Annihilation photon GAN source model for PET Monte Carlo simulation

David Sarrut, A Etxebeste, T Kaprelian, A Saporta, J M Létang

Annihilation photon GAN source model for PET Monte Carlo simulation

D. Sarrut[1], A. Etxebeste[1], T. Kaprelian[1], A. Saporta[1], JM. Létang[1]

[1] Université de Lyon, CREATIS; CNRS UMR5220; Inserm U1044; INSA-Lyon; Université Lyon 1; Centre Léon Bérard, France.

E-mail: david.sarrut@creatis.insa-lyon.fr

Abstract.

Objective. Following previous works on virtual sources model with GAN, we extend the proof of concept for generating back-to-back pairs of gammas with timing information, typically for Monte Carlo simulation of PET imaging.

Approach. A conditional GAN is trained once from a low statistic simulation in a given attenuation phantom and enables the generation of various activity source distributions. GAN training input is a set of gammas exiting a phantom, tracked from a source of positron emitters, described by position, direction and energy. A new parameterization that improves the training is also proposed. An ideal PET reconstruction algorithm is used to evaluate the quality of the GAN.

Main results. The proposed method is evaluated on NEMA IEC phantoms and with CT patient image showing good agreement with reference simulations. The proportions of 2-gammas, 1-gammas and absorbed-gammas are respected to within one percent, image profiles matched and recovery coefficients were close with less than 5% difference. GAN tends to blur gamma energy peak, e.g. 511 keV.

Significance. Once trained, the GAN generator can be used as input source for Monte Carlo simulations of PET imaging systems, decreasing the computational time with speedups up to $\times 400$ according to the configurations.

1. Introduction

In previous works [1–3], it has been shown that Generative Adversarial Network (GAN) can model phase-spaces such that the trained generator neural network approximates a distribution of particles, thereby being able to be used as a fast virtual source model (VSM,e) of particles. This approach was applied for modeling photon beams from Linac head in radiation therapy treatment [1] and for gammas exiting phantoms or patient CT during simulation of SPECT systems [2; 3]. Compared to conventional phase-space files, GAN generators are compact, few MB instead of few GB, and can generate particles at high speed (around 10^6 particles per second). Hence, splitting Monte Carlo simulation of imaging systems in two parts – (i) tracking emitted particles within the phantom/patient CT, (ii) tracking particles in detector, and modeling the first

part with the GAN generator – was shown to lead to computational time gain. GAN was proposed in medical physics Monte Carlo simulation for dosimetry, such as in [4], but also as source of particles in HEP community for particle showers in calorimeters (CaloGAN) [5] or jet simulation (DijetGAN) [6]. More generally, readers interested in machine learning methods for particle physics simulation, can refer to the following review papers [7–12] and this living review website.
However, previously proposed GAN-generated sources of particles cannot readily be used for the simulations of PET imaging system because they require particle time information and the knowledge of the initial emitting event to distinguish between true and scatter/random coincidences. In this work, we extend the method to consider pairs of particles and timing information. We also propose a new parameterization that improves the training and the use of conditional GAN in order for the generator to model a family of activity source distributions.

2. Material and methods

The goal is to train a GAN that can generate the correct distribution of gammas exiting a given phantom filled with any β^+ activity source distribution. The considered phantom may be a standard phantom to assess performance characteristics (e.g. NEMA) or a 3D CT image of a patient. Once trained, the GAN will be able to generate a distribution of (pairs of) gammas as they exit the phantom, sparing the needed computation time for the β^+ emission and the tracking of the gamma pairs. The GAN is trained specifically for a given phantom, but is conditional to the β^+ activity source distribution that can be arbitrarily chosen. It can be used as a fast source for another Monte Carlo simulation, like testing different PET devices, or for any other data analysis task.

Training dataset from low statistic Monte Carlo simulation. First, a training dataset is created by a low statistic Monte Carlo simulation which tracks particles from the β^+ emission to gammas exiting the phantom. It is composed of gammas reaching a given surface surrounding the phantom, typically a sphere surrounding the entire phantom. We called the output gammas the "exiting gammas". Each gamma is described by its exiting position, direction, energy and time. Those properties are the ones needed for a complete PET simulation. Moreover, we also store other information: the ID of the initial β^+ event (required to pair the gammas) and its position (later required by the conditional GAN). The dataset is thus composed of exiting energy + exiting time + exiting position + exiting direction + event position + event ID making a total of 12 dimensions. Note that the time is the time from the initial event emission, so the sum of travel times of the β^+ and the exiting gamma. The initial activity source for this training dataset may be simulated in several ways. We considered here the emission of β^+ particles with energy spectrum corresponding to a given β^+ emitter obtained from [13], hence including the range of positrons. The activity source is chosen to be uniform (no spatial a priori) over the whole phantom, each β^+ originating with the same probability

from any position in the phantom.

Pairing. From the obtained output phase-space, the exiting gammas are paired according to the ID of the initial β^+ emission. We consider here exiting pairs of back-to-back gammas, exiting single gammas and exiting events with 3 or more gammas originating from the same β^+. The pairs of events represent more than 91% of the events, while single events represent about 8%. Triplet and higher order multiple events about 1%. These ratios depend on the phantom and can vary. The single events correspond to back-to-back gammas where one gamma is absorbed in the phantom by photoelectric interaction. Triplet (and more) events correspond to additional (rare) gammas created either during disintegration (positronium) or by bremsstrahlung of β^+. Triplets are considered as a pair plus a single gamma, quadruplets as two pairs, etc. All incomplete pairs, such as totally absorbed pairs of gammas or single events, are still included in the dataset: every time a gamma is absorbed, a fake gamma is created with energy set to zero.
At the end of this step, we obtained a dataset with 16 dimensions (the position, the direction, the energy and the time for the two gammas) representative of the reference probability distribution of exiting gammas. It contains information about the ratio between exiting and absorbed gammas for the given attenuation phantom.

Parameterization. From this paired phase-space, the exiting gammas (resp. A and B) positions on the sphere are denoted $\mathbf{P_A}$ and $\mathbf{P_B}$, exiting directions $\mathbf{d_A}$ and $\mathbf{d_B}$, energies $\mathrm{E_A}$ and $\mathrm{E_B}$, and time $\mathrm{t_A}$ and $\mathrm{t_B}$. We parameterized this information with time-weighted LOR (Line of Response), similar to conventional parameterization of LOR in PET imaging. Indeed, we computed the time-weighted position $\mathbf{P_C}$ on the A to B LOR with $\mathbf{d_C} = \overrightarrow{\mathbf{P_A}\mathbf{P_B}}$ direction, such that $\mathbf{P_C} = \mathbf{P_A} + \frac{\mathrm{t_A}}{\mathrm{t_A}+\mathrm{t_B}}\mathbf{d_C}$, and stored the sum of times $\mathrm{t_C} = \mathrm{t_A} + \mathrm{t_B}$ to be able to retrieve both timestamps. Note that, if a gamma undergoes Compton scattering, this position $\mathbf{P_C}$ does not correspond to the true location of the annihilation. For the fake gammas (unpaired with energy set to zero), a fake B event is artificially created with the position $\mathbf{P_B}$, time $\mathrm{t_B}$ and direction $\mathbf{d_B}$ computed from the (known) emission point as if the gamma traveled in straight line, without interaction. Moreover, the energy $\mathrm{E_B}$ is set to zero to be identified as a single event. The figure 1 illustrates this parameterization in a situation where one of the two back-to-back gammas changes its direction due to a Compton interaction. The initial 16 dimensions were thus replaced by 15 dimensions, we will show in the experiments that it allows a better training.

