state, and decrease when the patient is exhaling. This finding appears to be consistent with the hypothesis that the counts increase during the lung expansion, because the density is decreasing and consequently also the projections are less attenuated.

With regards to the application of the *CorrWeights* method, the failure rates proved to be very high and with no clear pattern among the use of the different kinds of sinograms, therefore suggesting that this method cannot be utilised in the upper lung to correct the direction of the PCA signal. This was to be expected, as the assumption at the base of the method's functioning is that the biggest change in the data is generated by the axial motion of the organs, that in these datasets is presumed to be minimal. In addition, these findings can also be interpreted as an additional confirmation that the major changes that occur in these datasets, which are detected by PCA, are not generated by the same effect as in the

lower lung.

6.5 Conclusions

In the first part of this chapter, data were generated to specifically reproduce the effect of density-only changes in the upper lung, disregarding the contribution generated by motion. To make the simulation more realistic, this was achieved by using respiratory gated PET/CT patient data (as cine-CT was available), and applying local volume changes to the reconstructed PET end-inhalation gate. The utilised transformation was obtained making use of the Jacobian determinant of the deformation fields between CT gates. PCA was applied and the first PC analysed. The total counts in the simulated series of PET gates was also evaluated and compared to the weight of the first PC. The results show that the PC weight and the total counts present the same trend, suggesting that are both representative of the same effect. Furthermore, as by construction the only changes that occur in the data are those generated by local volume changes, these results indicate that PCA can detect density changes in the lung.

The application of the method on 15 PET/CT patients, selecting the upper lung bed position, showed that the generated respiratory signal is highly similar to the RPM, and that using TOF sinograms yields signals more highly correlated to the RPM when compared to using the non-TOF sinograms. The difference between using 5 or 11 TOF bins has not been found to be significant. The application of the cylindrical masking did not influence the performance of the method for this bed position. In any case it is advisable to apply it as standard procedure, as it has been shown to improve the results for the lower lung.

The sign-determination method of choice for the lower lung, *CorrWeights*, provided high failure rates in all cases (non-TOF and TOF, without and with the masking) for the upper lung. This is an indication that for the upper part of the thorax the assumptions required for the method to correctly function are not met, those being that the main cause of change in the data is given by axial motion of the organs. A possible solution for the sign-determination in the upper lung could involve the use of the signal generated by the total counts in time, which has been shown to always be positively correlated with the RPM (in the patients dataset utilised in this analysis). This signal is of lower quality, in terms of correlation with the RPM, similarly to the $u(t)$ signal obtained with the *CorrWeights* method (see Eq. 3.15), nevertheless it has a known relationship with the direction of motion. An automated decision should therefore be made, to select whether to use the *CorrWeights* method or the total counts signal. This needs further investigation.

In conclusion, this study provided preliminary insight into the potential of extracting other kinds of information by applying PCA on PET raw data, specifically the density changes in the upper lung data that occur during respiration, even when organ motion is

expected to be small in the considered FOV. More patients need to be investigated and also reconstruction studies need to be conducted, to evaluate the extent of the quantitative differences between respiratory gates.

Chapter 7

Conclusions

Respiratory motion in PET imaging of the thorax leads to artefacts in the reconstructed images, lowering the quantitative accuracy of this imaging modality. Numerous approaches and methodologies have been suggested in order to decrease this effect and improve the image quality, for which a respiratory motion signal is required. In recent years several groups have proposed DD methods for respiratory motion detection, which do not require external devices and only rely on the PET raw data. Among the available methods, PCA has shown promising results, producing respiratory signals highly correlated to those of external devices and also to the internal motion extracted from MRI images. For these reasons, PCA has been selected as the method of choice for our research.

The main aim of this thesis was to investigate the application of PCA on PET lung data for respiratory motion detection, and to provide additional methods to further its performance. Furthermore, the methodologies presented throughout this work could also be applied to evaluate other DD algorithms, therefore constituting a potential standard procedure for the assessment of a DD method's performance.

In this chapter, the main conclusions of each of the previous chapters are highlighted, followed by an outline of the novel contributions and by an overview of the possible topics of future work. Finally, a list of the publications that have arisen from the work presented here, and also from collaborations with other researchers, is given.

7.1 Summary of Main Conclusions

In Chapter 3, two sinogram-based methods were proposed for the sign-determination of PCA respiratory signal (*CorrWeights* and *CorrSino*) and their performance was compared to two versions of a previously published method, which relies on image registration. The main assumption of all the methods is that respiration generates motion that is mostly in the feet-head direction (axial direction). All methods successfully determined the correct

direction of the respiratory signal when applied to simulation data. When applied to clinical patient data, selecting the PET bed position that corresponds to the lower lung, the newly proposed methods distinctly outperformed the registration-based ones, producing lower failure rates. Among the two proposed methods, *CorrWeights* is not specifically related to PCA, therefore it could be utilised to determine the correct direction of the motion of respiratory signals obtained with any other DD method. In this chapter the methods were implemented for non-TOF data.

In Chapter 4, PCA and the sign-determination methods were extended to TOF data. The methods were also combined with a sinogram masking procedure, in order to select data likely affected by respiratory motion: a cylinder is created in image space, centrally placed in the FOV, then non-TOF and TOF forward-projection are applied and a binary masked is created, assigning a value of 1 to all non-zero elements and zero elsewhere. Multiplying the data with the generated mask effectively selects the central part of the patient's body, where motion is expected to occur. A simulation study was then performed to highlight the different results produced by PCA when using non-TOF and TOF sinograms without and with the proposed selection in sinogram space. The results showed that using TOF sinograms, and applying the sinogram masking process, increase the capability of PCA to capture changes that occur in the areas of interest. The same methodologies were then utilised on patient data in the following chapter.

