Physics and Imaging in Radiation Oncology 33 (2025) 100724

Contents lists available at ScienceDirect

Physics and Imaging in Radiation Oncology

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

Technical Note

Ultra-fast, one-click radiotherapy treatment planning outside a treatment planning system

Gerd Heilemann [a,*], Lukas Zimmermann [a], Tufve Nyholm [b], Attila Simkó [b], Joachim Widder [a], Gregor Goldner [a], Dietmar Georg [a], Peter Kuess [a]

[a] Department of Radiation Oncology, Medical University of Vienna, Waehringer Guertel 18-20 1090 Vienna, Austria
[b] Department of Diagnostics and Intervention, Umeå University 90185 Umeå, Sweden

ARTICLE INFO

Keywords:
Auto-planning
Automation
Prostate
VMAT
Deep learning
Treatment planning
Artificial intelligence

ABSTRACT

We present an automated radiation oncology treatment planning pipeline that operates between segmentation and plan review, minimizing manual interaction and reliance on traditional planning systems. Two AI models work in sequence: the first generates a dose distribution, and the second creates a deliverable DICOM-RT plan. Trained and validated on 276 plans, and tested on 151 datasets, the system produced clinically deliverable plans—complete with all VMAT parameters—in about 38 s. These plans met target coverage and most organ-at-risk constraints. This proof-of-concept demonstrates the feasibility of generating high-quality, deliverable DICOM plans within seconds.

1. Introduction

Radiation oncology is essential in cancer treatment, with over half of patients receiving radiotherapy. Enhancing workflow efficiency is critical to meet the increasing demand, enabling clinical teams to treat more patients effectively while managing costs and resources [1]. Automation and artificial intelligence (AI) are promising strategies in adaptive settings to improve access and quality in radiotherapy, potentially increasing the return on investment [2–5].

The complexity of radiotherapy treatment planning (TP)—which depends heavily on the planner's skill, experience, and time investment—poses significant challenges amid rising demand. Inefficient workflows create bottlenecks that hinder the advancements of personalized radiotherapy, delaying the integration of online treatment adaptation [5].

Deep learning (DL) has the potential to revolutionize TP by significantly enhancing speed, accuracy, and consistency. Various models have been developed to predict the dose or fluence of Volumetric Modulated Arc Therapy (VMAT) plans [6–10]. However, these models often rely on traditional treatment planning systems (TPS) to finally convert predictions into executable DICOM RT plans for linear accelerators (LINACs), highlighting a gap in achieving a fully automated planning process outside commercial software.

We introduce a pipeline that combines a dose prediction model [10] with another model that predicts machine parameters for VMAT plans [11]. This integration creates a novel fully automated pipeline capable of producing deliverable treatment plans without relying on a TPS in less than one minute in one click.

2. Materials and methods

This study was approved by the Institutional Review Board (1373/2023).

2.1. Autonomous treatment planning workflow

The pipeline integrated two DL models (Models A and B). It encompasses the TP workflow after segmentation and before plan review. This included preprocessing the input DICOM data (CT images and RT structures), transferring and upscaling the output from one model to the other, and ultimately generating machine parameters. These parameters were saved as a DICOM RT plan file, fully compatible with the original DICOM CT and structure datasets through matching unique identifiers (UID). The complete workflow is depicted in Fig. 1.

In Step 1 (Fig. 1), CT images were converted into a numerical array for the initial dose prediction (Model A) together with binary masks for the planning target volume (PTV) and organs-at-risk (OARs) obtained from the RT structure file, which was resampled to the same resolution

* Corresponding author.
E-mail address: gerd.heilemann@meduniwien.ac.at (G. Heilemann).

https://doi.org/10.1016/j.phro.2025.100724
Received 1 October 2024; Received in revised form 4 February 2025; Accepted 4 February 2025
Available online 9 February 2025

(128 × 128 × 128 voxels with a voxel size of 3.125 mm). For input of Model B, the CT and binary masks were resampled to 128 × 128 × 320 voxels as Model B required a higher resolution in the cranio-caudal direction (resulting in voxels size of $3.125 \times 3.125 \times 1.25\,m^3$).

