DTU Library

A machine learning method for fast and accurate characterization of depth-of-interaction gamma cameras
Paper

Pedemonte, Stefano; Pierce, Larry; Van Leemput, Koen

PAPER

A machine learning method for fast and accurate characterization of depth-of-interaction gamma cameras

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

A machine learning method for fast and accurate characterization of depth-of-interaction gamma cameras

Stefano Pedemonte[1], Larry Pierce[2] and Koen Van Leemput[1,3]

[1] Martinos Center for Biomedical Imaging, MGH, Harvard Medical School, Boston, MA, USA

[2] Imaging Research Lab, University of Washington, Seattle, WA, USA

[3] Department of Applied Mathematics and Computer Science, Technical University of Denmark, Lyngby, Denmark

E-mail: `spedemonte@mgh.harvard.edu`

Abstract. Measuring the depth-of-interaction (DOI) of gamma photons enables increasing the resolution of emission imaging systems. Several design variants of DOI-sensitive detectors have been recently introduced to improve the performance of scanners for Positron Emission Tomography (PET). However, the accurate characterization of the response of DOI detectors, necessary to accurately measure the DOI, remains an unsolved problem. Numerical simulations are, at the state of the art, imprecise, while measuring directly the characteristics of DOI detectors experimentally is hindered by the impossibility to impose the depth-of-interaction in an experimental set-up. In this article we introduce a machine learning approach for extracting accurate forward models of gamma imaging devices from simple pencil-beam measurements, using a nonlinear dimensionality reduction technique in combination with a finite mixture model. The method is purely data-driven, not requiring simulations, and is applicable to a wide range of detector types. The proposed method was evaluated both in a simulation study and with data acquired using a monolithic gamma camera designed for PET (the cMiCE detector), demonstrating the accurate recovery of the DOI characteristics. The combination of the proposed calibration technique with maximum-a-posteriori estimation of the coordinates of interaction provided a depth resolution of ≈ 1.14 mm for the simulated PET detector and ≈ 1.74 mm for the cMiCE detector. The software and experimental data are made available at `http://occiput.mgh.harvard.edu/depthembedding/`.

1. Introduction

Scintillation gamma cameras are composed of one or multiple scintillator crystals coupled to a set of photo-detectors (Fig. 1). The interaction of a gamma photon with the scintillating material determines, through a cascade of radiative and non-radiative processes, the emission of secondary photons. These are captured by the set of photo-detectors, producing electrical signals that are amplified by the front-end electronics and, in most designs, digitalized and streamed to a digital computer (Fiorini et al., 2008).

2

Figure 1. Schematic representation of a monolithic gamma camera. The scintillator is coupled with the array of photo-detectors indexed with $d = 1, \ldots, D$. The forward model of the camera is discretized in voxels indexed with (i, j, k) along the three axis (x, y, z). The interaction of a gamma photon in (i, j, k) determines the noisy measurement vector $\mathbf{u} = \{u_1, \ldots, u_D\}$. The discrete forward model of the camera is characterized by the set of expected measurements $\{l_{dijk}\}$ in detector element d for gamma interaction in (i, j, k).

The set of signals acquired by the photo-detectors bears information about the energy of the gamma photon and about the location of the interaction within the scintillator crystal. Event estimation algorithms then enable the reconstruction of the coordinates of interaction and energy of the gamma photon from the measurement of secondary photons.

In the original design introduced by Anger (1958), of which modern designs are a derivation, the gamma camera was conceived as a planar imaging device, providing information about the coordinates of interaction on a plane. Recently, there has been increasing interest in depth-of-interaction (DOI) detectors, which provide information about the coordinates of interaction of the gamma photon in three-dimensional space. In particular in PET imaging, DOI sensitive detectors eliminate the degradation due to parallax errors, promising to solve the long standing problem of the non-uniform resolution across the scanner field-of-view (Ito et al., 2011; MacDonald and Dahlbom, 1998). With the advancement of time-of-flight (TOF) PET, the DOI information has also shown to play a key role in the improvement of the time resolution by enabling a more accurate measurement of the length of the lines-of-response, therefore promising to further increase the spatial resolution of TOF PET scanners (Ito et al., 2011; Shibuya et al., 2007).

1.1. Existing methods for the estimation of the DOI

Various designs of DOI sensitive cameras have been proposed, falling widely in two categories: pixellated and monolithic systems. Here we focus on the monolithic design, although our results can also be applied to pixellated detectors (see the discussion at the end of the article in Section 5).

Several methods for estimating the DOI of photons in monolithic gamma cameras have been proposed in the literature. A first class of methods utilizes the centroid algorithm initially described by Anger (1958) for two-dimensional estimation. In these

3

techniques, the in-plane coordinates are first reconstructed using the standard centroid algorithm, after which the DOI is derived from a measurement of the width of the light distribution (Lerche et al., 2005). While these techniques are computationally efficient and can be implemented in hardware using simple resistive networks (Lerche et al., 2005), they suffer from the bias and random errors characteristic of the centroid method and assume an arbitrary relation (e.g., linear in Lerche et al. 2005) between the width of the light distribution and the DOI.

A second class of methods is based on the L-nearest-neighbors (L-NN) algorithm, introduced by Maas et al. (2009). Instead of *estimating* the DOI, these methods aim at *correcting* for it. Three versions have been described. The first version, described in Maas et al. 2009, is based on a calibration phase requiring the exhaustive illumination of the detector, using pencil beams covering the entire entrance plane, at multiple incidence angles. Such measurement requires a complex set-up, with a rotating stage and electronic collimation, and very long acquisition times; furthermore the estimation method requires an extremely large number of vector comparisons. To address these limitations, a second version was proposed in van Dam et al. 2011b, where the calibration time is reduced by acquiring photon beams only in the perpendicular direction. Such a simplified L-NN algorithm, however does not enable the estimation of the DOI. The third version of the algorithm, described in van Dam et al. 2011a, extends the simplified L-NN algorithm to include an estimation of the DOI. The two-dimensional intersection between the gamma ray and the entrance plane of the camera is first estimated using the two-dimensional method described in van Dam et al. 2011b, after which the DOI is estimated by comparing the measurement in the brightest photo-detector with values stored in a lookup-table, precomputed using a Monte Carlo simulator. Since this method relies on an accurate Monte Carlo simulation of the light transport process, it requires a complex software set-up and may be sensitive to the parameterization of the materials, unpredictable surface effects (such as optical coupling between the scintillator crystal and the array of detectors), and variations in manufacturing processes.

A third class of methods is based on training neural networks to predict the two-dimensional intersection between the gamma ray and the entrance plane of the camera form the array of photo-detector measurements (Marone et al., 2009; Bruyndonckx et al., 2008, 2006, 2004). While these methods lead to computationally efficient estimation algorithms amenable to real-time execution (Wang et al., 2011), similarly to Maas et al. 2009 they currently rely on the exhaustive illumination of the detector, using pencil beams covering the entire entrance plane of the camera, at multiple incidence angles. Such measurement requires a complex set-up, with a rotating stage and electronic collimation, and very long acquisition times.

A final class of DOI estimation methods, initially investigated by Gray and Macovski (1976) and further explored by Barrett et al. (2009a), is based on statistical models. In this approach, the estimation is posed as an inverse problem, where the calculation of the location of interaction and the energy of the gamma photon requires evaluating a forward problem that predicts the signals produced by the gamma camera

4

for a given location of interaction and energy of the gamma photon, effectively modeling the characteristics of the camera. While model-based statistical methods provide optimal estimates of the coordinates of interaction and energy if the forward model is exact (Barrett et al., 2009a), their accuracy is in practice dictated by the accuracy of the forward model.

Various numerical simulation techniques for forward modeling of gamma cameras have been described in the literature, including analytical approaches such as the simple solid-angle model based on L'Huilier's formula (Pedemonte et al., 2009; Li et al., 2010), and methods based on Monte Carlo integration of the light transport equations (Boschini and Fiorini, 1999). However, analytical models are too simplistic to capture the complexity of real-world gamma cameras; whereas Monte Carlo simulations are complex to set up and yet may not fully capture the real characteristics of the imaging devices.

An alternative to forward modeling using numerical simulation is to measure the characteristics of gamma imaging devices experimentally. Experimental characterization can be accomplished by exhaustively scanning the two-dimensional entrance surface of the detector with a perpendicular collimated pencil-beam: averaging the signals produced for all interactions at each given position of the beam then yields a two-dimensional estimate of the response of the camera (Barrett et al., 2009b). However, the depth at which the gamma photons interact cannot be imposed or directly observed, making the measurement of the response of the camera as a function of the three-dimensional coordinates a challenging missing-data problem.

Ling et al. (2007) proposed an algorithm for the experimental estimation of the depth-dependent forward model for DOI estimation using perpendicular pencil-beam measurements. The method proposed by Ling et al. estimates the depth-dependent forward model by clustering the events at a given beam location in depth bins according to the measurement value in the channel with the most intense signal, and averaging the data vectors in each cluster. However this method suffers from two limitations: it uses only the measurement of the brightest photo-detector to cluster the calibration data points in depth bins, making inefficient use of the noisy measurements; and the calibration data is clustered arbitrarily, without establishing a quantitative relation between the depth-of-interaction and the depth bins. Consequently, as reported in the experiments in Ling et al. (2007), the algorithm is in practice limited to the discrimination of 2 or 4 levels of DOI. Another method, based on maximum likelihood estimation and polynomial parametrization of the forward model, was proposed in Hunter et al. (2009); however this method requires an initial estimate of the depth-dependent forward model obtained from a Monte Carlo simulation.

1.2. Contribution of this paper

In this paper, we propose a novel, purely data-driven methodology for the estimation of the depth-dependent forward model of gamma cameras from simple pencil-beam

5

measurements, enabling the accurate estimation of the three-dimensional coordinates of interaction and energy of gamma photons using model-based statistical estimation. The technique that we describe does not require either the use of simulations, other pre-computed values, or the optical and the electrical characteristics of the system, simplifying the calibration of gamma cameras and enabling its application to a wide variety of detector types. It enables short calibration measurements by making use of the information from all available photo-detectors to estimate the DOI of each photon. We test the proposed methodology on both simulated and measured data from monolithic PET detectors and show that the method provides for accurate depth-dependent detector characterization using pencil-beam measurements enabling the consequent accurate estimation of the three-dimensional coordinates of interaction.

The rest of this article is organized as follows: section 2 describes the probabilistic model of the gamma detection process and reviews the model-based methods for estimating the coordinates of interaction and energy of the gamma photons. Section 3 presents the challenge of measuring experimentally the forward model of DOI gamma cameras, required by such model-based methods, and introduces our new machine learning method for estimating the forward model from simple pencil-beam measurements. Section 4 evaluates the algorithm in simulations and real-data experiments. Section 5 summarizes the proposed algorithm and the results of the experiments, and presents a discussion of the method in relation to previously published techniques for DOI estimation.

