



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Simulation Based Inference for PET iterative reconstruction

Bastien Bergere¹, Thomas Dautremer¹, and Claude Comtat^{1,2}

¹ CEA, LIST, F-91120, Palaiseau, France

² Université Paris-Saclay, Palaiseau, France

Abstract. The analytical projector (system matrix) used in most PET reconstructions does not incorporate Compton scattering and other important physical effects that affect the process generating the PET data, which can lead to biases. In our work, we define the projector from the generative model of a Monte-Carlo simulator, which already encompasses many of these effects. Based on the simulator’s implicit distribution, we propose to learn a continuous analytic surrogate for the projector by using a neural density estimator. This avoids the discretization bottleneck associated with direct Monte-Carlo estimation of the PET system matrix, which leads to very high simulation cost. We compare our method with reconstructions using the classical projector, in which corrective terms are factored into a geometrically derived system matrix. Our experiments were carried out in the 2D setting, which enables smaller-scale testing.

Keywords: PET reconstruction · Simulation-Based Inference (SBI) · Scatter correction · Conditional Density Estimation · Neural networks.

1 Introduction

Positron emission tomography (PET), whose aim is to map metabolic activity inside the patient via the administration of a radioactive tracer, is a standard imaging technique in oncology for diagnosis and therapy. The inverse problem of PET reconstruction involves obtaining an image of the spatial distribution of the radiotracer from the detections in coincidence of pairs of gamma photons.

Due to the complex nature of the physical phenomena that affect the emission, transport and detection of γ -photons, many biases impact the direct line-integral (Radon) model used to define the analytical projector in most iterative reconstruction algorithms. Notably, because of the low energy resolution of detectors, coincidences due to photons scattered inside the body may represent up to half of the total detected events. In model-based iterative reconstruction, ad hoc corrective terms must therefore be analytically estimated and factored into the PET system matrix to account for these effects. However these methods developed over time are complex and can be prone to specification errors.

In contrast, multiple simulators based on particle transport provide realistic generative models of a PET scan, to the extent that they are commonly used

in nuclear medicine as a reference. In our work we define the projector from the generative model of one of these Monte Carlo simulator [2], which incorporates realistic physical effects (mainly Compton scattering and positron physics).

A number of works [13, 17] that can be found in the extensive review [5] have already studied the feasibility and benefits of using a Monte-Carlo projector, since performance of iterative reconstructions depends on how faithfully the model describes the real PET process. However, all these studies have in common some direct Monte-Carlo estimation of the PET matrix, in which a fixed number of observations are simulated for each point source [7]. Due to the high resolution of modern PET scanners this leads inevitably to computation and memory shortage, as the system matrix can have up to 10^{12} elements.

In our proposal we instead consider a continuous Monte-Carlo projector, bypassing the bottleneck in the simulation phase due to discretization. The drawback is that the simulator’s distribution is implicit and does not provide a tractable representation of the projector, which is needed to compute the likelihood.

As in the general case where there is no likelihood, the main way to proceed is to use density estimation techniques [1] to construct a tractable surrogate. Our work is in line with this general idea, called simulation-based inference [3] (SBI), which has been growing due to the advent of high-fidelity simulators as well as the development of powerful neural density estimators [8]. Our proposal is to use the linearity in the Poisson model, so that learning a likelihood surrogate for PET imaging can be reduced to approximating the conditional distribution of the continuous projector. Due to the multimodal and low-dimensional nature of this law, we favour a neural density estimator in the form of a Gaussian mixture.

The principal objective of this article is to establish a proof of concept for the SBI approach in two-dimensional PET.

2 Continuous Surrogate Projector

For each positron emission occurring at $\mathbf{s}=(x, y)$, the simulator generates the trajectory of the annihilation photons and outputs their angular position on the detectors ring $\boldsymbol{\theta}=(\theta_1, \theta_2)$. Let p_{sim} denote the conditional density³ of (θ_1, θ_2) given an emission point (x, y) . Its discrete analogue is the PET matrix described by the conditional probabilities a_{ij} of detection in a line of response (LoR) i given a emission in pixel j . Conditionally on an emission map $\lambda(x, y)$, simulated observations follow a spatial Poisson process with intensity Λ over $[0, 2\pi]^2$, where

$$\Lambda(\theta_1, \theta_2) = \int p_{sim}(\theta_1, \theta_2 | x, y)\lambda(x, y)dxdy \quad (1)$$

Given the number of latent variables involved in generating photons trajectories within the simulator, the implicit density p_{sim} is intractable, as is the Poisson likelihood. Nevertheless, thanks to linearity in λ in (1), it is sufficient to approximate the 2D conditional distribution p_{sim} to then obtain a likelihood surrogate.

