



Generative adversarial networks (GAN) for compact beam source modelling in Monte Carlo simulations

David Sarrut, Nils Krah, Jean-Michel Létang

► To cite this version:

David Sarrut, Nils Krah, Jean-Michel Létang. Generative adversarial networks (GAN) for compact beam source modelling in Monte Carlo simulations. *Physics in Medicine and Biology*, 2019, 64 (21), 10.1088/1361-6560/ab3fc1 . hal-02276243

HAL Id: hal-02276243

<https://hal.science/hal-02276243>

Submitted on 8 Nov 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Generative Adversarial Networks (GAN) for compact beam source modelling in Monte Carlo simulations

D. Sarrut¹, N. Krah^{1,2}, JM. Létang¹

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

² University of Lyon, CNRS/IN2P3, Institute of Nuclear Physics, UMR 5822, Villeurbanne, France

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

Abstract.

A method is proposed and evaluated to model large and inconvenient phase space files used in Monte Carlo simulations by a compact Generative Adversarial Network (GAN). The GAN is trained based on a phase space dataset to create a neural network, called Generator (G), allowing G to mimic the multidimensional data distribution of the phase space. At the end of the training process, G is stored with about 0.5 million weights, around 10MB, instead of few GB of the initial file. Particles are then generated with G to replace the phase space dataset.

This concept is applied to beam models from linear accelerators (linacs) and from brachytherapy seed models. Simulations using particles from the reference phase space on one hand and those generated by the GAN on the other hand were compared. 3D distributions of deposited energy obtained from source distributions generated by the GAN were close to the reference ones, with less than 1% of voxel-by-voxel relative difference. Sharp parts such as the brachytherapy emission lines in the energy spectra were not perfectly modeled by the GAN. Detailed statistical properties and limitations of the GAN-generated particles still require further investigation, but the proposed exploratory approach is already promising and paves the way for a wide range of applications.

1. Introduction

Monte Carlo simulations are widely used to characterize sources of particles such as those of linac photon/electron beams, X-ray tubes, proton beam nozzles, brachytherapy radionuclide seeds, particles emerging from a voxelised patient geometry (to simulate a nuclear imaging process), etc. The computation time to perform such simulations is generally high and phase space files have been acknowledged as a necessary means to avoid repeated and redundant execution of part of the simulation. A typical example is the dose calculation in a patient CT image where the simulation is split into two parts [1]. A first, detailed, Monte Carlo simulation is performed to transport particles

through the accelerator treatment head elements (primary collimation, flattening filter, monitor chambers, mirrors, secondary collimation, etc.), up to a virtual plane. The properties (energy, position, direction) of all particles reaching the plane are stored in the phase space file and depend on the detailed properties of the treatment head components, such as its shape and materials. A second simulation tracks particles from the phase space plane through the multi-leaf collimator and the patient CT image to estimate the absorbed dose distribution.

However, phase space files are typically up to several tens of gigabytes large and inconvenient to use efficiently. Statistical limitations may also arise when more particles are required than stored in the phase space file. Several virtual source models for linac beam modelling have been proposed in the literature. For example in [2, 3], the authors describe the statistical properties of the phase space distribution by analytical functions, by evaluating the dependence of the parameters and by adapted sampling procedures. Chabert et al. [4] used 4D correlated histograms with different adaptive binning schemes to represent an Elekta Synergy 6 MV photon beam. Recently, Brualla et al. [5] proposed a method to extract light-weight spectra from large phase space files. This method may be sufficient for some applications, but may neglect correlations between energy, position and direction and require adaptive binning. Overall, the proposed methods were useful, but are specific to a simulation type and may not be applied in applications other than linac simulations. In this work, we propose the use of GAN as a generic beam modelling technique.