Conditional GAN training. Following previous works on source GAN [1–3], a GAN was developed to learn generating exiting gammas following the initial distributions. The two neural networks of the GAN, the Critic (C) and the Generator (G) are trained simultaneously and act as follows. G takes as input a vector of noise of size N and generates N pairs of gammas. We considered *conditional* GAN [14; 15] with the positions

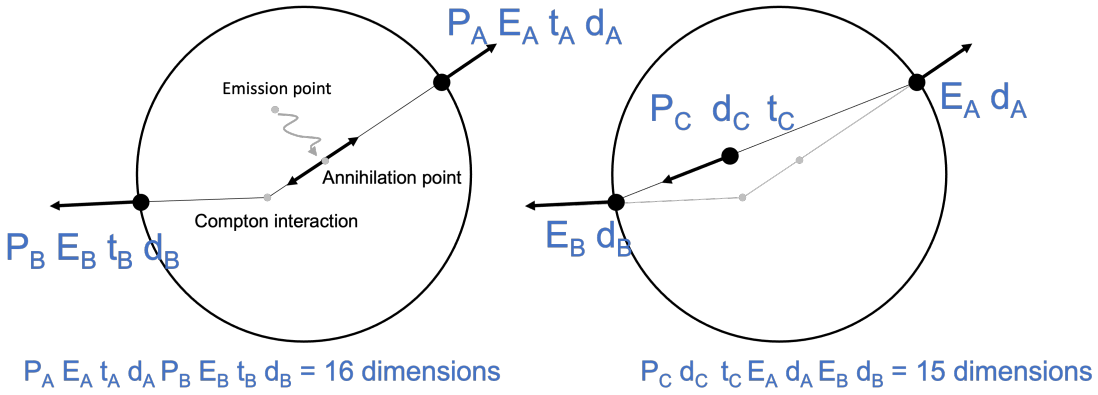

Figure 1. Illustration of the proposed 15-dimension parameterization of the pairs of exiting gammas with time-weighted LOR.

of β^+ emission $\mathbf{P_0}$ as the vector of conditions (the directions are considered isotropic). User of the GAN will then be able to consider any set of $\mathbf{P_0}$ to give as input to the GAN to obtain the corresponding exiting gammas. In order to train G for generating gamma distributions resembling the ones in the training dataset, C is trained to distinguish between gamma distribution coming from the real training dataset and distributions generated by G. The architectures of the C and G were both 4 layers of H = 600 fully connected neurons with leaky Relu non linearity. The dimension of the noise vector z was set 20. It leads to approximately 1.4×10^6 weights for each C and G network. The loss function was Wasserstein [16] with gradient penalty [17; 18] weighted by λ. We chose a Squared Hinge regularization [19] with $\lambda = 10$. The optimisation was performed by RMSProp gradient descend [20] with adaptive learning rates starting at 4×10^{-4} and decreasing by a factor of 0.8 every 5,000 iterations. Batch size was 10^4 particles and 60 epochs were performed (approximately 10^5 iterations). Note that, because the input dimensions were of multiple natures (position, energy, direction and time), an initial normalization (zero mean and unit variance) was required for all dimensions.
At the end of the training process, we obtain a conditional neural network G that takes as input a condition vector of source positions and a random sample z from the standard distribution. From this input, G will produce a vector of back-to-back exiting gammas, taking into account positron range, attenuation and scatter of the phantom. G can thus be considered as a "forward generative model". The (slow) β^+ and gamma tracking in the phantom can thus be replaced by (fast) GAN generated gammas that can then be tracked towards a given PET system with conventional Monte Carlo tracking.

Evaluation with ideal reconstruction. It is not straightforward to evaluate the ability of the GAN to produce gammas distribution. The main idea here is to compare GAN generated gammas to gammas obtained from an analog reference Monte Carlo simulation. While it is possible to compare the marginal distributions of those two distributions, it is only a partial assessment because there are correlations between the different dimensions

of the distribution. Another alternative is to compare detected gammas from a simulated PET acquisition, obtained from an analog Monte Carlo reference simulation and from GAN generated gammas tracked in the PET imaging system. However in that case, the imperfect detection of the PET system will impair the comparison and could lead to artificially good results, only due to the limited spatial and time resolution. Instead, we decided to consider a virtual ideal detection system with perfect capabilities: infinite time and spatial resolution. From all exiting pairs of gammas, a 3D image can be reconstructed taking into account the gammas time-of-flight to locate the interaction points in the phantom, considering the speed of light. The virtual detector is (arbitrarily) located at a given distance (for example 20 cm from the phantom), allowing to take into account potential uncertainty in the exiting directions of the generated gammas. This reconstruction method can be related to the reconstruction-free approach described in [21]. Such an "ideal reconstruction" can be considered as a worst case test as real-life PET detection systems will obviously depict lower resolution and hence be more tolerant to differences between the distributions. Finally, note that, as explained before, the single and absorbed gammas are also considered in the process.

3. Experiments

All Monte Carlo experiments were performed with an upcoming GATE version 10 [22; 23] using Geant4 version 11. The beta version of Gate 10 has been released in April 2023‡. All additional features and tools developed in this work are open-source and included in this version.

Several experiments were conducted using a numerical model of the "NEMA IEC Body Phantom" which is 180 mm high and composed of a lung insert and 6 spheres with various diameters: 10, 13, 17, 22, 28 and 37 mm (see figure 2). The GATE model of the phantom was developed with the exact dimensions and materials (plastic, lung ICRP material, etc). All elements were described with analytical shapes such as boxes, spheres, cylinders, etc. The six spheres can be filled with some given activity concentrations, for example with ^{18}F or ^{68}Ga. The background activity was simulated with an activity source in the large water compartment. In order to illustrate the ability of the GAN to learn the forward model including all physical processes (absorption, scattering), we artificially introduced a large tungsten insert ($3 \times 8 \times 10$ cm, see figure 2) in the IEC phantom. This insert generates additional scatter and absorption enhancing the capability of the GAN to reproduce the consequence of those interactions on the exiting gammas.

Sources were described as β^+ particles with energy spectrum obtained from [13], hence including range of positrons. All particles exiting the phantom were scored as soon as they reach a spherical surface embedding the entire phantom. The phase-space included position, direction, energy and time of all the gammas reaching the surface. The time was defined according to the initial β^+ creation at $t = 0$ and thus included the travel

‡ https://github.com/OpenGATE/opengate

time (time of flight) of both the positron and the annihilation gammas. The phase-space also stored the initial position of the emission point (to be used as the condition of the conditional GAN). Simulations were run such that the output phase-spaces contained approximately 3.2×10^7 gammas (3.1 GB), from 1.7×10^7 primaries.
All computation time analyses were performed on the same computer: Intel® Core™ i9-10900K CPU @ 3.70GHz with NVIDIA Quadro RTX (4000/PCIe/SSE2, 8 GB memory). For the sake of comparison, we considered single core computation but multi-cores could of course be used to further decrease computation time. We are using the recommended physics list named "EM Option4", with production cuts of 0.1 mm and are considering only the time spent in the tracking of the particles in the phantom.

