

Generalizing Monte Carlo dose reconstruction beyond the training anatomy using conditional diffusion models [★]

Hong-Phuong Dang¹ and Richard Gozan^{1,2}

¹ CentraleSupélec, IETR, UMR CNRS 6164, Cesson-Sévigné, France

² ENSAI, Bruz, France

`hong-phuong.dang@centralesupelec.fr`

`richard.gozan@centralesupelec.fr`

Abstract. Accurate Monte Carlo (MC) dose calculation is essential in radiotherapy, but obtaining high-sampling (HS) dose maps requires a very large number of particle histories and remains computationally expensive. Deep learning has recently been investigated as an asymptotic variance reduction strategy to reconstruct HS dose maps from low-sampling (LS) simulations. However, most existing approaches formulate this task as deterministic image-to-image regression and may be sensitive to anatomical distribution shifts. This limitation is particularly important when HS references are scarce for a new anatomical site. In this work, we investigate conditional denoising diffusion probabilistic models (DDPMs) for MC dose reconstruction beyond the training anatomy. The proposed model learns the conditional distribution of the HS dose map given the corresponding LS simulation and planning computed tomography (CT) image. It reconstructs the HS dose map through an iterative stochastic denoising process rather than a single deterministic prediction. The model is trained on abdominal data and evaluated on abdominal test data and five unseen pelvic and head-and-neck slices. On the abdominal test set, it performs close to a deterministic U-Net. On the limited cross-anatomy set, the DDPM produces visually more faithful reconstructions in the examples considered. Geometric alignment further improves its results. Although preliminary, these findings motivate larger-scale evaluation of conditional diffusion models for MC dose reconstruction.

Keywords: Monte Carlo dose reconstruction · diffusion models · cross-anatomy generalization · dosimetry

1 Introduction

Accurate dose calculation is essential in radiotherapy, as it directly affects both tumor control and healthy tissue preservation. Monte Carlo (MC) simulations are among the most accurate methods for radiotherapy dose calculation because

[★] This work was partially funded by the French National Research Agency through the ExOTics project (ANR-25-CE45-3443-01)




Fig. 1. Evolution of dose maps over time during MC simulation

they explicitly model particle transport and interactions within patient anatomy. Their main limitation is computational cost. Due to the stochastic nature of MC simulations, obtaining statistically reliable high-sampling (HS) dose maps requires a very large number of particle histories, resulting in long computation times. In contrast, low-sampling (LS) simulations are fast but exhibit strong statistical fluctuations, resulting in high uncertainty in the estimated dose. Fig. 1 illustrates the progressive reduction of statistical noise as the number of simulated particle histories increases.

Many variance reduction techniques (VRTs) have been proposed to improve the efficiency of MC simulations without introducing estimation bias [1, 4, 6]. Hardware acceleration strategies, including GPU implementations and massively parallel computing, have also considerably reduced computation times [2, 10]. Despite these advances, HS MC simulations remain computationally demanding.

Deep learning (DL) has recently emerged as a promising asymptotic VRT (aVRT). Instead of obtaining an HS dose map solely by increasing the number of simulated particles, a neural network is trained to reconstruct the HS map from a fast LS simulation, often using the planning CT image as anatomical context. Several approaches have shown substantial acceleration of MC dose calculations, in some cases approaching real-time inference [9, 7, 16, 15]. Among these methods, convolutional neural networks (CNNs), particularly U-Net-based architectures, have become the most widely adopted framework for supervised LS-to-HS reconstruction [8, 13, 9, 12, 11, 3]. Most of these methods remain deterministic. Given a pair (CT, LS), the network outputs a single estimate of the HS. This deterministic formulation is effective when training and testing anatomies follow the same distribution, but it may be sensitive to anatomical distribution shifts. This limitation is particularly critical in radiotherapy. Generating HS MC dose maps remains computationally expensive, and large paired LS/HS datasets are difficult to obtain for every anatomical site, treatment configuration, or clinical protocol. A model trained on one anatomy should therefore ideally transfer to other anatomies with minimal or no additional adaptation. DL-based surrogates also depend strongly on the representativeness of the training data. Explicitly incorporating prior information, such as physical constraints or medical expert knowledge, also remains challenging in purely data-driven architectures. Finally, deterministic regression does not directly provide a mechanism for sampling plausible reconstructions or estimating reconstruction uncertainty.