Generative Adversarial Networks (GAN) were recently reported [6, 7] as deep neural network architectures which allow mimicking a distribution of multidimensional data and have gained large popularity due to their success in realistic image synthesis. GAN learn representations of a training dataset by implicitly modeling high-dimensional distributions of data. Once trained, the resulting model is a neural network called Generator G that produces an element $\mathbf{x}$ that is supposed to belong to the underlying probability distribution of the training data. Based on this principle, several variants have been proposed such as the Wasserstein GAN (WGAN) [8] that will be used in this work. GAN and derivatives are still emerging concepts and an active field of research.

The idea to use machine learning techniques for Monte Carlo simulations and high energy physics is not new. Several works have been reported such as the use of neural networks in condensed matter physics [9, 10] or the use of GAN [11] for fast simulation of particle showers in electro-magnetic calorimeters. However, to our knowledge, GANs have never been proposed for representing large phase space distributions.

2. Materials & Methods

2.1. Training data

A phase space dataset generated by a MC simulation contains a set of particles described by position, direction cosines, energy, statistical weight, etc. The dataset is intended

to be as large as possible in order to be representative of the underlying probability distribution of the particle source, denoted as p_{real} . A phase space file is an extensive finite description of multidimensional data and thus may suffer from residual uncertainty known as latent variance [12]. It is used as a training dataset with samples $\mathbf{x} \in \mathbb{R}^d$ of dimension d . In practice, d is typically equal to 7, with 3 parameters for the particle position, 3 for the direction and 1 for the energy.

2.2. GAN optimisation

In the following, we first introduce the general concept behind a GAN and then summarize the aspects specific to a WGAN, a special category of GAN used in this work.

The goal is to train a generative function G that models a distribution $p_{\boldsymbol{\theta}}$. Parameters $\boldsymbol{\theta}$ are related to the distribution model approximating the target distribution p_{real} only known by samples from a training dataset [6]. The neural network architecture is composed of two multilayer perceptrons, D and G , competing against each other, hence the term adversarial. The *generator* $G(\mathbf{z}; \boldsymbol{\theta}_G)$ is trained to produce samples distributed similarly as the data distribution of $\mathbf{x}$. It takes $\mathbf{z}$ as input, sampled from a simple multidimensional normal prior distribution, $\mathcal{N}^d(0, 1)$, and produces a sample $\mathbf{x}$ as if it were drawn from p_{real} . The parameters $\boldsymbol{\theta}_G$ are the weights of the network G . The *discriminator* $D(\mathbf{x}; \boldsymbol{\theta}_D)$ is trained to distinguish between samples from the real data distribution and those generated by G . It takes $\mathbf{x}$ as input and yield a single scalar as output that represents the probability of $\mathbf{x}$ coming from the real data rather than from the generator. D is trained to maximize the probability of correctly identifying samples from the training data as real and those generated by G as fake. The parameters $\boldsymbol{\theta}_D$ are the weights of the network D .

The GAN training process is a zero-sum non-cooperative game which converges when the discriminator and the generator reach Nash equilibrium [13]. At such an equilibrium, one player (neural network) will not change its action (weights) regardless of what the opponent (the other network) may do. In the conventional GAN formulation [6], the considered cost function was the Binary Cross Entropy (BCE) for both G and D . $\text{BCE}(p, q)$ between two distributions p and q is related to the Kullback-Leibler divergence which measures the performance of a classification model whose output is a probability value between 0 and 1. The loss function of a GAN quantifies the similarity between the data distribution generated by G and the real sample distribution and it has been shown that this corresponds to the Jensen-Shannon divergence (JSD) when the discriminator is optimal [6]. JSD is a symmetric and smooth version of the Kullback-Leibler divergence.