³ By design, p_{sim} is symmetric under exchange of variables θ_1, θ_2

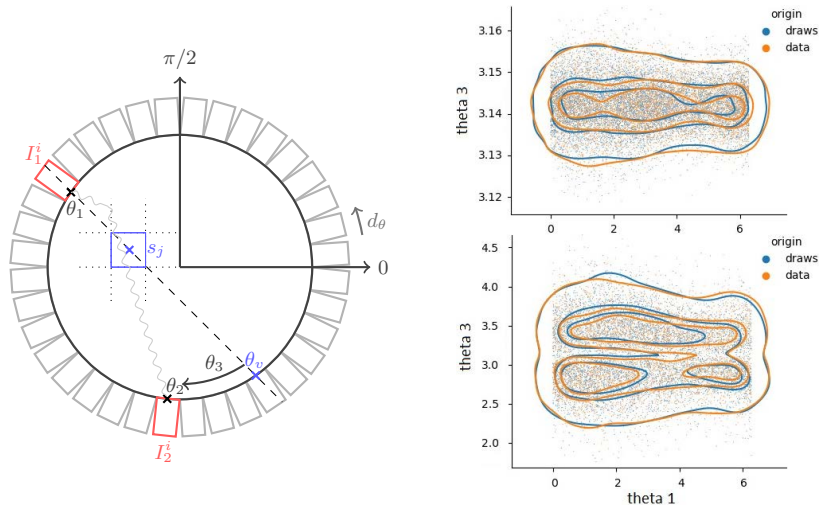


Fig. 1. Left : the new observations parameters (θ_1, θ_3) are computed conditionally to the emitting pixel $\mathbf{s}_j = (x_j, y_j)$. The angle θ_3 is the difference (centered in π) between θ_2 and the virtual angle $\theta_v(\theta_1, \mathbf{s}_j)$. Right : At pixel $\mathbf{s}_j = (-4, 2)$, the conditional distribution of (θ_1, θ_3) is sampled from the neural surrogate (orange) and the simulator (blue) for both scattered coincidences (below) and unscattered (above, different scale).

2.1 Conditional Density Estimation

Combining our prior knowledge about the multimodal structure of data with the learning capacity of neural nets, we choose to model the data with Mixture Density Networks (MDN) [1]. The MDN is constructed with two elements :

i) a conditional distribution written as a mixture of M Gaussian components⁴, where parameter M controls the model’s expressivity :

$$q(\boldsymbol{\theta} | \mathbf{s}) = \sum_{m=1}^M \pi_m(\mathbf{s}) \times \mathcal{N}(\boldsymbol{\theta} | \mu_m(\mathbf{s}), \Sigma_m(\mathbf{s})) \quad (2)$$

ii) a single neural net (MLP) with weights \mathbf{w} , that takes as an input the variable $\mathbf{s} = (x, y)$, and outputs all parameters $\pi_m(\mathbf{s}; \mathbf{w})$ and $\mu_m(\mathbf{s}; \mathbf{w}), \Sigma_m(\mathbf{s}; \mathbf{w})$ describing the mixture.

Choosing the best density estimator then becomes an optimization problem on \mathbf{w} , with the conditional likelihood (2) over the train set as the objective.

Finally, we define a new set of angles (θ_1, θ_3) parameterizing the independent observations, as shown in Figure 1. This representation, which is problem-specific, allows us to distinguish and learn separately (with two distinct MDN) the unscattered (q_1) and scattered (q_2) components of p_{sim} . For the first component, due to the small variance of θ_3 , we decompose $q_1(\theta_1, \theta_3 | \mathbf{s}) := \bar{q}_1(\theta_1 | \mathbf{s}) \times q_{GG}(\theta_3 | \theta_1, \mathbf{s})$, where \bar{q}_1 is a MDN in 1D and q_{GG} is another density network with

⁴ can be generalized to any choice of law

a single generalized Gaussian component. A classifier h is also trained with a neural net to predict the observation labels c_i associated with each component (cf. section 3). Conditionally to the weights $\mathbf{w} = [w_1, w_2, w_3, w_4]$ of \bar{q}_1, q_2, q_{GG} and h , the continuous surrogate is finally described by the following formula⁵ :

$$\hat{p}(\theta_1, \theta_3 | \mathbf{s}) = h(1 | \mathbf{s})q_1(\theta_1, \theta_3 | \mathbf{s}) + h(2 | \mathbf{s})q_2(\theta_1, \theta_3 | \mathbf{s}) \quad (3)$$

3 Stochastic Simulator

To generate PET data, we use the Phoebe software [2], developed at CEA List from Penelope [14], which enables Monte-Carlo simulation of photon and electron/positron transport. The main interactions photons undergo are the photoelectric effect (absorption) and Compton inelastic scattering. To reduce simulation time, the distributions of the positron path range⁶ and annihilation acolinearity are precomputed using Phoebe and sampled for each emission.