Diffusion models (DMs) offer a different paradigm. Rather than learning a direct deterministic mapping from (CT, LS) to HS, they learn a conditional denoising process that progressively transforms Gaussian noise into a plausible HS dose map. This iterative stochastic formulation may learn representations that are less tied to a single anatomical distribution than standard image-to-image regression. It also provides a natural framework for future posterior sampling and uncertainty estimation. Despite their success in natural and medical image generation, DMs remain largely unexplored for MC dose reconstruction.

In this work, we investigate conditional denoising diffusion probabilistic models (DDPMs) for MC dose reconstruction from LS simulations and CT images. Our primary objective is to investigate the behavior of DMs beyond the anatomical distribution observed during training while assessing their reconstruction accuracy on the training anatomy. We formulate MC dose reconstruction as a conditional DDPM problem, where HS dose maps are sampled conditionally on the corresponding LS dose map and CT image. We conduct a preliminary evaluation in a cross-anatomy setting. Training is performed exclusively on abdominal CT data. Testing includes abdominal data and previously unseen pelvic and head-and-neck anatomies. We also investigate the role of geometric consistency through test-time alignment. Our results show that conditional DMs achieve reconstruction performance close to that of a deterministic U-Net on the training anatomy. On the limited cross-anatomy test set, the qualitative results suggest promising behavior on previously unseen anatomies.

The paper is organized as follows. Section 2 formulates the MC dose reconstruction problem. Section 3 presents the conditional DDPM. Section 4 describes the experimental setup and reports the results on abdominal and cross-anatomy data. Section 5 concludes the paper and discusses future work.

2 Problem statement

The objective of this work is to reconstruct HS dose maps from LS simulations while exploiting the anatomical information provided by the planning CT image. Let $\mathcal{D} = \{(\text{CT}_i, \text{LS}_i, \text{HS}_i)\}_{i=1}^N$ denote a dataset of triplets. The LS and HS dose maps are generated using the same MC simulation setup. They differ only in the number of simulated particle histories. Most existing DL approaches formulate MC dose reconstruction as a deterministic image-to-image regression problem. Given a pair (CT, LS), a neural network f_θ is trained to produce a deterministic estimate of the corresponding HS dose map :

$$f_\theta : (\text{CT}, \text{LS}) \mapsto \widehat{\text{HS}} \quad (1)$$

The parameters θ are obtained by minimizing a supervised reconstruction loss ℓ , typically the mean absolute error (MAE) or the mean squared error (MSE),

$$\min_{\theta} \sum_{i=1}^N \ell(f_\theta(\text{CT}_i, \text{LS}_i), \text{HS}_i). \quad (2)$$

We instead formulate HS dose reconstruction as a stochastic iterative process conditioned on the LS dose map and the corresponding CT image. The CT




Fig. 2. Forward corruption process q and conditional backward denoising process p_θ used for MC dose reconstruction.

provides the anatomical geometry and tissue composition governing radiation transport. The LS provides an initial noisy estimate of the deposited dose. Both inputs therefore contain complementary information for reconstructing the HS.

We also consider a cross-anatomy setting. The model is trained exclusively on abdominal data and evaluated on both abdominal data and previously unseen pelvic and head-and-neck anatomies. This protocol provides a preliminary evaluation beyond the training anatomy.

3 Methodology

3.1 Conditional DDPM for Monte Carlo dose reconstruction

We formulate MC dose reconstruction as a conditional diffusion problem. Given an LS dose map and the corresponding planning CT image, the objective is to model the conditional distribution of the associated HS dose map. The proposed approach relies on denoising diffusion probabilistic models (DDPMs). A fixed forward process progressively corrupts an HS dose map by adding Gaussian noise. A neural network is then trained to approximate the corresponding backward process. During sampling, the backward diffusion process starts from Gaussian noise and progressively generates an HS dose map conditioned on the LS dose map and the corresponding CT image. The LS dose map provides a noisy observation of the deposited dose, while the CT image provides information about the anatomical geometry and tissue composition. Both conditioning variables remain fixed throughout the backward diffusion process. Fig. 2 illustrates the forward corruption process and the conditional backward denoising process.