2. Review of model-based DOI estimation

We start by modeling position-sensitive gamma imaging devices in the probabilistic framework, and by reviewing maximum-a-posteriori estimation of the coordinates of interaction and of the energy. We emphasize that the focus of this paper is on accurately characterizing gamma cameras, as detailed in section 3, so that these methods can be used in practice. Reconstructing tomographic images based on the provided coordinates of interaction is an entirely separate problem for which standard tomographic reconstruction algorithms can be used, and falls outside the scope of this work.

2.1. Probabilistic model of the interaction process

Let (x, y, z) be the coordinates of three-dimensional space with z being the direction normal to the entrance plane of the camera (Fig. 1) and let $\xi \left[\frac{photons}{KeV} \right]$ be the photon yield of the scintillator. In consequence of the interaction of a gamma photon with energy E $[KeV]$, the crystal emits on average $E\xi$ secondary photons. These propagate through the camera and are eventually detected by the photo-detectors, indexed with $d = 1, \ldots, D$, producing the measurement vector $\mathbf{u} = \{u_1, \ldots, u_D\}$. Let the geometry and the optical characteristics of the device be expressed by the probability $\tilde{s}_d(x, y, z)$

6

that a secondary photon emitted in consequence of a gamma interaction in (x, y, z) enters the active area of photo-detector d; and let $G_d \left[\frac{V}{photon} \right]$ express the characteristics of photo-detector d, encompassing the quantum efficiency and the gain of the charge amplifier‡. The function of three-dimensional space $\tilde{l}_d(x, y, z) = \xi G_d \tilde{s}_d(x, y, z)$ expresses the overall characteristics of the gamma camera, encompassing the geometry, the optical characteristics of the scintillator and of the detectors, the quantum efficiency of the detectors and the gain of the read-out electronics.

In order to discretize the forward model for numerical implementation, let us partition the volume of the scintillator into a grid of (I, J, K) not necessarily cubic voxels (Fig. 1), indexed in the three-dimensional space with (i, j, k), and assume that $\tilde{l}_d(x, y, z)$ is constant within each voxel. Let us indicate with l_{dijk} the discrete approximation of $\tilde{l}_d(x, y, z)$:

$$\mathbf{L} = \{l_{dijk}\}. \tag{1}$$

In consequence of the interaction of a gamma photon in voxel i, j, k, the signal expected at photo-detector d is, in the discrete approximation:

$$\bar{u}_d = E l_{dijk}. \tag{2}$$

The set of parameters $\mathbf{L}$ constitutes the discrete numerical representation of the characteristics of the gamma camera, expressing the signal expected at each read-out channel in consequence of the interaction of a gamma photon at a given location of interaction; per unit of energy of the gamma photon. Section 3 will describe our novel method to estimate $\mathbf{L}$ experimentally from pencil-beam measurements; the remainder of this section completes the description of the probabilistic model of the photon interaction process and reviews how $\mathbf{L}$ is utilized, in the model-based approach, to estimate the coordinates of interaction and energy.

Under the assumption that the coordinates of interaction and energy of the gamma photon are the unique cause of the measurements (e.g., no dark counts), the measurements are conditionally independent such that the conditional probability $p(\mathbf{u}|E, i, j, k)$ factorizes according to:

$$p(\mathbf{u}|E, i, j, k) = \prod_{d=1}^{D} p(u_d|E, i, j, k). \tag{3}$$

Assuming that the uncertainty of the measurement of each u_d is dominated by the intrinsic uncertainty due to the photon counting statistics, $p(u_d|E, i, j, k)$ is expressed by the Poisson mass function with expectation $E l_{dijk}$:

$$p(u_d|E, i, j, k) = \frac{e^{-E l_{dijk}} (E l_{dijk})^{u_d}}{u_d!}. \tag{4}$$

‡ The unit measures may differ depending on the type of detector and on the characteristics of the read-out electronics.

7

Figure 2. The prior probability of interaction is determined by integrating the Beer-Lambert law of absorption, depicted in panel (A) in the 1-D case, and depends on the geometry of the camera, the characteristics of the scintillating material and of the trajectories of the incoming gamma photons. (B) Prior probability of interaction for parallel photons of unknown (x, y) location, utilized in the event estimation experiments.

Equations (3) and (4) constitute the probabilistic model of the interaction process, parameterized by the forward model $\mathbf{L} = \{l_{dijk}\}$.

2.2. Event estimation

While often the event estimation problem is formulated as maximum-likelihood inference (see Barrett et al., 2009b), here we adopt maximum-a-posteriori (MAP) inference in order to include information about the location of interaction of the gamma photon available prior to the measurement. In general, knowledge of the geometry of the gamma camera, of the characteristics of the scintillating material and of the trajectories of the incoming gamma photons, determines a non-uniform prior probability distribution of the location of interaction, which we denote with $p(i, j, k)$. This is illustrated graphically in Fig. 2; for the case of collimated parallel photons with unknown (x, y) location (Fig. 2.B), we have that $p(i, j, k) = \frac{1}{IJ}p(k)$, where $p(k)$ will be derived from the Beer Lambert law of absorption in section 3.1.

Under the MAP criterion, the event estimation problem is formulated as the problem of finding the coordinates of interaction $\hat{i}, \hat{j}, \hat{k}$ and the energy $\hat{E}$ that maximize the posterior probability distribution of i, j, k, E, given the photo-detector measurement vector $\mathbf{u}$:

$$\hat{i}, \hat{j}, \hat{k}, \hat{E} = \arg \max_{(i,j,k,E)} p\left(E, i, j, k | \mathbf{u}\right), \tag{5}$$

with the posterior probability expressed by Bayes formula:

$$p(E, i, j, k | \mathbf{u}) = \frac{p(\mathbf{u} | E, i, j, k) p(E, i, j, k)}{p(\mathbf{u})}, \tag{6}$$

and with the model of the interaction process $p(\mathbf{u}|E, i, j, k)$ expressed by Eq. (3). The solution of (5) is given by (see Appendix A for derivation):

8

$$\hat{i}, \hat{j}, \hat{k} = \arg \min_{i,j,k} \left\{ \ln(\sum_{d=1}^{D} l_{dijk}) \sum_{d=1}^{D} u_d - \sum_{d=1}^{D} (u_d \ln l_{dijk}) - \ln p(i, j, k) \right\}. \quad (7)$$

$$\hat{E} = \frac{\sum_{d=1}^{D} u_d}{\sum_{d=1}^{D} l_{d\hat{i}\hat{j}\hat{k}}}; \quad (8)$$

A variety of algorithms can be utilized to maximize (7), such as simple exhaustive search or fast adaptive grid search methods (Pedemonte et al., 2009).

3. Estimating the forward model using machine learning

Before the MAP estimation of the coordinates of interaction and energy of section 2 - Eq. (7) and (8) can be performed, the forward model $\mathbf{L}$ of the gamma camera needs to be known. Here we propose a new method to estimate the forward model automatically and accurately from simple pencil-beam measurements. The method that we propose combines two techniques from the machine learning literature: nonlinear dimensionality reduction – which aims to simplify structured high-dimensional data by projecting it into a lower-dimensional space – and finite mixture modeling – a statistical method for characterizing models with unobserved, latent variables.

The schematic of an experimental set-up for estimating the characteristics of a gamma camera $\mathbf{L}$ is shown in Fig. 3. A collimated beam of gamma photons with known energy $\bar{E}$, perpendicular to the (x, y) plane, is positioned at locations displaced on a regular grid of $I \times J$ points, covering the entire entrance plane of the gamma camera. The use of a beam perpendicular to the (x, y) plane decouples the problem of estimating $\mathbf{L}$ in a set of $I \times J$ independent problems. The data acquired at each beam position $\bar{i}, \bar{j}$ bears information about a subset of $\mathbf{L}$: $\mathbf{L}_{\bar{i}\bar{j}} = \{l_{d\bar{i}\bar{j}k}\}$. In the remainder of this section, we will concentrate on the problem of estimating the parameters $\mathbf{L}_{\bar{i}\bar{j}}$ for a given location of the beam.

For each position of the beam $(\bar{i}, \bar{j})$, N interaction events are recorded; let us index with $n = 1, \ldots, N$ the interaction events and indicate with $\mathbf{U} = \{\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_N\}$ the set of N corresponding measurements $\mathbf{u}_n = \{u_{1n}, u_{2n}, \ldots, u_{Dn}\}$. The problem of estimating the characteristics of the gamma camera is a missing-data problem: if the depth-of-interaction of the gamma photons were known, it would be straightforward to compute $\mathbf{L}_{\bar{i}\bar{j}}$ by averaging the measurements associated to each discrete value of the depth-of-interaction. However, in reality both $\mathbf{L}_{\bar{i}\bar{j}}$ and the depth-of-interaction of the photons are unknown.

We investigate two approaches to address this missing-data problem. The first approach poses the problem of estimating $\mathbf{L}_{\bar{i}\bar{j}}$ as a finite mixture modeling problem, accounting for the actual noise properties of the measurement system when estimating the forward model; however it leads to a local optimization algorithm that can be sensitive to its initialization. The second approach addresses the problem from a
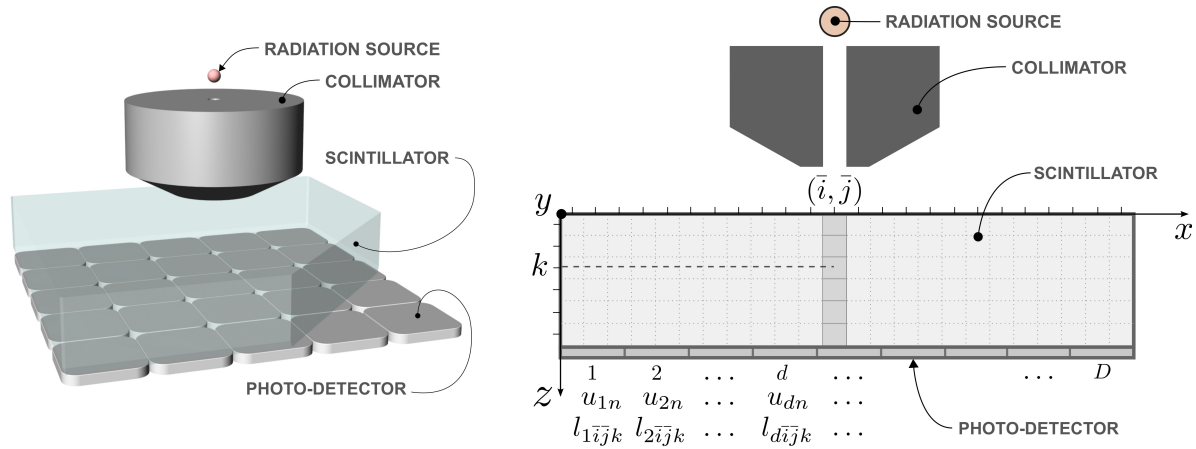
Figure 3. Schematic of an experimental set-up to learn the depth-dependent response of a monolithic gamma camera. A collimated photon beam is focused at a given location $(\bar{i}, \bar{j})$. Though the depth-of-interaction can not be imposed in the experiment, the depth-dependent response can be learned from the signals produced by multiple photon interactions at $(\bar{i}, \bar{j})$.

different perspective, by resorting to the concept of nonlinear dimensionality reduction. This leads to a globally convergent algorithm, that however does not account for photon counting noise. In the experiments, we will demonstrate that the combination of the two approaches provides a globally-convergent, fully-automated algorithm for the accurate estimation of the forward model.