From the Phoebe code, which is in 3D, we generate 2D particle trajectories by conditioning the photons to remain in the plane during annihilation and later during Compton interactions. Time-dependent effects that introduce bias and non-linearities, such as radiotracer kinetics and detector dead time, as well as random coincidences, are not taken into account. Detector physics is not taken into account either, with the exception of a Gaussian blur to the energy values (20 % energy resolution). A low-energy threshold of 420 keV is then applied, leading to a scatter fraction of around 35-40 %.

Given an emission map λ , emission points are sampled in the simulator according to a Poisson process. The energy values (e_1, e_2) of the two photons when entering the detectors are used to label the observations : if $e_1 = e_2 = 511$ keV, they are labeled as unscattered ($c = 1$), otherwise as scattered ($c = 2$). Simulated events that do not result in a coincidence (at least one of the photons is not detected) are recovered and are labeled accordingly ($c = 3$) .

4 Experiments Setup

Attenuation map : Volumes of the materials in the Field of View (FOV) are described in Phoebe by quadratic surfaces. In our experiments, we chose a head-like scene with an outer ring of bone and water inside (cf. Figure 2).

Training data : For $k = 1, \dots, N_{unif} = 3 \times 10^6$, list-mode data $(x_k, y_k, \theta_1^k, \theta_3^k)$ along with the corresponding labels c_k are generated by the simulator using an uniform emission map over the FOV (simulation time $t \approx 5$ min). The data are then splitted into train and validation sub-samples. Given the grid that will be used for reconstruction (see section 4.1), the (x_k, y_k) values are then replaced by the coordinates of the pixel containing them, enabling the networks to learn the conditional density of (θ_1, θ_3) given the emission pixel.

⁵ conditional on the event being detected ($h \neq 3$). The sensitivity at \mathbf{s} is equal to the probability $h(1|\mathbf{s}) + h(2|\mathbf{s}) = 1 - h(3|\mathbf{s})$

⁶ using the positron emission energy spectrum of fluor F18.

Networks training : The four neural networks describing \bar{q}_1, q_2, q_{GG} and h are trained over our training sample (for q_2 only on the scattered data and for \bar{q}_1, q_{GG} only on the unscattered). Each of them is defined with four 256-width hidden layers, which theoretically allow to approximate any function of interest [9]. The Adam optimizer is used for gradient descent in order to minimize the negative log-likelihood or the categorical cross-entropy in the specific case of the classifier. Tunable hyperparameters include the activation functions, the Adam learning rate and the batch size. They are adjusted using the validation set. In contrast the number M of components of the MDN is fixed and set to 20. The training is implemented using the deep-learning Python library TensorFlow. Computations were done on our hardware comprised of a single Nvidia RTX 3080 GPU kernel and AMD Ryzen 64-cores CPU (total training time $t \approx 15$ min).

4.1 Reconstruction settings

As in most cases in PET, the image λ is reconstructed on a finite basis of pixels, in our case of total dimension $J=200 \times 200$. To evaluate our proposal, we use an image of a brain obtained by MRI, to which a tumor has been added (cf. Figure 2) to generate a realistic fluorodeoxyglucose [^{18}F]FDG PET phantom. With the attenuation map shown in Figure 2, which in practice must be obtained by a joint CT scan, this "true image" $\lambda^* \in \mathbf{R}^J$ is then used in Phoebe to generate a total number of $N_1 = 10^6$ unscattered-equivalent⁷ observations $(\theta_1^k, \theta_2^k)_{1 \leq k \leq N}$. By dividing the outside ring uniformly we also define $n_D = 500$ "virtual" detectors (cf. Figure 1), which leads to a total number $I = 4.4 \times 10^4$ of valid detectors pairs (LoR) : $(I_1^i, I_2^i)_{1 \leq i \leq I}$. An histogram $\mathbf{Y} \in \mathbf{R}^I$ is then obtained after binning the continuous observations (θ_1, θ_2) over the intervals (I_1^i, I_2^i) . The histogram is subsequently transformed into a sinogram.