Test case #1. The first test case was performed with a constant activity concentration of 2 kBq·mL^{-1} in all 6 spheres and no background activity.

Test case #2. The second test case was performed with several (relative) concentration activities in spheres and background: 100 for 10 mm, 17 mm and 28 mm diameter spheres, 50 for 37 mm diameter sphere and the central insert, 20 for 13 and 22 mm diameter spheres and 0.1 for the background making a total of about 1.8 MBq. The tests were performed with both analog Monte Carlo simulation (reference) and with GAN-generated back-to-back gammas. Images were reconstructed with the ideal reconstruction method and compared in terms of Recovery Coefficient computed for all 6 spheres: $RC = \frac{A_s - A_{bg}}{A_s^r - A_{bg}^r}$ with A_s and A_{bg} the obtained activity (with GAN) in a sphere s and in the background, and A_s^r and A_{bg}^r the reference activities (from Monte Carlo simulation). The noise levels [24] in different ROI (the hot spheres) were also compared: $NL = \frac{A_{SD}}{A_{\text{mean}}}$ with A_{SD} the standard deviation of the activity values of all voxels in the ROI, and A_{mean} the mean activity.

Test case #3. A third test case was also performed with a CT image as phantom and ^{68}Ga, having a larger positron range than ^{18}F. The training dataset was generated with a uniform ^{68}Ga activity source inside the patient contour. Then, to illustrate the versatility of the trained GAN, two generated conditional activity distribution of ^{68}Ga were chosen, the first with 3 spheres of activity inside the patient (see image in figure 6), and the second as a voxelized source of a ^{68}Ga PET image (see image in figure 2-right). In addition, in order to illustrate the reliability of the method, two other artificial test cases were generated: one with several point-sources and one with a star-shape source in order to highlight the differences from the reference simulation. In all cases, CT and activity images were resampled to (4 mm)^3 voxel size.

4. Results

The figure 3 illustrates the interest of the proposed parameterization for the test #1 (IEC phantom with activity in all six spheres, no background). The two profiles extracted from

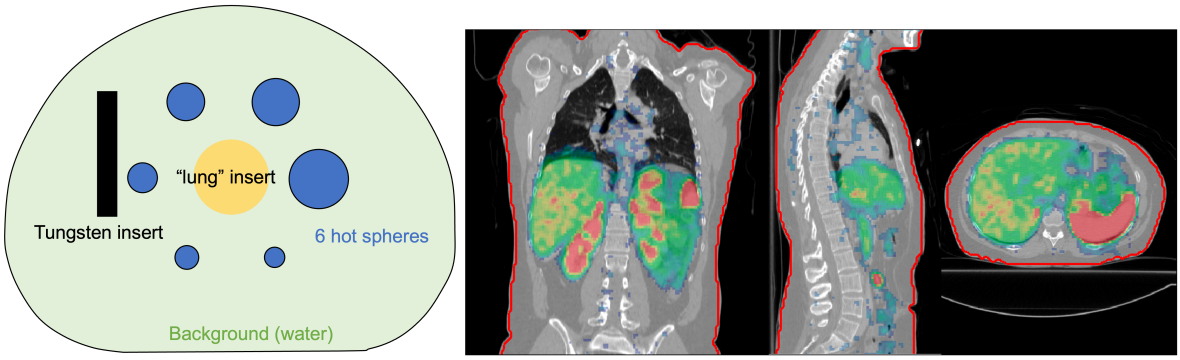

Figure 2. Schematic of the IEC 6 spheres phantom with additional tungsten insert (Left). CT image with patient contour (in red) and PET activity superimposed (Right).

the ideally reconstructed images show differences when using or not the parameterization. Recovery coefficient were between 90 and 102% for all spheres with the proposed parameterization and from 70% to 98% without. Noise levels were lower than 4% for the five larger spheres and around 9.5% for the smallest one, while they were between 7 and 11% when the proposed parameterization is not used. The increase attenuation for the smallest sphere near the tungsten insert is visible.

Similarly, the figure 4 depicts the results for test #2, with spheres, central insert and background having various activity concentrations (profile lines are depicted in both images). Recovery coefficients with respect to the reference Monte Carlo simulation were in average on all spheres of 5.5%, with noise level difference of 5.2%. Note that it is possible to perform ideal reconstruction with an energy threshold (for example discard all gammas with energy lower than 500 keV). In that case, comparison of reference and GAN generated gamma leads to very close images (less than 5% difference).

Several physical characteristics of the gammas distribution were compared between reference and GAN-generated. We computed the numbers of paired/unpaired/absorbed gammas in table 1. As explained before, absorbed gammas are detected with an energy strictly equal to zero in the reference dataset. Since GAN tends to blur the values, a threshold of 20 keV was applied to the generated distribution: generated gammas with energy below this threshold were considered absorbed. The figure 5 displays the energy spectra of the exiting gammas. We can notice the energy exactly equal to zero for the absorbed gammas in the reference dataset, while GAN-generated gammas lead to a slightly spread distribution centered around zero (even with negative values). Also, the 511 keV peak is spread for GAN-generated gammas.

The CT slices with reconstructed ^{68}Ga events superimposed for the two first test-#3 cases (with hot spheres and with realistic activity distribution) are presented in Figures 6 and 7. The upper rows show the results of the reference Monte Carlo simulation, while the bottom rows show the corresponding images generated by the GAN. It is worth noting that the noise levels in the two image versions were slightly different (within 5%) which can be observed around the liver region in Figure 7.

The versatility of the proposed method is illustrated in Figures 8 and 9 for the two other tests, namely the point-sources and star-shape sources. The reconstructed intensity profiles for both methods (reference and GAN) were in good agreement: the pixel-wise mean of absolute differences between the reference and GAN-generated images was computed relative to the maximum value and observed to be lower than 3% for both tests. It is worth noting that in the point-sources example, all sources had the same intensity but showed different reconstructed values due to attenuation, and were well-reproduced.