3.1. Prior probability distribution of depth

We first derive the expression of the prior distribution $p(k)$ in the case of the experimental set-up of Fig. 3, as this will be used in both the finite mixture modeling and the nonlinear dimensionality reduction approach. With collimated photons traveling in the direction perpendicular to the camera plane, the Beer-Lambert exponential law of absorption expresses the probability of interaction as a function of z:

$$p(z) = \alpha_{\bar{E}} e^{-\alpha_{\bar{E}} z}, \tag{9}$$

with $\alpha_{\bar{E}}$ $[cm^{-1}]$ being the narrow-beam attenuation coefficient of the scintillator crystal for the gamma photons of energy $\bar{E}$ utilized for the experimental characterization. Letting Δz be the size of the voxels along direction z, the probability that a detected gamma photon has interacted in voxel k is, integrating (9):

$$p(k) = \frac{\int_{(k-1)\Delta z}^{k\Delta z} p(z)\mathrm{d}z}{\sum_{k'=1}^{K} \int_{(k'-1)\Delta z}^{k'\Delta z} p(z)\mathrm{d}z} = \frac{e^{-\alpha_{\bar{E}}(k-1)\Delta z} - e^{-\alpha_{\bar{E}} k\Delta z}}{1 - e^{-\alpha_{\bar{E}} K\Delta z}}. \tag{10}$$

This is easily calculated, given the parameter $\alpha_{\bar{E}}$ of the scintillator crystal.

10

3.2. Estimating the forward model using finite mixture modeling

The missing-data problem of estimating the depth-dependent forward model can be addressed by means of an iterative procedure: using a provisional estimate of $\mathbf{L}_{\bar{i}\bar{j}}$, one may estimate the depth-of-interaction k of each gamma photon, and then update $\mathbf{L}_{\bar{i}\bar{j}}$ by averaging the signals produced by interactions associated to each value of k. An iterative algorithm of this kind indeed emerges by formulating the problem of estimating $\mathbf{L}_{\bar{i}\bar{j}}$ as maximum-likelihood estimation. The generative probabilistic model of the photon detection process, described in section 2.1, can be utilized to express the likelihood of the parameters $\mathbf{L}_{\bar{i}\bar{j}}$ when observing the set of N gamma interactions with known $\bar{i}, \bar{j}$ location:

$$p(\mathbf{u}_n|\mathbf{L}_{\bar{i}\bar{j}}, k) = \prod_{d=1}^{D} p(u_d|E, \bar{i}, \bar{j}, k),$$

with $p(u_d|E, \bar{i}, \bar{j}, k)$ given by Eq. (4). Marginalizing over the depth-of-interaction k, the probability distribution associated to the measurement vector $\mathbf{u}_n$ is a finite mixture model:

$$p(\mathbf{u}_n|\mathbf{L}_{\bar{i}\bar{j}}) = \sum_{k=1}^{K} p(k)p(\mathbf{u}_{dn}|\mathbf{L}_{\bar{i}\bar{j}}, k),$$

Due to the independence of the N events of interaction, the probability of $\mathbf{U}$ conditional to the forward model parameters is given by the product of terms $p(\mathbf{u}_n|\mathbf{L}_{\bar{i}\bar{j}})$, so that finally:

$$p(\mathbf{U}|\mathbf{L}_{\bar{i},\bar{j}}) = \prod_{n=1}^{N} \sum_{k=1}^{K} p(k) \prod_{d=1}^{D} \frac{e^{-El_{dijk}}(El_{dijk})^{u_{dn}}}{u_{dn}!}. \tag{11}$$

While there is no closed form solution for the maximizer of Eq. (11) with respect to $\mathbf{L}_{\bar{i}\bar{j}}$, we can utilize the Expectation Maximization (EM) algorithm to update $l_{\bar{i}\bar{j}k}$ iteratively. The EM algorithm provides a gradient-type iterative update formula with implicit line search that guarantees an increase of the likelihood value at each new iteration. The EM update formula consists of two steps (Bishop, 2007):

$$h_{nk}^{(q)} = \frac{p(k) \prod_{d=1}^{D} \left[e^{-\bar{E}l_{d\bar{i}\bar{j}k}^{(q-1)}} \left(\bar{E}l_{d\bar{i}\bar{j}k}^{(q-1)} \right)^{u_{dn}} \right]}{\sum_{k'=1}^{K} p(k') \prod_{d=1}^{D} \left[e^{-\bar{E}l_{d\bar{i}\bar{j}k'}^{(q-1)}} \left(\bar{E}l_{d\bar{i}\bar{j}k'}^{(q-1)} \right)^{u_{dn}} \right]},$$

$$l_{d\bar{i}\bar{j}k}^{(q)} = \frac{\sum_{n=1}^{N} h_{nk}^{(q)} u_{dn}}{\bar{E} \sum_{n=1}^{N} h_{nk}^{(q)}}. \tag{12}$$

At each iteration q, the weights $h_{nk}^{(q)}$ associating each measurement n to depth k are computed from $\mathbf{L}_{\bar{i}\bar{j}}^{(q-1)}$, and $\mathbf{L}_{\bar{i}\bar{j}}^{(q)}$ is then computed, according to Eq. (12), as the sum of the measurement vectors weighted by $h_{nk}^{(q)}$. Although EM guarantees an increase of the

likelihood at each new iteration, it is a greedy algorithm that can get trapped in local optima, making the solution dependent on the initialization. In the remainder of the paper, we will refer to this first algorithm for the estimation of the depth-of-interaction of a set of photons with the acronym *MLE-EM*: maximum-likelihood estimation by means of the Expectation Maximization algorithm.

3.3. Estimating the forward model using nonlinear dimensionality reduction

In this section we present the problem of estimating the forward model from a different perspective that leads to a globally convergent algorithm. The formulation that follows is based on two observations:

- As depicted in Fig. 2.A, for a given position of the beam, according to the Beer-Lambert law of absorption, the number of interactions per unit volume is maximum at the entrance location of the beam and decreases exponentially along z.
- For a given position of the beam, two measurement vectors are expected to be similar if they are produced by interactions at near depth locations, and less similar if they are produced by two interactions at distant depths.

The second statement is true if the detector indeed provides information about the depth-of-interaction. In consequence of these two considerations, amongst N photon interaction measurements, there is an abundance of measurements produced by interactions that have occurred near the surface and scarcity of measurements produced by deep interactions; the latter will be similar to one another and different from the measurements originated near the surface. Since the photons from a pencil beam have essentially only one degree of freedom in their interaction within the crystal (depth), we expect the D-channel beam data produced by the D photodetectors to lie near a one-dimensional manifold (Saul and Roweis, 2003) embedded within the D-dimensional signal space. The method that we will describe in the following utilizes concepts of dimensionality reduction in conjunction with the observations above to determine the depth-of-interaction of each event.

One degree of freedom, depth: Nonlinear dimensionality reduction algorithms provide a low dimensional representation $\mathbf{R} = \{\mathbf{r}_1, \ldots, \mathbf{r}_n, \ldots, \mathbf{r}_N\}$, $\mathbf{r}_n \in \mathcal{R}^{D'}$ of a dataset of higher dimensionality, $\mathbf{U} = \{\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_n, \ldots, \mathbf{u}_N\}$, $\mathbf{u}_n \in \mathcal{R}^D$, under the assumption that $\mathbf{U}$ lies on a manifold embedded in $\mathcal{R}^D$ with intrinsic dimensionality $D' < D$. We consider here the Locally Linear Embedding (LLE) algorithm (Saul and Roweis, 2003), which constructs a low dimensional representation with invariant self-similarity between the data points. Let us give a self-contained description of the principles underlying the LLE algorithm. The first step of LLE is to find a matrix of weights $\mathbf{W}$ that expresses each point $\mathbf{u}_n$ as a linear combination of its η nearest neighbors in $\mathcal{R}^D$. This step is performed simultaneously for every point $\mathbf{u}_n \in \mathbf{U}$ by first measuring the distance between all pairs of data points; selecting the η nearest neighbors of each point $\mathbf{u}_n$; and solving:

12

$$\hat{\mathbf{W}} = \arg\min_{\mathbf{W}} \left\{ \sum_{n=1}^{N} \left| \mathbf{u}_n - \sum_{m=1}^{N} W_{nm}\mathbf{u}_m \right|^2 \right\}, \tag{13}$$

subject to the two constraints that (1) each data point $\mathbf{u}_n$ is reconstructed only from its neighbors, enforcing $W_{nm} = 0$ if $\mathbf{u}_m$ does not belong to the set of η neighbors of $\mathbf{u}_n$; and (2) that the rows of the weight matrix sum to one for every n: $\sum_m W_{nm} = 1$. In the implementation of LLE that we adopt here, the metric of distance for the selection of the η neighbors is the Euclidean norm. Given such choice, the algorithm has only one free parameter, the number of neighbors η. As described in detail in Saul and Roweis (2003), the constrained minimization of (13) has a closed-form solution. Once the reconstruction weights $\hat{\mathbf{W}}$ have been computed, the LLE algorithm maps each high-dimensional observation $\mathbf{u}_n$ to a low-dimensional vector $\mathbf{r}_n$, which represents the global internal coordinate of the point on the manifold. This is done by finding the set $\hat{\mathbf{R}}$ of D'-dimensional vectors such that, according to the weight matrix $\hat{\mathbf{W}}$, $\mathbf{r}_n$ is near $\mathbf{r}_m$ if $\mathbf{u}_n$ is near $\mathbf{u}_m$:

$$\hat{\mathbf{R}} = \arg\min_{\mathbf{R}} \left\{ \sum_{n=1}^{N} \left| \mathbf{r}_n - \sum_{m=1}^{N} \hat{W}_{nm}\mathbf{r}_m \right|^2 \right\}. \tag{14}$$

Defining a quadratic form in the vectors $\mathbf{r}_n$, the embedding cost (14) is minimised globally by solving a sparse $N \times N$ eigenvalue problem; details are provided in Saul and Roweis (2003).