4.2 Projector binning (system matrix)

Since observations are sampled from a Poisson process with intensity Λ , count data over the disjoint sets of LoR $L_i = (I_1^i \times I_2^i) \sqcup (I_2^i \times I_1^i)$ have the following distribution :

$$\mathbf{Y} \sim \bigotimes_{i=1}^I \mathcal{P}_{oisson}(\Lambda_i) \quad , \quad \Lambda_i = \iint_{L_i} \Lambda(\theta_1, \theta_2) d\theta_1 d\theta_2 \quad (4)$$

Defining $a_{ij} = \iint_{L_i} p_{sim}(\theta_1, \theta_2 | \mathbf{s}_j) d\theta_1 d\theta_2$ and remembering the definition of Λ in (1), we then have $\Lambda_i \approx \sum_j \lambda_j a_{ij}$. The mean value of the count data $\mathbf{\Lambda} \in \mathbf{R}^I$ is thus described by the linear system $\mathbf{\Lambda} = \mathbf{A}\lambda$, where $\mathbf{A} = (a_{ij}) \in \mathbf{R}^{I \times J}$ corresponds to the PET matrix. In our work, coefficients (a_{ij}) are estimated on

⁷ The total number of observations needed to get the mentioned number if we remove all the less-informative scattered observations.

the fly at each iteration by integrating the continuous surrogate in the 2D data space $[0, 2\pi] \times [0, 2\pi]$:

$$a_{ij} \approx \iint_{L_i} \hat{p}(\theta_1, \theta_3 \mid \mathbf{s}_j; \hat{\mathbf{w}}) \times 1 \, d\theta_1 d\theta_2 \quad (5)$$

where \mathbf{s}_j is the coordinates of pixel j and θ_3 is a function⁸ of θ_1, θ_2, s_j (cf. Fig. 1). Here $\hat{\mathbf{w}}$ denotes our solution to the optimization problem (cf. section 2.1). For the scattered component, the surrogate q_2 in equation (3) is deemed sufficiently regular and a only one value is used to approximate (5). Whereas for unscattered counts, the integral over $d\theta_1$ is estimated by importance sampling by directly calling the generative model $\bar{q}_1(\cdot | s_j)$.

For comparison, we will also use the factorized matrix $\mathbf{A} = \mathbf{H}_{dat} \mathbf{A}_{att} \mathbf{A}_{geo} \mathbf{H}_{img}$ associated to the analytical projector commonly used in PET reconstruction [12]. The diagonal matrix \mathbf{A}_{att} is obtained from the attenuation map and the sparse matrix \mathbf{A}_{geo} , which corresponds to the discretized Radon transform is efficiently computed using Siddon’s algorithm [16]. Finally, the matrix \mathbf{H}_{img} (resp. \mathbf{H}_{dat}) corresponds to a 2D convolution in the image (resp. histogram data) space.

4.3 Iterative reconstruction (EM)

The image is reconstructed by maximizing the concave Poisson log-likelihood

$$L(\boldsymbol{\lambda}) \stackrel{c}{=} \sum_i (y_i \log(\mathbf{a}_i^T \boldsymbol{\lambda}) - \mathbf{a}_i^T \boldsymbol{\lambda}) \quad (6)$$

To do that, we use the deterministic (start from a constant image) EM algorithm [6], which is the reference method for PET statistical reconstruction [11].

Since the inverse problem is ill-posed, it is necessary to regularize the EM estimates. The first method used is early Stopping with Gaussian post-smoothing⁹. Regularization is also achieved by adding to the log-likelihood objective a prior term that penalizes non-smooth images. Using a Gibbs Field prior with the Fair potential function [18], we maximize the new objective using the iterative MAP-EM algorithm of De Pierro [10, 18].

Unlike our surrogate model, scatter is not taken into account in the analytical projector and must therefore be corrected. From a rough estimate $\tilde{\boldsymbol{\lambda}}$, the number b_i of scattered counts in each LoR i is estimated beforehand¹⁰, and then the Poisson model’s mean value of counts in each LoR i is changed to $\mathbf{a}_i^T \boldsymbol{\lambda} + b_i$.