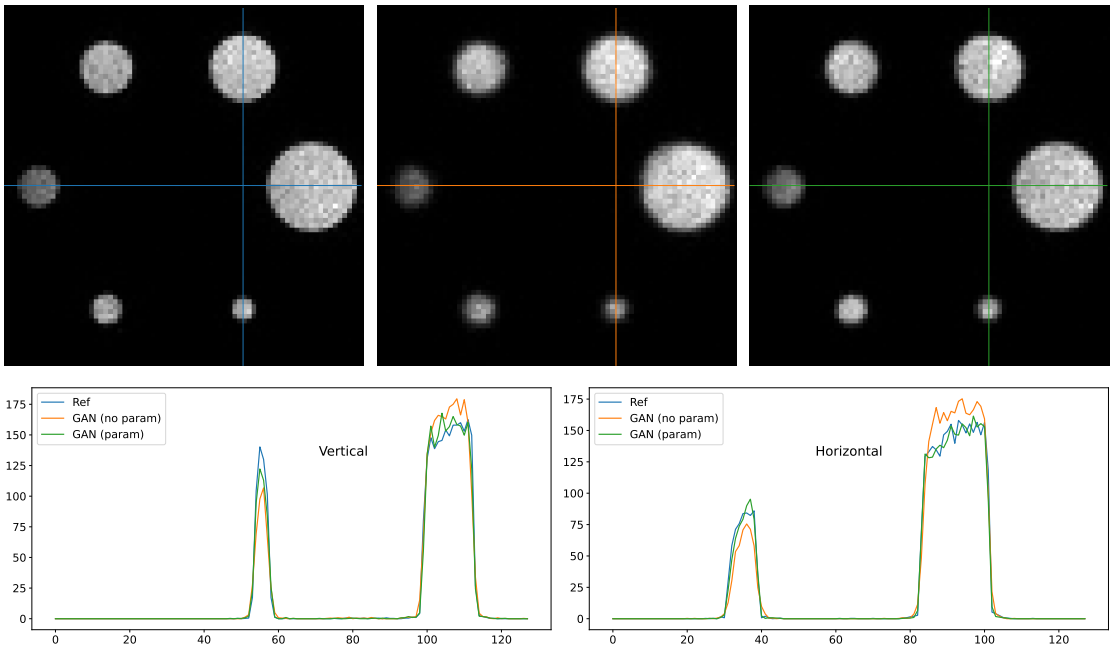

Figure 3. First row shows slices of the reconstructed images for test #1, for 1) reference simulation, 2) GAN-based simulation without parameterization, 3) GAN-based simulation with parameterization. Second row depicts horizontal and vertical profiles of the three slices.

	Test #1			Test #3		
	ref	**cGAN**	**ratio**	**ref**	**cGAN**	**ratio**
Exiting 2-gammas	90.5%	90.3%	*100.2%*	91.2%	90.8%	*100.3%*
Exiting 1-gammas	9.1%	9.2%	*98.4%*	8.5%	8.8%	*96.5%*
Absorbed gammas	0.4%	0.5%	*89.1%*	0.4%	0.4%	*102.4%*

Table 1. Analysis of exiting gammas in reference and GAN-generated distributions for tests #1 and #3: proportions of 2-gammas, 1-gammas and absorbed gammas. The third columns show the differences.

Computation time analysis. The speed for a reference Monte Carlo simulation greatly varies according to the setup. It is mostly related to the number of intermediate steps

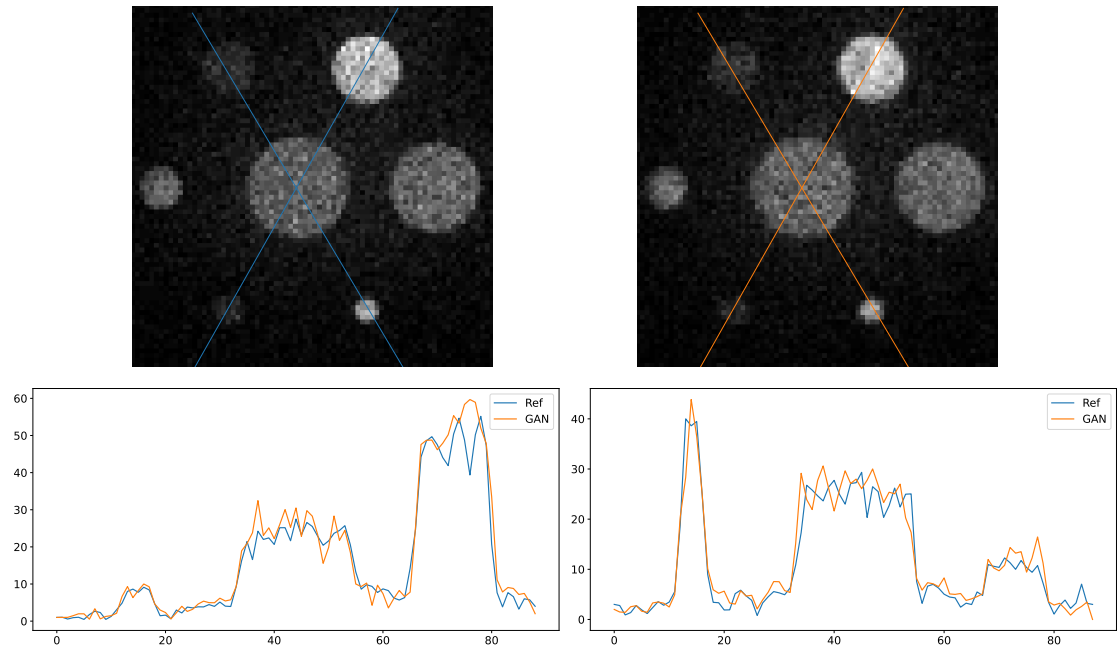

Figure 4. Slices images for test #2: reference at left and GAN generated at right. The plots on the second row show profiles extracted on both images for both methods along the given diagonal lines.

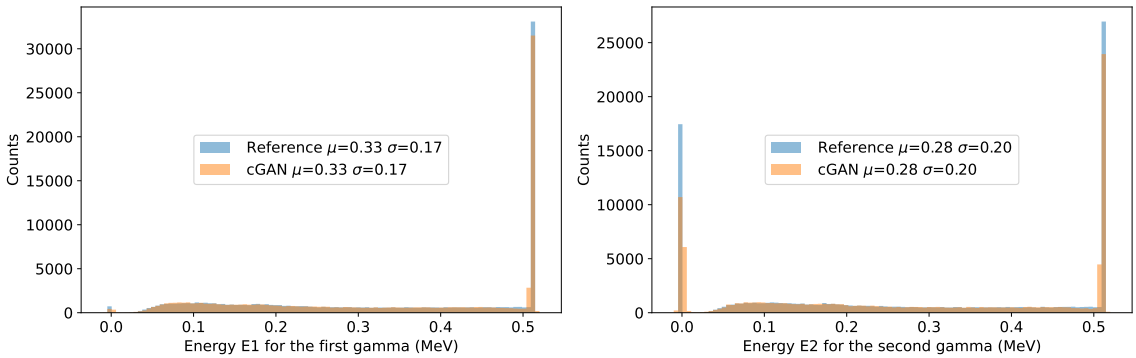

Figure 5. Analysis of exiting gammas energy spectrum in reference and GAN-generated distributions for tests #1. Gamma with energy equal to zero (reference) or close to zero (GAN-generated) are considered as absorbed gammas. Left is the energy of the first gamma and right the energy of the second gamma in all pairs.

that the Geant4 engine will have to process. With Geant4 several parameters impact the computation time: the physics list (e.g. option1 around twice faster than option4), the tracking cuts or the type of volume (voxelized or analytical). For example, tracking ^{68}Ga β^+ is slower than tracking ^{18}F β^+, approximately 20%, because the higher energy of the former leads to more steps to track. With ^{68}Ga β^+, when the phantom is described by analytical shapes (spheres, cylinders, etc.), PPS (Particle Per Second) speed was about 6×10^3. With voxelized volume, the speed drops to 2×10^3 PPS, considering