For illustration, the scatter-plot of Fig. 4 reports the coordinates of 5000 simulated interaction events, projected onto the manifold of dimensionality $D' = 3$, using $\eta = 12$ (the set-up of the simulation will be described in detail in the experiments in section 4). Note that the choice $D' = 3$ here is purely for the purpose of visualization – a value of $D' = 1$ will be used in actual computations, as detailed below. The color of each point represents the depth-of-interaction of the gamma photon, which is known in the simulation experiment. In Fig. 4 one observes that the structure of the data is intrinsically 1-dimensional. Following this observation, in order to capture the 1-dimensional degree-of-freedom of the data, which appears to be associated with the variable depth-of-interaction, we utilize the LLE algorithm with $D' = 1$. The scalar measure r_n therefore denotes the location of data point n on the 1-dimensional manifold. Although the manifold coordinates r_n are related to the depth-of-interaction, the exact mapping between the manifold coordinates and the depth remains at this point unknown. In order to create a mapping between the low-dimensional representation of the data points and depth, in the following we utilize the notion that the photon-flux is an exponential function of depth.

Depth from manifold density:

Page 13 of 32

AUTHOR SUBMITTED MANUSCRIPT - PMB-104764.R2

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Figure 4. Simulated monolithic gamma camera (simulation geometry in Fig. 6-II): scatter-plot of the data points, generated by a collimated beam, projected onto the manifold of dimension $D' = 3$. Data points are colored according to the ground-truth depth-of-interaction label k. From this visualization of the data projected onto the three-dimensional manifold it can be observed that the data live in a 1-dimensional manifold. The density of data points decreases exponentially with increasing depth-of-interaction. Fig. 5 represents the projection for $D = 1$.

Let us indicate with θ_k the cumulative mass function that expresses the probability that the gamma photon has interacted at a depth smaller or equal to k:

$$\theta_k = \sum_{k'=1}^{k} p(k'). \tag{15}$$

The values of θ_k can be computed precisely using (10) and (15), given the narrow-beam attenuation coefficient $\alpha_{\bar{E}}$ of the scintillating material. A mapping between the coordinates of the manifold and the depth-of-interaction can now be found by considering that the number of data points expected to interact at depth k is $p(k)N$. The mapping between the 1-dimensional manifold and the discretized depth values can be expressed by means of $K + 1$ scalar values $t_h, h = 0, \ldots, K$; $t_h \in \mathbb{R}$, such that $t_{h+1} > t_h$ and such that a measurement with low dimensional representation r_n is assigned to depth k if $t_{k-1} < r_n \leq t_k$, as depicted in Fig. 5.

The algorithm that we propose to compute each $\{t_k\}$ in turn, selects the scalar values t_k by searching through the set of low dimensional representations of the experimental data points. The algorithm proceeds as follows: the first scalar is set
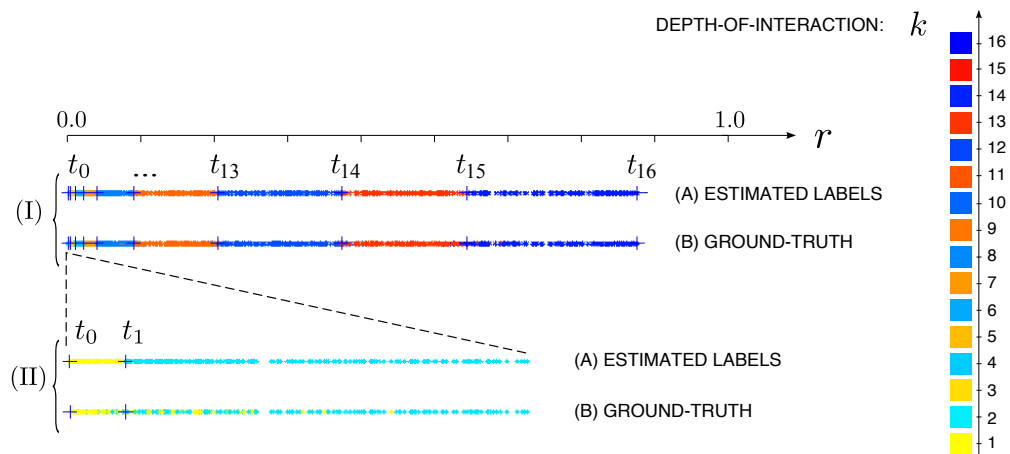
14

Figure 5. Simulated monolithic gamma camera (simulation geometry in Fig. 6-II): scatter-plot of the data points, generated by a collimated beam, projected onto the manifold of dimension $D = 1$. The location of each data-point n on the manifold is characterized by scalar r_n. (I): All N training data points. (II): Close up for r near 0, corresponding to depth-of-interaction $k = 1$ and $k = 2$ voxels. $t_h = \{t_0, t_1, \ldots, t_K\}$ partition the manifold in $K = 16$ regions. In plot (A) the data points are colored according to the value of r_n (estimated depth-of-interaction). In plot (B) data points are colored according to the true depth-of-interaction.

to the minimum value in the set: $t_0 = \min(r_1, r_2, \ldots)$; the remaining scalar values are chosen, for each value of k, as the minimum value in the search set for which the number of data points such that $r_n \leq t_k$ is equal or larger to $\theta_k N$. It has to be noted that the LLE algorithm does not guarantee that increasing values of r correspond to increasing values of z or vice versa. Therefore before searching for the values $\{t_h\}$, the sign of $\{r\}$ is inverted if the average distance between the η data points with the highest r values is smaller than the average distance between the η data points with the lowest r values. This guarantees that the region of the manifold with low values of r corresponds to the entrance of the camera, where the density of data points is higher.

The resulting map $\{t_h\}$ assigns each measurement n to a DOI k according to the value r_n of the low dimensional representation of the measurement vector. The forward model $l_{d\bar{i}\bar{j}k}$ is then calculated as the average of the measurement vectors assigned to each value of k, normalized by the energy of the gamma source. In the remainder of the paper, we will name this second algorithm for the estimation of the depth-of-interaction of a set of photons *Depth Embedding*.

4. Experiments and Results

We evaluated the algorithms presented in the previous section with two experiments: (A) a simulation study of a monolithic PET gamma camera, replicating the Continuous Miniature Crystal Element (cMiCE) PET detector (Miyaoka et al., 2011) developed at the Imaging Research Laboratory at the University of Washington; and (B) a study

15

based on real data acquired with the cMiCE PET detector. The simulation study highlights the importance of accurate 3D characterization; investigates the accuracy of the estimates of the forward model obtained with the Depth Embedding and MLE-EM algorithms; and investigates the effect of the parameters η (the number of neighbors in the dimensionality reduction algorithm) and N (the number of training gamma photons per beam location). The real-data study utilizes the Depth Embedding and MLE-EM algorithms to extract a 3-D characterization of the cMiCE camera, quantifying the improvement of the spatial resolution obtained in comparison to 2-D estimation and to the standard centroid method (Anger, 1958).

The algorithms were implemented in the Python programming language, using the version of the LLE algorithm included in the scikit-learn (Pedregosa et al., 2011) machine learning module (version 0.16). The average execution time for the computation of the elements of the forward model for a single beam position, for the real data experiment, was 5.3 sec (3.1 sec for Depth Embedding and 2.2 sec for MLE-EM), using an Apple MacBook Pro with 2.7 GHz Intel Core i7 CPU and 16 GB RAM.

4.1. Simulation study

We utilized a numerical simulation to generate a ground-truth for the evaluation of the recovery of the parameters of the forward model **L**. This also allowed us to investigate the optimal selection of the parameters of the Depth Embedding algorithm and to analyze the effect of inaccuracies of the forward model on the estimation of the coordinates of interaction, highlighting the importance of estimating the 3-D model in a precise fashion.

Simulated forward model: We simulated a monolithic scintillation gamma camera using the SimSET photon-tracking software (Harrison and Lewellen, 2012). The simulation parameters were set to replicate the characteristics of the cMiCE PET detector, which consists of a $50 \times 50 \times 15$ mm^3 LYSO crystal ($\alpha_{511KeV} = 0.83 \; [cm^{-1}]$, $\xi = 25 \; \frac{photons}{KeV}$) coupled to a 8×8 (64-channel) array of photo-detectors. Photons were generated from a single point source $40mm$ from the entrance face of the LYSO crystal. SimSET was utilized to sample the trajectories of the gamma photons. Collimation was simulated by emitting photons in a right circular cone so that the circular base of the cone is a $1mm$ diameter circle on the entrance face of the crystal. This results in the cone having a diameter of $\sim 1.4mm$ at the photosensor plane and incident photons that are parallel to within 0.1°. This is similar to our real-data experimental setup of the cMiCE PET camera (section 4.2). Realistic photon scatter was simulated within the crystal, and the position and energy of each gamma photon interaction by photoelectric absorption and scatter was recorded. Transport of the secondary photons was then modeled with two components in order to simulate a perfectly reflective crystal entrance face and perfectly absorbent side faces: the solid angle from the point of the gamma photon interaction to each element of the photosensor and the reflection on the entrance face of the crystal, which was obtained by computing a second solid angle from a virtual point

16

source that is the reflection of the original point source through the entrance face. The photodetectors were considered noiseless and were each assigned a quantum efficiency of 0.35 , to replicate the average characteristics of the detector modules of the CMiCE camera, resulting in an average of ≈ 5000 photons per event of interaction; noise in the detection of the secondary photons was obtained using a Poisson noise generator. The simulator was utilized to generate both the ground truth forward model and the simulated events of interaction. The ground truth forward model $\mathbf{L}^{[true]}$ was generated by averaging the noiseless signals produced by the non-scattered interactions.

Characterization: The set-up described in section 3 for the characterization of the gamma camera was replicated in the simulation: the pencil-beam, perpendicular to the (x, y) plane, was positioned on a grid of $(I = 48 \times J = 48)$ points with stride $1\ mm \times 1\ mm$. For each position, $N = 40000$ events were simulated and an energy window was applied to reject scatter, resulting in approximately 10000 qualified events per beam position. The energy window was set adaptively with a lower threshold of 5/6 of the photo-peak, halfway to the Compton edge and an upper threshold of 1.25 times the full width at tenth max of the photo-peak (the value of the thresholds depends on the energy resolution of the detector at each beam location; the average values of the two thresholds were $426KeV$ and $550KeV$). Once the measurements $\mathbf{U}$ were simulated this way, we aimed to estimate the characteristics of the gamma camera $\mathbf{L}$ with the three following methods:

(i) *2-D characterization*;

(ii) *3-D characterization with Depth Embedding*;

(iii) *3-D characterization with MLE-EM*.