4.4 Evaluation methodology

All figures of merits are computed and averaged over $N = 20$ replicates, each obtained with an independent set of simulated observations. In addition to Mean

⁸ the Jacobian resulting from this change of variable is equal to 1

⁹ We choose a FWHM value of 2mm

¹⁰ The $\tilde{\boldsymbol{\lambda}}$ estimate is obtained with the analytical model and using all counts, and from there \mathbf{b} is estimated by simulation (can also be done analytically).

Square Error (MSE), we compute the relative bias (B) and the relative standard deviation (σ) values, which are defined as follows

$$B = \frac{\sum_{j \in R} (\bar{\lambda}_j^N - \lambda_j^*)}{\sum_{j \in R} \lambda_j^*} \quad \sigma^2 = \frac{\frac{1}{N} \sum_{n=1}^N \sum_{j \in R} (\lambda_j^{(n)} - \bar{\lambda}_j^N)^2}{\sum_{j \in R} (\lambda_j^*)^2} \quad (7)$$

with $\bar{\lambda}_j^N = \frac{1}{N} \sum_{n=1}^N \lambda_j^{(n)}$ denoting the mean reconstruction at pixel j . These relative values are computed for two regions (R) of the brain : grey matter¹¹ and lesion (tumor) and are plotted in Fig.3 for different choice of hyperparameters.

5 Results and Discussion




Fig. 2. Reconstructions are displayed in the center and rightmost columns, with those obtained using the SBI (resp. factorized) matrix in the first (resp. second) row. Left-most column : Ground Truth (top) vs Attenuation at 511 keV (down). Center-left : EM reconstruction (one replica). Center-right : MAPEM reconstruction (one replica). Rightmost column : mean EM reconstruction (10 replicates).

Figure 2 compares PET images reconstructions obtained with $N_1 = 10^6$ counts. The images are obtained using the EM and MAPEM algorithms, each with the two different discretized projectors : the simulation-based projector (SBI) versus the analytical projector (factorized matrix). The values of the regularization parameters used are $t = 80$ for EM early stopping and $\beta = 0.005$ for the weight of the Gibbs prior (they corresponds to the larger marks in Fig.3). They are chosen to minimize the MSE, which averages to 2.45×10^4 with the SBI matrix and 2.50×10^4 with the factorized matrix using EM algorithm (resp. 2.29×10^4 with SBI matrix and 2.48×10^4 with factorized matrix using MAPEM).

¹¹ the grey matter region includes the entire green colored level in the ground truth

We observe similar bias values for both projectors, but noisier reconstructions with the geometrical one. This result can be partly explained by the fact that the detector response weights (rows of the PET matrix) are more sparse in the geometric model than in the SBI model, which simulates a number of degrading physical effects including scattering. These intuitions are confirmed by the more detailed bias - standard deviations plots presented in Figure 3. For a given bias value (these are negative largely due to partial volume effects), we observe that reconstructions with SBI systematically have a lower standard deviation.




Fig. 3. Bias-Standard deviation plots computed across 20 replicates. Plotted points correspond to different regularization hyperparameters (ranging from $t = 40$ to $t = 120$ for EM and $\beta = 5 \times 10^{-5}$ to $\beta = 10^{-3}$ for MAPEM).

5.1 Discussion

Our goal is to propose a new conceptual approach that allows to efficiently incorporate a realistic generative model of the PET process into an iterative model-based reconstruction (cf. Supplementary Material). As in the case of reconstruction with accurate resolution modeling, we observe less pixel variance using our simulation-based projector (a slower convergence is also observed). Using a different simulator such as GATE [15] to generate test data should be the main way to attest the validity of our proposal.

The advantage of the SBI approach is that it is much easier to add complex physical effects (such as annihilation non-collinearity) to the simulator’s generative model than to an explicit analytical model. These effects can be added naturally into the simulator by defining and drawing latent variables. They are then automatically integrated into the surrogate model during the supervised learning phase as the estimator implicitly marginalizes over all latent variables. As many software packages [4, 15] already offer accurate stochastic simulation of the PET process (including crystal detection), this is a promising field of study. Clinical feasibility ultimately depends on the scalability of our method to 3D,

which seems achievable without major changes. Indeed, learning a continuous model remains a density estimation problem in low-dimensional spaces (only three coordinates are needed to define an emission source and four for a line of response). In terms of computational resources, the main limitations are likely to be related to the discretization of the continuous model during reconstruction. Another important question concerns the management of the possible bias introduced by the simulator and the regression.

Acknowledgments. This work is supported by CEA-List and FOCUS program.

Disclosure of Interests. Authors have no competing interests.

