

PyDoseRT: A physics-informed, plug-and-play dose engine for gradient-based radiotherapy treatment planning

Attila Simkó¹
Matthias Kronsteiner²
Simon Glatzer²
Minh Vu¹
Josef A. Lundman¹
Joakim Jonsson¹
Jörgen Olofsson¹
Kristina Sandgren¹
Wolfgang Lechner²
Dietmar Georg²
Tommy Löfstedt³
Tufve Nyholm¹
Anders Garpebring¹
Gerd Heilemann²

ATTILA.SIMKO@UMU.SE

¹ Department of Diagnostics and Intervention, Umeå University, Umeå, Sweden

² Department of Radiation Oncology, Medical University of Vienna, Vienna, Austria

³ Department of Computing Science, Umeå University, Umeå, Sweden

Editors: Under Review for MIDL 2026

Abstract

We present PyDoseRT, a GPU-accelerated, differentiable dose engine implemented in PyTorch that computes 3D dose distributions from machine-deliverable parameters (MLC leaf positions, jaw settings, gantry angles, monitor units) while preserving gradients throughout. The engine was validated on 181 clinical VMAT prostate plans from two institutions, achieving mean gamma pass rates of 99.6% and 97.5% (2%/2 mm). We further trained a deep learning model on the LUND-PROBE dataset to predict delivered dose from patient anatomy, using PyDoseRT as a differentiable layer for end-to-end training. The model inherently produced deliverable plans with Dice coefficients of 0.87 ± 0.02 and 0.92 ± 0.03 for the 50% and 95% isodose volumes on a held-out validation set. PyDoseRT enables TPS-independent, gradient-based radiotherapy optimization and provides a platform for deep learning-based treatment planning.

Keywords: Radiotherapy, dose engine, differentiable physics, treatment planning, deep learning

1. Introduction

Radiotherapy treatment planning involves optimizing thousands of parameters—multileaf collimator (MLC) positions, jaw positions, gantry angles, monitor units—to achieve clinically acceptable dose distributions. This optimization currently requires expert manual adjustments in commercial treatment planning systems (TPS). The rising global cancer

burden (Abdel-Wahab et al., 2024; Aggarwal et al., 2023) and the adoption of image-guided radiotherapy (Qiu et al., 2023) demand automated workflows to replace manual planning.

Deep learning (DL) has shown promise for dose prediction (Zimmermann et al., 2021; Gao et al., 2025), but predicted dose maps are not deliverable plans and must be imported into a TPS for aperture optimization (Mekki et al., 2025; Chang et al., 2025). Parameter prediction methods bypass the TPS but do not model how parameter changes affect dose (Heilemann et al., 2023, 2025). A differentiable dose engine bridges this gap: it enables gradient-based optimization of delivery parameters and can be embedded as a differentiable layer within DL models for end-to-end training. While prior work explored PyTorch-based optimization for CyberKnife (Liang et al., 2022) and GPU-based proton dose computation (Bhattacharya et al., 2025), neither offered full VMAT physics with gradient access.

We present PyDoseRT, a modular, gradient-enabled dose engine in PyTorch, and demonstrate its use as a differentiable training component for DL-based treatment parameter prediction on the LUND-PROBE dataset.

2. Materials and Methods

2.1. Dose Engine

PyDoseRT is a physics-based pencil-beam dose engine that maps delivery parameters (MLC leaf/jaw positions, monitor units) to 3D dose distributions for a given patient CT and beam geometry. The pipeline consists of six layers: (1) fluence map generation from aperture parameters with physics corrections; (2) fluence projection into 3D with beam divergence and inverse-square law; (3) radiological depth computation via ray-tracing; (4) depth-dependent pencil-beam kernel construction (Nyholm et al., 2006); (5) grouped 2D convolution of fluence with kernels; (6) coordinate transformation and dose accumulation. Layers on optimizable parameters preserve gradients; fixed computations (radiological depth, kernels) run outside the gradient graph. Beam models were commissioned against water-tank measurements for an Elekta Versa HD and a Varian TrueBeam via an automated fitting pipeline.

2.2. Dose Engine Validation

The engine was validated on two prostate VMAT cohorts: 19 patients (Umeå, Gold Atlas (Nyholm et al., 2018), Varian TrueBeam, 42.7 Gy/7 fx) and 162 patients (Vienna, Elekta Versa HD, 60 Gy/20 fx). Clinical DICOM RTPLAN parameters were recalculated in PyDoseRT and compared against TPS doses using gamma analysis (2%/2 mm, 10% threshold, excluding 1 cm from external contour), isodose Dice coefficients, and mean absolute dose difference (MADD).