The 2-D characterization (method (i)), which we will indicate in the remainder using $\mathbf{L}^{[2D]} = \left\{ l_{dij}^{[2D]} \right\}$, was obtained simply by averaging all the signals acquired for a given position of the beam, disregarding the depth-of-interaction: $l_{d\bar{i}\bar{j}}^{[2D]} = \frac{1}{\overline{E}N} \sum_{n=1}^{N} u_{dn}$, for $\mathbf{U} = \{u_{dn}\}$ acquired at position $\overline{ij}$. It was found that the MLE-EM algorithm (method (iii)), which is a greedy optimizer that can get trapped in local optima, fails to converge to a meaningful solution when initialized randomly. In the following we therefore present the results of the MLE-EM algorithm initialized with the model produced by the Depth Embedding algorithm (method (ii)), along with the results obtained using the Depth Embedding algorithm alone.

In order to determine the optimal number of neighbors for the Depth Embedding algorithm (method (ii)), the algorithm was first utilized to reconstruct $N = 5000$ events per beam position and executed multiple times, with η ranging between 5 and 20; the number of depth bins was set to $K = 16$, determining a voxels thickness of $0.94mm$. For each value of η, the estimated model $\hat{\mathbf{L}} = \{\hat{l}_{dijk}\}$ was compared with the ground-truth model $\mathbf{L}^{[true]} = \{l_{dijk}^{[true]}\}$ by measuring the mean percentage error:
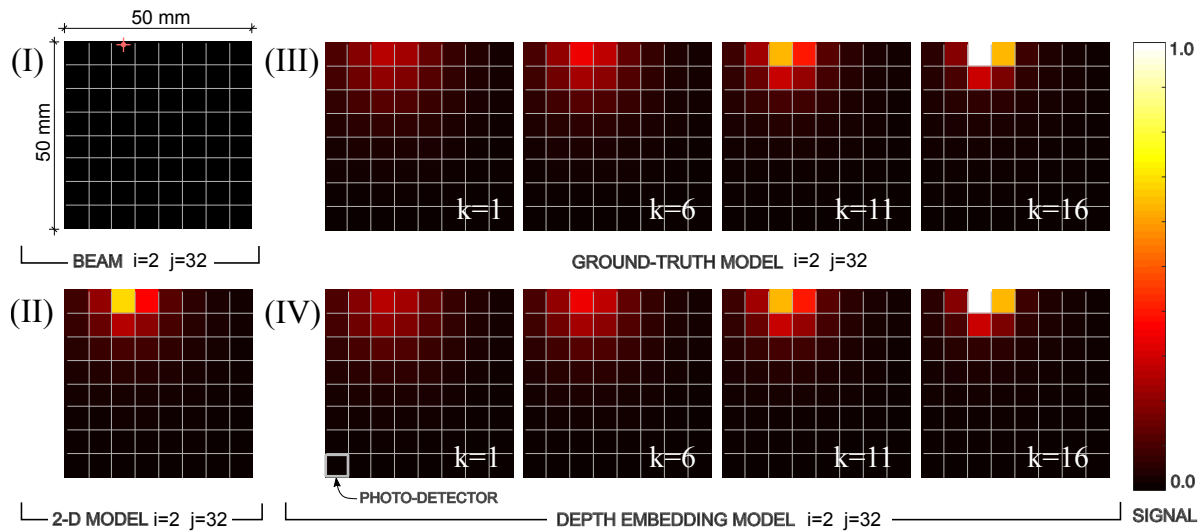
Figure 6. Simulation of a monolithic gamma camera consisting of a $50 \times 50 \times 15$ mm scintillator crystal coupled with an array of $D = 64$ photo-detectors of size 6.25×6.25 mm each. (I) Representative position of the beam $i = 2, j = 32$. (II) 2-D forward model for $i = 2, j = 32$, estimated by simple averaging; $N = 5000$. (III) Ground truth 3-D forward model for beam location $i = 2, j = 32$ and depth $k = 1, 6, 11, 16$. (IV) 3-D forward model for estimated with Depth Embedding; $\eta = 12$, $N = 5000$; $i = 2, j = 32$, $k = 1, 6, 11, 16$.

$$MPE = \frac{100\%}{DIJK} \sum_{dijk}^{DIJK} \left| \frac{l_{dijk}^{[true]} - \hat{l}_{dijk} + \delta}{l_{dijk}^{[true]} + \delta} \right|, \tag{16}$$

with $\delta = 10^{-9}$ to avoid division by zero. The MPE of the estimated parameters, reported in the plot in Fig. 7.I, was found to have a minimum of 6.18% for $\eta = 11$ and to present a plateau in the range of η comprised between 7 and 17; in subsequent experiments, the number of neighbors was set to the value at the center of the plateau: $\eta = 12$. Fig. 4, to which we have previously referred in section 3-B to describe the rationale of the Depth Embedding algorithm, shows the scatter-plot of the data points projected, onto the three-dimensional manifold, depicting the variable density of data points along the manifold. The projection of the data into the three-dimensional manifold is for visualization purpose only. Fig. 5 reports the scatter-plot of the data points projected on the manifold of dimension $D' = 1$, colored according to the true depth-of-interaction and to the depth-of-interaction estimated using the Depth Embedding algorithm. The images in Fig. 6.III and 6.IV display, respectively, the ground-truth and the parameters of the forward model estimated with the Depth Embedding algorithm (method (ii)) for a representative position of the beam $i = 2, j = 32$, depicted in Fig. 6.I, and for four representative depths $k = 1, 6, 11, 16$, visually confirming that the algorithm correctly captures the depth-dependent response. For visual comparison, Fig. 6.II displays the 2-D forward model obtained with method (i), i.e., averaging the measurement vectors.

In order to evaluate the effect of the MLE-EM refinement (method (iii)) and to

18

Figure 7. Accuracy of the forward model estimates in the simulation of a monolithic gamma camera. (I) Mean percentage error of the forward model estimated with Depth Embedding, with multiple settings of the parameter η; $N = 5000$. (II) Mean percentage error, with variable number of data points N, of the forward model estimated using Depth Embedding and of the forward model refined with MLE-EM.

determine the sensitivity of the parameter N, the characterization experiment was then repeated for 15 values of N, equally spaced between $N = 200$ and $N = 10000$; for each value of N, the forward model was first estimated using Depth Embedding (method (ii)) and then refined using MLE-EM. MLE-EM was executed until convergence (in average 15 iterations). The MPE for the models produced by Depth Embedding alone and for the models refined using MLE-EM is reported in Fig. 7.II. In both cases, the MPE decreases with increasing N, stabilizing for $N \gtrsim 1500$. Values of N larger than 1500 are therefore advisable to minimize the error of the model estimates. The MLE-EM algorithm was found to consistently reduce the MPE, reducing the MPE to 2.82% (from the initial value of 6.18% produced by Depth Embedding), in the case of $N = 5000$.

The forward models estimated with the methods (ii) and (iii) were up-sampled on a grid $I = 200 \times J = 200 \times K = 60$ and the forward model estimated with method (i) was up-sampled on a grid $I = 200 \times J = 200$ using tri-linear interpolation.

Event estimation: Once the three forward models were estimated with the methods (i), (ii) and (iii) described in the previous section and up-sampled, a second simulation was set up to evaluate and compare the effect of the forward models on the estimation of the coordinates of interaction. The detector was flooded using 25 pencil-beams, displaced on a regular grid with stride 12.0 mm × 12.0 mm. The interaction of $N = 5000$ photons per beam location was simulated. The coordinates of interaction of each simulated photon were then reconstructed using five different methods:

(A) The centroid method described in Anger (1958);

(B) MAP with the 2-D forward model;

(C) MAP with the Depth Embedding 3-D forward model;

(D) MAP with the MLE-EM 3-D forward model;

(E) MAP with the 3-D ground-truth forward model.

The 2-D event estimation (B) was performed using a 2-D version of Eq. (7):

$$\hat{i}, \hat{j} = \arg \min_{i,j} \left\{ \sum_{d=1}^{D} u_d - \sum_{d=1}^{D} \left(u_d \ln l_{dij}^{[2D]} \right) \right\}. \tag{17}$$

Figure 8. Estimation of simulated events of interaction for the set-up of Fig. 6-II. The images report the 2-D histogram of the estimations of 37500 events of interaction uniformly distributed across 25 discrete locations (for the 3-D estimation methods, the histograms are obtained by disregarding the depth information). (A) Estimation with the centroid algorithm (Anger, 1958); (B) MAP estimation with the 2-D forward model obtained by experimental characterization; (C) MAP estimation with the 3-D forward model obtained with the Depth Embedding algorithm; (D) MAP estimation with the 3-D forward model obtained with the Depth Embedding algorithm and refined with MLE-EM; (E) MAP estimation using the ground-truth 3-D forward model.

The 3-D MAP event estimation (C)(D)(E) were performed using Eq. (7) with uniform prior probability of i, j and probability of k given by Eq. (10), parameterized with the parameter α_{511Kev} of the simulated scintillator crystal (Fig. 2.B):

$$p(i, j, k) = \frac{1}{IJ} \frac{e^{-\alpha_{511Kev}(k-1)\Delta z} - e^{-\alpha_{511Kev}k\Delta z}}{1 - e^{-\alpha_{511Kev}K\Delta z}}, \tag{18}$$

with $K = 60$ and $\Delta z = 0.25$ mm. For all the MAP methods (B)(C)(D)(E), simple exhaustive search was used to solve (8), (7) and (17).

The results of the estimations are reported in Fig. 8, where the intensity of the images is proportional to the number of events associated to each of the 200×200 pixels on the (x, y) plane of the camera (in the case of 3-D estimation, the depth was simply discarded when creating these 2-D histograms). Fig. 8.E reports the best achievable result, obtained by reconstructing the events using MAP estimation with the ground-truth forward model (method (E)); here the events are, on average, correctly positioned near the beam locations. Note that even in this idealized scenario, which provides an upper bound on the achievable performance in practice, the errors are distributed in a slightly anisotropic way near the edges of the device: this is because the detector geometry intrinsically provides less information about the location of interaction in directions perpendicular to the edge orientation there. The 2-D estimation using the centroid method (Fig. 8.A) was found to introduce large systematic errors in the peripheral regions of the camera, a problem known as barrel effect. MAP estimation using the estimated 2-D forward model (Fig. 8.B) was also found to mis-locate the events of interaction in the (x, y) plane, introducing a systematic error, visible at all beam locations, due to the lack of depth information. The estimations obtained with the two estimated 3-D forward models (Fig. 8.C-8.D)) are visually undistinguishable from the results obtained using the ground-truth forward model, appearing to eliminate

20

Table 1. Estimation of simulated events of interaction for the set-up of Fig. 6-II: mean bias (Eq.(19)) and standard deviation (Eq.(20)) of the reconstructed (x, y) beam locations and root mean squared error of the reconstructed DOI (Eq.(21)) for the five methods described in the simulation study section.