Step 2 represents **Model A**, the dose prediction model, that was based on a U-Net architecture [10]. Its output was a 128 × 128 × 128 voxel-sized dose distribution.

In Step 3, this predicted dose distribution was up-sampled to 128 × 128 × 320 voxels, agreeing with the CT input for Model B. Additionally, the dose within the PTV was homogenized by adjusting each voxel's dose (D):

$$D_{new} = (1-\alpha)D + \alpha\overline{D}_{PTV}$$

where $\overline{D}_{PTV}$ is the mean dose within the PTV and α controls the weighting and was set to 0.5. This reduces variation and increases ho-

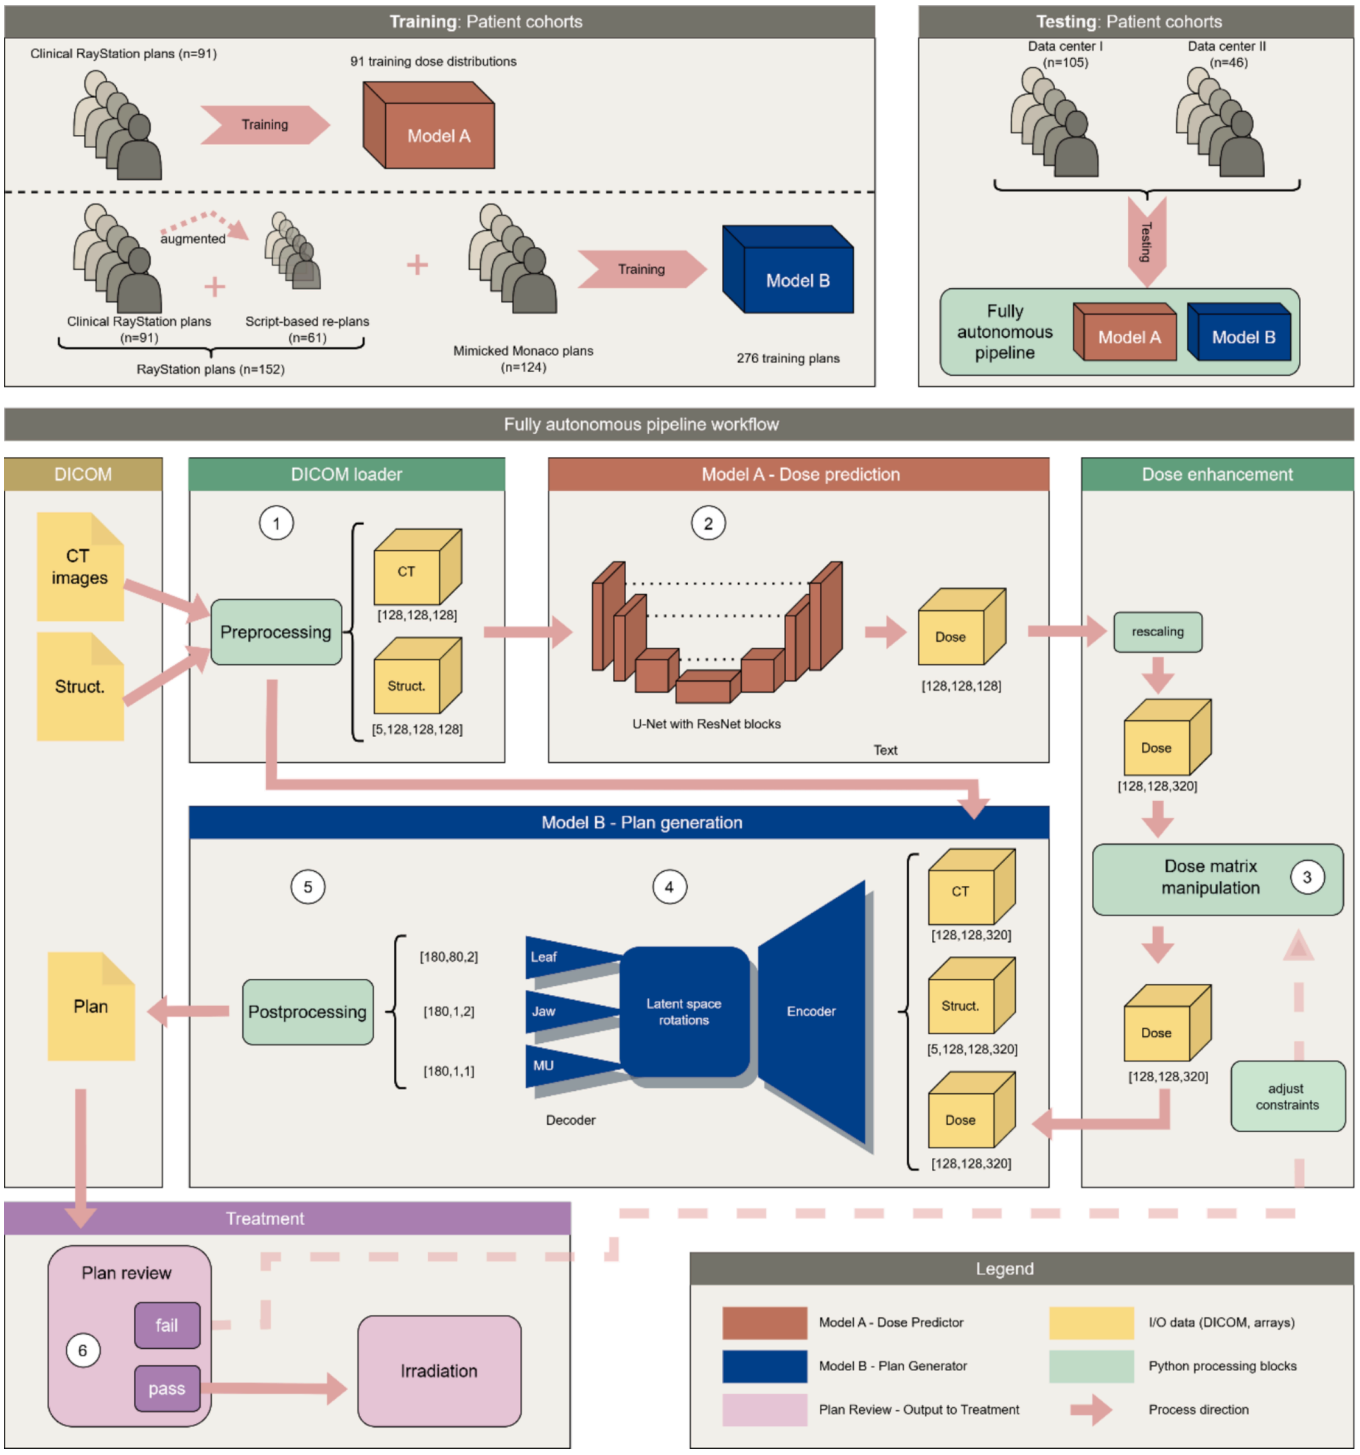

Fig. 1. Top: Patient cohorts used for training (left) and testing (right). Bottom: Fully autonomous pipeline workflow, starting with image and structure preprocessing in step 1 (steps symbolized as numbered white circles), followed by dose prediction in Model A (step 2). An adjustment of the dose is performed in step 3, before the dose acts as input for Model B, the plan generator, in step 4. In the postprocessing (step 5), potential machine violations are corrected to ensure the plan can be irradiated at the LINAC. Plan review (step 6) is the final external quality check.

mogeneity in the PTV.

In Step 4, this predicted dose distribution, together with the CT and structure masks, was used as input for **Model B,** generating machine parameters—leaf/jaw positions and monitor units (MU)— for all control points of a VMAT arc [11].

