Physics and Imaging in Radiation Oncology 30 (2024) 100575

Contents lists available at ScienceDirect

Physics and Imaging in Radiation Oncology

journal homepage: www.sciencedirect.com/journal/physics-and-imaging-in-radiation-oncology

Original Research Article

A deep learning based dynamic arc radiotherapy photon dose engine trained on Monte Carlo dose distributions

Marnix Witte *, Jan-Jakob Sonke

Department of Radiation Oncology, The Netherlands Cancer Institute, Amsterdam, The Netherlands

ARTICLE INFO

Keywords:
Radiotherapy
Dose engine
Deep learning

ABSTRACT

Background and purpose: Despite hardware acceleration, state-of-the-art Monte Carlo (MC) dose engines require considerable computation time to reduce stochastic noise. We developed a deep learning (DL) based dose engine reaching high accuracy at strongly reduced computation times.

Materials and methods: Radiotherapy treatment plans and computed tomography scans were collected for 350 treatments in a variety of tumor sites. Dose distributions were computed using a MC dose engine for ~30,000 separate segments at 6 MV and 10 MV beam energies, both flattened and flattening filter free. For dynamic arcs these explicitly incorporated the leaf, jaw and gantry motions during dose delivery. A neural network was developed, combining two-dimensional convolution and recurrence using 64 hidden channels. Parameters were trained to minimize the mean squared log error loss between the MC computed dose and the model output. Full dose distributions were reconstructed for 100 additional treatment plans. Gamma analyses were performed to assess accuracy.

Results: DL dose evaluation was on average 82 times faster than MC computation at a 1 % accuracy setting. In voxels receiving at least 10 % of the maximum dose the overall global gamma pass rate using a 2 % and 2 mm criterion was 99.6 %, while mean local gamma values were accurate within 2 %. In the high dose region over 50 % of maximum the mean local gamma approached a 1 % accuracy.

Conclusions: A DL based dose engine was implemented, able to accurately reproduce MC computed dynamic arc radiotherapy dose distributions at high speed.

1. Introduction

Deep learning (DL) techniques have previously been applied to accelerate the radiotherapy planning process [1–3]. Accurate and fast dose computation is a requirement for Intensity Modulated Radiation Therapy (IMRT). Monte Carlo (MC) techniques provide highest accuracy, and may efficiently be applied in clinical treatment planning thanks to graphics processing unit (GPU) acceleration [4,5]. Nonetheless, accurate dose estimation requires computation times in the order of minutes, which may hamper application in on-line adaptive settings where sub-minute computation times would be essential.

To improve the balance between accuracy and runtime of dose computation, trained up-sampling of low resolution two-dimensional (2D) dose distribution slices was investigated [6]. Volumetric modulated arc therapy (VMAT) distributions for prostate cancer patients were reconstructed from pencil beam computations in water using a knowledge distillation framework with trained teacher and student networks [7]. Many contributions to the literature have used a three-dimensional (3D) U-Net architecture for dose computation. A dose verification tool for a novel cobalt IMRT machine was developed [8]. De-noising of fast MC computations with high stochastic noise was implemented [9]. Fluence maps for fixed beam IMRT of prostate plans were pre-processed using ray tracing to feed the network [10]. Similarly, individual multi-leaf collimator (MLC) shapes were projected into 3D parameter sets, and then used to compute fixed beam IMRT doses for prostate cancer patients using conventional and magnetic resonance (MR) guided linacs [11,12]. This method was extended to VMAT under MR guidance for various treatment sites [13]. Similarly, a transformer based U-Net solution was recently presented for MR image guided treatment [14]. Apart from the 3D U-Net approach, 3D convolutions were combined with 2D sequences connected using a transformer network for proton dose calculation and to reconstruct VMAT dose distributions, taking 3D projected shape information as input [15,16]. A 2D convolutional approach combined with a long-short term memory (LSTM) network

* Corresponding author.
E-mail address: m.witte@nki.nl (M. Witte).

https://doi.org/10.1016/j.phro.2024.100575
Received 11 October 2023; Received in revised form 3 April 2024; Accepted 3 April 2024
Available online 7 April 2024

was used to model proton dose deposition [17].