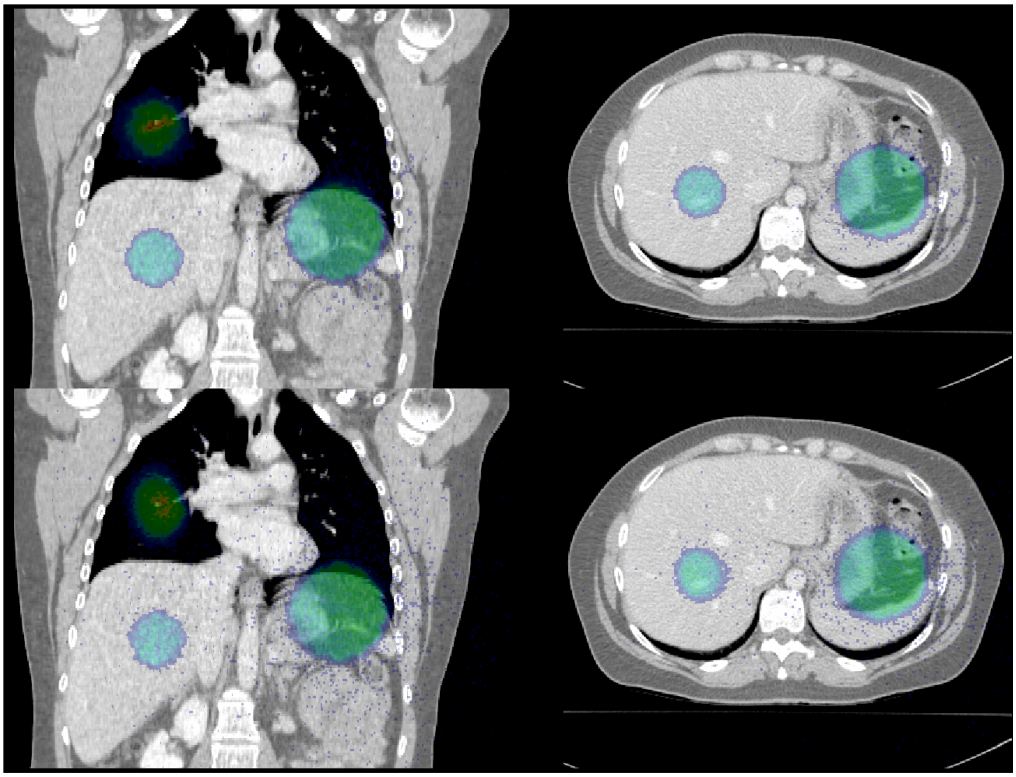

Figure 6. CT slices with ideally reconstructed events superimposed for experiment #3 with three spherical sources. Top: from reference Monte Carlo simulation (around 15k particles per second); Bottom: from GAN generated events (around 830k particles per second).

voxel's side size of 2 mm. If the phantom is larger than the IEC phantom, like the complete CT image of the patient in previous tests, speed can drop to 1.5×10^3 PPS.
The GAN-based method is split into three stages. The first stage is dedicated to the generation of the training dataset. Around 2×10^7 exiting particles were needed, which leads to about 50 minutes for analytical volume, or 10 hours for voxelized volumes. The second stage is the training, which lasts less than 2 hours. Those two first stages are only needed once.
The raw generation of particles with GAN is close to 8.3×10^5 PPS and is independent of the volume type (voxelized or not). Hence the speedups for particles tracking in a volume are around $\times 130$ (analytical) or $\times 400$ (voxelized) times faster than the analog version. This speed should be mitigated if the generated particles are re-inserted into a subsequent Monte Carlo simulation. In that case, the Geant4 kernel engine must create new particles from the properties generated by the GAN. We implemented such a GAN source in GATE/Geant4 and the computation time was around 1.8×10^5 PPS, about 90 times faster than the reference (voxelized phantom). Note that the comparison is not fully fair as part of the GAN generation uses GPU while everything else is CPU. We computed the speedups on another architecture without GPU (Apple M1 Pro, 2021) and found speedup factors around 40. The drop in performance between the raw GAN

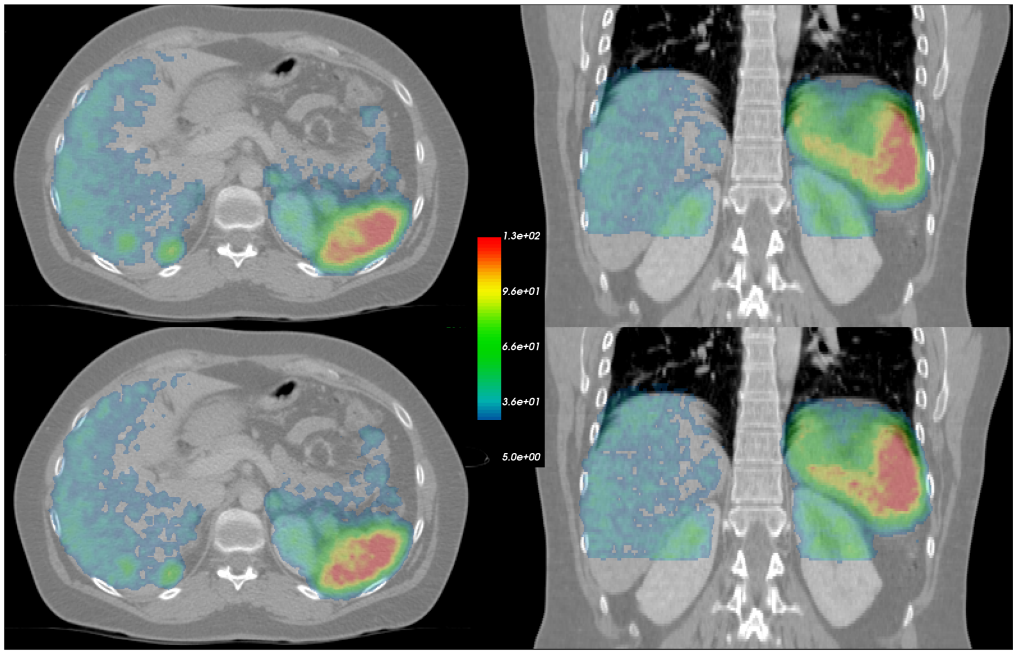

Figure 7. CT slices with ideally reconstructed events superimposed for experiment #3 with realistic ^{68}Ga distribution. Top: from reference Monte Carlo simulation (around 15k particles per second); Bottom: from GAN generated events (around 830k particles per second).