2.3. Deep Learning-Based Parameter Prediction

An attention-based DL model was trained on the LUND-PROBE dataset to predict machine parameters from patient CT, target, and organ-at-risk masks. Predicted MLC positions, jaw settings, and monitor units are passed through PyDoseRT during training; losses on the resulting dose backpropagate through the engine to the network weights, enabling end-to-end optimization without a commercial TPS.

3. Results

3.1. Dose Engine Validation

Water phantom depth-dose curves and lateral profiles showed close agreement with measurements ($MAE \leq 1.5\%$). Clinical recalculation results are summarized in Table 1. On an NVIDIA A40, a VMAT forward pass took 2.1 s and a forward-backward pass 5.1 s.

Table 1: Dose engine validation: agreement between PyDoseRT and TPS (mean [range]).

Metric	Umeå cohort	Vienna cohort	Unit
Gamma pass (2%/2 mm)	99.64 [98.79–100.00]	97.52 [89.64–99.76]	%
Dice (50% isodose)	0.98 [0.98–0.99]	0.98 [0.94–0.99]	–
Dice (95% isodose)	0.96 [0.95–0.98]	0.95 [0.84–0.98]	–
MADD	0.07 [0.05–0.11]	0.05 [0.02–0.11]	Gy

3.2. Deep Learning-Based Parameter Prediction

Table 2 summarizes the DL model performance on the held-out LUND-PROBE validation set (83 patients). The model predicted deliverable machine parameters that, when recalculated with PyDoseRT, produced dose distributions in close agreement with clinical reference plans. Predictions can also be exported as DICOM RTPLANS.

Table 2: DL-predicted plan quality on the LUND-PROBE validation set (mean \pm std).

Metric	Value	Unit
Dice (50% isodose)	0.87 \pm 0.02	–
Dice (95% isodose)	0.92 \pm 0.03	–
MADD	0.93 \pm 0.13	Gy

4. Discussion and Conclusions

PyDoseRT demonstrates that physics-informed radiotherapy planning can be performed entirely within automatic differentiation frameworks. Unlike existing open-source engines (Wieser et al., 2017; Bhattacharya et al., 2025) that lack differentiability or GPU acceleration, and unlike RL-based approaches (Mekki et al., 2025; Achlatis et al., 2025) that require surrogate dose models, PyDoseRT enables direct gradient-based optimization of delivery parameters through physical dose computation. With forward-backward pass times of ~ 5 s per VMAT plan, it supports iterative optimization within clinically relevant timeframes (Qiu et al., 2023). Current limitations include 2D pencil beam kernels, a focus on prostate treatments, and fixed gantry angles. The LUND-PROBE experiment demonstrates PyDoseRT’s utility as a differentiable training layer for end-to-end DL-based planning. Code is available at <https://github.com/UMU-DDI/PyDoseRT>, with examples.

Acknowledgments

We are grateful for the financial support obtained from the Cancer Research Foundation in Northern Sweden (AMP 24-1151). The computations were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS), partially funded by the Swedish Research Council through grant agreement no. 2022-06725.

References

May Abdel-Wahab, C Norman Coleman, Jesper Grau Eriksen, Peter Lee, Ryan Kraus, Ekaterina Harsdorf, Becky Lee, Adam Dicker, Ezra Hahn, Jai Prakash Agarwal, Pataje G S Prasanna, Michael MacManus, Paul Keall, Nina A Mayr, Barbara Alicja Jereczek-Fossa, Francesco Giammarile, In Ah Kim, Ajay Aggarwal, Grant Lewison, Jiade J Lu, Douglas Guedes de Castro, Feng-Ming (Spring) Kong, Haidy Afifi, Hamish Sharp, Verna Vanderpuye, Tajudeen Olasinde, Fadi Atrash, Luc Goethals, and Benjamin W Corn. Addressing challenges in low-income and middle-income countries through novel radiotherapy research opportunities. *The Lancet Oncology*, 25, 6 2024. URL <https://linkinghub.elsevier.com/retrieve/pii/S147020452400038X>.

Stefanos Achlatis, Efstratios Gavves, and Jan-Jakob Sonke. Physics-guided radiotherapy treatment planning with deep learning. 6 2025.