An MC dose engine can account for the continuous gantry and leaf motions during VMAT dose delivery, which so far none of the DL models in literature have considered. Furthermore, so far all DL based photon dose engines used some form of 3D data pre-processing, which limits the efficiency at inference time.

We implemented a dose engine using a recurrent 2D convolutional neural network architecture, directly taking as input the Digital Imaging and Communications in Medicine (DICOM) planning computed tomography (CT) image and the Radiotherapy Plan (RTPLAN) parameters at pairs of control points specifying gantry rotation and leaf and jaw motions during dynamic arc delivery, and producing a 3D segment dose distribution on the patient grid as output. This allows the model to be trained end-to-end on high quality MC dose distributions for dynamic as well as static beams, while the limited memory footprint and computational burden favor application under time critical conditions. The aim of this work was to demonstrate the feasibility of a 2D recurrent network for VMAT dose computation, and to compare its speed and accuracy with existing solutions.

2. Materials and methods

2.1. Patient data

Following institute review board approval, 350 DICOM treatment plans were collected from our institutional database, filtering to ensure a planning CT and RTPLAN were present, in which static (IMRT) or dynamic (VMAT) irradiation had been planned for any site of treatment. A total of 181 plans were static, containing 7460 control points in 863 beams, while 169 plans used 447 dynamic beams consisting of 26842 control points.

The plans were recomputed at a 2 mm resolution using the GPUMCD standalone dose calculation library (Version 1, Elekta AB, Stockholm, Sweden) at a 1 % statistical uncertainty setting. As only models for an Elekta accelerator using a multileaf collimator (MLC) with 160 leaves with a width of 5 mm at isoc depth were available, interpolated leaf positions were added to plans created for an (older) 80 leaf MLC. The GPUMCD engine separately models 6 MV and 10 MV energies, and uses separate models for flattened and flattening filter free (FFF) beams. Irrespective of the original intent, each plan was recomputed for the four cross combinations, storing each segment dose distribution separately. This led to four sets of ∼30000 segment dose distributions to be used for model training. A separate model was trained for each combination of energy and filter setting, resulting in four sets of trained model weights. A validation set of six treatment plans were held back to aid model selection and hyper-parameter tuning, five of which were VMAT for the treatment sites brain, head-and-neck (H&N), lung, breast and rectum, and one IMRT for a lung case including the mediastinum.

The GPUMCD engine consists of a Microsoft Windows dynamic link library, and the functions contained within it expect the data as derived from the RTPLAN to be presented in a particular form, for each segment providing a full state of the linear accelerator at both the beginning and ending control points. The MC engine then simulates particles at states in between, for a VMAT plan at interpolated MLC and gantry positions. The processed RTPLAN information used to evaluate each segment was captured and saved in a separate JavaScript Object Notation (JSON) file accompanying each segment dose distribution, to be used as model input (see Supplement).

Additionally, 100 plans were retrieved from the database to serve as test set, consisting of five sets of twenty plans for the major tumor sites breast, H&N, lung, prostate, and rectum. These were collected from a more recent time frame and were all VMAT plans using the 160 leaf MLC. All plans used two treatment arcs, except the breast plans which had between two and eight arcs. Full patient dose distributions were evaluated using the GPUMCD engine at 1 % accuracy with the beam energy, FFF status and dose grid resolution as used for clinical planning.

2.2. Dose engine based on deep learning

Assuming dose only depends on CT values encountered down to a given depth, the model was constructed on a diverging coordinate system (Fig. 1a) as a 2D convolutional neural network on the (x, z) plane perpendicular to the beam axis, while using recurrence to model dose variation along the beam direction y (Fig. 1b).