In practice, we found a conventional GAN difficult to train and subject to mode collapse [8]. Instead, we used WGAN, proposed by Arjovsky et al. [8], which uses the Wasserstein (or Earth Mover's) distance as an alternative loss function. The Wasserstein distance between two distributions p and q is the cost of the optimal transport needed

to deform p into q . It has been shown that it helps stabilizing the learning process, because it is less subject to vanishing gradients than a conventional GAN. In practice, there are few changes compared to the original GAN. First, the loss functions become the following expressions:

$$J_D(\boldsymbol{\theta}_D, \boldsymbol{\theta}_G) = \mathbb{E}_{\mathbf{z}} [D(G(\mathbf{z}))] - \mathbb{E}_{\mathbf{x}} [D(\mathbf{x})] \quad (1)$$

$$J_G(\boldsymbol{\theta}_D, \boldsymbol{\theta}_G) = -\mathbb{E}_{\mathbf{z}} [D(G(\mathbf{z}))] \quad (2)$$

Second, after every gradient update, the weights $\boldsymbol{\theta}_D$ are clamped to a small fixed range (e.g. $[-0.01, 0.01]$) in order to constrain weights to a compact space. Finally, the authors [8] also recommend using the RMSProp optimizer [14] instead of the conventional Adam optimizer [15] because the latter uses momentum processes that may cause instability in the model training. In a WGAN, D does not act as a explicit discriminator, but is a helper for approximating the Wasserstein metric between real and generated data distributions. D is called the “critic”. The training is no more performed until Nash equilibrium, but until loss convergence.

2.3. GAN architecture and parameters

A GAN requires setting several interconnected hyperparameters that influence results in different ways. We describe here the optimal set of parameters we empirically found. We comment on their influence we observed in the result section. The architecture of both D and G was the following: we used $H = 400$ neurons in each of the 3 hidden linear fully connected layers. The values were set empirically, based on experimental results. As advocated in [8], the activation function was a Rectified Linear Unit (ReLU) $r(x) = \max(0, x)$ for all layers except for the last one of G where instead a sigmoid function was used. We set the dimension of $\mathbf{z}$ to 6. The total number of weights was around 5×10^5 for both D and G . The learning rate was chosen empirically and set to 10^{-5} . Stochastic batches of 10^4 samples were used at each iteration. The discriminator was updated more frequently than the generator, four times per iteration versus once, as advocated in [8]. We set the number of iterations (or *epochs* in the deep learning community) to 80 000.

2.4. Implementation

All simulations were implemented in Gate version 8.0 [16], using Geant4 version 10.3 [17]. Neural network operations (training, samples generation) were implemented in Python with the PyTorch framework [18] using CUDA GPU acceleration. Once the network is trained, it can be read and used within Gate/Geant4 thanks to a newly developed Gate module exploiting the PyTorch C++ API [19]. Source code is open-source and will be available in the next Gate release. Computations were performed on an Intel Xeon CPU E5-2640 v4 @ 2.40GHz and an NVIDIA Titan Xp (GP102-450-A1) with 12 GB memory.

2.5. Experiments

The proposed method was evaluated with phase space files of two linear accelerators provided by IAEA [20] and of brachytherapy ^{125}I seeds [21, 16]. The linacs are a 6 MV photon Elekta and a Cyberknife with IRIS collimator of 60 mm. For the linac phase space files, the particles were parameterized by energy E , position x, y, z , and direction cosines dx, dy, dz . Both phase space files were recorded in a plane, so the z coordinate was constant and ignored in the following, leading to 6 dimensions. The phase space distribution are available as two files (see table 1): one file was used to train the GAN (PHSP₁); the other file was used for evaluation (PHSP₂). To generate a purely continuous energy distribution, the phase space datasets were pre-processed to remove the 511 keV peak. The phase space files of the ^{125}I brachytherapy seed (59.49 days half life, maximum 35 keV) were obtained from simulations of a seed capsule model composed of a double-wall made of titanium surrounding a tungsten x-ray marker coated with an organic carbon layer [16] (The Best Medical model 2301 source, Best Medical International, Springfield, VA). The energy distribution of gammas exiting the capsule is mainly composed of three emission lines (around 4.47, 4.9 and 27.4 keV). In that sense, the brachytherapy example is complementary to the linac examples. Two files of 2.5 GB were generated, each containing 1.04×10^8 gammas, described by 7 parameters each (energy, position, direction). One file was used to train the GAN and the other one for evaluation.