In Chapter 5, the performance of PCA and the sinogram-based sign-determination methods were thoroughly evaluated, utilising patient datasets of the lower part of the lung. In the first part, PET/MR data were utilised to compare the PCA respiratory signal to the internal motion of the diaphragm, extracted from specific MR sequences. The acquisitions also included the use of a respiratory pneumatic belt. The results show that the PCA signal, the internal motion extracted from the MR data and the pneumatic belt signal are very similar to each other, with mean correlation values above 0.9. These results prove that the PCA output is highly reliable, and furthermore that the external motion of the chest and the internal motion of the organs are both good measures of the breathing pattern of the patient. Furthermore, the analysis of the shape of the respiratory signals showed that no evident drift was present in the PCA signals compared to the other two (MR-derived and belt). The second part of the chapter concerns the evaluation of the methods on a large cohort of PET patients (28 datasets, monitored with the RPM), with respect to using non-TOF and TOF data, and with respect to the duration of the interval taken into account and the level of noise. The results show that utilising TOF sinograms, together with the application of the masking process in sinogram space, improves the quality of the PCA signal and the reliability of the sign-determination methods, the impact being more significant when using lower counts (i.e. higher level of noise) as opposed to shorter interval durations. The analysis

of the sinogram-based sign-determination methods confirmed the *CorrWeights* method as the best performing among the two, as shown in the preliminary results of Chapter 3. In the third part, the RPM and PCA respiratory signals were compared by evaluating the detected motion on gated images reconstructed without attenuation correction. Image registration was used to obtain the maximum detected displacement, and also to extract the slope of the linear fit of the displacement per gate. The use of the PCA signal for gating resulted in most cases in similar results to when using the RPM signal, and for those patients where differences were found, the PCA gating produced the higher detected displacement and higher slope value. The closer analysis of those patients showed that the RPM signal revealed a drift in its values, not present in the PCA case, that affected the gate assignment. Therefore, with the utilised gating parameters, using the PCA signal for gating has proved to be best at discriminating the patient's motion.

In Chapter 6, the methods presented throughout the thesis were applied to PET data of the upper chest. The aim was to investigate whether PCA is able to detect respiratory-induced changes in the data of the upper lungs, where motion is expected to be small and the cause of the biggest change in the data is presumed to be the variation in lung density. Simulation data were generated to simulate density changes only: 5 respiratory gates were produced and PCA was applied. The results show that PCA is able to detect the simulated density changes in the data. Furthermore, the results of a ROI analysis on the attenuated and non-attenuated projections support the hypothesis that the dominant effect on PET data due to volume changes is given by the variation in the attenuation, and not by the variation in the activity concentration. PCA and *CorrWeights* were then applied on 15 clinical patients' datasets. The PCA respiratory signals showed high correlation with the RPM, and the use of TOF sinograms proved to be beneficial, with mean correlation values going from 0.73 with non-TOF to 0.82 with 11 TOF bins. In contrast to the findings of Chapter 5, for this data of the upper chest, the sinogram masking process introduced in Chapter 4 did not influence the results. Furthermore, the total detected counts were evaluated as a temporal signal, which was found to be positively correlated with the RPM. This supports the finding, obtained from the simulation data, that for the upper chest PET counts increase in inspiration and decrease in expiration. *CorrWeights* provided high failure rates that imply this method cannot be utilised with these datasets. It was suggested that the total counts signal could be utilised for sign-determination, as its correlation with the RPM has been found to always be positive, however a more targeted analysis is required.

In summary, throughout this thesis the application of PCA to detect changes in the PET data has been investigated for the purpose of respiratory motion detection. Methodologies to improve its performance, and to better exploit the information available in TOF-PET data, have been proposed and implemented, making it a robust, fast and easily implementable

tool. All the methods have been thoroughly evaluated, both on simulation data and clinical patient data. The extensive comparison with external devices' signal has provided proof that this DD method constitutes a reliable tool for respiratory motion management in lung PET imaging. Furthermore, the utilised evaluation process constitutes and effective strategy that could be used for the analysis of any other DD method. Finally, the potential to utilise PCA to extract other information from the PET data of the upper chest has been shown.

7.2 Original Contributions

The following points summarise the main novel contributions produced by the author, among the research presented in this work:

- Implementation of the sign-determination method *CorrWeights*, and the comparison between its performance and that of the *CorrSino* method (invented by Dr. Kris Thielemans) and the registration-based methods.

- Investigation of the sign-determination methods' performance with respect to duration of data interval taken into account.

- Implementation of the TOF unlisting process: in `C` language within the GE proprietary source code for the unlisting of PET data generated by GE PET/CT scanners (in this thesis Discovery STE, Discovery 690 and Discovery 710); in `C++` within the STIR code for the unlisting of the PET data generated by the GE SIGNA PET/MR scanner (source code available at `https://github.com/UCL/STIR/tree/GE_Signa`).

- Extension of PCA and sign-determination methods to TOF sinograms, and the introduction of the practice of selecting the total number of TOF bins of choice (depending on the total number of TOF bins provided by the dataset and on the possible mashing factors).

- Automatic cylindrical masking process in sinogram space, which does not require any manual input.

- Simultaneous comparison of PCA respiratory signal with the internal organ motion, extracted from dynamic MR images, and an external device signal (pneumatic belt), including temporal synchronisation of the three signals, providing proof of high similarity between the three signals.

- Evaluation of the proposed DD method, in comparison to the RPM, using several metrics: the maximum detected motion between gates, and the slope of the linear fit of the motion per gate with respect to the mean motion within each gate.

- Investigation and comparison of the gate assignment generated when using RPM or PCA signals, aiming to discover the reason for the considerable differences that had been found for some patients.

- Introduction of the use of PCA on the upper part of the lungs, investigating the possibility to detect density changes, and evaluation on simulation and patient data.

7.3 Suggested Future Work

There are several areas, arising from the research presented in this thesis, which would require further investigation.

Regarding the proposed sign-determination method *CorrWeights* presented in Chapter 3, it would be useful to define a threshold value, or a specific metric, to allow deciding whether to trust its sign choice, aiming to reduce the failure rate when applying the method. More generally, it should provide an indication of the robustness of the method's performance in each specific case, as it has been shown that results can be patient dependent. A study could therefore be performed with a large cohort of patients in order to provide such a decision-making tool.

A refinement of the masking process in sinogram space, presented in Chapter 4, could be implemented. When the CT image is available, a mask that is tailored to the patient's body shape could be created, therefore selecting more specifically the areas of interest. The analysis of the impact on the PCA respiratory signal and sign-determination method performance would then have to follow.