In Step 5, a vendor-specific post-processing step was implemented to ensure no predicted machine parameters violated specific constraints, securing conformity within the machine's capabilities (e.g., minimum leaf gaps and travel speeds). The final parameters were then compiled into a DICOM plan, including patient metadata, achieving compatibility with the CT and structure set.

2.2. Data and model training

Models A and B were trained separately and specifically for this study with datasets not used in previous publications. The clinical TP protocol is described in the supplementary materials.

Model A was a U-Net with additional ResNet blocks interspersed between the up- and down-sampling convolutions. The training was done with an AdamW optimizer coupled with the One Cycle scheduler and an initial learning rate of 0.001. We combined L1 loss with a feature loss [12], which extracts higher-order information using a pre-trained model [13]. Specifically, this study introduced an additional gradient loss function, designed to compute the dose gradient across the dose distribution. This loss function was incorporated to improve the model's ability to predict accurate dose distributions. The batch size was 2 for 40 k iterations.

For training (see Fig. 1 top left) the dose prediction model (**Model A**), 91 clinically approved datasets (CT, structures, dose) from patients treated for localized prostate cancer at the Department of Radiation Oncology at the Medical University of Vienna (MUW) between 01/2021 and 12/2023 were used.

Model B employed an encoder with three decoders to accurately predict the necessary parameters for creating a compliant DICOM RT plan [11]. The encoder employed a ResNet-18 with 3D convolutions as its backbone, using pre-trained MedicalNet weights [14] and processed the original plan dose. The latent space output was manipulated by rotating it along the patient's longitudinal axis, creating beams-eye-views for each control point by resampling the data at the beam's isocenter. Three separate decoders for the leaf and jaw positions and the monitor units predict all machine parameters for a DICOM RT plan. Training was done with the Lion optimizer [15] for 40 k iterations and a batch size of 5. Separate L1 losses were computed for leaf position, jaw positions, and MUs.

For training **Model B** (plan generation), the same 91 clinically approved prostate plans as for **Model A** were used (CT, structures, dose, and plan). Additionally, the training data was augmented through a semi-automated TP workflow. This workflow utilized a knowledge-based scripting approach, developed in-house (unpublished) and implemented via the Python scripting interface in RayStation (RaySearch Laboratories, Sweden). The database stored a range of anatomical features, including OAR and target volumes, and proximity measures (e.g., surface-to-surface and center-of-mass distances). Anatomically similar patients were identified by comparing segmented CT data with stored anatomical profiles. Their corresponding optimization parameters were then applied to initiate the initial optimization in the TPS. In addition, mean dose-volume histograms (DVHs) from these similar patients were calculated and used to inform a projected target DVH. This added 61 datasets.

To further expand the training data, the dose-mimicking module in RayStation was employed. This approach utilized prostate treatment plans previously generated in the TPS Monaco (Elekta, Sweden) which were incompatible with RayStation plans, thereby contributing 124 additional datasets to the training pool, leading to a final training dataset size of 276.

2.3. Evaluation of the generated plans

A total of 151 independent patients (see Fig. 1 top right), not part of the training, were used for testing (see Fig. 1). The DICOM CT images and RT structure files were used as input. Of these, 105 patients were prostate patients from MUW treated between 2021 and 2024 and 46 from Umeå University (Department of Diagnostics and Intervention) from 2022 to 2023. Different TPSs were used in both centers, as well as treatment machines and prescription levels (e.g., ultra-hypofractionated). Although most original clinical plans could not be directly compared with those generated in this study, a direct comparison for two plans with compatible prescription levels is included in the supplementary data.

International OAR guidelines, with bladder and rectum dose constraints based on CHHiP [16] recommendations were used to quantify plan quality. In addition, constraints developed from local clinical experience at MUW were applied (Table 1). Target evaluation included PTV coverage, reporting $D_{95\%}$, $D_{1\%}$ consistent with established standards [17,18]. For plan review (Fig. 1, step 6) the generated plans were recalculated with a commercial dose engine, i.e., collapsed-cone in RayStation or Monte Carlo in SciMoCa (ScientificRT GmbH, Germany) [19]. The MUs were then normalized to match the minimum requirements for tumor coverage.