For the examples considered in this work, the quantity of clinical interest is the accuracy of the dose calculation. For the linac tests, MC simulations were run in order to score the deposited energy in a $20 \times 20 \times 20 \text{ cm}^3$ water box with voxels of $4 \times 4 \times 4 \text{ mm}^3$ (Elekta) and $2 \times 2 \times 2 \text{ mm}^3$ (Cyberknife). As source, we used either a phase space file or GAN-generated particles (PHPS_{GAN}). 10^8 primary photons were used for the Elekta and 4×10^7 for the Cyberknife. Particles in the phase space files were not used multiple times. For the brachytherapy test, 79 seeds were evenly placed in the prostate region of a CT image, emitting a total of 10^8 gammas, in one case taken from a phase space file and in the other generated by the GAN. The deposited energy was scored in $2 \times 2 \times 2 \text{ mm}^3$ voxels.

In all cases, the MC relative statistical uncertainty $\sigma(k) = S(k)/D(k)$ of the energy deposited in a voxel was computed with the history-by-history method [22], with $S(k)$ the statistical uncertainty in voxel with index k (eq. 3), N the total number of primary events in the simulation, $d_{k,i}$ the energy deposited in voxel k by event i , and $D(k) = \sum_i d_{k,i}$ the total deposited energy in voxel k . The statistical uncertainty is the standard error of the mean of the scored quantity, here the mean dose, and thus converges to zero. The obtained values for the dose map in the three tests was less than 3% for all voxels with more than 10% of the maximum dose.

PHSP	Size	Nb of particles
Elekta PRECISE 6MV	2 files of 3.9 GB	1.3×10^8 photons each file
CyberKnife IRIS 60mm	2 files of 1.6 GB	5.8×10^7 photons each file
^{125}I brachy seed	2 files of 2.5 GB	1.04×10^8 photons each file

Table 1. Characteristics of the three used datasets

$$S(k) = \sqrt{\frac{1}{N-1} \left(\frac{\sum_i^N d_{k,i}^2}{N} - \left(\frac{\sum_i^N d_{k,i}}{N} \right)^2 \right)} \quad (3)$$

In order to compare particles from a phase space file with those generated by the GAN, marginal distributions of all parameters were plotted. Simulations to compute the deposited energy in water using phase space and GAN-generated particles were compared by analyzing the voxel-by-voxel differences of the deposited energy. The distribution of voxel differences naturally contains uncertainty. We evaluated the similarity of this uncertainty between phase space and GAN generated data. If the GAN produces realistic a phase space distribution, the uncertainty should be similar in both cases. We thus compared the distribution of voxel-wise differences between two simulations performed with two different phase space files (Δ_{PHPS}), and between two simulations using particles from the GAN and a phase space file, respectively (Δ_{GAN}). The differences were normalised by the maximum value in the image, as a proxy for the prescribed dose, denoted $\hat{D}_{\text{PHPS}_2} = \max_{\{k\}} D(k)$, see equation 5.

$$\Delta_{\text{PHPS}}(k) = \frac{D_{\text{PHPS}_2}(k) - D_{\text{PHPS}_1}(k)}{\hat{D}_{\text{PHPS}_2}} \quad (4)$$

$$\Delta_{\text{GAN}}(k) = \frac{D_{\text{PHPS}_2}(k) - D_{\text{GAN}}(k)}{\hat{D}_{\text{PHPS}_2}} \quad (5)$$

Moreover, in every voxel, we computed the ratio between voxel difference and uncertainty. If the statistical error were normally distributed, the distributions of those ratios should have a zero mean and unit standard deviation. Finally, for the linac experiments, we computed depth dose curves (along z) and transversal dose profiles at 20mm depth. For the brachytherapy example, we focused on the energy distribution and visually inspected the distribution of deposited dose.