3.2 Forward and backward diffusion processes

Let x_0 denote a reference HS Monte Carlo dose map. During the forward diffusion process, x_t , with $t \in \{1, \dots, T\}$, denotes the noisy sample obtained after t diffusion steps. Gaussian noise is progressively added according to a predefined variance schedule $\{\beta_t\}_{t=1}^T$. After T steps, x_T approximately follows a standard Gaussian distribution,

$$x_T \sim \mathcal{N}(0, \mathbf{I}). \quad (3)$$

The forward diffusion process is defined as

$$q(x_t | x_{t-1}) = \mathcal{N}\left(x_t; \sqrt{1 - \beta_t} x_{t-1}, \beta_t \mathbf{I}\right). \quad (4)$$

Algorithm 1 Training of the conditional DDPM

-
- 1: **repeat**
 - 2: Sample $(x_0, \text{LS}, \text{CT})$ from the training dataset $\mathcal{D}$
 - 3: Sample $t \sim \text{Uniform}(\{1, \dots, T\})$
 - 4: Sample $\epsilon \sim \mathcal{N}(0, \mathbf{I})$
 - 5: Compute

$$x_t = \sqrt{\bar{\alpha}_t} x_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$$

- 6: Take a gradient descent step on

$$\|\epsilon - \epsilon_\theta(x_t, t, \text{LS}, \text{CT})\|_2^2$$

- 7: **until** convergence
-

Algorithm 2 Sampling with the conditional DDPM

-
- 1: Sample $x_T \sim \mathcal{N}(0, \mathbf{I})$
 - 2: **for** $t = T, \dots, 1$ **do**
 - 3: Sample $z \sim \mathcal{N}(0, \mathbf{I})$ if $t > 1$, otherwise set $z = 0$
 - 4: Compute

$$x_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, t, \text{LS}, \text{CT}) \right) + \sigma_t z$$

- 5: **end for**
 - 6: **return** x_0
-

It admits the closed-form expression

$$q(x_t | x_0) = \mathcal{N}(x_t; \sqrt{\bar{\alpha}_t} x_0, (1 - \bar{\alpha}_t) \mathbf{I}), \quad (5)$$

where $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$.

The backward diffusion process progressively reconstructs x_0 from x_T . In the conditional formulation, the backward transition is defined as

$$p_\theta(x_{t-1} | x_t, \text{LS}, \text{CT}) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t, \text{LS}, \text{CT}), \sigma_t^2 \mathbf{I}), \quad (6)$$

where

$$\begin{aligned} \mu_\theta(x_t, t, \text{LS}, \text{CT}) &= \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, t, \text{LS}, \text{CT}) \right), \\ \sigma_t^2 &= \tilde{\beta}_t = \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \beta_t. \end{aligned}$$

The network ϵ_θ predicts the Gaussian noise added during the forward process. The noisy sample x_t evolves throughout the backward process. The LS dose map and CT image remain fixed and serve as conditioning variables.

Algorithms 1 and 2 summarize the training and sampling procedures.

3.3 Conditional DDPM architecture

The denoising network is based on a U-Net architecture. It consists of a five-level encoder-decoder with skip connections. Group normalization and SiLU activa-

tion functions are used throughout the network. At each diffusion step, the noisy sample x_t , the corresponding LS, and the CT are concatenated along the channel dimension and jointly processed by the U-Net. The diffusion timestep t is represented by a sinusoidal embedding ϕ_t , projected by a two-layer multilayer perceptron (MLP) with 128 and 256 hidden units, and injected additively into each convolutional block. Fig. 3 illustrates the overall conditioning mechanism.




Fig. 3. Conditional DDPM architecture. At each diffusion step, the noisy sample x_t , the LS dose map, and the corresponding CT image are concatenated and processed by the U-Net. The diffusion timestep is represented by a sinusoidal embedding ϕ_t , projected by a two-layer MLP, and injected additively into each residual block. The network predicts the Gaussian noise $\hat{\epsilon}$ added during the forward diffusion process.

3.4 Training objective

Following [5], the network is trained to predict the Gaussian noise added during the forward diffusion process. We use the simplified noise-prediction objective

$$\mathcal{L}_{\text{simple}} = \mathbb{E}_{(x_0, \text{LS}, \text{CT}) \sim \mathcal{D}, t, \epsilon} \left[\|\epsilon - \epsilon_{\theta}(x_t, t, \text{LS}, \text{CT})\|_2^2 \right], \quad (7)$$

where $t \sim \text{Uniform}(\{1, \dots, T\})$ and $\epsilon \sim \mathcal{N}(0, \mathbf{I})$. x_t is obtained from the forward diffusion process. Minimizing this objective trains the network to estimate the noise contained in x_t . This estimate is then used at each step of the backward diffusion process to reconstruct the corresponding HS dose map during sampling.