The MLC shape was rendered on this plane at a grid spacing in the x direction corresponding to the leaf separation of 5 mm at the isoc depth, and using 2 mm spacing in the leaf travel direction z. Based on the leaf and jaw positions at the beginning and ending control point of each segment, the fractional exposure was computed for each grid cell. For each segment the size of the grid was chosen to fit the MLC shape with a 4 cm margin to allow accurate modeling of the scattered dose in e.g. lung tissue. Zero padding was applied symmetrically to allow concatenation of multiple segments into a batch, equalizing the x and z dimensions while keeping each segment shape centered on the grid. While sampling CT values onto the 2D plane, the table, gantry, and collimator rotations were incorporated in the transformation matrix. For a VMAT segment the CT was transformed to the mid gantry position, while the amount of gantry rotation (zero in the case of IMRT) and the collimator angle were passed as input variables to the network. A 2.5 mm spacing was used between plane positions in the y direction, while skipping planes only containing air in front of and behind the patient (in beam's eye view). The coordinates x, y and z were used as input variables to the network, as well as a magnification factor relating the size of a 2D pixel to the nominal size at the isoc, as it expands with depth i. Explicit incorporation of coordinates as network inputs enabled the convolutional layers to encode the varying voxel geometries at different treatment depths and positions relative to the beam axis. Collimator angle θ_{coll} and gantry rotation $\Delta\theta_{\text{gantry}}$ were included as continuous variables by expressing these using pairs of sin and cos functions.

A gated recurrent unit (GRU) architecture with 64 hidden channels was deployed, in which the fully connected units were substituted with 2D convolutions using a 3×3 kernel size [18]. The 2D input variables (CT, MLC and x and z coordinates) were combined with the scalar inputs (collimator angle, gantry motion, y coordinate and magnification factor) by addition rather than concatenation, as this would have required expansion of the latter to 2D, introducing many redundant values in GPU memory. At each depth i the output of the GRU was transformed using a fully connected layer to produce the coefficients for a tri-cubic spline interpolation to sample the final output dose distribution on the patient's dose grid [19]. This sampling of dose onto the final grid was part of the gradient based training.

2.3. Model training

Model training was performed on a Linux GPU cluster with 24 GB memory per node, using PyTorch with its Adam optimizer [20]. Training loss was evaluated as the Mean Squared Logarithmic Error (MSLE) between the MC generated segment dose distribution and the model output, evaluated at 10000 dose grid voxels sampled randomly for each segment. The logarithm in the MSLE function helped to penalize absolute dose differences in low dose regions more heavily, leading to a more consistent relative accuracy across dose levels.

Using a learning rate of 10^{-4}, training ran for five days using 10^5 parameter updates (approximately 6 epochs) with single segment batches, followed by 10^5 parameter updates using batches of three segments and a final 10^4 parameter updates using a reduced learning rate of 10^{-5} (see Supplement).

2.4. Model evaluation

Model evaluation and testing as well as the MC dose computations were done on a Microsoft Windows personal computer with Nvidia A4000 GPU with 16 GB memory, using PyTorch's C++ interface to

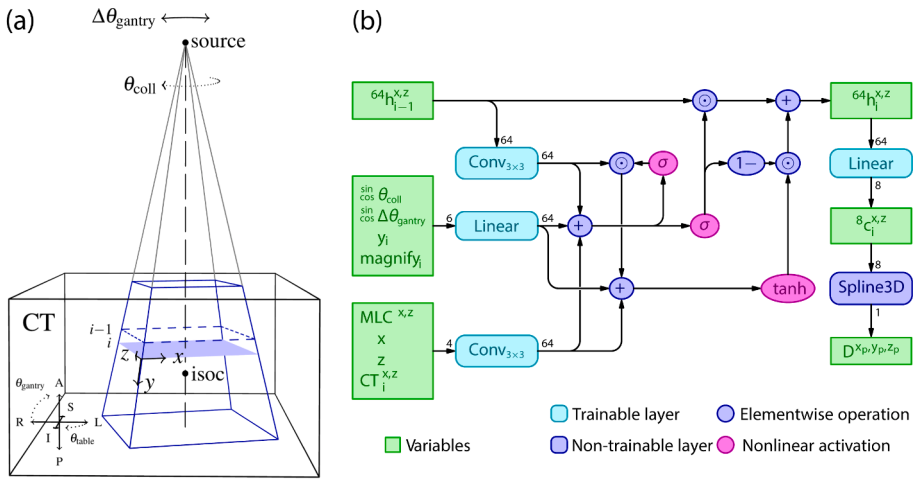

Fig. 1. (a) Geometry of the deep learning dose engine model. Computation takes place on a 2D grid perpendicular to the beam axis, accounting for beam divergence. Table rotation, mid-gantry position and collimator rotation are incorporated as a transformation of the CT. Gantry motion and collimator angle are used as input variables in the model. (b) The network architecture. A 2D convolutional recurrent neural network processes MLC, CT and geometrical data and produces cubic spline coefficients to interpolate dose on the patient dose grid. 2D and scalar input data are combined by addition.

integrate with our local software environment.