Results

As an example, figure 1 depicts the evolution of the loss J_D on training and validation datasets and J_G as a function of iterations (J_G does not depend on the kind of dataset). For visual clarity, we only depict data for the Elekta test case, subsampled to every 100 iterations (not subsampled in the magnified view). J_D (equation 1) is negative at

the beginning because the generator is not yet sufficiently trained, so that D applied to real data is larger than D applied to generated data. Once G is trained, J_D is close to zero. Furthermore, J_D is slightly less negative when evaluated on the validation dataset rather than training dataset because D retains the validation dataset as slightly less likely. The G-loss J_G depicted larger variation than the other losses but we did not investigate this further. Similar behavior was obtained for the other tests. Figures 2 and 4 display the marginal distributions of the 6 parameters (E, x, y, dx, dy, dz ; z was fixed) extracted from the phase space file compared with the ones obtained from the GAN. Figure 8 displays a closeup of the energy distribution. Note that the datasets used for training the GAN were always different from the ones used for validation. Figures 3 and 5 also illustrate the correlations (covariance normalised by the product of their standard deviations) between the 6 parameters for the two linac tests. The left panel in figure 6 shows the distribution of the relative differences Δ_{PHSP} and Δ_{GAN} for all three tests. The mean differences are indicated with vertical lines. The right hand panels show the distribution of the ratio between differences and uncertainty, which should ideally depict a mean value of zero and a standard deviation of one. Figures 7 show transversal and depth profiles of deposited energy for both tests. Figures 9 illustrate the deposited energy obtained for the brachytherapy test case from simulations with particles from phase space files and through GAN, respectively. The training process took around 2 hours for all tests and generation of 10^6 samples from the GAN took about one second. The final GAN model requires less than 10 MB of storage space.




Figure 1. J_D loss computed on training and validation datasets and J_G as a function of iterations, for the Elekta test case (subsamped every 100 iterations, excepted in the magnified view).

Training GAN is notoriously difficult: the models may not converge and mode collapses are common [8]. This was also our experience and, in the following, we




Figure 2. Marginal distributions of the 6 parameters obtained from the reference phase space file (PHSP) and from the GAN, for Elekta 6 MV linac.




Figure 3. Correlation matrices between all 6 parameters for phase space file (PHSP, left) and GAN (right), for Elekta 6 MV linac.

summarize our observations regarding the hyperparameters of the training process.

- We observed that the Wasserstein version of the GAN is required. With the conventional GAN formulation, we did not achieve acceptable results.
- The number of neurons ($H = 400$) and layers (3 hidden layers for each G and D)




Figure 4. Marginal distributions of the 6 parameters obtained from the reference phase space file (PHSP) and from the GAN, for the Cyberknife linac with IRIS collimator.




Figure 5. Correlation matrices between all 6 parameters for the phase space file (PHSP, left) and GAN (right), for Cyberknife linac.

were empirically set. Fewer neurons (around 300) lead to inferior results, while a larger number of neurons took more time to converge and did not seem to improve results.

- The dimension of the latent variable $\mathbf{z}$ is representative of the underlying latent

space structure of the multidimensional distribution. In this work, it was therefore fixed to 6 for linac tests and 7 for the brachytherapy test. We observed few differences with 5 or 7 dimensions. However, a too low value (lower than 5) lead to degraded results.

- According to [8], it is suggested to perform more discriminator updates than generator updates. Here, we selected four discriminator updates for one generator update. A lower ratio (2:1) degraded results.
- The learning rate was set to 10^{-5} and we did not find any clear improvement with other values. The batch size was empirically set to 10^4 . Smaller numbers of samples per batch lead to inaccurate probability densities in the generated distributions (which are 6/7-dimensional). This is probably because the multidimensional phase space is too sparsely represented by such samples to be faithfully learned. Larger batches did not really improve the results. Figure 1 shows a typical training process. The values of the loss function strongly oscillated during the first few thousands of iterations and then tended to slowly converge and oscillate around a fixed value. Repeated training with the same set of parameters lead to slightly different results because of the stochastic nature of the learning process.