Regarding the PET data generated by the GE Signa PET/MR scanner, utilised in the first part of Chapter 5, the small TOF bin size could be exploited in order to choose a variety of TOF mashing factors. This would allow to generate sinograms with different number of TOF bins (in this case 3, 9, 13, 27, 39 or 117) therefore a comparison could be performed between the PCA results obtained with each of the sinograms. This would allow one to better appreciate the influence of the total number of TOF bins on the PCA ability to detect respiratory motion.

The analysis on reconstructed PET gates, presented in the third part of Chapter 5, could be performed choosing different gating parameters, e.g. using phase gating instead of displacement gating, or performing the latter with a threshold on the value of the respiratory signal (to disregard potential outliers). The conclusions are not expected to vary considerably, but it would be informative to gain insight into the impact of such parameters on the results of the gating process with the PCA respiratory signal.

Following the results shown in Chapter 6, in order to determine the influence that density changes in the PET data could have on the quantitation process, a study involving

reconstructed images should be considered. The data would be gated with the PCA signal and then reconstructed. The SUV of detectable lesions, and the value of ROIs placed in various parts of the lung, could be compared between gates. Additionally, the analysis of the application of the *CorrWeights* method on the upper chest datasets, together with the total counts signal, should be addressed more in detail, aiming at providing a reliable tool for sign-determination in the whole thorax.

Finally, the impact of using PCA as a motion detection tool for respiratory motion management in clinical practice should be investigated. The patients that have been acquired at the Institute of Nuclear Medicine at UCLH (see Section 5.2.1) are the first group of a larger cohort of patient that could to be utilised for the application of respiratory motion correction, performed with proprietary software provided by the vendor (in this case General Electric). The aim of the study would be to evaluate the feasibility of utilising PCA instead of an external device for motion detection in PET lung imaging (in the above-mentioned department, the available device is the RPM). The study could also include the use of newer algorithms for motion compensation with matched attenuation correction, presented in Bousse et al. (2016a), and the comparison of their results with those obtained with the vendor software.

7.4 List of Published Papers and Presentations

7.4.1 Publications

A. Bousse, **O. Bertolli**, D. Atkinson, S. Arridge, S. Ourselin, B. F. Hutton, and K.Thielemans. "Maximum-likelihood joint image reconstruction/motion estimation in attenuation-corrected respiratory gated PET/CT using a single attenuation map." *IEEE Transactions on Medical Imaging*, vol. 35, no. 1 (2016): 217-228.

A. Bousse, **O. Bertolli**, D. Atkinson, S. Arridge, S. Ourselin, B. F. Hutton, and K.Thielemans. "Maximum-likelihood joint image reconstruction and motion estimation with misaligned attenuation in TOF-PET/CT." *Physics in Medicine and Biology*, vol. 61, no. 3 (2016): L11.

O. Bertolli, S. Arridge, S. D. Wollenweber, C. W. Stearns, B. F. Hutton, and K. Thielemans. "Sign determination methods for the respiratory signal in data-driven PET gating." *Physics in Medicine and Biology*, vol.62, no. 8 (2017): 3204.

7.4.2 Published Conference Proceedings

B. F. Holman, V. Cuplov, **O. Bertolli**, A. M. Groves, B. F. Hutton, and K. Thielemans. "Density variation during respiration affects PET quantitation in the lung." In Nuclear Science

Symposium and Medical Imaging Conference (NSS/MIC), pp. 1-3. IEEE, 2015.

O. Bertolli, S. Arridge, S. D. Wollenweber, C. W. Stearns, B. F. Hutton, and K. Thielemans. "Sign determination methods for the respiratory signal in data-driven PET gating." In *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, pp. 1-3. IEEE, 2015.

O. Bertolli, S. Arridge, S.D. Wollenweber, C.W. Stearns, B. F. Hutton, and K. Thielemans. "Data driven respiratory signal detection in PET taking advantage of time-of-flight data." In *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, pp. 1-3. IEEE, 2016.

O. Bertolli, S. Arridge, S. D. Wollenweber, C. W. Stearns, B. F. Hutton, and K. Thielemans. "Improvement of the Sign Determination Method for Data-Driven respiratory signal in TOF-PET." In *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, 2017 (in press)

O. Bertolli, V. Cuplov, S. Arridge, S. D. Wollenweber, C. W. Stearns, B. F. Hutton, and K. Thielemans. "Detection of Lung Density Variations With Principal Component Analysis in PET." In *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, 2017 (in press)

E. Ovtchinnikov, D. Atkinson, C. Kolbitsch, B. A. Thomas, **O. Bertolli**, C. O. da Costa-Luis, N. Efthimiou, R. Fowler, E. Pasca, P. Wadhwa, E. Emond, J. Matthews, A. J. Reader, C. Tsoumpas, C. Prieto, M. Turner,K. Thielemans. "SIRF: Synergistic Image Reconstruction Framework." In *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, 2017 (in press)

7.4.3 Published Conference Abstracts

O. Bertolli, P. Schleyer, S. Arridge, R. Manjeshwar, B. F. Hutton, and K. Thielemans. "Data driven respiratory gating in PET: determination of motion direction." *Journal of Nuclear Medicine*, vol.56, no. supplement 3 (2015): 1808-1808.