Lastly, a patient-specific quality assurance (PSQA) process was employed to confirm the clinical viability of our generated plans. A random set of 15 plans was tested on a PSQA phantom (Delta4 Scandidos, Sweden), with acceptance criteria aligned with those applied to clinical plans, namely a gamma-index threshold (>90 %). This allows verification of the automatically generated plans to ensure they meet delivery and safety standards.

3. Results

In total 151 prostate treatment plans were generated with the pipeline. The average time required to produce a treatment plan from a segmented CT without any user interaction was 38 ± 6 s, while about half of that time was spent on pre- and post-processing.. The inference was done on a workstation (AMD Ryzen 7 3700-X 8 Core 3.6 GHz, RTX 2080Ti).

The results of these 151 predicted plans concerning the clinical goals are presented in Table 1 and in the supplementary materials. The rectum

Table 1
Results of the clinical goal evaluation based on CHHiP [16] and the Medical University of Vienna (MUW) internal clinical goals for 151 generated plans.

ROI	Source	Criteria	Goal	Plans that achieved the goal
Bladder	CHHiP	V40.8 Gy	50 %	100 %
		V48.6 Gy	25 %	97 %
		V60Gy	5 %	76 %
	MUW	V54Gy	20 %	100 %
		V48Gy	40 %	100 %
		V44Gy	50 %	100 %
		V30Gy	80 %	100 %
		D1%	66 Gy	92 %
Rectum	CHHiP	V60Gy	3 %	58 %
		V57Gy	15 %	100 %
		V52.8 Gy	30 %	100 %
		V48.6 Gy	50 %	100 %
		V40.8 Gy	60 %	100 %
	MUW	V54Gy	20 %	93 %
		V48Gy	40 %	100 %
		V40Gy	45 %	100 %
		V30Gy	60 %	100 %
		V20Gy	54 %	100 %
		D1%	66 Gy	99 %
FemoralHead	MUW	V46Gy	5 %	100 %
PTV Prostate	MUW	D95%	57 Gy	100 %
		D1%	66 Gy	51 %

criteria were fulfilled in nearly all 151 cases, except V_{60Gy} (CHHiP), which was exceeded in 42 % of the evaluated plans. V_{60Gy} was also the most challenging criterion for the bladder and was not fulfilled in 24 %. The MUW criteria focused on the maximum dose substitute ($D_{1\%}$), which was exceeded in 7 % of the plans. While target coverage ($D_{95\%}$) was achieved in all generated plans, dose heterogeneity ($D_{1\%}$) was more challenging, with success in only 51 %.

The PSQA measurements of the randomly selected 15 treatment plans demonstrated that the plans could be delivered at a clinical LINAC. All 15 plans met the clinical criteria (2 %/2mm, >90 % gamma pass rates) for acceptance with pass rates of 95.4 ± 0.7 %.

4. Discussion

The pipeline efficiently generated VMAT treatment plans directly from segmented CT images, autonomously producing DICOM RT plan files ready for treatment without the need for a traditional TPS in 38 ± 6 s.

A recent ESTRO challenge explored the extent of automation possible in radiotherapy planning [8]. Our study addressed the process from segmented CT images to plan review and differed from previous studies in several aspects; other approaches included script-based solutions within a TPS [9,20,21] or dose prediction models [10,22–24], which rely on TPS to generate plan files. Fluence prediction [25,26], eliminates the need for mathematical optimization to convert dose distributions into fluence maps but TPS sequencing is required to produce the plan file. Recently, reinforcement learning (RL) agents were explored for automatic plan optimization [27,28]. RL shows potential, however, further studies are required to describe the limitations and challenges of this technique in the context of radiation oncology [5,29].

Some studies presented plan generation outside a commercial TPS – such as interfacing with matRAD [8] or Hyperion [30]. However, these studies use conventional plan optimization or sequencing within a TPS, a shortcoming surpassed by our approach by combining two AI models in sequence to directly predict machine parameters.