	(A) CENTROID	(B) 2-D MAP	(C) DEPTH-EMB	(D) MLE-EM	(E) TRUE
$\overline{BIAS}_{xy}$ [mm]	3.16	0.88	0.36	0.32	0.33
$\overline{STD}_{xy}$ [mm]	1.80	1.34	1.13	1.10	1.05
$\overline{Error}_z$ [mm]			1.16	1.14	1.13
max($BIAS_{xy}$) [mm]	6.69	2.62	1.36	1.27	1.24
max(STD_{xy}) [mm]	2.95	2.05	1.66	1.62	1.57
max($Error_z$) [mm]			1.57	1.52	1.48

Figure 9. Depth-dependence of STD_{xy} and of $Error_z$ in the ideal case of MAP estimation with known forward model (E) for a representative beam position near the edge of the detector.

the bias observed in the 2-D estimations.

In order to assess quantitatively the accuracy of the estimations, we measured the bias (systematic error) and standard deviation (random error) of the reconstructed (x, y) beam locations and the root mean squared error of the depth-of-interaction estimates:

$$BIAS_{xy}^{[p]} = \left[\left(x^{[T,p]} - \bar{x}^{[R,p]} \right)^2 + \left(y^{[T,p]} - \bar{y}^{[R,p]} \right)^2 \right]^{\frac{1}{2}}, \quad (19)$$

$$STD_{xy}^{[p]} = \left[\frac{1}{N} \sum_{n=1}^{N} \left(x_n^{[R,p]} - \bar{x}^{[R,p]} \right)^2 + \left(y_n^{[R,p]} - \bar{y}^{[R,p]} \right)^2 \right]^{\frac{1}{2}}, \quad (20)$$

$$Error_z^{[p]} = \left[\frac{1}{N} \sum_{n=1}^{N} \left(z_n^{[T,p]} - z_n^{[R,p]} \right)^2 \right]^{\frac{1}{2}}, \quad (21)$$

with $x^{[T,p]}$ and $y^{[T,p]}$ indicating the true locations of the beam in position p; $z_n^{[T,p]}$ indicating the true depth of the n-th interaction; $x_n^{[R,p]}$, $y_n^{[R,p]}$ and $z_n^{[R,p]}$ indicating the coordinates of interaction reconstructed with each method in turn; and with

21

$\bar{x}^{[R,p]} = \frac{1}{N}\sum_{n=1}^{N} x_n^{[R,p]}$, $\bar{y}^{[R,p]} = \frac{1}{N}\sum_{n=1}^{N} y_n^{[R,p]}$. The mean bias and deviation of the beam and the mean error of the depth estimates across the entire camera surface were then computed by averaging over the 25 beam locations: $\overline{BIAS}_{xy} = \frac{1}{25}\sum_{p=1}^{25} BIAS_{xy}^{[p]}$; $\overline{STD}_{xy} = \frac{1}{25}\sum_{p=1}^{25} STD_{xy}^{[p]}$; $\overline{Error}_z = \frac{1}{25}\sum_{p=1}^{25} Error_z^{[p]}$ to characterize globally the systematic and random error of the beam coordinates and the error of the measurement of depth. The maximum values were also recorded in order to observe the characteristics of the camera in the regions with the lowest performance. The numerical values are reported in Table 1. Fig. 9 provides an insight of the depth-dependence of STD_{xy} and $Error_z$ by plotting their values as a function of the depth-of-interaction for the case of ideal estimation with known forward model. One can notice that, due to the intrinsic characteristics of the camera, the (x, y) error is larger near the surface, which the error in estimating z is larger near the detectors. The 3-D estimation using the forward model estimated with the Depth Embedding algorithm produced a considerable reduction of the systematic errors of the beam positions and of the random errors, when compared to the standard centroid method and to 2-D MAP estimation. The bias of the beam positions, in particular, is virtually eliminated ($\overline{BIAS}_{xy} = 0.36$ mm), producing therefore a more uniform effective resolution across the (x, y) plane. The MLE-EM refinement improved the depth resolution, reducing the depth error (from $\overline{Error}_z = 1.16$ mm to $\overline{Error}_z = 1.14$ mm) and producing a further slight reduction of the bias ($\overline{BIAS}_{xy} = 0.32$ mm) and deviation of the beam (from $\overline{STD}_{XY} = 1.13$ mm to $\overline{STD}_{XY} = 1.10$ mm). The resulting bias, beam deviation and depth errors are virtually identical to the optimal values obtained using the ground-truth 3-D forward model: $\overline{BIAS}_{xy} = 0.33$ mm; $\overline{STD}_{XY} = 1.05$ mm, ; $\overline{Error}_z = 1.13$ mm.

4.2. Characterization of the cMiCE PET camera

The Depth Embedding and MLE-EM algorithms were applied to data collected from the Continuous Miniature Crystal Element (cMiCE) PET detector (Miyaoka et al., 2011), which consists of a $50 \times 50 \times 15$ mm³ LYSO crystal ($\alpha_{511KeV} = 0.83$ $[cm^{-1}]$) coupled to a 8×8 (64-channel) position-sensitive photomultiplier tube (H8500, Hamamatsu Photonics K.K., Japan).

Characterization: In order to speed-up the acquisition, the single beam described in the previous sections was replaced with a 4×4 array of collimated beams formed using 16 ^{22}Na point sources with activity between 15 and 20 μCi (Fig. 11.I). Each beam had a FWHM of 0.52 mm at the crystal entrance face. The beam array was utilized to scan the surface of the detector on a 12×12 square grid with stride 1 $mm \times 1$ mm, producing in total $I = 48 \times J = 48$ beam positions over the entire surface of the camera. For each detected photon, the source beam was then recovered using an initial estimate of the coordinates of interaction based on the centroid algorithm (Anger, 1958). In order to discard scattered photons, the events were filtered using an energy window and a likelihood-based filter. The energy windowing procedure was identical to the one used in the simulation experiments (see section 4.1); after applying it, approximately 2000

22

Figure 10. cMiCE PET camera: scatter-plot of the data points projected onto the manifold of dimension $D' = 3$ for a representative beam position.

interactions per beam position were registered. The second filter discarded events with abnormal light distribution by first computing the 2-D forward model by averaging all the events at a given location of interaction, and then selecting $N \approx 1500$ events with the highest likelihood associated to the 2-D forward model (expressed by the formula in the curly brackets in Eq. (17)). This second filter is a heuristic aimed at further removing scattered photons when constructing the forward model, motivated by the idea that non-scattered events at all depths are expected to be closer than scattered events to the mean of the non-scattered events - this is obtained by averaging all the measurements at one location, under the assumption that non-scattered events dominate when computing the mean. The decision to further remove $\frac{1}{4}$ of the events at all beam locations was motivated by the aim of removing as many scattered photons as possible, while guaranteeing a sufficient number of events. The acquisition experiment was performed twice, in order to generate two separate data sets: one for the characterization of the camera and one to test the estimation of the coordinates of interaction using the estimated forward model.

The forward model **L** was estimated by applying the Depth Embedding algorithm and the MLE-EM algorithm to the first data set. The number of neighbors for Depth Embedding was set to the value found in the simulation experiment ($\eta = 12$) and the number of points of the 3-D grid for the model estimation was set to $I = 48, J = 48, K = 16$, with voxel size $1.00 \times 1.00 \times 0.94$ mm. The MLE-EM algorithm was initialized with
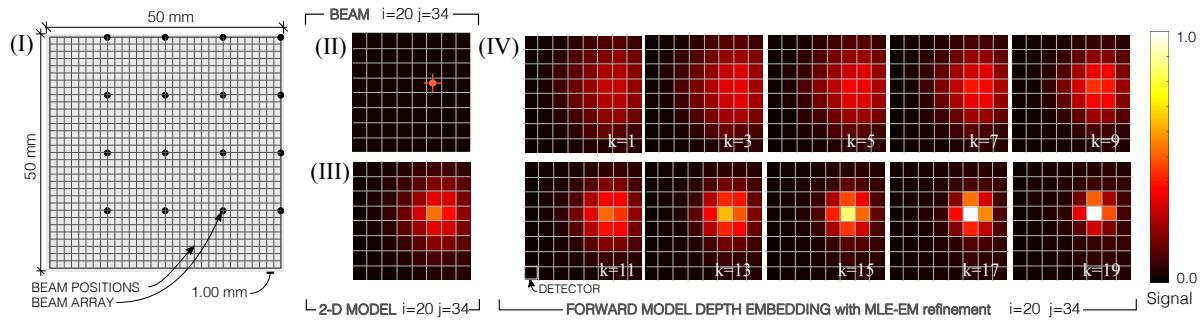
Figure 11. Characterization of the cMiCE PET camera. (I) Measurement set-up: in order to speed up the acquisition, an array of 4×4 beams was used to scan exhaustively the camera surface. The beam position was then recovered using an initial estimation based on the centroid algorithm. (II) A representative beam position $i = 20, j = 34$. (III) 2-D forward model computed by averaging the measurement vectors, displayed for $i = 20, j = 34$; $N = 1822$. (IV) 3-D forward model obtained with the Depth Embedding algorithm with MLE-EM refinement; $i = 20, j = 34$; $\eta = 12, N = 1412$.

the forward model produced by the Depth Embedding algorithm and iterated until convergence (in average 18 iterations).

The low dimensional visualization ($D' = 3$) of the data acquired for a representative beam location (i=20, j=34) is reported in Fig. 10, with color labels representing the estimated depth-of-interaction. At visual inspection, the 3-D forward model recovered using the Depth Embedding algorithm, displayed in Fig. 11.IV for the representative beam location displayed in Fig. 11.II, presents a more peaked response for large values of k; indicating that the depth information is correctly captured. The forward model obtained with the MLE-EM refinement did not present visible differences from the model obtained with Depth Embedding alone. The 2-D forward model obtained by simple averaging is reported in Fig. 11.III for visual comparison.