E. Ovtchinnikov, D. Atkinson, **O. Bertolli**, C. O. da Costa-Luis, N. Efthimiou, R. Fowler, C. Kolbitsch, J. Matthews, C. Prieto, A. J. Reader, B. A. Thomas, C. Tsoumpas, M. Turner and K. Thielemans. "SIRF: Synergistic Image Reconstruction Framework". In *6th Conference on PET-MRI and SPECT-MRI (PSMR)*, 2017

7.4.4 Submitted Conference Abstracts

O. Bertolli, B. Fernandez, S. Arridge, B. F. Hutton, C. Comtat, M. Soussan, and Kris Thielemans. "Comparison of motion surrogate signal derived from MR images, respiratory belt and PCA-PET for motion correction in PET/MR". For *ISMRM Conference, Paris, France*, 2018

O. Bertolli, T. Sanderson, S. Arridge, A. Alnaim, S. Wan, S. D. Wollenweber, C. W. Stearns, B. F. Hutton, S. Arridge, and Kris Thielemans. "Evaluation of data-driven respiratory gating in PET with PCA". For *SNMMI Conference, Philadephia, USA*, 2018

7.4.5 Other Presentations

O. Bertolli, "Detecting motion from the data", in *UCL PET/MRI Methodology Symposium*, London 2016.

O. Bertolli, "Improvements in Data-Driven respiratory signal extraction with the use of TOF-PET data", in *Workshop on the Technology Frontier for Single Photon Detection and Advanced Scintillator Timing*, at the *Nuclear Science Symposium and Medical Imaging Conference*, 2016

O.Bertolli, "Unlisting GE Signa PET/MR listmode data with STIR" at the STIR User's Meeting, at the *Nuclear Science Symposium and Medical Imaging Conference*, 2016

7.4.6 Filed patents

O.Bertolli and K. Thielemans, "Data driven motion detection in PET taking advantage of TOF data"
Application number: GB1618487.1
Filing date: 14 December 2016 (Journal 6656)
See `https://www.ipo.gov.uk/p-journal/p-pj/p-pj?startYear=2016&startMonth=December&startDay=14th+-+6656&endYear=2016&endMonth=December&endDay=14th+-+6656&filter=ucl&perPage=10&sort=Publication+Date`

Bibliography

Adam Alessio and Paul Kinahan. PET image reconstruction. *Nuclear Medicine*, 2, 2006.

Dale L Bailey. Data acquisition and performance characterization in PET. In *Positron emission tomography*, pages 41–62. Springer, 2005.

Dale L Bailey, Joel S Karp, and Suleman Surti. Physics and instrumentation in PET. In *Positron emission tomography*, pages 13–39. Springer, 2005a.

Dale L Bailey, David W Townsend, Peter E Valk, and Michael N Maisey. *Positron emission tomography*. Springer, 2005b.

Harrison H Barrett and Kyle J Myers. *Foundations of Image Science*. John Wiley & Sons, 2013.

Bernard Bendriem and David W Townsend. *The Theory and Practice of 3D PET*, volume 32. Springer Science & Business Media, 2013.

Ottavia Bertolli, Simon Arridge, Scott D Wollenweber, Charles W Stearns, Brian F Hutton, and Kris Thielemans. Sign determination methods for the respiratory signal in data-driven PET gating. *Physics in Medicine and Biology*, 62(8):3204, 2017.

V Bettinardi, L Presotto, E Rapisarda, M Picchio, L Gianolli, and MC Gilardi. Physical performance of the new hybrid PET/CT Discovery-690. *Medical Physics*, 38(10):5394–5411, 2011.

Thomas Beyer, David W Townsend, Tony Brun, Paul E Kinahan, et al. A combined PET/CT scanner for clinical oncology. *Journal of Nuclear Medicine*, 41(8):1369, 2000.

Thomas Beyer, Gerald Antoch, Stefan Müller, Thomas Egelhof, Lutz S Freudenberg, Jörg Debatin, and Andreas Bockisch. Acquisition protocol considerations for combined PET/CT imaging. *Journal of Nuclear Medicine*, 45(1 suppl):25S–35S, 2004.

M. Blume, A. Martinez-Moller, A. Keil, N. Navab, and M. Rafecas. Joint reconstruction of image and motion in gated positron emission tomography. *IEEE Transactions on Medical Imaging*, 29(11):1892–1906, 2010.

Luc Boucher, Serge Rodrigue, Roger Lecomte, and François Bénard. Respiratory gating for 3-dimensional PET of the thorax: feasibility and initial results. *Journal of Nuclear Medicine*, 45(2):214–219, 2004.

Alexandre Bousse, Ottavia Bertolli, David Atkinson, Simon Arridge, Sébastien Ourselin, Brian F Hutton, and Kris Thielemans. Maximum-likelihood joint image reconstruction/motion estimation in attenuation-corrected respiratory gated PET/CT using a single attenuation map. *IEEE Transactions on Medical Imaging*, 35(1):217–228, 2016a.

Alexandre Bousse, Ottavia Bertolli, David Atkinson, Simon Arridge, Sébastien Ourselin, Brian F Hutton, and Kris Thielemans. Maximum-likelihood joint image reconstruction and motion estimation with misaligned attenuation in TOF-PET/CT. *Physics in Medicine and Biology*, 61(3):L11, 2016b.

Lisa Gottesfeld Brown. A Survey of Image Registration Techniques. *ACM Comput. Surv.*, 24(4):325–376, December 1992.

Ralph A Bundschuh, Axel Martínez-Moeller, Markus Essler, María-José Martínez, Stephan G Nekolla, Sibylle I Ziegler, and Markus Schwaiger. Postacquisition detection of tumor motion in the lung and upper abdomen using list-mode PET data: a feasibility study. *Journal of Nuclear Medicine*, 48(5):758–763, 2007.

Florian Büther, Mohammad Dawood, Lars Stegger, Frank Wübbeling, Michael Schäfers, Otmar Schober, and Klaus P Schäfers. List mode–driven cardiac and respiratory gating in PET. *Journal of Nuclear Medicine*, 50(5):674–681, 2009.

Florian Büther, Iris Ernst, Mohammad Dawood, Peter Kraxner, Michael Schäfers, Otmar Schober, and Klaus P Schäfers. Detection of respiratory tumour motion using intrinsic list mode-driven gating in positron emission tomography. *European Journal of Nuclear Medicine and Molecular Imaging*, 37(12):2315–2327, 2010.

Florian Büther, Thomas Vehren, Klaus P. Schäfers, and Michael Schäfers. Impact of Data-driven Respiratory Gating in Clinical PET. *Radiology*, 281(1):229–238, 2016. doi: 10.1148/radiol.2016152067. PMID: 27092660.

Delphine L Chen, Joseph Cheriyan, Edwin R Chilvers, Gourab Choudhury, Christopher Coello, Martin Connell, Marie Fisk, Ashley M Groves, Roger N Gunn, Beverley F Holman, et al. Quantification of lung PET images: challenges and opportunities. *Journal of Nuclear Medicine*, 58(2):201–207, 2017.