Average MC and DL computation times were computed per treatment site and overall, and plotted against planning target volume (PTV) size.

Model accuracy was assessed by computation of global gamma values between the MC and DL dose distributions with criteria of 2 and 1 (% and mm) for voxels above 10 % of dose maximum, similar to previously published studies. Voxels at a CT number less than −800 were disregarded, unless these were part of the delineated lungs. For each treament site, orthogonal dose and gamma plots were produced for the case with highest (worst) gamma score.

Given the large numbers of voxels at low doses whose accuracy is compared to the maximum dose, global gamma evaluation tends to lead to high pass rates. As a more critical test we therefore also evaluated local gamma using a 2 % and 1 mm criterion for voxels with at least 10 % and 50 % of dose maximum, and reported mean gamma values.

3. Results

3.1. Computation speed

The duration of dose computation was found to depend strongly on the size of the high dose region, so on PTV size. This applied both to the MC computations and the DL evaluations, see Fig. 2. Rectum cases which in general have low plan complexity still computed relatively quickly despite large PTV sizes and are concentrated in the lower right hand side of the diagram, while the more heavily modulated plans for head-and-neck targets could take long even for smaller PTVs, and are found towards the upper left hand side.

The speedup factor of DL compared to the MC computation depended on treatment site, and for breast cases was two times higher than for prostate (Table 1). While breast and rectum PTVs were similar, the speedup was higher for the breast cases.

3.2. Computation accuracy

Global gamma pass rates in the test were similarly high for all treatment sites at 2 %, 2 mm criterion, while at 1 %, 1 mm the H&N cases showed a somewhat reduced pass rate on average (Table 2).

In the plan with lowest global gamma pass rate for each of the tumor sites, higher gamma values can be seen at anomalies near the patient external (Fig. 3b, d, f), at contrast agent (3d) or fiducial markers (3h), and in bone (3d, j) (Fig. 3). Also, dose build-up regions in the rectum case show high gamma (3j), while in the prostate case intermediate gamma values reside in the high dose region (3h).

From local gamma evaluations, it followed that for voxels above 10 % of maximum dose the DL model was accurate well within 2 %, while for voxels above 50 % it approached a 1 % accuracy on average (Table 3).

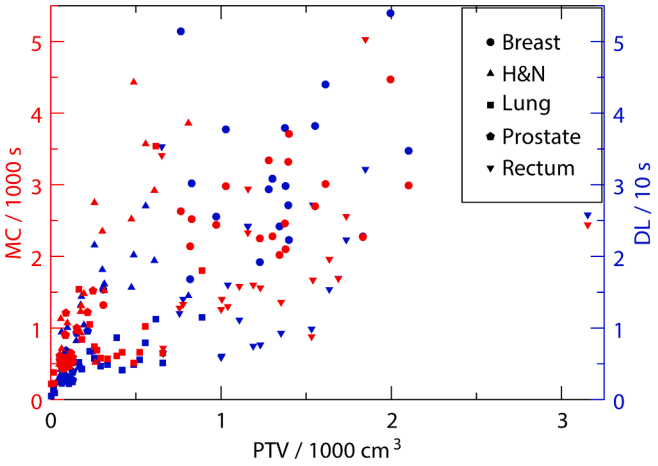

Fig. 2. Dose evaluation run times for MC at 1 % accuracy (red) and DL (blue) against PTV size. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 1
Planning target volumes, DL run times and speedup factors of DL over MC at 1 % accuracy. Mean value and standard deviation (SD) were computed per treatment site and overall.