Event estimation: In the case of real data, the true characteristics of the camera **L** are unknown; the accuracy of the estimates of the 3-D forward model of the camera can therefore not be quantified directly. However, the photon beam measurements provide an indirect form of ground-truth: the (x, y) locations of interaction are known. As we have observed in the previous experiment, the lack of 3-D information determines a mis-location of the coordinates of interaction in the (x, y) plane; utilizing the photon-beam measurements, we therefore quantified the systematic and random errors of the measured beam positions in order to assess indirectly the accuracy of the forward model. The test data set was acquired with the same set-up utilized for the characterization of the camera, described in the previous section. For each event of interaction in the test data set, the coordinates of interaction were reconstructed using the four methods described in the simulation experiment section:

(A) The centroid method described in (Anger, 1958);

(B) MAP with the 2-D forward model;

24

(C) MAP with the Depth Embedding 3-D forward model;

(D) MAP with the MLE-EM 3-D forward model.

Prior to the estimation, the forward models were up-sampled to a grid of $200 \times 200 \times 60$ by tri-linear interpolation. As in the simulation experiments, the MAP estimations were performed by exhaustive search of the coordinates of interaction.

For visual assessment, a 2-D histogram of the reconstructed events was formed by selecting events emitted from a subset of the beam locations displaced on a 5×5 grid. Fig. 12.I reports the 2-D histogram for each estimation method and displays, for reference, the 5×5 beam locations (Fig. 12-left). Both the centroid method (Fig. 12.A) and 2-D MAP estimation (Fig. 12.B) were found to introduce visible errors in the estimated (x, y) coordinates of interaction; the errors appear reduced in both the cases of 3-D estimation using Depth Embedding alone (Fig. 12.C) and in the case of Depth Embedding with the MLE-EM refinement (Fig. 12.D).

For each estimation method, the accuracy of the reconstructed coordinates of interaction was quantified by the bias (systematic error) and standard deviation (random error) of the (x, y) coordinates of the beams. The bias and standard deviation were computed according to Eq. (19) and Eq. (20) for each beam position ($\bar{i} = 1, 2, \ldots 48$, $\bar{j} = 1, 2, \ldots 48$). Table 2 reports their means, representing a global measure of bias and of deviation of the beams, and their maximum values. The 3-D estimation methods were found to reduce considerably the systematic errors, producing a small bias across the surface of the gamma camera: $\overline{BIAS}_{xy} = 0.64$ mm in the case of Depth Embedding and $\overline{BIAS}_{xy} = 0.62$ mm in the case of Depth Embedding and MLE-EM refinement, compared to $\overline{BIAS}_{xy} = 3.16$ mm in the case of the centroid algorithm and $\overline{BIAS}_{xy} = 0.72$ in the case of 2-D estimation. The standard deviation of the (x, y) coordinates decreased from $\overline{STD}_{xy} = 2.85$ mm for the centroid algorithm and $\overline{STD}_{xy} = 2.90$ mm for the 2-D estimation to $\overline{STD}_{xy} = 2.54$ mm for the Depth Embedding algorithm alone and $\overline{STD}_{xy} = 2.51$ mm for Depth Embedding with the MLE-EM refinement.

In order to evaluate the recovery of the DOI, histograms of the reconstructed DOI values were computed form the reconstructed coordinates of interaction, utilizing the events from all beam locations. The visual comparison with the true distribution of the DOI, reported in Fig. 13, indicates the correctness of the forward models and the accurate recovery of depth.

For further evaluation of the recovery of the DOI, the camera was illuminated with a beam incident at 45 degrees (45 degrees rotation of the calibration beam around the x axis). $N = 12412$ events were recorded after applying the energy threshold. Fig. 14 reports the (y, z) histogram of the reconstructed coordinates of interaction and the full-width at half-maximum (FWHM) of the reconstructed beam. The FWHM, calculated at three levels of depth (near the entrance surface, at the center of the crystal and near the detector plane) with the methodology described in the caption of Fig. 14, did not present large variations as a function of depth. The low value of the FWHM (1.76 mm

Figure 12. Estimation of events of interaction measured with the cMiCE PET camera. (I) Left: Positions of the beams. (B) Estimation with the centroid algorithm. (C) MAP estimation with 2-D forward model obtained by experimental characterization of the camera. (D) MAP estimation with the 3-D forward model obtained with Depth Embedding. (E) MAP estimation with the 3-D forward model obtained with Depth Embedding and refined with the MLE-EM algorithm. (II) Mean bias and mean standard deviation of the reconstructed (x, y) beam positions.

Table 2. Estimation of events of interaction measured with the cMiCE PET camera: mean bias and mean standard deviation of the reconstructed (x, y) beam positions.

	(A) CENTROID	(B) 2-D MAP	(C) DEPTH-EMB	(D) MLE-EM
$\overline{\text{BIAS}}_{xy}$ [mm]	3.16	0.72	0.64	0.62
$\overline{\text{STD}}_{xy}$ [mm]	2.85	2.90	2.54	2.51
$\max(\text{BIAS}_{xy})$ [mm]	6.32	1.72	1.70	1.69
$\max(\text{STD}_{xy})$ [mm]	3.92	3.62	3.12	3.09
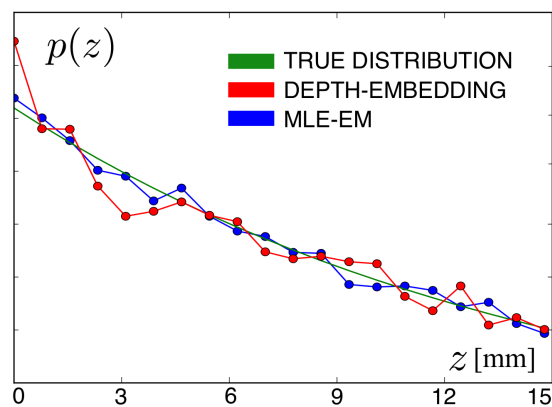
Figure 13. cMiCE PET camera, perpendicular beams: Beer-Lambert prior depth probability (green); histogram of the reconstructed DOI using Depth Embedding (red) and Depth Embedding with MLE-EM refinement (blue).

26

Figure 14. cMiCE PET camera: estimation of a pencil beam incident at 45 degrees. The images in the top row display the integral sum of the 3-D histogram along the x-axis for the estimations obtained using Depth Embedding (left) and Depth Embedding with MLE-EM refinement (right). The two images report the full-width-at-half-maximum (FWHM) of the reconstructed beam, computed by rotating the 3-D histogram by 45 degrees by bi-linear interpolation, by averaging the rotated histogram over three 5-mm slabs, depicted by the dotted lines, and fitting three Gaussian functions to the resulting 1-dimensional histograms (reported in the bottom row). The FWHM of the 45 degrees beam represents an indirect approximate measurement of the DOI resolution.

for Depth Embedding and 1.74 mm for Depth Embedding with MLE-EM refinement at the center of the crystal) and the absence of distortions demonstrate the accurate estimation of the depth-of-interaction. A previous publication reported a DOI resolution of the cMiCE detector with the 15 mm crystal of 4.8 mm (Miyaoka et al., 2009).

5. Discussion and conclusion

In this paper we have addressed the problem of estimating the depth-dependent forward model of gamma cameras from simple pencil-beam measurements. We have introduced two complementary methods based on machine learning. The first method, which we named MLE-EM, is based on finite mixture modeling. It accounts for the Poisson noise characteristics of the detection process, but yields an iterative algorithm that is sensitive to initialization. The second method, which we named Depth Embedding, is based on nonlinear dimensionality reduction. It yields a robust algorithm suitable for fully automated use, but does not take into account the Poisson noise characteristics. In order to combine the best of both worlds, we proposed to run the two algorithms in tandem, initializing the MLE-EM algorithm with the solution obtained with the Depth Embedding method.

The proposed algorithms were tested with simulated data of a compact monolithic gamma camera, and with data produced by the Continuous Miniature Crystal Element (cMiCE) PET detector (Miyaoka et al., 2011) developed at the Imaging Research Laboratory, University of Washington. It was found that the use of the Depth Embedding method alone to estimate the 3-D characteristics of a camera enables the

27

accurate estimation of the depth-of-interaction and improves the in-plane estimation accuracy when compared to 2-D estimation and to the standard centroid method, eliminating systematic errors, such as the barrel effect, and reducing random errors. The refinement of the 3-D model using the MLE-EM algorithm produces a further improvement of the estimation accuracy.

In contrast to previously proposed methods, which required acquisition of calibration data at multiple angles (Maas et al., 2009; Wang et al., 2011; Marone et al., 2009; Bruyndonckx et al., 2008, 2006, 2004) or the use of a simulator (Hunter et al., 2009; van Dam et al., 2011b), the method that we introduced enables the estimation of the depth-dependent forward model of gamma imaging devices from pencil-beam data acquired in a single direction without the need for a simulator or of heuristics specific to a camera geometry. The method is purely data-driven, not only circumventing the imprecision of simulations, but also greatly simplifying the use of the algorithm, which can be applied to a new device without any modeling effort.

By enabling simple, fast and accurate estimation of the depth-dependent forward model of gamma cameras, the method presented here makes it possible to fully exploit the advantages offered by model-based estimation. First, the estimation of the coordinates of interaction and energy are optimal, producing the minimum achievable bias and random errors (Barrett et al., 2009a). Second, the model-based approach is orders of magnitude more computationally efficient than nearest neighbor methods (Maas et al., 2009; van Dam et al., 2011b,a), requiring a single vector comparison per location of interaction, instead of thousands. Efficient implementations of the model-based MAP estimation algorithm achieve millions of events per second on commodity hardware (Pedemonte et al. (2009), `https://sourceforge.net/projects/oree/`).

The calibration procedure that we described, conversely to methods based on neural networks (Wang et al., 2011; Marone et al., 2009; Bruyndonckx et al., 2008, 2006, 2004) and to the original L-NN method (Maas et al., 2009), does not require coincidence information. Therefore the set-up for the calibration is simpler and the method is amenable for use, without modification, in applications other than PET, such as for the calibration of detectors for Single Photon Emission Computed Tomography (SPECT).

Although in this paper we have focused, for simplicity of exposure, on monolithic detectors alone, we believe our results translate directly to pixellated detectors and hybrid designs (Huber et al., 2001). Although in pixellated systems, the DOI is often computed in hardware, model-based estimation can be utilized to estimate the DOI in pixellated systems. For instance, the simple ratio between the measurements at opposite sides of the crystal in pixellated cameras with two layers of photo-detectors (Du et al., 2008) can be replaced with model-based estimation. We expect that the calibration procedure that we describe here will enable the accurate recovery of the depth-of-interaction in a variety of detector designs using simple calibration procedures, including the pixellated detector types described in Huber et al. (2001); Du et al. (2008, 2009); Orita et al. (2005).

The source code of the algorithms and the data of the experiments are

28

made publicly available under BSD licence at - `http://occiput.mgh.harvard.edu/depthembedding/`.