Sanghee Cho, Sangtae Ahn, Quanzheng Li, and Richard M Leahy. Exact and approximate fourier rebinning of PET data from time-of-flight to non time-of-flight. *Physics in Medicine and Biology*, 54(3):467, 2009.

Maurizio Conti. State of the art and challenges of time-of-flight PET. *Physica Medica*, 25(1): 1–11, 2009.

Maurizio Conti. Focus on time-of-flight PET: the benefits of improved time resolution. *European Journal of Nuclear Medicine and Molecular Imaging*, 38(6):1147–1157, 2011.

Vesna Cuplov, Beverley Fiona Holman, Jamie R McClelland, Marc Modat, Brian F Hutton, and Kris Thielemans. Issues in quantification of registered respiratory gated PET/CT in the lung. *Physics in Medicine and Biology*, in press.

Joël Daouk, Marie Leloire, Loïc Fin, Pascal Bailly, Julie Morvan, Isabelle El Esper, Lazhar Saidi, Véronique Moullart, Géraldine François, Vincent Jounieaux, et al. Respiratory-gated 18F-FDG PET imaging in lung cancer: effects on sensitivity and specificity. *Acta Radiologica*, 52(6):651–657, 2011.

Margaret E. Daube-Witherspoon and Gerd Muehllehner. Treatment of Axial Data in Three-Dimensional PET. *Journal of Nuclear Medicine*, 28(11):1717–1724, nov 1987.

M. Dawood, N. Lang, Xiaoyi Jiang, and K.P. Schafers. Lung motion correction on respiratory gated 3-D PET/CT images. *IEEE Transactions on Medical Imaging*, 25(4):476–485, 2006.

Mohammad Dawood, Florian Büther, Norbert Lang, Otmar Schober, and Klaus P Schäfers. Respiratory gating in positron emission tomography: a quantitative comparison of different gating schemes. *Medical Physics*, 34(7):3067–3076, 2007.

Michel Defrise. A factorization method for the 3D X-ray transform. *Inverse Problems*, 11(5): 983, 1995.

Alberto Del Guerra. *Ionizing radiation detectors for medical imaging.* World Scientific, 2004.

Gaspar Delso, Sebastian Fürst, Björn Jakoby, Ralf Ladebeck, Carl Ganter, Stephan G Nekolla, Markus Schwaiger, and Sibylle I Ziegler. Performance measurements of the Siemens mMR integrated whole-body PET/MR scanner. *Journal of nuclear medicine*, 52(12):1914–1922, 2011.

A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B*, 39:1–38, 1977.

Yusuf E Erdi, Sadek A Nehmeh, Tinsu Pan, Alexander Pevsner, Kenneth E Rosenzweig, Gikas Mageras, Ellen D Yorke, Heiko Schoder, Wendy Hsiao, Olivia D Squire, et al. The CT motion quantitation of lung lesions and its impact on PET-measured SUVs. *Journal of Nuclear Medicine*, 45(8):1287–1292, 2004.

Hadi Fayad, Tinsu Pan, Jean François Clement, and Dimitris Visvikis. Technical note: Correlation of respiratory motion between external patient surface and internal anatomical landmarks. *Medical Physics*, 38(6):3157–3164, 2011.

Tao Feng, Jizhe Wang, Youjun Sun, Wentao Zhu, Yun Dong, and Hongdi Li. Self-gating: An adaptive center-of-mass approach for respiratory gating in pet. *IEEE Transactions on Medical Imaging*, 2017a.

Tao Feng, Jizhe Wang, Wentao Zhu, Yun Dong, and Hongdi Li. Robust Device-less Cardiac Gating Method for Cardiac PET Applications. *Journal of Nuclear Medicine*, 58(supplement 1):700–700, 2017b.

Loïc Fin, Joël Daouk, Julie Morvan, Pascal Bailly, Isabelle El Esper, Lazhar Saidi, and Marc-Etienne Meyer. Initial clinical results for breath-hold CT-based processing of respiratory-gated PET acquisitions. *European Journal of Nuclear Medicine and Molecular Imaging*, 35(11):1971–1980, 2008.

Murray F. Freeman and John W. Tukey. Transformations related to the angular and the square root. *The Annals of Mathematical Statistics*, 21(4):607–611, 12 1950.

Karl Friston, John Ashburner, Christopher D Frith, J-B Poline, John D Heather, Richard SJ Frackowiak, et al. Spatial registration and normalization of images. *Human Brain Mapping*, 3(3):165–189, 1995.

David P Gierga, Johanna Brewer, Gregory C Sharp, Margrit Betke, Christopher G Willett, and George TY Chen. The correlation between internal and external markers for abdominal tumors: implications for respiratory gating. *International Journal of Radiation Oncology* Biology* Physics*, 61(5):1551–1558, 2005.

Olivier Guivarc'h, Alexandre Turzo, Dimitris Visvikis, and Yves Bizais. Synchronization of pulmonary scintigraphy by respiratory flow and by impedance plethysmography. In *Proc. SPIE Med. Imaging*, pages 1166–1175, 2004.

Dongfeng Han, John Bayouth, Qi Song, Aakant Taurani, Milan Sonka, John Buatti, and Xiaodong Wu. Globally optimal tumor segmentation in PET-CT images: a graph-based co-segmentation method. In *Biennial International Conference on Information Processing in Medical Imaging*, pages 245–256. Springer, 2011.

Jianfeng He, Graeme JO Keefe, Sylvia Jie Gong, Gareth Jones, Tim Saunder, Andrew M Scott, and Moshi Geso. A novel method for respiratory motion gated with geometric sensitivity of the scanner in 3D PET. *IEEE Transactions on Nuclear Science*, 55(5): 2557–2565, 2008.

Andre Heremans, Johny A Verschakelen, Maurits Demedts, et al. Measurement of lung density by means of quantitative CT scanning: a study of correlations with pulmonary function tests. *Chest*, 102(3):805–811, 1992.

Mirco Heß, Florian Büther, Fabian Gigengack, Mohammad Dawood, and Klaus P. Schäfers. A dual-Kinect approach to determine torso surface motion for respiratory motion correction in PET. *Medical Physics*, 42(5):2276–2286, 2015.