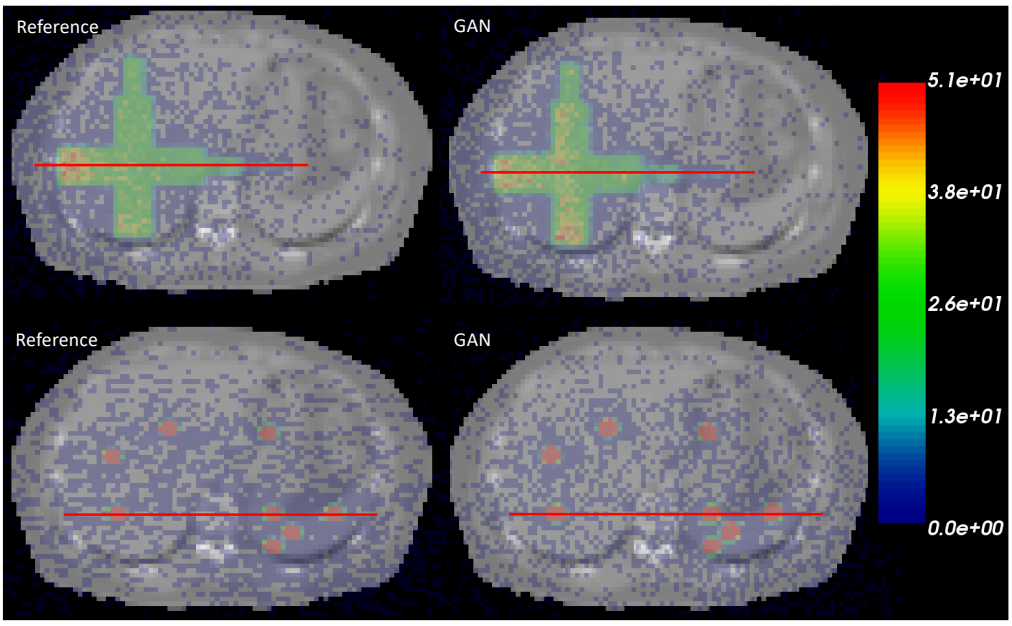

Figure 8. Experiment #3 with artificial sources (star-shape and point-sources). Left are computed with the reference method and right with the proposed GAN-method. Red lines correspond to the profiles extracted in figure 9.

generation and the Geant4 particles creation is not fully understood and may be improved

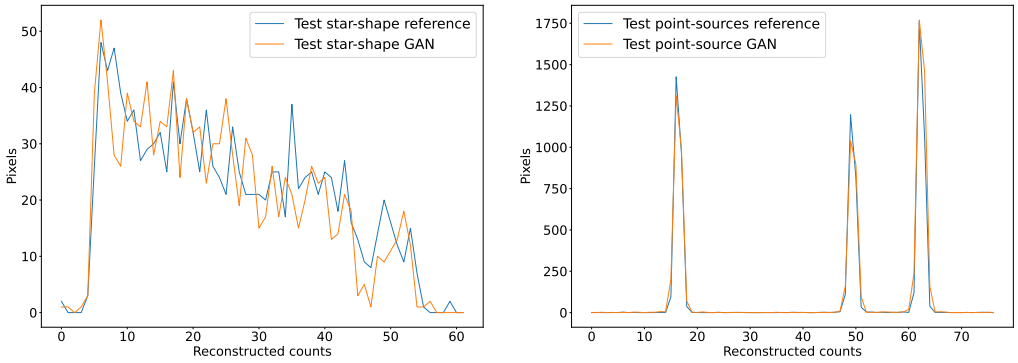

Figure 9. Extracted image profiles from figure 8, corresponding to star-shape and points-sources tests.

with a better implementation.
In addition, particles are usually tracked not only in phantom but also in a PET detector. In that case, the time needed to track the particle in the PET head is not modified. With a standard simulation of PET system in GATE, we computed that without GAN, 40% of the time is spent in tracking the particle in the phantom and 60% in the PET head. With GAN, only 7% of the time is spent in the phantom.

5. Discussion and conclusion

We proposed a neural network forward model that generates exiting pairs of back-to-back gammas from a given attenuation phantom. The model is compact (around 10 MB) and can quickly generate millions of gammas closely following the initial probability distributions, including scattering and absorption. The model is specific for a given attenuation phantom, so the dataset and the training must be redone (only once) if it is changed. However, thanks to the conditional part, the model is valid for an arbitrary source distribution, chosen by the user. For example, we show that one single model for a given patient CT image can be used to generate PET images for several activity distributions.
While we obtained results close to the reference Monte Carlo simulations (to within a few percent), it is difficult to test the complete reliability of the method (see star-shape and point-sources test cases). To date, it is not known if the GAN-based simulations are biased or not.
The model can be used either independently of a Monte Carlo simulation for data analysis or within (as a source of particles). For the later, this GAN-based approach is more suited for complex and large simulations when detailed tracking in the phantom is important, i.e. for studying the impact of ^{68}Ga positron range on PET images. With small CT voxel (2 mm), the speedup reach $\times 90$ and the time spent in tracking particles in the phantom become negligible compared to other parts of the simulation. Outside GATE/Geant4, the generation speed of particles is at about 8.3×10^5 PPS.

The G and D networks in our study were originally designed as fully connected neural networks or Densely connected Neural Network (DNN), which is a simple architecture when compared to other modern deep learning approaches. We also explored using convolutional networks (CNN) to replace the deep fully connected layers with 1D convolution. Our best results were achieved using a hybrid architecture consisting of 200 neurons for the two initial fully connected layers, followed by a UNet shape of 8-16-16-8 layers with 3D feature maps and two final fully connected layers. While this approach led to results that were close in terms of recovery coefficients, the training time increased (with a factor of 1.5× or 2×), and the generation time was significantly slower with CNN than with DNN (2×-3× times slower with our computer and GPU). Moreover, sharp situations (such as the point source test case) depicted smoothed intensity and were less accurate. Although other architectures, such as VAE or generative diffusion processes [25], could potentially improve the results, we decided to stick with the DNN approach due to its simplicity and efficiency.

The limitations are the following. First, the information about triple (or higher order multiple gammas) gammas, even if the total number of gammas is correct, is lost and cannot be retrieved. The GAN tends to blur the distributions, hence peak energy values such as zero energy to detect absorbed gammas or 511 keV of annihilation photons are slightly spread.

The code, including both the additional functionalities inserted in GATE and the training of the model, is open-source and is included in GATE version 10.

Acknowledgments

This work was performed within the framework of the MOCAMED project (ANR-20-CE45-0025), the SIRIC LYriCAN Grant INCa-INSERM-DGOS-12563, the LABEX PRIMES (ANR-11-LABX-0063) of Université de Lyon, within the program "Investissements d'Avenir" (ANR- 11-IDEX-0007) operated by the ANR, and the POPEYE ERA PerMed 2019 project (ANR-19-PERM-0007-04). We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research. This work was granted access to the HPC resources of IDRIS under the allocation 2019-101203 made by GENCI (Jean Zay computing center).

References

[1] D. Sarrut, N. Krah, and J. M. Létang, "Generative adversarial networks (GAN) for compact beam source modelling in Monte Carlo simulations," *Physics in Medicine and Biology*, vol. 64, p. 215004, Oct. 2019.

[2] D. Sarrut, A. Etxebeste, N. Krah, and J.-M. Létang, "Modeling complex particles

phase space with GAN for Monte Carlo SPECT simulations: A proof of concept," *Physics in Medicine and Biology*, 2021.