4 Experimental validation

4.1 Data and experimental setup

MC simulations were performed using the open-source GPU-based GGEMS platform to generate paired LS and HS dose maps. For each patient, a voxelized phantom was constructed from the planning CT image, and a conical photon beam source was placed behind the patient. Each simulation configuration was run at two sampling levels. LS dose maps were generated using 10^6 photons, while HS reference dose maps were generated using 3×10^9 photons. LS simulations are obtained rapidly, whereas generating a single HS dose map of size 64×64 requires approximately 10 hours.

The abdominal dataset was built from the public dataset introduced in [14], which contains 82 patients. After preprocessing and slice extraction, 6 560 axial 2D samples were obtained. Each sample consists of a triplet (LS, HS, CT). A patient-level split was used to avoid data leakage. The training, validation, and

Table 1. Reconstruction performance on the abdominal test set.

Model Conditioning inputs	PSNR (dB)				SSIM			
	Min	Mean	Max	Std	Min	Mean	Max	Std
U-Net LS	30.85	38.67	44.27	2.14	0.920	0.974	0.989	0.009
U-Net CT + LS	33.60	43.44	50.61	3.43	0.960	0.991	0.998	0.006
DDPM –	18.56	23.43	30.17	1.75	0.344	0.721	0.893	0.089
DDPM CT	24.29	35.79	45.30	4.18	0.491	0.905	0.991	0.079
DDPM CT + LS	29.56	42.42	50.07	3.25	0.897	0.981	0.997	0.014

test sets contain 58, 16, and 8 patients, respectively, corresponding to 4640, 1280, and 640 slices.

Cross-anatomy generalization was evaluated on MC dose maps from two anatomical regions not observed during training. The evaluation set contains four pelvic slices and one head-and-neck slice. The limited number of cases results from the high computational cost of generating the corresponding HS MC references. These samples therefore provide a preliminary proof-of-concept evaluation of cross-anatomy reconstruction.

The conditional DDPM was trained for 400 epochs using the Adam optimizer with a learning rate of 10^{-4} and a batch size of 10. The diffusion process used $T = 1000$ timesteps and a linear noise schedule ranging from $\beta_1 = 10^{-4}$ to $\beta_T = 0.02$. An exponential moving average of the model parameters was maintained during training with a decay rate of 0.9999. Experiments were conducted on a workstation equipped with an Intel Core i7-10700F CPU, 32 GB of RAM, and an NVIDIA GeForce RTX 2070 SUPER GPU. On this hardware, sampling a 64×64 dose map required approximately 3 seconds.

Reconstruction quality was evaluated using the peak signal-to-noise ratio (PSNR) and the structural similarity index measure (SSIM). The current evaluation is limited to image-similarity metrics because clinically relevant volumetric dose metrics require 3D dose distributions and segmented anatomical structures, which are not available in the present 2D proof-of-concept setting.

4.2 In-distribution performance

We first evaluate the models on the abdominal test set, which corresponds to the anatomical region used during training. Table 1 reports the performance obtained with several independently trained models and different input configurations. For the DDPM, the noisy sample x_t and the diffusion timestep t are always provided to the network and are therefore omitted from the conditioning column. This column indicates whether CT and/or LS are additionally used as conditioning inputs.

The results show the importance of both the LS dose map and the planning CT image. For the U-Net, adding the CT image to the LS input increases the mean PSNR from 38.67 dB to 43.44 dB and the mean SSIM from 0.974 to 0.991. For the DDPM, CT conditioning increases the mean PSNR from 23.43 dB to 35.79 dB compared with the unconditional model. Joint conditioning on CT and LS further increases the mean PSNR to 42.42 dB and the mean SSIM to 0.981.

These results confirm that the two inputs provide complementary information. The LS dose map provides a noisy estimate of the deposited dose. The CT image provides information about anatomical geometry and tissue composition.

In the remainder of the experiments, both models receive the CT image and the LS dose map. Under this common input configuration, the conditional DDPM performs slightly below the deterministic U-Net on the abdominal test set, while remaining close in terms of both PSNR and SSIM. We next conduct a preliminary evaluation of the two models under cross-anatomy distribution shifts.

4.3 Preliminary cross-anatomy evaluation

Here, we examine the behavior of the proposed model beyond the anatomical region observed during training. Fig. 4 compares the reconstructions produced by the deterministic U-Net and the conditional DDPM on selected pelvic and head-and-neck slices, without additional adaptation.