	PTV (cm^3)		T_{DL} (s)		T_{MC}/ T_{DL}	
	mean	SD	mean	SD	mean	SD
Total	662	630	16	11	82	34
Breast	1291	430	27	7	118	31
H&N	272	211	18	12	75	28
Lung	327	239	9	7	70	32
Prostate	107	50	7	3	59	14
Rectum	1311	564	19	10	87	30

Table 2
Pass rates per plan of global gamma using a 2 %, 2 mm and a 1 %, 1 mm criterion for voxels receiving at least 10 % of the maximum dose. Mean value and standard deviation (SD) were computed per treatment site and overall.

	$\gamma^{G}_{10}(2^{\%}, 2^{mm})$ %^{pass}			$\gamma^{G}_{10}(1^{\%}, 1^{mm})$ %^{pass}		
	mean	SD	min	mean	SD	min
Total	99.6	0.3	97.6	86.2	7.3	62.3
Breast	99.5	0.3	98.9	86.4	4.5	80.0
H&N	99.6	0.2	99.1	82.4	9.3	65.3
Lung	99.7	0.3	98.7	87.2	7.4	62.3
Prostate	99.6	0.3	98.9	87.3	7.6	73.4
Rectum	99.4	0.5	97.6	87.8	5.9	70.0

4. Discussion

We demonstrated a novel dose engine using a 2D convolutional recurrent neural network on a computational grid accounting for beam divergence. The model does not depend on time consuming 3D preprocessing steps, and includes the leaf and gantry motions in VMAT delivery to accurately reproduce high quality MC generated dose distributions. The model was shown to perform well on lower cost mid-range GPU hardware.

The model architecture using 2D convolutions and recurrence in the third dimension resembles previous work on proton dose distributions, in which a LSTM was used [17]. Similarly, others used a transformer network to connect 2D slices [16]. Transformers have demonstrated

Table 3
Mean voxel values per plan of local gamma using a 2 %, 1 mm criterion for voxels receiving at least 10 % and 50 % of the maximum dose. Mean value and standard deviation (SD) were computed per treatment site and overall.

	$\gamma^{L}_{10}(2^{\%}, 1^{mm})$		$\gamma^{L}_{50}(2^{\%}, 1^{mm})$	
	mean	SD	mean	SD
Total	0.75	0.13	0.53	0.20
Breast	0.79	0.11	0.55	0.08
H&N	0.78	0.14	0.52	0.14
Lung	0.67	0.07	0.44	0.07
Prostate	0.81	0.13	0.57	0.30
Rectum	0.71	0.12	0.40	0.09