There still are several limitations in our proof of concept. The model does not generalize well outside the training dataset, which means localized prostate cancer patients, moderately hypofractionated (20 × 3 Gy). All patients were treated with a rectal balloon and no spacer. Most datasets were from low-risk patients, with only few including extensive seminal vesicles. Thus, handling more complex target-normal tissue configurations will require larger training datasets. Additionally, the pipeline currently produces files for one specific LINAC vendor, necessitating future generalization. For several cases, the plan quality did not meet clinical expectations, highlighting the need for more refined model architectures. Potential improvements include dividing the problem into 2D projections to reduce model size and increase batch size, or leveraging transformer networks to better capture the sequential nature of the MLC sequence [31]. Note that the pipeline starts after the segmentation process, thus tasks such as inserting a couch model were not addressed in this study. Moreover, the typical iterative replanning process familiar to conventional TPS users remains underdeveloped in our pipeline. Although we suggest a possible approach—indicated by the dashed line in Fig. 1 linking step 6 back to step 3—this remains a conceptual solution, and additional research is needed to fully implement and validate a robust replanning mechanism.

Finally, our pipeline does not inherently provide dose calculations using a commercial dose engine. However, this can be advantageous: since our pipeline rapidly produces deliverable plan files, these can be reviewed using more lightweight, web-based secondary plan review platforms—such as SciMoCa (ScientificRT GmbH, Germany), DOSIsoft (DOSIsoft, France), VERIQA RT MonteCarlo (PTW, Germany), DoseCHECK (Sun Nuclear, USA), or RadCalc (LAP GmbH Laser Applications, Germany)—without the overhead of an expensive commercial TPS. A future with rapid AI-based replanning presents significant challenges for quality assurance. Given the strict regulations in some countries, a fundamental re-evaluation of QA requirements may be necessary, as we see with dedicated adaptive radiotherapy machines today [32,33].

5. Declaration on Generative AI

Generative AI (OpenAI ChatGPT v-o1) was utilized for proofreading this manuscript. All AI-generated suggestions were carefully reviewed and edited by the authors to ensure accuracy and clarity.

Funding

None.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: The Department of Radiation Oncology at the Medical University of Vienna has institutional research contracts with RaySearch Laboratories (Sweden), ELEKTA (Sweden), and Brainlab (Germany).

Acknowledgments

None.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.phro.2025.100724.