Figure 6. Distributions of relative differences between PHSP1 and PHSP2 and between PHSP1 and GAN. Vertical lines indicate the mean differences (in black for PHSP1-PHSP2 and in red for PHSP1-GAN). Right images show the difference relative to the statistical uncertainty, this distribution should have zero mean and standard deviation of one. Top row for Elekta machine, center row for Cyberknife device, and bottom row for brachytherapy test. For this last plot, the peaks structure correspond to the seeds in the CT image.




Figure 7. Transverse profiles at 20mm depth (left) and depth profiles (right) deposited energy for Elekta (top) and Cyberknife (bottom) machines. Curves were obtained for PHSP2 and GAN-based simulations.




Figure 8. Closeup on the energy distribution obtained from the reference brachytherapy phase space (PHSP) and from the GAN.




Figure 9. Slices of CT prostate image with deposited energy overlay (in MeV), computed by phase space (PHSP, left) and GAN generated (center) particles. The right image shows the dose difference Δ_{GAN} relative to the maximum dose (the maximum difference was below 4%).

Discussion

The distributions generated by the GAN were close to those represented the phase space files as shown in figures 2, 3, 4, and 5. Note that the figures only display marginal distributions and correlation between pairs of parameters, but high order correlations were also modeled by the GAN. In particular, X-dX and Y-dY correlations are related to the cone geometry of the linac beam.

When the GAN was used to compute deposited energy in a water box, overall results were good and could not be distinguished visually or by looking at the transversal or depth profiles (figure 7). Still, differences between phase space and GAN based simulation results are slightly larger than differences between two simulations using phase space files. For example, with the Elekta linac, the distributions of differences show a small shift towards negative values, visible in the top-left panel of figure 6. This shift is however not dosimetrically relevant (less than 0.04% of the maximum dose). A similar effect is also visible in the distribution of the differences relative to the statistical uncertainty. The top and bottom right panels in figure 6 depict zero mean and unity standard deviation for Δ_{PHSP} , while small shifts (0.02 and 0.06) are present for Δ_{GAN} . We also verified that the difference distribution between two simulations, both using GAN generated particles, was very close to the one between two phase space based simulations (not shown in the figure because curves completely overlapped with the Δ_{PHSP} distribution). The last plots in figures 6 and 9 show the almost unnoticeable differences between the deposited energy distributions. As a conclusion, the experiments illustrate that a GAN generated particle distribution is close to the distribution represented by the phase space file, but not exactly equivalent: differences between GAN and phase space files are larger than between two phase space files. The difference can hardly be seen in the deposited energy here, but may be important for other applications.

For the brachytherapy case, figure 8 displays a close-up around the main energy peaks, illustrating how the GAN only approximates the peak values. It seems that sharp features in the distributions are still difficult to model with this method. Nevertheless, the dose maps were comparable. If an exact representation of the energy spectrum is required, a workaround could consist in training several GANs: one for each peak (ignoring the energy parameter) and one for the continuous part. We note that for the Elekta linac, the 511 keV peak represents less than 0.4% of the total number of photons.

As already mentioned in papers related to GAN [6, 8], training a GAN is still a difficult process, involving manual tuning of hyperparameters and empirical decisions. In this work, the same set of hyperparameters, given in section 2.3, led to consistent results for all three tested phase space files. The hyperparameter space is large and other more optimal hyperparameters than the ones used in the presented work might exist. Pruning techniques may also be employed in order to trim the network size by removing unimportant neurons to improve speed and results. Moreover, better results may be achieved with a different network architecture or training process. Findings

described here may serve as a first guide to further investigations.