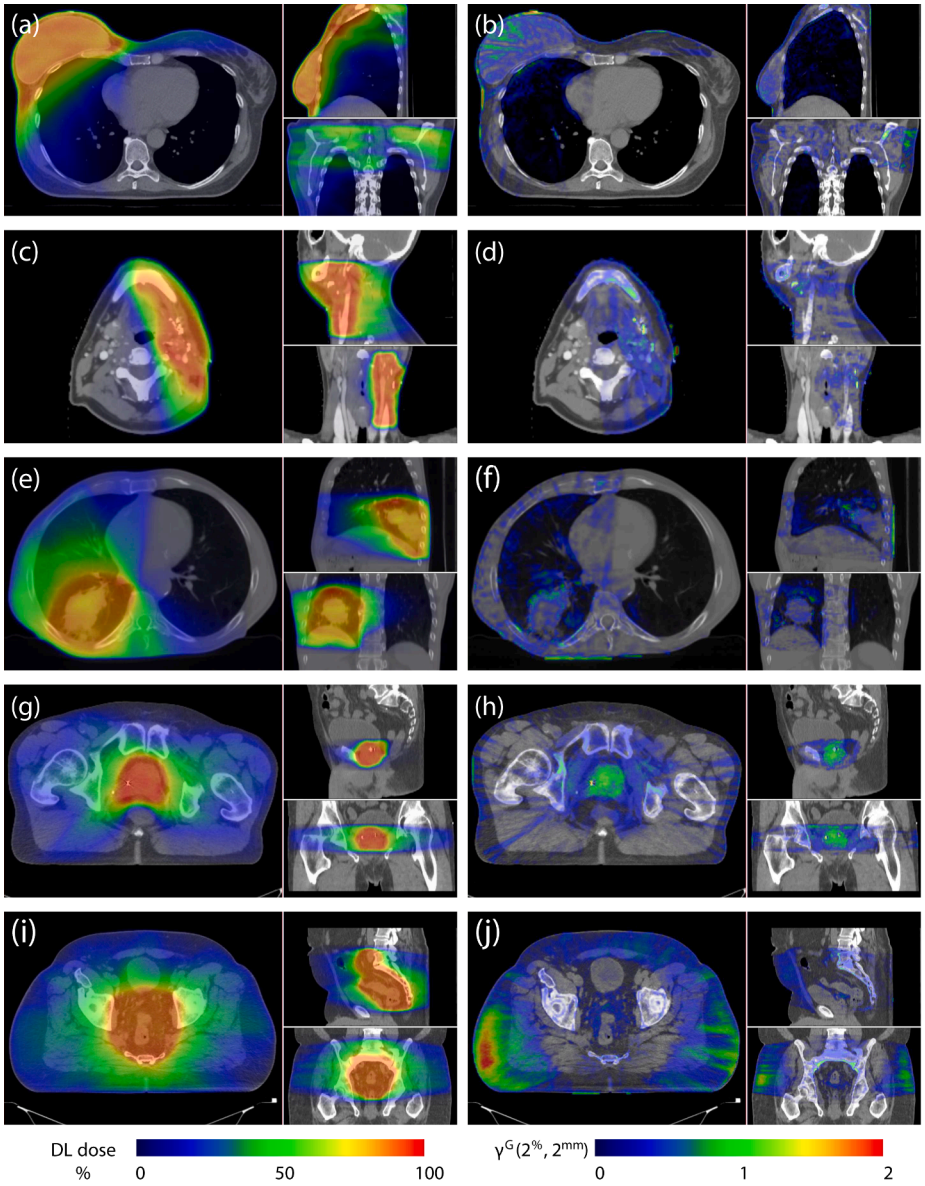

Fig. 3. DL dose distributions (left) and global gamma values using a 2 %, 2 mm criterion (right) for the plan with the lowest pass rate in each of the tumor sites.

unparalleled power to encode time dependencies, while the LSTM represents a more complex, potentially more powerful version of a recurrent network than the GRU we deployed. As the depth dependency of dose is determined by localized radiation diffusion and cumulative effects such as beam hardening, we still expected the GRU to be appropriate.

Directly comparing computational speed with previously published implementations is difficult because of the dependency on PTV size and differences in GPU hardware. Run times around one minute for full evaluation of a prostate step & shoot IMRT plan were reported, although using an older GPU [11]. When generalizing their methods to VMAT irradiation in the presence of a 1.5 T external magnetic field, these authors struggled to reach the necessary computational efficiency to handle the large numbers of segments [13]. The average run time of eight seconds for a set of prostate, lung, H&N and brain cases previously reported is consistent with our results, although these authors used a higher end GPU with larger memory (Nvidia A40 with 48 GB) [16]. As our model was implemented end to end as GPU accelerated code and does not depend on pre-processing steps, it is expected to scale well with hardware capabilities.

Gamma pass rates at a 2 % and 2 mm criterion were also comparable with those reported by [16]. These authors reported reduced pass rates in H&N cases, possibly related to air cavities. Our pass rates were more consistent across treatment sites, however we explicitly filtered out air cavities in the gamma analysis. Using a 1 % and 1 mm criterion our pass rates were lower for prostate and lung cases, but higher for H&N.

Run times for the DL model were observed to depend on the size of the high dose region, but also on plan complexity. In a heavily modulated plan, numerous irregularly shaped segments contribute to the high dose region. Areas blocked by leaves still receive scattered dose, contributing both to low and high dose regions. Margins around the segment shapes should ensure these regions are included in the computations. For a given volume of high dose, a more modulated plan will thus require longer computation time.

In constructing the model, simplifying assumptions were made. While the dose deposited at any point should mainly depend on the CT values that were encountered down to that depth, this is not true for the contribution of back-scattered radiation. A similar issue arises due to the rotation of the computational plane relative to the isoc in a VMAT segment. This causes one end of the plane to rise towards CT values that were already seen and could be accommodated for in the network's hidden state, but on the other end of the plane this rotation means the dose actually depends on CT values that are yet to be sampled. A possible solution for these issues would be to use a bidirectional network, however preliminary results using such a setup showed little difference, and the additional complication was deemed not worthwhile.