References

1. Bishop, C.M.: Pattern Recognition and Machine Learning (Information Science and Statistics). Springer-Verlag, Berlin, Heidelberg (2006)
2. Champciaux, V.: Calcul accéléré de la dose périphérique en radiothérapie. Ph.D. thesis (01 2021)
3. Cranmer, K., Brehmer, J., Louppe, G.: The frontier of simulation-based inference. Proceedings of the National Academy of Sciences **117**, 201912789 (05 2020). <https://doi.org/10.1073/pnas.1912789117>
4. España, S., Lopez Herraiz, J., Vicente, E., Vaquero, J.J., Desco, M., Udías, J.: PeneloPET, a Monte Carlo PET simulation tool based on PENELOPE: features and validation. Physics in medicine and biology **54**, 1723–42 (03 2009). <https://doi.org/10.1088/0031-9155/54/6/021>
5. Iriarte, A., Marabini, R., Matej, S., Sorzano, C., Lewitt, R.: System models for PET statistical iterative reconstruction: a review. Computerized Medical Imaging and Graphics **48** (12 2015). <https://doi.org/10.1016/j.compmedimag.2015.12.003>
6. Lange, K.L.: EM reconstruction algorithms for emission and transmission tomography. Journal of computer assisted tomography **8** **2**, 306–16 (1984)
7. Ortuño, J.E., Guerra-Gutiérrez, P., Rubio, J.L., Kontaxakis, G., Santos, A.: 3D-OSEM iterative image reconstruction for high-resolution PET using precalculated system matrix. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **569**(2), 440–444 (2006)
8. Papamakarios, G., Nalisnick, E., Rezende, D.J., Mohamed, S., Lakshminarayanan, B.: Normalizing flows for probabilistic modeling and inference. Journal of Machine Learning Research **22**(57), 1–64 (2021), <http://jmlr.org/papers/v22/19-1028.html>
9. Park, S., Yun, C., Lee, J., Shin, J.: Minimum width for universal approximation. In: International Conference on Learning Representations (2021)
10. Pierro, A.R.D.: A modified expectation maximization algorithm for penalized likelihood estimation in emission tomography. IEEE Trans. Medical Imaging **14**, 132–137 (1995)
11. Qi, J., Leahy, R.: Iterative reconstruction techniques in emission computed tomography. physics in medicine biology, 51, r541-r578. Physics in medicine and biology **51**, R541–78 (09 2006). <https://doi.org/10.1088/0031-9155/51/15/R01>

12. Qi, J., Leahy, R.M., Cherry, S.R., Chatzioannou, A., Farquhar, T.H.: High-resolution 3D Bayesian image reconstruction using the microPET small-animal scanner. *Physics in medicine & biology* **43**(4), 1001 (1998)
13. Rannou, F., Chatzioannou, A.: Fully 3D system model estimation of OPET by Monte Carlo simulation. In: *IEEE Symposium Conference Record Nuclear Science 2004*. vol. 6, pp. 3433–3436 Vol. 6 (2004)
14. Salvat, F., Fernández-Varea, J., Sempau, J.: PENELOPE. a code system for Monte Carlo simulation of electron and photon transport. *NEA Data Bank, Workshop Proceeding, Barcelona* pp. 4–7 (01 2007)
15. Sarrut, D., Bała, M., Bardiès, M., Bert, J., Chauvin, M., Chatzipapas, K., Dupont, M., Etxebeste, A., Fanchon, L., Jan, S., Kayal, G., Kirov, A., Kowalski, P., Krzemien, W., Labour, J., Schöneck, M., Loudos, G., Mehadji, B., Ménard, L., Roncali, E.: Advanced Monte Carlo simulations of emission tomography imaging systems with GATE. *Physics in Medicine Biology* **66** (05 2021). <https://doi.org/10.1088/1361-6560/abf276>
16. Siddon, R.L.: Fast calculation of the exact radiological path for a three-dimensional CT array. *Medical Physics* **12**(2), 252–255 (1985). <https://doi.org/https://doi.org/10.1118/1.595715>
17. Southekal, S., Purschke, M.L., Schlyer, D.J., Vaska, P.: Quantitative PET imaging using a comprehensive monte carlo system model. *IEEE Transactions on Nuclear Science* **58**(5), 2286–2295 (2011). <https://doi.org/10.1109/TNS.2011.2160094>
18. Wang, G., Qi, J.: Edge-preserving PET image reconstruction using trust optimization transfer. *IEEE transactions on medical imaging* **34** (11 2014). <https://doi.org/10.1109/TMI.2014.2371392>