[3] A. Saporta, A. Etxebeste, T. Kaprelian, J. M. Létang, and D. Sarrut, "Modeling families of particle distributions with conditional GAN for Monte Carlo SPECT simulations," *Physics in Medicine & Biology*, vol. 67, p. 234001, Dec. 2022.

[4] F. Mentzel, K. Kroninger, M. Lerch, O. Nackenhorst, J. Paino, A. Rosenfeld, A. Saraswati, A. C. Tsoi, J. Weingarten, M. Hagenbuchner, and S. Guatelli, "Fast and accurate dose predictions for novel radiotherapy treatments in heterogeneous phantoms using conditional 3D-UNet generative adversarial networks," *Medical Physics*, vol. n/a, no. n/a, 2022.

[5] M. Paganini, L. de Oliveira, and B. Nachman, "CaloGAN: Simulating 3D High Energy Particle Showers in Multi-Layer Electromagnetic Calorimeters with Generative Adversarial Networks," *Physical Review D*, vol. 97, p. 014021, Jan. 2018.

[6] R. Di Sipio, M. F. Giannelli, S. K. Haghighat, and S. Palazzo, "DijetGAN: A Generative-Adversarial Network approach for the simulation of QCD dijet events at the LHC," *Journal of High Energy Physics*, vol. 2019, p. 110, Aug. 2019.

[7] D. Guest, K. Cranmer, and D. Whiteson, "Deep Learning and its Application to LHC Physics," *Annual Review of Nuclear and Particle Science*, vol. 68, pp. 161–181, Oct. 2018.

[8] A. Radovic, M. Williams, D. Rousseau, M. Kagan, D. Bonacorsi, A. Himmel, A. Aurisano, K. Terao, and T. Wongjirad, "Machine learning at the energy and intensity frontiers of particle physics," *Nature*, vol. 560, pp. 41–48, Aug. 2018.

[9] J. Albrecht, A. A. Alves, G. Amadio, G. Andronico, N. Anh-Ky, L. Aphecetche, J. Apostolakis, M. Asai, L. Atzori, M. Babik, G. Bagliesi, M. Bandieramonte, S. Banerjee, M. Barisits, L. A. T. Bauerdick, S. Belforte, D. Benjamin, C. Bernius, W. Bhimji, R. M. Bianchi, I. Bird, C. Biscarat, J. Blomer, K. Bloom, T. Boccali, B. Bockelman, T. Bold, D. Bonacorsi, A. Boveia, C. Bozzi, M. Bracko, D. Britton, A. Buckley, P. Buncic, P. Calafiura, S. Campana, P. Canal, L. Canali, G. Carlino, N. Castro, M. Cattaneo, G. Cerminara, J. Cervantes Villanueva, P. Chang, J. Chapman, G. Chen, T. Childers, P. Clarke, M. Clemencic, E. Cogneras, J. Coles, I. Collier, D. Colling, G. Corti, G. Cosmo, D. Costanzo, B. Couturier, K. Cranmer, J. Cranshaw, L. Cristella, D. Crooks, S. Crépé-Renaudin, R. Currie, S. Dallmeier-Tiessen, K. De, M. De Cian, A. De Roeck, A. Delgado Peris, F. Derue, A. Di Girolamo, S. Di Guida, G. Dimitrov, C. Doglioni, A. Dotti, D. Duellmann, L. Duflot, D. Dykstra, K. Dziedziniewicz-Wojcik, A. Dziurda, U. Egede, P. Elmer, J. Elmsheuser, V. D. Elvira, G. Eulisse, S. Farrell, T. Ferber, A. Filipcic, I. Fisk, C. Fitzpatrick, J. Flix, A. Formica, A. Forti, G. Franzoni, J. Frost, S. Fuess, F. Gaede, G. Ganis, R. Gardner, V. Garonne, A. Gellrich, K. Genser, S. George, F. Geurts, A. Gheata, M. Gheata, F. Giacomini, S. Giagu, M. Giffels, D. Gingrich, M. Girone, V. V. Gligorov, I. Glushkov, W. Gohn, J. B. Gonzalez Lopez, I. González Caballero, J. R. González Fernández, G. Govi, C. Grandi, H. Grasland, H. Gray, L. Grillo,

W. Guan, O. Gutsche, V. Gyurjyan, A. Hanushevsky, F. Hariri, T. Hartmann, J. Harvey, T. Hauth, B. Hegner, B. Heinemann, L. Heinrich, A. Heiss, J. M. Hernández, M. Hildreth, M. Hodgkinson, S. Hoeche, B. Holzman, P. Hristov, X. Huang, V. N. Ivanchenko, T. Ivanov, J. Iven, B. Jashal, B. Jayatilaka, R. Jones, M. Jouvin, S. Y. Jun, M. Kagan, C. W. Kalderon, M. Kane, E. Karavakis, D. S. Katz, D. Kcira, O. Keeble, B. P. Kersevan, M. Kirby, A. Klimentov, M. Klute, I. Komarov, D. Konstantinov, P. Koppenburg, J. Kowalkowski, L. Kreczko, T. Kuhr, R. Kutschke, V. Kuznetsov, W. Lampl, E. Lancon, D. Lange, M. Lassnig, P. Laycock, C. Leggett, J. Letts, B. Lewendel, T. Li, G. Lima, J. Linacre, T. Linden, M. Livny, G. Lo Presti, S. Lopienski, P. Love, A. Lyon, N. Magini, Z. L. Marshall, E. Martelli, S. Martin-Haugh, P. Mato, K. Mazumdar, T. McCauley, J. McFayden, S. McKee, A. McNab, R. Mehdiyev, H. Meinhard, D. Menasce, P. Mendez Lorenzo, A. S. Mete, M. Michelotto, J. Mitrevski, L. Moneta, B. Morgan, R. Mount, E. Moyse, S. Murray, A. Nairz, M. S. Neubauer, A. Norman, S. Novaes, M. Novak, A. Oyanguren, N. Ozturk, A. Pacheco Pages, M. Paganini, J. Pansanel, V. R. Pascuzzi, G. Patrick, A. Pearce, B. Pearson, K. Pedro, G. Perdue, A. Perez-Calero Yzquierdo, L. Perrozzi, T. Petersen, M. Petric, A. Petzold, J. Piedra, L. Piilonen, D. Piparo, J. Pivarski, W. Pokorski, F. Polci, K. Potamianos, F. Psihas, A. Puig Navarro, G. Quast, G. Raven, J. Reuter, A. Ribon, L. Rinaldi, M. Ritter, J. Robinson, E. Rodrigues, S. Roiser, D. Rousseau, G. Roy, G. Rybkine, A. Sailer, T. Sakuma, R. Santana, A. Sartirana, H. Schellman, J. Schovancová, S. Schramm, M. Schulz, A. Sciabà, S. Seidel, S. Sekmen, C. Serfon, H. Severini, E. Sexton-Kennedy, M. Seymour, D. Sgalaberna, I. Shapoval, J. Shiers, J.-G. Shiu, H. Short, G. P. Siroli, S. Skipsey, T. Smith, S. Snyder, M. D. Sokoloff, P. Spentzouris, H. Stadie, G. Stark, G. Stewart, G. A. Stewart, A. Sánchez, A. Sánchez-Hernández, A. Taffard, U. Tamponi, J. Templon, G. Tenaglia, V. Tsulaia, C. Tunnell, E. Vaandering, A. Valassi, S. Vallecorsa, L. Valsan, P. Van Gemmeren, R. Vernet, B. Viren, J.-R. Vlimant, C. Voss, M. Votava, C. Vuosalo, C. Vázquez Sierra, R. Wartel, and The HEP Software Foundation, "A Roadmap for HEP Software and Computing R&D for the 2020s," *Computing and Software for Big Science*, vol. 3, p. 7, Mar. 2019.