The rendering of the MLC shape onto the computational grid does not fully encode the complexity of VMAT dose delivery. In reality the simultaneous motions of MLC leaves and gantry lead to a complicated 3D pattern of dose deposition; the MC engine used for the ground truth dose in this study indeed computes VMAT dose in this realistic fashion. The DL model does not represent this complexity, but rather averages out the effects of leaf motions and gantry rotation. As our model shows good gamma evaluation results these simplifications seem justified, however if yet higher accuracy would be required this could offer a point of improvement.

In the training data set, segment shapes were re-used across beam energies and flattening filter settings. Also, MLC160 segments were constructed from MLC80 data by interpolating leaf positions. Accumulation of such segments would not lead to valid, conformal patient dose distributions. As the DL model was set up to reproduce per segment dose distributions, this still leads to an accurate and valid model.

Visual inspection of the gamma value maps that were presented revealed mainly clinically irrelevant discrepancies at anomalies such as contrast agent or at the patient contour. Still, differences in the gamma distributions observed in the prostate and rectum cases were remarkable, as these share a seemingly similar patient anatomy. However, segments for the rectum case were much larger, with MLC positions further away from the beam axis contributing to the dose deposition. Different MLC positions can be expected to be associated with different residual modeling errors.

Our model was written as a C++ module depending on LibTorch, allowing for easy integration with existing software solutions on Linux and Microsoft Windows platforms. In an on-line adaptive treatment setting the high computational speed allows to predict dose variations as a result of daily anatomical changes to high precision, potentially enabling real-time computation that can keep up with the dose delivery. For an on-line MR-linac work flow, GPUMCD at a 3 % accuracy was previously used, which should approximately be nine times faster than at 1 %, while our model showed 82 times speedup [13]. However, our model does not apply to dose delivery in presence of an external magnetic field. This will be the subject of our next analyses.

In summary, we have presented for the first time a GPU accelerated DL dose engine built using 2D network layers, capable of reproducing high quality MC dose distributions involving dynamic leaf and gantry motions. With its high speed and accuracy, it could become a valuable asset in a future online adaptive treatment setting.

CRediT authorship contribution statement

Marnix Witte: Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Visualization. **Jan-Jakob Sonke:** Conceptualization, Writing – review & editing, Supervision.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: The Netherlands Cancer Institute department of Radiation Oncology has entered a strategic research alliance with Elekta AB, Stockholm, Sweden.

Acknowledgements

Research at the Netherlands Cancer Institute is supported by institutional grants of the Dutch Cancer Society and of the Dutch Ministry of Health, Welfare and Sport. We acknowledge Elekta Solutions AB, Stockholm, Sweden for providing research software tools, and would like to thank Jacques Craenmehr for the collection of patient data.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, athttps://doi.org/10.1016/j.phro.2024.100575.

References

[1] Zeverino M, Piccolo C, Wuethrich D, Jeanneret-Sozzi W, Marguet M, Bourhis J, et al. Clinical implementation of deep learning-based automated left breast simultaneous integrated boost radiotherapy treatment planning. Phys Imaging Radiat Oncol 2023;28:100492. https://doi.org/10.1016/j.phro.2023.100492.
[2] Kusters M, Miki K, Bouwmans L, Bzdusek K, Van Kollenburg P, Smeenk RJ, et al. Evaluation of two independent dose prediction methods to personalize the automated radiotherapy planning process for prostate cancer. Phys Imaging Radiat Oncol 2022;21:24–9. https://doi.org/10.1016/j.phro.2022.01.006.
[3] Van De Sande D, Sharabiani M, Bluemink H, Kneepkens E, Bakx N, Hagelaar E, et al. Artificial intelligence based treatment planning of radiotherapy for locally advanced breast cancer. Phys Imaging Radiat Oncol 2021;20:111–6. https://doi.org/10.1016/j.phro.2021.11.007.
[4] De Martino F, Clemente S, Graeff C, Palma G, Cella L. Dose calculation algorithms for external radiation therapy: an overview for practitioners. Appl Sci 2021;11:6806. https://doi.org/10.3390/app11156806.
[5] Hissoiny S, Raaijmakers AJE, Ozell B, Després P, Raaymakers BW. Fast dose calculation in magnetic fields with GPUMCD. Phys Med Biol 2011;56:5119–29. https://doi.org/10.1088/0031-9155/56/16/003.