Beverley F Holman, Vesna Cuplov, Brian F Hutton, Ashley M Groves, and Kris Thielemans. The effect of respiratory induced density variations on non-TOF PET quantitation in the lung. *Physics in Medicine and Biology*, 61(8):3148, 2016.

Hui Hu. Multi-slice helical CT: Scan and reconstruction. *Medical Physics*, 26(1):5–18, 1999.

Malcolm Hudson and Richard Larkin. Accelerated image reconstruction using ordered subsets of projection data. *IEEE Transactions on Medical Imaging*, 13:601–609, 1994.

Dan Ionascu, Steve B Jiang, Seiko Nishioka, Hiroki Shirato, and Ross I Berbeco. Internal-external correlation investigations of respiratory induced motion of lung tumors. *Medical Physics*, 34(10):3893–3903, 2007.

M.W. Jacobson and J.A. Fessler. Joint estimation of image and deformation parameters in motion-corrected PET. In *IEEE Nuclear Science Symp. and Medical Imaging Conf. Record*, volume 5, pages 3290–3294 Vol.5, 2003.

Hossein Jadvar and J Anthony Parker. *Clinical PET and PET/CT*. Springer Science & Business Media, 2006.

Xiao Jin. *Event-by-Event Motion Correction in Positron Emission Tomography: Development, Evaluation, and Applications*. Yale University, 2013.

E Kanoulas, J A Aslam, G C Sharp, R I Berbeco, S Nishioka, H Shirato, and S B Jiang. Derivation of the tumor position from external respiratory surrogates with periodical updating of the internal/external correlation. *Physics in Medicine and Biology*, 52(17): 5443, 2007.

Chien-Min Kao. Windowed image reconstruction for time-of-flight positron emission tomography. *Physics in Medicine and Biology*, 53(13):3431, 2008.

Paul J Keall, Gig S Mageras, James M Balter, Richard S Emery, Kenneth M Forster, Steve B Jiang, Jeffrey M Kapatoes, Daniel A Low, Martin J Murphy, Brad R Murray, et al. The management of respiratory motion in radiation oncology report of aapm task group 76a). *Medical Physics*, 33(10):3874–3900, 2006.

Adam L Kesner, Paul J Schleyer, Florian Büther, Martin A Walter, Klaus P Schäfers, and Phillip J Koo. On transcending the impasse of respiratory motion correction applications in routine clinical imaging–a consideration of a fully automated data driven motion control framework. *EJNMMI Physics*, 1(1):8, 2014.

Adam Leon Kesner and Claudia Kuntner. A new fast and fully automated software based algorithm for extracting respiratory signal from raw PET data and its comparison to other methods. *Medical Physics*, 37(10):5550–5559, 2010.

P.E. Kinahan and J.G. Rogers. Analytic 3D image reconstruction using all detected events. *IEEE Transactions on Nuclear Science*, 36(1):964–968, 1989.

PE Kinahan, DW Townsend, T Beyer, and D Sashin. Attenuation correction for a combined 3D PET/CT scanner. *Medical Physics*, 25(10):2046–2053, 1998.

Glenn F Knoll. *Radiation detection and measurement*. John Wiley & Sons, 2010.

Jeffrey A Kolthammer, Kuan-Hao Su, Anu Grover, Manoj Narayanan, David W Jordan, and Raymond F Muzic. Performance evaluation of the Ingenuity TF PET/CT scanner with a focus on high count-rate conditions. *Physics in Medicine and Biology*, 59(14):3843, 2014.

F Lamare, T Cresson, J Savean, C Cheze Le Rest, AJ Reader, and D Visvikis. Respiratory motion correction for PET oncology applications using affine transformation of list mode data. *Physics in Medicine and Biology*, 52(1):121, 2006.

Kenneth Lange, Richard Carson, et al. EM reconstruction algorithms for emission and transmission tomography. *Journal of Computer Assisted Tomography*, 8(2):306–16, 1984.

Joseph Larmor. LXIII. On the theory of the magnetic influence on spectra; and on the radiation from moving ions. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 44(271):503–512, 1897.

Craig S Levin, Sri Harsha Maramraju, Mohammad Mehdi Khalighi, Timothy W Deller, Gaspar Delso, and Floris Jansen. Design features and mutual compatibility studies of the time-of-flight PET capable GE SIGNA PET/MR system. *IEEE Transactions on Medical Imaging*, 35(8):1907–1914, 2016.

Tom K Lewellen. Time-of-flight PET. In *Seminars in nuclear medicine*, volume 28, pages 268–275. Elsevier, 1998.

Tianfang Li, Brian Thorndyke, Eduard Schreibmann, Yong Yang, and Lei Xing. Model-based image reconstruction for four-dimensional PET. *Medical physics*, 33(5):1288–1298, 2006.

Martin J Lipton. Cine computerized tomography. *The International Journal of Cardiac Imaging*, 2(4):209–221, 1987.

Chi Liu, Larry A Pierce II, Adam M Alessio, and Paul E Kinahan. The impact of respiratory motion on tumor quantification and delineation in static pet/ct imaging. *Physics in Medicine and Biology*, 54(24):7345, 2009.

Val J Lowe, JW Fletcher, Lisa Gobar, Michael Lawson, Peter Kirchner, Peter Valk, John Karis, Karl Hubner, Dominique Delbeke, Elizabeth V Heiberg, et al. Prospective investigation of positron emission tomography in lung nodules. *Journal of Clinical Oncology*, 16(3): 1075–1084, 1998.

Richard Manber, Kris Thielemans, Brian F Hutton, Anna Barnes, Sébastien Ourselin, Simon Arridge, Celia OâĂŹMeara, Simon Wan, and David Atkinson. Practical pet respiratory motion correction in clinical pet/mr. *Journal of Nuclear Medicine*, 56(6):890–896, 2015.

Ravindra Manjeshwar, Xiaodong Tao, Evren Asma, and Kris Thielemans. Motion compensated image reconstruction of respiratory gated PET/CT. In *Biomedical Imaging: Nano to Macro, 2006. 3rd IEEE International Symposium on*, pages 674–677. IEEE, 2006.