[10] D. Bourilkov, "Machine and deep learning applications in particle physics," *International Journal of Modern Physics A*, vol. 34, p. 1930019, Dec. 2019.

[11] K. Albertsson, P. Altoe, D. Anderson, J. Anderson, M. Andrews, J. P. A. Espinosa, A. Aurisano, L. Basara, A. Bevan, W. Bhimji, D. Bonacorsi, B. Burkle, P. Calafiura, M. Campanelli, L. Capps, F. Carminati, S. Carrazza, Y.-f. Chen, T. Childers, Y. Coadou, E. Coniavitis, K. Cranmer, C. David, D. Davis, A. De Simone, J. Duarte, M. Erdmann, J. Eschle, A. Farbin, M. Feickert, N. F. Castro, C. Fitzpatrick, M. Floris, A. Forti, J. Garra-Tico, J. Gemmler, M. Girone, P. Glaysher, S. Gleyzer, V. Gligorov, T. Golling, J. Graw, L. Gray, D. Greenwood, T. Hacker, J. Harvey, B. Hegner, L. Heinrich, U. Heintz, B. Hooberman, J. Junggeburth, M. Kagan, M. Kane, K. Kanishchev, P. Karpiński, Z. Kassabov, G. Kaul, D. Kcira, T. Keck, A. Klimentov, J. Kowalkowski, L. Kreczko, A. Kurepin, R. Kutschke, V. Kuznetsov,

N. Köhler, I. Lakomov, K. Lannon, M. Lassnig, A. Limosani, G. Louppe, A. Mangu, P. Mato, N. Meenakshi, H. Meinhard, D. Menasce, L. Moneta, S. Moortgat, M. Neubauer, H. Newman, S. Otten, H. Pabst, M. Paganini, M. Paulini, G. Perdue, U. Perez, A. Picazio, J. Pivarski, H. Prosper, F. Psihas, A. Radovic, R. Reece, A. Rinkevicius, E. Rodrigues, J. Rorie, D. Rousseau, A. Sauers, S. Schramm, A. Schwartzman, H. Severini, P. Seyfert, F. Siroky, K. Skazytkin, M. Sokoloff, G. Stewart, B. Stienen, I. Stockdale, G. Strong, W. Sun, S. Thais, K. Tomko, E. Upfal, E. Usai, A. Ustyuzhanin, M. Vala, J. Vasel, S. Vallecorsa, M. Verzetti, X. Vilasís-Cardona, J.-R. Vlimant, I. Vukotic, S.-J. Wang, G. Watts, M. Williams, W. Wu, S. Wunsch, K. Yang, and O. Zapata, "Machine Learning in High Energy Physics Community White Paper," *arXiv:1807.02876 [hep-ex, physics:physics, stat]*, May 2019.

[12] D. Sarrut, A. Etxebeste, E. Muñoz, N. Krah, and J. M. Létang, "Artificial Intelligence for Monte Carlo Simulation in Medical Physics," *Frontiers in Physics*, vol. 9, p. 601, 2021.

[13] X. Mougeot, "Reliability of usual assumptions in the calculation of beta and nu spectra," *Physical Review C*, vol. 91, p. 055504, May 2015.

[14] M. Mirza and S. Osindero, "Conditional Generative Adversarial Nets," *arXiv:1411.1784 [cs, stat]*, Nov. 2014.

[15] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-Image Translation with Conditional Adversarial Networks," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5967–5976, July 2017.

[16] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML'17, (Sydney, NSW, Australia), pp. 214–223, JMLR.org, Aug. 2017.

[17] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, "Improved Training of Wasserstein GANs," in *Advances in Neural Information Processing Systems*, vol. 30, Curran Associates, Inc., 2017.

[18] A. Jolicoeur-Martineau and I. Mitliagkas, "Connections between Support Vector Machines, Wasserstein distance and gradient-penalty GANs," *arXiv:1910.06922 [cs, stat]*, Oct. 2019.

[19] H. Petzka, A. Fischer, and D. Lukovnicov, "On the regularization of Wasserstein GANs," *arXiv:1709.08894 [cs, stat]*, Mar. 2018.

[20] T. Tieleman and G. Hinton, "Lecture 6.5-RMSprop: Divide the gradient by a running average of its recent magnitude," *COURSERA: Neural networks for machine learning*, vol. 4, no. 2, pp. 26–31, 2012.

[21] S. I. Kwon, R. Ota, E. Berg, F. Hashimoto, K. Nakajima, I. Ogawa, Y. Tamagawa, T. Omura, T. Hasegawa, and S. R. Cherry, "Ultrafast timing enables reconstruction-free positron emission imaging," *Nature Photonics*, pp. 1–5, Oct. 2021.

[22] D. Sarrut, M. Bala, M. Bardiès, J. Bert, M. Chauvin, K. Chatzipapas, M. Dupont, A. Etxebeste, L. M. Fanchon, S. Jan, G. Kayal, A. S. Kirov, P. Kowalski, W. Krzemien, J. Labour, M. Lenz, G. Loudos, B. Mehadji, L. Ménard, C. Morel, P. Papadimitroulas, M. Rafecas, J. Salvadori, D. Seiter, M. Stockhoff, E. Testa, C. Trigila, U. Pietrzyk, S. Vandenberghe, M.-A. Verdier, D. Visvikis, K. Ziemons, M. Zvolský, and E. Roncali, "Advanced Monte Carlo simulations of emission tomography imaging systems with GATE," *Physics in Medicine and Biology*, vol. 66, p. 10TR03, May 2021.

[23] D. Sarrut, N. Arbor, T. Baudier, D. Borys, A. Etxebeste, H. Fuchs, J. Gajewski, L. Grevillot, S. Jan, G. C. Kagadis, H. G. Kang, A. Kirov, O. Kochebina, W. Krzemien, A. Lomax, P. Papadimitroulas, C. Pommranz, E. Roncali, A. Rucinski, C. Winterhalter, and L. Maigne, "The OpenGATE ecosystem for Monte Carlo simulation in medical physics," *Physics in Medicine and Biology*, vol. 67, Sept. 2022.

[24] K. Ramonaheng, J. A. van Staden, and H. du Raan, "The effect of calibration factors and recovery coefficients on 177Lu SPECT activity quantification accuracy: A Monte Carlo study," *EJNMMI Physics*, vol. 8, p. 27, Mar. 2021.

[25] S. Bond-Taylor, A. Leach, Y. Long, and C. G. Willcocks, "Deep Generative Modelling: A Comparative Review of VAEs, GANs, Normalizing Flows, Energy-Based and Autoregressive Models," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, pp. 7327–7347, Nov. 2022.