
Demo: Generative AI helps Radiotherapy Planning with User Preference

Riqiang Gao¹ Simon Arberet¹ Martin Kraus¹ Han Liu¹ Wilko F.A.R. Verbakel²
Dorin Comaniciu¹ Florin C. Ghesu¹ Ali Kamen¹

¹Digital Technology and Innovation, Siemens Healthineers

²Varian Medical Systems, Siemens Healthineers

Abstract

Radiotherapy planning is a highly complex process that often varies significantly across institutions and individual planners. Most existing deep learning approaches for 3D dose prediction rely on reference plans as ground truth during training, which can inadvertently bias models toward specific planning styles or institutional preferences. In this study, we introduce a novel generative model that predicts 3D dose distributions based solely on user-defined preference “flavors”. These customizable preferences enable planners to prioritize specific trade-offs between organs-at-risk (OARs) and planning target volumes (PTVs), offering greater flexibility and personalization. Designed for seamless integration with clinical treatment planning systems, our approach assists users in generating high-quality plans efficiently. Comparative evaluations demonstrate that our method can surpasses the Varian RapidPlanTM model in both adaptability and plan quality in some scenarios. Demo video: <https://huggingface.co/Jungle15/DoseProposerDemo>.

1 Introduction

Radiotherapy (RT) plays a pivotal role in cancer treatment, being applicable for about half of all cancer cases [18]. In recent years, deep learning has been integrated into numerous stages of the RT planning workflow, including tasks such as dose prediction [19, 5, 11, 30], fluence map generation [28, 3], end-to-end including multileaf collimator (MLC) leaf sequencing [10, 16], and dose calculation [29]. Among these, 3D dose prediction has gained particular attention due to its potential to streamline the planning process. By leveraging inputs like CT images, delineated structures such as planning target volumes (PTVs) and organs at risk (OARs), and configuration data, deep models can estimate full 3D dose distributions without relying on time-consuming optimization [5, 11]. This capability has broad applications: it can inform objective functions during plan optimization [4], facilitate quality assurance [13, 14], and serve as a foundation for fully automated, AI-driven planning systems [8, 10].

RapidPlanTM [25] of Varian is a knowledge-based treatment planning system that uses conventional machine learning and a library of prior high-quality plans to model the relationship between anatomy and dose distribution. It generates personalized dose-volume histograms (DVHs) [7] and optimization objectives, improving plan consistency, quality, and efficiency across various treatment sites.

Despite its clinical popularity, RapidPlan has notable limitations: (1) it relies on DVH predictions, which can miss critical spatial dose details, reducing the accuracy of personalized optimization for OARs, especially in anatomically complex cases; (2) its PCA-based regression pipeline is typically trained on dozens of plans, limiting generalizability across institutions and patient populations. Users need to train institution-specific models to reflect guidelines and preferences.

Recent deep learning models enable direct 3D dose prediction from anatomy and contours, offering voxel-level accuracy and spatial awareness with large-scale data [5, 9]. This enhances personalization,




Figure 1: Our flexible dose proposer pipeline. Stage I pre-trains a foundational dose decoder, regularizing the dose prediction to be realistic. The second stage trains a flexible dose prediction model with heterogeneous conditions.

interpretability, and supports adaptive, automated planning. However, user interaction with a single model remains unexplored. Moreover, dose prediction alone is not a deliverable plan, and while dose mimicking [4] and end-to-end AI pipelines [10] have been studied in research settings, integration into commercial clinical platforms is still limited.

To address these gaps, we contribute the following in this work:

- We propose a two-stage training framework with a foundational dose decoder, improving the stability of training in complex scenarios.
- To our knowledge, this is the first dose prediction model with interactive sliders, enabling real-time customization of trade-offs between target homogeneity and OAR sparing.
- We integrate our AI model with a widely used clinical treatment planning system and demonstrate its effectiveness in generating high-quality radiotherapy plans.

2 Method

2.1 Problem Definition and Intuition

The homogeneity index (HI) and conformity index (CI) are two main metrics to evaluate plan quality on PTVs. The CI evaluates how well the prescribed dose conforms to the shape and size of the PTV. The HI measures the uniformity of the dose distribution within the PTV. $HI = (D_{05} - D_{95}) / D_{50}$, and $CI = V_{covered} / V_{PTV}$. D_t represents dose received by $t\%$ of the PTV (i.e., $(100-t)$ -th percentile). V_{PTV} is volume size of the PTV, and $V_{covered}$ is volume size of the PTV receiving at least prescribed dose. On the other hand, OAR sparing (e.g., mean doses of OARs), is often inversely related to PTV coverage and uniformity. Traditional planning models often fail to capture diverse user preferences, requiring separate models for different styles.

Generative AI has demonstrated strong capabilities in producing images and videos based on user prompts [22, 21, 26]. However, current AI-based dose prediction models largely overlook individual planning preferences. This work aims to reduce that gap. After dose prediction, we convert the 3D dose to optimization objectives inspired by [9] to generate deliverable plans (more in Appendix 11).