6. Acknowledgements

Appendix A. Model-based event estimation with Poisson model

Here we derive the maximum-a-posteriori (MAP) estimation algorithm for the coordinates of interaction and energy i, j, k, E of gamma photons under the assumption of Poisson noise. The maximum-a-posteriori (MAP) estimates of i, j, k, E maximize the posterior probability (6), or equivalently minimize its negative logarithm.

Although the energy E and the location of interaction i, j, k are a-priori mutually dependent due to Beer-Lambert law, in practical scenarios they can be considered approximately independent a priori; therefore we use $p(E, i, j, k) = p(E)p(i, j, k)$. We further consider the prior probability distribution associated to the energy $p(E)$ uninformative (i.e., uniform within a sufficiently wide range). Dropping the constants $\ln p(\mathbf{u})$ and $\ln p(E)$:

$$\ln p(E, i, j, k | \mathbf{u}) = E \sum_{d=1}^{D} l_{dijk} - \sum_{d=1}^{D} u_d \ln l_{dijk} - \ln(E) \sum_{d=1}^{D} u_d + \sum_{d=1}^{D} \ln(u_d!) - \ln p(i, j, k). \tag{A.1}$$

Minimizing the expression with respect to E, we obtain:

$$\frac{\partial}{\partial E} \ln p(E, i, j, k | \mathbf{u}) = -\sum_{d=1}^{D} l_{dijk} + \frac{\sum_{d=1}^{D} u_d}{E} = 0. \tag{A.2}$$

This has a closed-form solution that depends on the location of interaction $\hat{i}, \hat{j}, \hat{k}$:

$$\hat{E} = \frac{\sum_{d=1}^{D} u_d}{\sum_{d=1}^{D} l_{d\hat{i}\hat{j}\hat{k}}}. \tag{A.3}$$

The minimum of $\ln p(E, i, j, k | \mathbf{u})$ with respect to i, j, k, does not have a closed-form solution. By replacing (A.3) in (A.1), one observes that the MAP estimate of the location of interaction does not depend on the energy:

$$\hat{i}, \hat{j}, \hat{k} = \arg \min_{i,j,k} \left\{ \ln(\sum_{d=1}^{D} l_{dijk}) \sum_{d=1}^{D} u_d - \sum_{d=1}^{D} (u_d \ln l_{dijk}) - \ln p(i, j, k) \right\} ; \tag{A.4}$$

therefore the estimation consists in first finding the location of interaction by maximizing (A.4) and then computing the energy in closed-form using (A.3). A variety of algorithms can be utilized to maximize (A.4), such as simple exhaustive search or more computationally-efficient adaptive grid search methods (Pedemonte et al., 2009).

REFERENCES   30

References

Anger, H. O. (1958). Scintillation camera, *Review of Scientific Instruments* **29**(1): 27–33.

Barrett, H., Hunter, W., Miller, B., Moore, S., Chen, Y. and Furenlid, L. (2009a). Maximum-likelihood methods for processing signals from gamma-ray detectors, *IEEE Transactions on Nuclear Science* (56): 725–735.

Barrett, H., Hunter, W., Miller, B., Moore, S., Chen, Y. and Furenlid, L. (2009b). Maximum-Likelihood Methods for Processing Signals From Gamma-Ray Detectors, *IEEE Transactions on Nuclear Science* **56**(3): 725–735. Symposium on Radiation Measurements and Applications, Berkeley, CA, JUN 02-05, 2008.

Bishop, C. (2007). *Pattern Recognition and Machine Learning*, first edition edn, Springer.

Boschini, L. and Fiorini, C. (1999). Scidra: a monte carlo simulator for the study and optimization of gamma cameras based on a scintillator coupled to a silicon photodetector array, *IEEE Nuclear Science Symposium Conference Record*, Vol. 1, pp. 464–469.

Bruyndonckx, P., Lemaitre, C., van der Laan, D., Maas, M., Schaart, D., Devroede, O., Krieguer, M. and Tavernier, S. (2006). Comparison of nonlinear position estimators for continuous scintillator detectors in PET, *IEEE Nuclear Science Symposium Conference Record*, Vol. 4, pp. 2518–2522.

Bruyndonckx, P., Lemaitre, C., van der Laan, D., Maas, M., Schaart, D., Wang, Y., Li, Z., Krieguer, M. and Tavernier, S. (2008). Evaluation of machine learning algorithms for localization of photons in undivided scintillator blocks for PET detectors, *IEEE Transactions on Nuclear Science* **55**(3, 1): 918–924.

Bruyndonckx, P., Leonard, S., Tavernier, S., Lemaitre, C., Devroede, O., Wu, Y. and Krieguer, M. (2004). Neural network-based position estimators for PET detectors using monolithic LSO blocks, *IEEE Transactions on Nuclear Science* **51**(5, 2): 2520–2525.

Du, H., Yang, Y. and Cherry, S. (2008). Comparison of four depth-encoding PET detector modules with wavelength shifting (WLS) and optical fiber read-out, *Physics in Medicine and Biology* (53): 1829–1842.

Du, H., Yang, Y., Glodo, J., Wu, Y., Shah, K. and Cherry, S. (2009). Continuous depth-of-interaction encoding using phosphor-coated scintillators, *Physics in Medicine and Biology* (54): 1757–1771.

Fiorini, C., Gola, A., Peloso, R., Geraci, A., Longoni, A., Padovini, G., Lechner, P., Struder, L., Hutton, B., Erlandsson, K., Mahmood, S., Van Mullekom, P., Pedretti, A., Moretti, R., Poli, G. and Lucignani, G. (2008). HICAM: Development of a high-resolution Anger camera for nuclear medicine, *IEEE Nuclear Science Symposium Conference Record*, pp. 3961–3964.

Gray, R. and Macovski, A. (1976). Maximum a posteriori estimation of position in scintillation cameras, *IEEE Transactions on Nuclear Science* (23): 849–852.

REFERENCES 31

Harrison, R. and Lewellen, T. (2012). The SimSET Program, *in* M. Ljungberg, S. Strand and M. King (eds), *Monte Carlo Calculations in Nuclear Medicine, Second Edition*, CRC Press, Boca Raton, Florida.

Huber, J., Moses, W., Andreaco, M. and Petterson, O. (2001). An LSO Scintillator Array for a PET Detector Module with Depth of Interaction Measurement, *IEEE Transactions on Nuclear Science* **48**(3): 1757–1771.

Hunter, W., Barrett, H. and Furenlid, L. (2009). Calibration method for ml estimation of 3d interaction position in a thick gamma-ray detector, *IEEE Transactions on Nuclear Science* **56**(1): 189–196.

Ito, M., Hong, S. and Jae Sung, L. (2011). Positron Emission Tomography (PET) Detectors with Depth-ofInteraction (DOI) Capability, *Biomedical Engineering Letters* **1**: 70–81.

Lerche, C., Benlloch, J. M., Snchez, F., Pavn, N., Escat, B., Gimenez, E., Fernndez, M., Torres, I., Gimnez, M., Sebasti, A. and Martnez, J. (2005). Depth of Gamma-Ray Interaction Within Continuous Crystals From the Width of Its Scintillation Light-Distribution, *IEEE Transactions on Nuclear Science* **52**(3): 560–572.

Li, Z., Wedrowski, M., Bruyndonckx, P. and Vandersteen, G. (2010). Nonlinear least-squares modeling of 3D interaction position in a monolithic scintillator block, *Physics in Medicine and Biology* **55**(21): 6515–6532.

Ling, T., Lewellen, T. and Miyaoka, R. (2007). Depth of interaction decoding of a continuous crystal detector module, *Physics in Medicine and Biology* (52): 2213–2228.

Maas, M., Schaart, D., van der Laan, D., Bruyndonckx, P., Lemaitre, C., Beekman, F. and van Eijk, C. (2009). Monolithic scintillator PET detectors with intrinsic depth-of-interaction correction, *Physics in Medicine and Biology* (54): 1893–1908.

MacDonald, L. and Dahlbom, M. (1998). Parallax correction in PET using depth of interaction information, *IEEE Transactions on Nuclear Science* **45**(4): 2232–7.

Marone, A., Fiorini, C., Baraldi, P., Cadini, F., Zio, E. and Camera, F. (2009). Employing neural network to determine the position of interaction of medium-high energy gamma rays, *IEEE Nuclear Science Symposium Conference Record*, pp. 645–649.

Miyaoka, R., Li, X., Hunter, W., Pierce, II, L., McDougald, W., Kinahan, P. and Lewellen, T. (2011). Resolution Properties of a Prototype Continuous Miniature Crystal Element (cMiCE) Scanner, *IEEE Transactions on Nuclear Science* **58**(5, 1): 2244–2249.

Miyaoka, R., Li, X., Lockhart, C. and Lewellen, T. (2009). New continuous miniature crystal element (cmice) detector geometries, *IEEE Nuclear Science Symposium Conference Record* pp. 3639–3642.

Orita, N., Murayama, H., Kawai, H., Inadama, N. and Tsuda, T. (2005). Three-dimensional array of scintillation crystals with proper reflector arrangement for a depth of interaction detector, *IEEE Transactions on Nuclear Science* **52**: 8–14.

REFERENCES 32

Pedemonte, S., Gola, A., Abba, A. and Fiorini, C. (2009). Optimum real-time reconstruction of gamma events for high resolution anger camera with the use of gpgpu, *IEEE Nuclear Science Symposium Conference Record*, pp. 3388–3394.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M. and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python, *Journal of Machine Learning Research* **12**: 2825–2830.

Saul, L. and Roweis, S. (2003). Think Globally, Fit Locally: Unsupervised Learning of Low Dimensional Manifolds, *Journal of Machine Learning Research* **4**: 119–155.

Shibuya, K., Nishikido, F., Inadama, N., Yoshida, E., Lam, C., Tsuda, T., Yamaya, T. and H., M. (2007). Timing resolution improved by DOI information in an LYSO TOF-PET detector, *IEEE Nuclear Science Symposium Conference Record* **M19-7**.

van Dam, H., Seifert, S., Vinke, R., Dendooven, P., Lohner, H., Beekman, F. and Schaart, D. (2011a). A practical method for depth of interaction determination in monolithic scintillator PET detectors, *Physics in Medicine and Biology* (56): 4135–4145.

van Dam, H., Seifert, S., Vinke, R., Dendooven, P., Löhner, H., Beekman, F. and Schaart, D. (2011b). Improved Nearest Neighbor Methods for Gamma Photon Interaction Position Determination in Monolithic Scintillator PET Detectors, *IEEE Transactions on Nuclear Science* **58**(5): 2139–2147.

Wang, Y., Du, J., Zhou, Z., Yang, Y., Zhang, L. and Bruyndonckx, P. (2011). FPGA Based Electronics for PET Detector Modules With Neural Network Position Estimators, *IEEE Transactions on Nuclear Science* **58**(1, 1): 34–42.