References

[1] Baumann M, Bacchus C, Aznar MC, Coppes RP, Deutsch E, Georg D, et al. Clinical research for global needs of radiation oncology. Radiother Oncol 2024;190:110076. https://doi.org/10.1016/j.radonc.2023.110076.
[2] Bray F, Parkin DM, Gnangnon F, Tshisimogo G, Peko JF, Adoubi I, et al. Cancer in sub-Saharan Africa in 2020: a review of current estimates of the national burden, data gaps, and future needs. Lancet Oncol 2022;23:719–28. https://doi.org/10.1016/S1470-2045(22)00270-4.
[3] Atun R, Jaffray DA, Barton MB, Bray F, Baumann M, Vikram B, et al. Expanding global access to radiotherapy. Lancet Oncol 2015;16:1153–86. https://doi.org/10.1016/S1470-2045(15)00222-3.
[4] Court L, Aggarwal A, Burger H, Cardenas C, Chung C, Douglas R, et al. Addressing the global expertise gap in radiation oncology: the radiation planning assistant. JCO Glob Oncol 2023. https://doi.org/10.1200/go.22.00431.
[5] Qiu Z, Olberg S, den Hertog D, Ajdari A, Bortfeld T, Pursley J. Online adaptive planning methods for intensity-modulated radiotherapy. Phys Med Biol 2023;68. https://doi.org/10.1088/1361-6560/accdb2.
[6] Babier A, Zhang B, Mahmood R, Moore KL, Purdie TG, McNiven AL, et al. OpenKBP: The open-access knowledge-based planning grand challenge and dataset. Med Phys 2021;48:5549–61. https://doi.org/10.1002/mp.14845.
[7] Lee H, Kim H, Kwak J, Kim YS, Lee SW, Cho S, et al. Fluence-map generation for prostate intensity-modulated radiotherapy planning using a deep-neural-network. Sci Rep 2019;9:1–11. https://doi.org/10.1038/s41598-019-52262-x.
[8] Gooding MJ, Aluwini S, Guerrero Urbano T, McQuinlan Y, Om D, Staal FHE, et al. Fully automated radiotherapy treatment planning: A scan to plan challenge. Radiother Oncol 2024;200. https://doi.org/10.1016/j.radonc.2024.110513.
[9] Church C, Yap M, Bessrour M, Lamey M, Granville D. Automated plan generation for prostate radiotherapy patients using deep learning and scripted optimization. Phys Imaging. Radiat Oncol 2024;32. https://doi.org/10.1016/j.phro.2024.100641.
[10] Zimmermann L, Faustmann E, Ramsl C, Georg D, Heilemann G. Technical Note: Dose prediction for radiation therapy using feature-based losses and One Cycle Learning. Med Phys 2021;48:5562–6. https://doi.org/10.1002/mp.14774.
[11] Heilemann G, Zimmermann L, Schotola R, Lechner W, Peer M, Widder J, et al. Generating deliverable DICOM RT treatment plans for prostate VMAT by predicting MLC motion sequences with an encoder-decoder network. Med Phys 2023;50:5088–94. https://doi.org/10.1002/mp.16545.
[12] Johnson J, Alahi A, Fei-Fei L. Perceptual losses for real-time style transfer and super-resolution. In: Lecture notes in computer science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics); 2016. p. 694–711. vol. 9906 LNCS. https://doi.org/10.1007/978-3-319-46475-6_43.
[13] Tran D, Wang H, Torresani L, Ray J, Lecun Y, Paluri M. A closer look at spatiotemporal convolutions for action recognition. In: Proceedings of the IEEE