2.2 Flexible Dose Proposer

Our flexible dose proposer (FDP) is shown in Figure 1. Stage I trains a foundational dose decoder using a VQ-VAE architecture [24]. Stage II introduces multi-conditional inputs to enable context-aware dose prediction. For fast inference on 3D images, we deliberately exclude computationally intensive or iterative models such as diffusion-based approaches [15], although the core concept of FDP remains adaptable to other generative frameworks.

Stage I. The VQ-VAE pretraining is inspired by the success of latent diffusion models [22, 26], which use VAEs primarily for data compression to reduce inference time and computational cost. In contrast, our goal of Stage I pretraining is to stabilize Stage II training, especially under complex conditions like user preferences. Since we adopt one-step generation via GANs [12], our inference time is

Cohort Index	Cohort 0	Cohort 1	Cohort 2	Cohort 3	Cohort 4	Cohort 5
Train/Valid/Test	370/48/54	147/17/19	128/15/17	103/14/12	52/8/7	20/1/4

Table 1: Data distribution used in Stage II training. The Stage I is pre-trained with 31K doses.

minimal (shown in the demo) compared to diffusion-based approaches. Denote real/reference data and prediction as x and $\hat{x}$ respectively, the adversarial loss is defined as $L_{adv}(x, \hat{x}) = \mathbb{E}_x[(D(x) - 1)^2] + \mathbb{E}_{\hat{x}}[D(\hat{x})^2]$ and L_{vq} is the vector quantisation loss following [24]. To stabilize the training, we further introduce a uniform loss [27] to regularize the latent space:

$$\mathcal{L}_{\text{stage1}} = \underbrace{\mathbb{E}_i[\|x_i - \hat{x}_i\|]}_{\text{Reconstruction}} + \beta L_{vq} + L_{adv}(x, \hat{x}) + \underbrace{\lambda \cdot \log \left(\mathbb{E}_{i < j} \left[\exp(-t\|\hat{z}_i - \hat{z}_j\|^2) \right] \right)}_{\text{Uniformity}}. \quad (1)$$

Model structure following MONAI [6] for VQVAE. Network parameters are in Appendix 7.

Stage II. We utilize the foundational dose decoder for Stage II to stabilize the training. The CT image and RT structures have been concatenated as multi-channel input for the image encoder [11]. User preferences together with beam/angle plates, are encoded with adaptive instance normalization [17]. The encoder structure follows the MedNext [23] style, and network parameters are in Appendix 7.

Loss functions for Stage II include reconstruction losses of image space $\{x, \hat{x}\}$ and latent space $\{z, \hat{z}\}$, adversarial loss of the image space, and objective loss L_{obj} , which are defined as follows:

$$\mathcal{L}_{\text{stage2}}^{(i)} = \underbrace{\|x_i - \hat{x}_i\|}_{\text{Image Reconstruction}} + \underbrace{\|z_i - \hat{z}_i\|}_{\text{Latent Reconstruction}} + L_{adv}(x_i, \hat{x}_i) + \mathcal{L}_{\text{obj}}^{(i)}. \quad (2)$$

The L_{obj} is the inconsistency penalty between the user preference prompts and the prediction metrics:

$$\mathcal{L}_{\text{obj}}^{(i)} = \underbrace{\|\tilde{h} - \hat{h}\|}_{\text{ptv HI preference}} + \underbrace{\|p - \hat{p}\|}_{\text{ptv dose alignment}} + \underbrace{\|\tilde{w} \cdot u_{\text{oar}} - \hat{u}_{\text{oar}}\|}_{\text{oar-sparing preference}}, \quad (3)$$

where $\tilde{h}$ and $\hat{h}$ denote the homogeneity indices of user preference (i.e., slider value) and prediction, respectively; p and $\hat{p}$ represent the PTV mean dose of the reference and predicted plans; $\tilde{w}$ is the user-defined preference for OAR, and u_{oar} is the OAR mean dose of the reference plan. For simplicity, we omit writing out $\mathcal{L}_{\text{obj}}^{(i)}$ for each structure.

During training, the sliding bar values $\{\tilde{h}, \tilde{w}\}$ are randomly sampled within predefined ranges. This sampling strategy enhances the alignment between user inputs and the model’s predictions. In parallel, various loss regularizations—such as reconstruction loss and adversarial loss—are employed to guide learning. Additionally, Stage I pre-training introduces an extra layer of regularization. Together, these strategies encourage the model to integrate both patient-specific information (e.g., spatial distribution of PTV/OAR) and user interactions via the sliding bars into its dose predictions.

3 Experiments

Dataset. We have included six cohorts of head-and-neck cancer cases, as shown in Table 1. More details will be disclosed after the double-blind review. The Stage I pre-training decoder has used a much larger scale dataset with 31K doses (not necessarily from high-quality plans).