Jamie R McClelland, David J Hawkes, Tobias Schaeffter, and Andrew P King. Respiratory motion models: a review. *Medical Image Analysis*, 17(1):19–42, 2013.

Steven R Meikle and Ramsey D Badawi. Quantitative techniques in PET. *Positron Emission Tomography*, pages 93–126, 2005.

M Modat et al. *Computer Methods Prog. Bio.*, 98(3):278–284, 2010.

Sascha Moehrs, Michel Defrise, Nicola Belcari, Alberto Del Guerra, Antonietta Bartoli, Serena Fabbri, and Gianluigi Zanetti. Multi-ray-based system matrix generation for 3D PET reconstruction. *Physics in Medicine and Biology*, 53(23):6925, 2008.

Siti Salasiah Mokri, MI Saripan, MH Marhaban, and AJ Nordin. Lung segmentation in CT for thoracic PET-CT registration through visual study. In *Biomedical Engineering and Sciences (IECBES), 2012 IEEE EMBS Conference on*, pages 550–554. IEEE, 2012.

SA Nehmeh, YA Erdi, CC Ling, K.E. Rosenzweig, O.D. Squire, L.E. Braban, E. Ford, K. Sidhu, G.S. Mageras, S.M. Larson, et al. Effect of respiratory gating on reducing lung motion artifacts in PET imaging of lung cancer. *Medical Physics*, 29(3):366–371, 2002b.

Sadek A Nehmeh and Yusuf E Erdi. Respiratory motion in positron emission tomography/computed tomography: a review. In *Seminars in nuclear medicine*, volume 38, pages 167–176. Elsevier, 2008.

Sadek A Nehmeh, Yusuf E Erdi, Clifton C Ling, Kenneth E Rosenzweig, Heiko Schoder, Steve M Larson, Homer A Macapinlac, Olivia D Squire, and John L Humm. Effect of respiratory gating on quantifying PET images of lung cancer. *Journal of Nuclear Medicine*, 43(7):876–881, 2002a.

Sadek A Nehmeh, Yusuf E Erdi, Kenneth E Rosenzweig, Heiko Schoder, Steve M Larson, Olivia D Squire, and John L Humm. Reduction of respiratory motion artifacts in PET imaging of lung cancer by respiratory correlated dynamic PET: methodology and comparison with respiratory gated PET. *Journal of Nuclear Medicine*, 44(10):1644–1648, 2003.

Sadek A Nehmeh, Yusuf E Erdi, Gustavo SP Meirelles, Olivia Squire, Steven M Larson, John L Humm, and Heiko Schöder. Deep-inspiration breath-hold PET/CT of the thorax. *Journal of Nuclear Medicine*, 48(1):22–26, 2007.

Philip J Noonan, Jon Howard, Deborah Tout, Ian Armstrong, Heather A Williams, Tim F Cootes, William A Hallett, and Rainer Hinz. Accurate markerless respiratory tracking for gated whole body PET using the Microsoft Kinect. In *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2012 IEEE*, pages 3973–3974. IEEE, 2012.

John M Ollinger. Model-based scatter correction for fully 3D PET. *Physics in Medicine and Biology*, 41(1):153, 1996.

Yuki Otani, Ichirou Fukuda, Nobuhiro Tsukamoto, Yu Kumazaki, Hiroshi Sekine, Etsuko Imabayashi, Osamu Kawaguchi, Takayuki Nose, Teruki Teshima, and Takushi Dokiya. A comparison of the respiratory signals acquired by different respiratory monitoring systems used in respiratory gated radiotherapy. *Medical Physics*, 37(12):6178–6186, 2010.

Sébastien Ourselin, Alexis Roche, Gérard Subsol, Xavier Pennec, and Nicholas Ayache. Reconstructing a 3d structure from serial histological sections. *Image and Vision Computing*, 19(1):25–31, 2001.

Cihat Ozhasoglu and Martin J Murphy. Issues in respiratory motion compensation during external-beam radiotherapy. *International Journal of Radiation Oncology* Biology* Physics*, 52(5):1389–1399, 2002.

Tinsu Pan. Comparison of helical and cine acquisitions for 4D-CT imaging with multislice CT. *Medical Physics*, 32(2):627–634, 2005.

K Pearson. On lines and planes of closest fit to system of points in space. *Philiosophical Magazine*, 2(6):559–572, 1901.

Audrey Pepin, Joel Daouk, Pascal Bailly, Sebastien Hapdey, and Marc-Etienne Meyer. Management of respiratory motion in PET/computed tomography: the state of the art. *Nuclear Medicine Communications*, 35(2):113, 2014.

Yani Picard and Christopher J Thompson. Motion correction of PET images using multiple acquisition frames. *IEEE Transactions on Medical Imaging*, 16(2):137–144, 1997.

Jennifer Prekeges et al. *Nuclear Medicine Instrumentation*. Jones & Bartlett Publishers, 2012.

Feng Qiao, Tinsu Pan, John W Clark Jr, and Osama R Mawlawi. A motion-incorporated reconstruction method for gated PET studies. *Physics in Medicine and Biology*, 51(15): 3769, 2006.

Arman Rahmim, Olivier Rousset, and Habib Zaidi. Strategies for motion tracking and correction in PET. *PET Clinics*, 2(2):251–266, 2007.

Joseph M Reinhardt, Kai Ding, Kunlin Cao, Gary E Christensen, Eric A Hoffman, and Shalmali V Bodas. Registration-based estimates of local lung tissue expansion compared to xenon CT measures of specific ventilation. *Medical Image Analysis*, 12(6):752–763, 2008.

Silin Ren, Xiao Jin, Chung Chan, Yiqiang Jian, Tim Mulnix, Chi Liu, and Richard E Carson. Data-driven event-by-event respiratory motion correction using TOF PET list-mode centroid of distribution. *Physics in Medicine and Biology*, 62(12):4741, 2017.

Michael Riedel, N Navab, and A Moller. Respiratory motion estimation: tests and comparison of different sensors. *Interdisciplinary Project. Technische Universität München Fakultät für Informatik, Germany*, 2006.

Eike Rietzel, Tinsu Pan, and George TY Chen. Four-dimensional computed tomography: Image formation and clinical protocol. *Medical Physics*, 32(4):874–889, 2005.