The two models achieve close reconstruction performance on the abdominal test set (Table 1), but exhibit different qualitative behavior on the selected unseen anatomies. In the examples shown, the U-Net reconstructions exhibit stronger deviations from the HS reference, whereas the conditional DDPM better preserves the spatial dose patterns.

Given the limited number of cross-anatomy samples, these observations do not establish systematic generalization. They nevertheless provide preliminary evidence motivating further evaluation of the diffusion-based formulation under anatomical distribution shifts. We next examine whether geometric alignment between training and testing anatomies affects this behavior.

4.4 Impact of geometric alignment before reconstruction

The cross-anatomy data are represented in different orientations. We therefore assess whether geometrically aligning the test inputs to a common orientation affects reconstruction quality. The models are trained on the original abdominal data without geometric augmentation. Before reconstruction, each LS dose map and its corresponding CT image are rotated by the closest multiple of 90 degrees required to match the orientation of the abdominal training images. The same transformation is applied to the HS reference for evaluation. No retraining or parameter update is performed. Fig. 5 shows the reconstructions obtained from the geometrically aligned inputs.

Compared with the unaligned results shown in Fig. 4, geometric alignment improves the DDPM reconstructions in terms of PSNR and SSIM, but degrades the U-Net reconstructions. For the examples considered, these results suggest that the conditional DDPM benefits more than the U-Net from geometric consistency with the training data. A systematic assessment of geometric transformations remains necessary.




Fig. 4. Cross-anatomy reconstruction without geometric alignment. Both the U-Net and the conditional DDPM use the LS dose map and the corresponding CT image as conditioning inputs. They are trained exclusively on abdominal data and evaluated on previously unseen pelvic and head-and-neck anatomies. For these selected examples, the DDPM achieves higher PSNR and SSIM values than the deterministic U-Net and produces reconstructions closer to the HS reference. These examples are intended as a qualitative proof of concept rather than a systematic demonstration of cross-anatomy generalization.

4.5 Sampling-based reconstruction variability

Unlike deterministic image-to-image regression, the conditional DDPM can generate multiple reconstructions for the same conditioning pair (LS, CT). Each sampling run starts from an independent Gaussian realization x_T and follows a stochastic reverse diffusion trajectory. We repeat this procedure K times and compute the pixel-wise mean and standard deviation of the resulting reconstructions. The mean is used as the dose estimate, while the standard deviation quan-




Fig. 5. Effect of geometric alignment before reconstruction. The LS dose map and the corresponding CT image are rotated to a common orientation, while the trained models remain unchanged. The same transformation is applied to the HS reference for evaluation. In the examples shown, geometric alignment improves the DDPM reconstruction in terms of PSNR and SSIM, but degrades the U-Net reconstruction.

tifies the variability across samples. This variability should not be interpreted as a calibrated estimate of reconstruction error.

A standard U-Net produces one deterministic output, although predictive variability can be introduced through Monte Carlo dropout and repeated forward passes. Unlike this additional mechanism, stochastic sampling is intrinsic to the DDPM.

Fig. 6 shows the LS input, the mean DDPM reconstruction, the corresponding pixel-wise standard deviation, and the HS reference. This observation remains qualitative. Assessing the relationship between sampling variability and the actual reconstruction error requires a dedicated calibration study.




Fig. 6. Sampling-based reconstruction variability with the conditional DDPM. For fixed LS and CT inputs, the model is sampled $K = 5$ times from independent Gaussian noise realizations. The mean and pixel-wise standard deviation are shown with the HS MC reference.

5 Conclusion and perspectives

This work presented a conditional DDPM for MC dose reconstruction from LS simulations and planning CT images. The proposed method reconstructs HS dose maps through an iterative stochastic denoising process conditioned on both anatomical and dosimetric information.

On the abdominal test set, the conditional DDPM achieves reconstruction performance close to that of a deterministic U-Net. On a limited set of previously unseen pelvic and head-and-neck slices, the DDPM produces visually more faithful dose maps than the U-Net in the examples considered. These preliminary observations motivate a larger cross-anatomy evaluation but do not yet establish systematic generalization. The DDPM produces visually more faithful dose maps than the U-Net in these cases. Test-time geometric alignment further improves the DDPM reconstructions, suggesting that the model can exploit geometric consistency between training and testing anatomies.

The stochastic sampling process also provides a direct way to assess reconstruction variability. Repeated sampling for fixed LS and CT inputs yields multiple plausible dose maps, from which pixel-wise mean and standard deviation maps can be computed. These preliminary results illustrate the ability of the conditional DDPM to quantify sampling-based reconstruction variability, although further experiments are required to assess its calibration and its relationship with the actual reconstruction error.