Ajay Aggarwal, Laurence Edward Court, Peter Hoskin, Isabella Jacques, Mariana Kroiss, Sarbani Laskar, Yolande Lievens, Indranil Mallick, Rozita Abdul Malik, Elizabeth Miles, Issa Mohamad, Claire Murphy, Matthew Nankivell, Jeannette Parkes, Mahesh Parmar, Carol Roach, Hannah Simonds, Julie Torode, Barbara Vanderstraeten, and Ruth Langley. Archery: a prospective observational study of artificial intelligence-based radiotherapy treatment planning for cervical, head and neck and prostate cancer - study protocol. *BMJ Open*, 13, 12 2023.

Mahasweta Bhattacharya, Calin Reamy, Heng Li, Junghoon Lee, and William T. Hrinivich. A python package for fast gpu-based proton pencil beam dose calculation. *Journal of Applied Clinical Medical Physics*, 26, 6 2025.

Ho-hsin Chang, Joseph Harms, Rex Alexander Cardan, John B. Fiveash, Richard A. Popple, and Carlos E. Cardenas. nndosenet: Intuitive and flexible deep learning framework to train and evaluate radiotherapy dose prediction models. *medRxiv*, 2025.

Riqiang Gao, Simon Arberet, Martin Kraus, Han Liu, Wilko FAR Verbakel, Dorin Comaniciu, Florin-Cristian Ghesu, and Ali Kamen. Demo: Generative ai helps radiotherapy planning with user preference, 2025. URL <https://arxiv.org/abs/2512.08996>.

Gerd Heilemann, Lukas Zimmermann, Raphael Schotola, Wolfgang Lechner, Marco Peer, Joachim Widder, Gregor Goldner, Dietmar Georg, and Peter Kuess. Generating deliverable dicom rt treatment plans for prostate vmat by predicting mlc motion sequences with an encoder-decoder network. *Medical Physics*, 50, 2023.

- Gerd Heilemann, Lukas Zimmermann, Tufve Nyholm, Attila Simkó, Joachim Widder, Gregor Goldner, Dietmar Georg, and Peter Kuess. Ultra-fast, one-click radiotherapy treatment planning outside a treatment planning system. *Physics and Imaging in Radiation Oncology*, 33, 1 2025.
- Bin Liang, Ran Wei, Jianghu Zhang, Yongbao Li, Tao Yang, Shouping Xu, Ke Zhang, Wenlong Xia, Bin Guo, Bo Liu, Fugen Zhou, Qiuwen Wu, and Jianrong Dai. Applying pytorch toolkit to plan optimization for circular cone based robotic radiotherapy. *Radiation Oncology*, 17, 2022.
- Lina Mekki, William T Hrinivich, and Junghoon Lee. Dual-arc vmat machine parameter optimization for localized prostate cancer using deep reinforcement learning. *Physics in Medicine & Biology*, 70, 2025.
- Tufve Nyholm, Jörgen Olofsson, Anders Ahnesjö, and Mikael Karlsson. Photon pencil kernel parameterisation based on beam quality index. *Radiotherapy and Oncology*, 78, 2006.
- Tufve Nyholm, Stina Svensson, Sebastian Andersson, Joakim Jonsson, Maja Sohlin, Christian Gustafsson, Elisabeth Kjellén, Karin Söderström, Per Albertsson, Lennart Blomqvist, Björn Zackrisson, Lars E. Olsson, and Adalsteinn Gunnlaugsson. Mr and ct data with multiobserver delineations of organs in the pelvic area-part of the gold atlas project:. *Medical Physics*, 45, 2018.
- Zihang Qiu, Sven Olberg, Dick den Hertog, Ali Ajdari, Thomas Bortfeld, and Jennifer Pursley. Online adaptive planning methods for intensity-modulated radiotherapy. *Physics in Medicine & Biology*, 68, 2023. URL <https://iopscience.iop.org/article/10.1088/1361-6560/acceb2>.
- Hans Peter Wieser, Eduardo Cisternas, Niklas Wahl, Silke Ulrich, Alexander Stadler, Henning Mescher, Lucas Raphael Muller, Thomas Klinge, Hubert Gabrys, Lucas Burigo, Andrea Mairani, Swantje Ecker, Benjamin Ackermann, Malte Ellerbrock, Katia Parodi, Oliver Jakel, and Mark Bangert. Development of the open-source dose calculation and optimization toolkit matrad. *Medical Physics*, 44, 6 2017.
- Lukas Zimmermann, Erik Faustmann, Christian Ramsel, Dietmar Georg, and Gerd Heilemann. Technical note: Dose prediction for radiation therapy using feature-based losses and one cycle learning. *Medical Physics*, 48, 2021.