Paul J Schleyer, Michael J O'Doherty, Sally F Barrington, and Paul K Marsden. Retrospective data-driven respiratory gating for PET/CT. *Physics in Medicine and Biology*, 54(7):1935, 2009.

Paul J Schleyer, Michael J O'Doherty, and Paul K Marsden. Extension of a data-driven gating technique to 3D, whole body PET studies. *Physics in Medicine and Biology*, 56 (13):3953, 2011.

Matthias Schmand, Ziad Burbar, James Corbeil, Nan Zhang, Christian Michael, Lary Byars, Lars Eriksson, Ron Grazioso, Matthew Martin, Andrew Moor, et al. BrainPET: first human tomograph for simultaneous (functional) PET and MR imaging. *Journal of Nuclear Medicine*, 48(supplement 2):45P–45P, 2007.

W. P. Segars, G. Sturgeon, S. Mendonca, Jason Grimes, and B. M. W. Tsui. 4D XCAT phantom for multimodality imaging research. *Medical Physics*, 37(9):4902–4915, 2010.

Lawrence A Shepp and Yehuda Vardi. Maximum likelihood reconstruction for emission tomography. *IEEE Transactions on Medical Imaging*, 1(2):113–122, 1982.

D. L. Snyder and D. G. Politte. Image Reconstruction from List-Mode Data in an Emission Tomography System Having Time-of-Flight Measurements. *IEEE Transactions on Nuclear Science*, 30(3):1843–1849, June 1983. ISSN 0018-9499. doi: 10.1109/TNS.1983.4332660.

M Teräs, T Tolvanen, JJ Johansson, JJ Williams, and J Knuuti. Performance of the new generation of whole-body PET/CT scanners: Discovery STE and Discovery VCT. *European Journal of Nuclear Medicine and Molecular Imaging*, 34(10):1683–1692, 2007.

Kris Thielemans, Shailendra Rathore, Fredrik Engbrant, and Pasha Razifar. Device-less gating for PET/CT using PCA. In *IEEE Nuclear Science Symp. and Medical Imaging Conf. Record*, pages 3904–3910, 2011.

Kris Thielemans, Charalampos Tsoumpas, Sanida Mustafovic, Tobias Beisel, Pablo Aguiar, Nikolaos Dikaios, and Matthew W Jacobson. STIR: software for tomographic image reconstruction release 2. *Physics in Medicine and Biology*, 57(4):867, 2012.

Kris Thielemans, Paul Schleyer, Paul K Marsden, Ravindra M Manjeshwar, Scott D Wollenweber, and Alexander Ganin. Comparison of different methods for data-driven respiratory gating of PET data. In *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2013 IEEE*, pages 1–4. IEEE, 2013.

Kris Thielemans, Paul Schleyer, Paul K Marsden, Jarmo Teuho, M Ter, V Bettinardi, L Menezes, RM Manjeshwar, CW Stearns, et al. Data-driven dual-gating for cardiac PET. In *IEEE Nuclear Science Symp. and Medical Imaging Conf. Record*, pages 1–4. IEEE, 2014.

Brian Thorndyke, Eduard Schreibmann, Albert Koong, and Lei Xing. Reducing respiratory motion artifacts in positron emission tomography through retrospective stacking. *Medical Physics*, 33(7):2632–2641, 2006.

T. Tomitani. Image Reconstruction and Noise Evaluation in Photon Time-of-Flight Assisted Positron Emission Tomography. *IEEE Transactions on Nuclear Science*, 28(6):4581–4589, Dec 1981. ISSN 0018-9499. doi: 10.1109/TNS.1981.4335769.

S Vandenberghe, E Mikhaylova, E DâĂŹHoe, P Mollet, and JS Karp. Recent developments in time-of-flight PET. *EJNMMI physics*, 3(1):1, 2016.

SS Vedam, VR Kini, PJ Keall, V Ramakrishnan, H Mostafavi, and R Mohan. Quantifying the predictability of diaphragm motion during respiration with a noninvasive external marker. *Medical Physics*, 30(4):505–513, 2003.

JA Verschakelen, L Van Fraeyenhoven, Genevieve Laureys, M Demedts, and AL Baert. Differences in CT density between dependent and nondependent portions of the lung: influence of lung volume. *American Journal of Roentgenology*, 161(4):713–717, 1993.

D Visvikis, F Lamare, P Bruyant, N Boussion, and C Cheze Le Rest. Respiratory motion in positron emission tomography for oncology applications: Problems and solutions. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 569(2):453–457, 2006.

Dimitris Visvikis, Olivier Barret, Tim Fryer, Alexandre Turzo, Frederic Lamare, Le Rest, Catherine Cheze, and Yves Bizais. A posteriori respiratory motion gating of dynamic PET images. In *IEEE Nuclear Science Symp. and Medical Imaging Conf. Record*, volume 5, pages 3276–3280, 2003.

Christian Wachinger, Mehmet Yigitsoy, Erik-Jan Rijkhorst, and Nassir Navab. Manifold learning for image-based breathing gating in ultrasound and MRI. *Medical Image Analysis*, 16(4):806–818, 2012.

OL Wade. Movements of the thoracic cage and diaphragm in respiration. *The Journal of Physiology*, 124(2):193–212, 1954.

Mengdie Wang, Ning Guo, Hui Zhang, Georges Elfhakri, Guangshu Hu, and Quanzheng Li. Retrospective data-driven respiratory gating for PET using TOF information. In *Engineering in Medicine and Biology Society (EMBC), 2015 37th Annual International Conference of the IEEE*, pages 4520–4523. IEEE, 2015.

Charles C Watson. New, faster, image-based scatter correction for 3D PET. *IEEE Transactions on Nuclear Science*, 47(4):1587–1594, 2000.

Jingyan Xu and Benjamin MW Tsui. Improved Intrinsic Motion Detection Using Time-of-Flight PET. *IEEE Transactions on Medical Imaging*, 34(10):2131–2145, 2015.

Barbara Zitova and Jan Flusser. Image registration methods: a survey. *Image and Vision Computing*, 21(11):977–1000, 2003.