Since the initial GAN proposal [6], numerous variants, among them WGAN, have been investigated in the literature with more than 500 papers per month during the end of 2018 [23]. Further works are still needed to evaluate the interest of those variants, in particular to investigate if the difficulty to precisely model sharp features of the distributions that have been described earlier could be overcome. Also transfer learning may be well adapted here: a first network trained for a given phase space could be used as a starting point for the training of another one. Note also that GANs are usually employed for problems with a higher number of dimensions and less smooth distributions (natural images, speech) than phase space data. Other methods, such as Gaussian mixture models [24], may also be useful to model phase space files.

Overall, the proposed approach has several advantages. It allows modelling a large file of several GB by one of about 10 MB (486406 neuron weights in the G model). The generation of particles from a GAN is a very fast process (1 second for 10^6 particles) and an arbitrary number of particles can be generated from the GAN while the phase space dataset is finite. The use of a generator instead a phase space file also greatly simplifies the simulation workflow. The infrastructure for training a GAN on a new phase space is generic and convenient to setup. Users may only need to set hyperparameters values. The learning process in this work involved only 10^8 particles, which means that computation time is preserved compared to a larger phase space. Further investigations might explore how accurately a GAN can be trained based on a smaller training dataset.

3. Conclusion

In this work, we proposed a GAN architecture to learn large phase space distributions used in Monte Carlo simulations of linacs and brachytherapy. The advantages of this approach lie in the compact size of the model, the ability to quickly generate a large number of particles, and the very generic nature of the process that may potentially be applied to a large range of sources of particles. Indeed, this approach could be extended to other types of simulations where complex distributions of particles are involved. We think that the exploratory work presented here is a first of a kind involving advanced machine learning methods for Monte Carlo simulations and can be applied to a large class of applications in medical physics.

Acknowledgments

This work was performed within the framework of the SIRIC LYriCAN Grant INCa-INSERM-DGOS-12563, and 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. We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research.