The present study is limited to low-resolution 2D dose maps, a simplified irradiation geometry, and a very small cross-anatomy test set. It also relies on image-similarity metrics rather than clinical dosimetric endpoints. Consequently, the cross-anatomy and uncertainty results should be regarded as preliminary.

We are currently investigating posterior-guided diffusion strategies to further exploit the physical information available from MC simulations during reconstruction. Beyond this ongoing work, future research will systematically assess training-time data augmentation strategies and explore more efficient generative formulations to accelerate sampling. The proposed framework will also be ex-

tended to three-dimensional dose reconstruction and evaluated on larger datasets covering a broader range of anatomical sites and clinical applications.

References

1. Behloul, A., Visvikis, D., Bert, J.: Improved woodcock tracking on monte carlo simulations for medical applications. *Phys Med Biol* **63** (2018)
2. Bert, J., Perez-Ponce, H., Bitar, Z.E., Jan, S., Boursier, Y., Vintache, D., Bonissent, A., Morel, C., Brasse, D., Visvikis, D.: Geant4-based monte carlo simulations on gpu for medical applications. *Phys Med Biol* **58** (2013)
3. Dang, H.P., Wojdacki, T., Visvikis, D., Bert, J.: Apprentissage profond pour améliorer la qualité statistique et le temps de calcul de carte dosimétrique en radiothérapie. In: *Conférence Extraction et Gestion des Connaissances (EGC) - IA pour la Santé* (2024)
4. García-Pareja, S., Lallena, A., Salvat, F.: Variance-reduction methods for monte carlo simulation of radiation transport. *Frontiers in Physics* **9** (2021)
5. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: *Proceedings of the 34th International Conference on Neural Information Processing Systems. NIPS '20*, Curran Associates Inc., Red Hook, NY, USA (2020)
6. Kawrakow, I., Fippel, M.: Investigation of variance reduction techniques for Monte Carlo photon dose calculation using XVMC. *Phys Med Bio* **45**(8), 2163–2183 (Aug 2000)
7. Kontaxis, C., Bol, G., Legendijk, J., Raaymakers, B.: Deepdose: Towards a fast dose calculation engine for radiation therapy using deep learning. *Phys Med Biol* **65** (2020)
8. Lee, M.S., Hwang, D., Kim, J.H., Lee, J.S.: Deep-dose: a voxel dose estimation method using deep convolutional neural network for personalized internal dosimetry. *Scientific Reports* **9** (2019)
9. Liu, Z., Fan, J., Li, M., Han, Y., Hu, Z., Huang, P., Tian, Y., Miao, J., Dai, J.: A deep learning method for prediction of three-dimensional dose distribution of helical tomotherapy. *Med Phys* **46** (2019)
10. Maneval, D., Ozell, B., Després, P.: pgpumcd: an efficient gpu-based monte carlo code for accurate proton dose calculations. *Phys Med Biol* **64** (2019)
11. Mao, X., Pineau, J., Keyes, R., Enger, S.A.: RapidBrachyDL: Rapid radiation dose calculations in brachytherapy via deep learning. *Int. J. Radiat. Oncol. Biol. Phys.* **108**(3), 802–812 (nov 2020)
12. Nguyen, D., Jia, X., Sher, D., Lin, M.H., Iqbal, Z., Liu, H., Jiang, S.: 3d radiotherapy dose prediction on head and neck cancer patients with a hierarchically densely connected u-net deep learning architecture. *Phys Med Biol* **64**, 065020 (2019)
13. Roser, P., Zhong, X., Birkhold, A.I., Strobel, N., Kowarschik, M., Fahrig, R., Maier, A.K.: Physics-driven learning of x-ray skin dose distribution in interventional procedures. *Med Phys* (2019)
14. Roth, H., Farag, A., Turkbey, E.B., Lu, L., Liu, J., Summers, R.M.: Data from pancreas-ct. *The Cancer Imaging Archive* (2016)
15. Sarrut, D., Etxebeeste, A., Krah, N., Létang, J.: Modeling complex particles phase space with gan for monte carlo spect simulations: a proof of concept. *Phys Med Biol* **66** (2021)
16. Villa, M., Nasr, B., Benoit, D., Padoy, N., Visvikis, D., Bert, J.: Fast dose calculation in x-ray guided interventions by using deep learning. *Phys Med Biol* **68** (2023)