[6] Dong P, Xing L. Deep DoseNet: a deep neural network for accurate dosimetric transformation between different spatial resolutions and/or different dose calculation algorithms for precision radiation therapy. Phys Med Biol 2020;65:035010. https://doi.org/10.1088/1361-6560/ab652d.
[7] Tseng W, Liu H, Yang Y, Liu C, Lu B. An ultra-fast deep-learning-based dose engine for prostate VMAT via knowledge distillation framework with limited patient data. Phys Med Biol 2023;68:015002. https://doi.org/10.1088/1361-6560/aca5eb.
[8] Oh K, Gronberg MP, Netherton TJ, Sengupta B, Cardenas CE, Court LE, et al. A deep-learning-based dose verification tool utilizing fluence maps for a cobalt-60 compensator-based intensity-modulated radiation therapy system. Phys Imaging Radiat Oncol 2023;26:100440. https://doi.org/10.1016/j.phro.2023.100440.
[9] Neph R, Lyu Q, Huang Y, Yang YM, Sheng K. DeepMC: a deep learning method for efficient Monte Carlo beamlet dose calculation by predictive denoising in magnetic resonance-guided radiotherapy. Phys Med Biol 2021;66:035022. https://doi.org/10.1088/1361-6560/abca01.
[10] Xing Y, Nguyen D, Lu W, Yang M, Jiang S. Technical Note: a feasibility study on deep learning-based radiotherapy dose calculation. Med Phys 2020;47:753–8. https://doi.org/10.1002/mp.13953.
[11] Kontaxis C, Bol GH, Lagendijk JJW, Raaymakers BW. DeepDose: towards a fast dose calculation engine for radiation therapy using deep learning. Phys Med Biol 2020;65:075013. https://doi.org/10.1088/1361-6560/ab7630.
[12] Tsekas G, Bol GH, Raaymakers BW, Kontaxis C. DeepDose: a robust deep learning-based dose engine for abdominal tumours in a 1.5 T MRI radiotherapy system. Phys Med Biol 2021;66:065017. https://doi.org/10.1088/1361-6560/abe3d1.
[13] Tsekas G, Bol GH, Raaymakers BW. Robust deep learning-based forward dose calculations for VMAT on the 1.5T MR-linac. Phys Med Biol 2022;67:225020. https://doi.org/10.1088/1361-6560/ac97d8.
[14] Xiao F, Cai J, Zhou X, Zhou L, Song T, Li Y. TransDose: a transformer-based UNet model for fast and accurate dose calculation for MR-LINACs. Phys Med Biol 2022;67:125013. https://doi.org/10.1088/1361-6560/ac7376.
[15] Pastor-Serrano O, Perkó Z. Millisecond speed deep learning based proton dose calculation with Monte Carlo accuracy. Phys Med Biol 2022;67:105006. https://doi.org/10.1088/1361-6560/ac692e.
[16] Pastor-Serrano O, Dong P, Huang C, Xing L, Perkó Z. Sub-second photon dose prediction via transformer neural networks. Med Phys 2023;50:3159–71. https://doi.org/10.1002/mp.16231.
[17] Neishabouri A, Wahl N, Mairani A, Köthe U, Bangert M. Long short-term memory networks for proton dose calculation in highly heterogeneous tissues. Med Phys 2021;48:1893–908. https://doi.org/10.1002/mp.14658.
[18] Ballas N, Yao L, Pal CJ, Courville AC. Delving deeper into convolutional networks for learning video representations. CoRR 2015. abs/1511.06432.
[19] Press WH, Teukolsky SA, Vetterling WT, Flannery BP. Numerical Recipes 3rd Edition: The Art of Scientific Computing. 3rd ed. Cambridge University Press; 2007.
[20] Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, et al. PyTorch: an imperative style, high-performance deep learning library. In: Advances in Neural Inf Processing Systems 32. Curran Associates, Inc.; 2019. p. 8024–35.