References

- [1] Andreo P. “The physics of small megavoltage photon beam dosimetry”. *Radiotherapy and Oncology*, 126(2):205–213. 2018.
- [2] Fix M.K., Stampanoni M., Manser P., Born E.J., Mini R. and Rügsegger P. “A multiple source model for 6 MV photon beam dose calculations using Monte Carlo.” *Physics in medicine and biology*, 46(5):1407–27. May 2001.
- [3] Grevillot L., Bertrand D., Dessy F., Freud N. and Sarrut D. “A Monte Carlo pencil beam scanning model for proton treatment plan simulation using GATE/GEANT4.” *Phys Med Biol*, 56(16):5203–5219. Aug 2011.
- [4] Chabert I., Barat E., Dautremer T., Montagu T., Agelou M., de Suray A.C., Garcia-Hernandez J., Gempp S., Benkreira M., De Carlan L. et al. “Development and implementation in the Monte Carlo code PENELOPE of a new virtual source model for radiotherapy photon beams and portal image calculation”. *Physics in Medicine & Biology*, 61(14):5215. 2016.
- [5] Brualla L., Rodriguez M., Sempau J. and Andreo P. “PENELOPE/PRIMO-calculated photon and electron spectra from clinical accelerators”. *Radiation Oncology*, 14(1):6. 2019.
- [6] Goodfellow I., Pouget-Abadie J., Mirza M., Xu B., Warde-Farley D., Ozair S., Courville A. and Bengio Y. “Generative adversarial nets”. In *Advances in neural information processing systems*, pages 2672–2680. 2014.
- [7] Creswell A., White T., Dumoulin V., Arulkumaran K., Sengupta B. and Bharath A.A. “Generative adversarial networks: An overview”. *IEEE Signal Processing Magazine*, 35(1):53–65. 2018.
- [8] Arjovsky M., Chintala S. and Bottou L. “Wasserstein GAN”. *arXiv preprint arXiv:1701.07875*. 2017.
- [9] Carrasquilla J. and Melko R.G. “Machine learning phases of matter”. *Nature Physics*, 13(5):431. 2017.
- [10] Shen H., Liu J. and Fu L. “Self-learning Monte Carlo with deep neural networks”. *Physical Review B*, 97(20):205140. 2018.
- [11] Paganini M., de Oliveira L. and Nachman B. “CaloGAN: Simulating 3D high energy particle showers in multilayer electromagnetic calorimeters with generative adversarial networks”. *Physical Review D*, 97(1):014021. 2018.
- [12] Sempau J., Sanchez-Reyes A., Salvat F., ben Tahar H.O., Jiang S. and Fernández-Varea J. “Monte Carlo simulation of electron beams from an accelerator head using PENELOPE”. *Physics in Medicine & Biology*, 46(4):1163. 2001.
- [13] Fedus W., Rosca M., Lakshminarayanan B., Dai A.M., Mohamed S. and Goodfellow I. “Many Paths to Equilibrium: GANs Do Not Need to Decrease aDivergence At Every Step”. *arXiv preprint arXiv:1710.08446*. 2018.
- [14] Tieleman T. and Hinton G. “Lecture 6.5-RMSprop: Divide the gradient by a running average of its recent magnitude”. *COURSERA: Neural networks for machine learning*, 4(2):26–31. 2012.
- [15] Kingma D.P. and Ba J. “Adam: A method for stochastic optimization”. *arXiv preprint arXiv:1412.6980*. 2014.
- [16] Sarrut D., Bardies M., Bousson N., Freud N., Jan S., Létang J.M., Loudos G., Maigne L., Marcatili S., Mauxion T. et al. “A review of the use and potential of the GATE Monte Carlo simulation code for radiation therapy and dosimetry applications”. *Medical physics*, 41(6Part1):064301. Jun 2014.
- [17] Allison J., Amako K., Apostolakis J., Arce P., Asai M., Aso T., Bagli E., Bagulya A., Banerjee S., Barrand G., Beck B., Bogdanov A., Brandt D., Brown J., Burkhardt H., Canal P., Ott D., Chauvie S., Cho K. and Yoshida H. “Recent developments in GEANT4”. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 835:186–225. 2016.
- [18] Paszke A., Gross S., Chintala S., Chanan G., Yang E., DeVito Z., Lin Z., Desmaison A., Antiga L. and Lerer A. “Automatic differentiation in PyTorch”. 2017.

- [19] PyTorch. “PyTorch C++ API”. Technical report, <https://pytorch.org/cppdocs>. 2019.
- [20] Capote R., Jeraj R., Ma C., Rogers D., Sánchez-Doblado F., Sempau J., Seuntjens J. and Siebers J. “Phase-space database for external beam radiotherapy. Summary report of a consultants’ meeting”. Technical report, International Atomic Energy Agency. 2006.
- [21] Thiam C.O., Breton V., Donnarieix D., Habib B. and Maigne L. “Validation of a dose deposited by low-energy photons using GATE/GEANT4”. *Phys Med Biol*, 53(11):3039–3055. Jun 2008.
- [22] Chetty I.J., Curran B., Cygler J.E., DeMarco J.J., Ezzell G., Faddegon B.a., Kawrakow I., Keall P.J., Liu H., Ma C.M.C., Rogers D., Seuntjens J., Sheikh-Bagheri D. and Siebers J.V. “Report of the AAPM Task Group No. 105: Issues associated with clinical implementation of Monte Carlo-based photon and electron external beam treatment planning”. *Medical Physics*, 34(12):4818. 2007.
- [23] Hindupur A. “The GAN Zoo”. Technical report, <https://github.com/hindupuravinash/the-gan-zoo>. (consulted June 2019).
- [24] Doerner E. and Caprile P. “Implementation of a double Gaussian source model for the BEAMnrc Monte Carlo code and its influence on small fields dose distributions.” *Journal of applied clinical medical physics*, 17:212–221. ISSN 1526-9914. September 2016.