Computer Society Conference on Computer Vision and Pattern Recognition; 2018. https://doi.org/10.1109/CVPR.2018.00675.
[14] Chen S, Ma K, Zheng Y. Med3D: Transfer Learning for 3D Medical Image Analysis. ArXiv Preprint 2019;https://arxiv.org/abs/1904.00625.
[15] Chen X, Liang C, Huang D, Real E, Wang K, Liu Y, et al. Symbolic Discovery of Optimization Algorithms. ArXiv Preprint 2023;https://arxiv.org/abs/2302.06675.
[16] Dearnaley D, Syndikus I, Mossop H, Khoo V, Birtle A, Bloomfield D, et al. Conventional versus hypofractionated high-dose intensity-modulated radiotherapy for prostate cancer: 5-year outcomes of the randomised, non-inferiority, phase 3 CHHiP trial. Lancet Oncol 2016;17:1047–60. https://doi.org/10.1016/S1470-2045(16)30102-4.
[17] Dunlop A, Mitchell A, Tree A, Barnes H, Bower L, Chick J, et al. Daily adaptive radiotherapy for patients with prostate cancer using a high field MR-linac: Initial clinical experiences and assessment of delivered doses compared to a C-arm linac. Clin Transl Radiat Oncol 2020;23:35–42. https://doi.org/10.1016/j.ctro.2020.04.011.
[18] Catton CN, Lukka H, Gu C-S, Martin JM, Supiot S, Chung PW, et al. Randomized trial of a hypofractionated radiation regimen for the treatment of localized prostate cancer. J Clin Oncol 2017;35:1884–90. https://doi.org/10.1200/JCO.
[19] Hoffmann L, Alber M, Söhn M, Elstrøm UV. Validation of the Acuros XB dose calculation algorithm versus Monte Carlo for clinical treatment plans. Med Phys 2018;45:3909–15. https://doi.org/10.1002/mp.13053.
[20] Harms J, Pogue JA, Cardenas CE, Stanley DN, Cardan R, Popple R. Automated evaluation for rapid implementation of knowledge-based radiotherapy planning models. J Appl Clin Med Phys 2023;24. https://doi.org/10.1002/acm2.14152.
[21] Lou Z, Cheng C, Mao R, Li D, Tian L, Li B, et al. A novel automated planning approach for multi-anatomical sites cancer in Raystation treatment planning system. Phys Med 2023;109. https://doi.org/10.1016/j.ejmp.2023.102586.
[22] Babier A, Mahmood R, Zhang B, Alves VGL, Barragán-Montero AM, Beaudry J, et al. OpenKBP-Opt: An international and reproducible evaluation of 76 knowledge-based planning pipelines. Phys Med Biol 2022;67. https://doi.org/10.1088/1361-6560/ac8044.
[23] Nguyen D, Jia X, Sher D, Lin M-H, Iqbal Z, Liu H, et al. 3D radiotherapy dose prediction on head and neck cancer patients with a hierarchically densely connected U-net deep learning architecture. Phys Med Biol 2019;64:065020. https://doi.org/10.1088/1361-6560/ab039b.
[24] Fan J, Wang J, Chen Z, Hu C, Zhang Z, Hu W. Automatic treatment planning based on three-dimensional dose distribution predicted from deep learning technique. Med Phys 2019;46:370–81. https://doi.org/10.1002/mp.13271.
[25] Li X, Zhang J, Sheng Y, Chang Y, Yin FF, Ge Y, et al. Automatic IMRT planning via static field fluence prediction (AIP-SFFP): A deep learning algorithm for real-time prostate treatment planning. Phys Med Biol 2020;65. https://doi.org/10.1088/1361-6560/aba5eb.
[26] Yuan Z, Wang Y, Hu P, Zhang D, Yan B, Lu HM, et al. Accelerate treatment planning process using deep learning generated fluence maps for cervical cancer radiation therapy. Med Phys 2022;49:2631–41. https://doi.org/10.1002/mp.15530.
[27] Hrinivich WT, Bhattacharya M, Mekki L, McNutt T, Jia X, Li H, et al. Clinical VMAT machine parameter optimization for localized prostate cancer using deep reinforcement learning. Med Phys 2024;51:3972–84. https://doi.org/10.1002/mp.17100.
[28] Shen C, Chen L, Gonzalez Y, Jia X. Improving efficiency of training a virtual treatment planner network via knowledge-guided deep reinforcement learning for intelligent automatic treatment planning of radiotherapy. Med Phys 2021;48:1909–20. https://doi.org/10.1002/mp.14712.
[29] Li C, Guo Y, Lin X, Feng X, Xu D, Yang R. Deep reinforcement learning in radiation therapy planning optimization: A comprehensive review. Phys Med 2024;125. https://doi.org/10.1016/j.ejmp.2024.104498.
[30] Künzel LA, Leibfarth S, Dohm OS, Müller AC, Zips D, Thorwarth D. Automatic VMAT planning for post-operative prostate cancer cases using particle swarm optimization: A proof of concept study. Phys Med 2020;69:101–9. https://doi.org/10.1016/j.ejmp.2019.12.007.
[31] Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is all you need. IEEE Ind Appl Mag 2017;8:8–15. https://doi.org/arXiv:1706.03762.
[32] Xu Y, Xia W, Ren W, Ma M, Men K, Dai J. Is it necessary to perform measurement-based patient-specific quality assurance for online adaptive radiotherapy with Elekta Unity MR-Linac? J Appl Clin Med Phys 2024;25. https://doi.org/10.1002/acm2.14175.
[33] Zhao X, Stanley DN, Cardenas CE, Harms J, Popple RA. Do we need patient-specific QA for adaptively generated plans? Retrospective evaluation of delivered online adaptive treatment plans on Varian Ethos. J Appl Clin Med Phys 2023;24. https://doi.org/10.1002/acm2.13876.