OAR	SpinalCord05	Larynx-PTV	Lips	Mandible-PTV	OCavity-PTV	ParotidCon-PTV	ParotidIps-PTV	Esophagus
RP std	3.02	3.49	2.54	2.93	4.04	3.38	2.96	3.21
ours std	1.18	1.78	1.09	1.27	1.82	1.22	1.78	1.27
RP mean	3.25	6.29	4.71	6.41	7.37	4.51	5.54	2.43
ours mean	1.69	2.94	2.06	2.01	4.13	1.92	2.92	1.18

OAR	SubmandL-PTV	Shoulders	SubmandR-PTV	PosteriorNeck	PharConst-PTV	BrainStem 03	Trachea	<i>better count</i>
RP std	3.25	1.38	2.70	4.30	3.62	0.92	3.40	0
ours std	1.38	0.51	1.45	0.92	1.78	0.86	1.68	15
RP mean	4.43	0.97	6.29	8.15	4.36	1.49	3.42	1
ours mean	4.51	0.62	5.37	1.19	3.37	1.08	2.01	14

Table 2: Statistics of intra-patient differences of expected and achieved DVHs ($\downarrow$). RP: RapidPlan [20]. ours: the predicted dose from FDP and objectives extracted from AI dose for Eclipse planning.

OAR	SpinalCord 05	Larynx-PTV	Lips	Mandible-PTV	OCavity-PTV	ParotidCon-PTV	ParotidIps-PTV	Esophagus
RP std	0.98	2.49	2.69	3.03	6.34	1.29	1.61	0.31
ours std	0.73	2.76	1.92	1.05	3.61	0.56	0.90	0.32
RP mean	0.43	4.24	2.02	2.22	6.00	1.66	2.25	0.23
ours mean	0.45	1.55	1.14	0.83	3.19	0.64	0.84	0.23

OAR	SubmandL-PTV	Shoulders	SubmandR-PTV	Posterior Neck	PharConst-PTV	BrainStem 03	Trachea	better count
RP std	7.12	0.14	5.70	3.76	3.73	1.76	6.51	3
ours std	3.93	0.27	4.17	2.53	1.71	0.90	1.86	12
RP mean	7.53	0.12	5.05	3.92	4.30	1.13	2.49	2
ours mean	4.38	0.13	4.81	2.10	1.48	0.61	0.99	12

Table 3: Statistics of inter-patient differences of expected and achieved DVHs ($\downarrow$). RP: RapidPlan [20]. ours: the predicted dose from FDP and objectives extracted from AI dose for Eclipse planning.




Figure 2: Estimated vs. achieved DVHs of RapidPlan and our FDP model. RapidPlan estimations are not well aligned with achieved plans, may due to its limited generalization to unseen patients.

Evaluation Metrics. Dose-volume histograms (DVHs) [7] are graphical representations in RT planning, providing a quantitative assessment of dose distribution across target volumes and surrounding organs at risk. Given DVHs of a specific organ (DVH_i has been normalized with same index of $Volume$, we compare the differences of $Dose$), intra-patient differences are defined as $D = DVH_{\text{expect}} - DVH_{\text{achieved}}$, where DVH_{expect} and DVH_{achieved} are the expected and achieved DVHs for the organ. We define the metrics (mean and standard deviation) and show illustration about intra- and inter-patient differences at Appendix 8. Lower metric values represent the expected plan and achieved plan are closer, indicating better DVH estimations.

We use a high-quality RapidPlan [20, 1] as baseline. The RT structure creation of all test plans followed the RapidPlan requirements. All the reference plans for AI are based on this RapidPlan.

Results. Table 2 and Table 3 shows the statistics of intra-patient and inter-patient differences, respectively. In both scenarios, lower values especially lower standard deviation (std) indicate better estimations of DVHs. Our FDP model outperforms RapidPlan in both scenarios.

Figure 2 shows some examples about estimated and achieved DVHs of RapidPlan and our FDP model. We notice that RapidPlan model are less robust than the proposed method, indicated by the less alignment between the blue lines and orange lines.

Table 4 shows the quality comparison of achieved plans of RapidPlan and our FDP. For OAR metrics (mean dose), FDP has achieved more larger percent on “better” than RapidPlan for OAR sparing count (14 over 0), and 1 over 0 for PTV homogeneity & conformity count.

OAR	SpinalCor 05	Larynx-PTV	Lips	Mandible-PTV	OCavity-PTV	ParotidCon-PTV	ParotidIps-PTV	Esophagus
better	47.50	30.00	47.50	31.25	64.56	32.50	48.68	30.23
worse	5.00	21.43	0.00	1.25	1.27	6.25	5.26	4.65
similar	47.50	48.57	52.50	67.50	34.18	61.25	46.05	65.12

OAR	SubmandL-PTV	Shoulders	SubmandR-PTV	Posterior Neck	PharConst-PTV	BrainStem 03	Trachea	OAR count
better	59.57	0.00	71.11	12.50	56.16	7.50	51.16	14
worse	2.13	0.00	6.67	6.25	2.74	3.75	0.00	0
similar	38.30	100.00	22.22	81.25	41.10	88.75	48.84	-

PTV	HI (PTVHigh)	CI (PTVHigh)	HI (PTVMid)	CI (PTVMid)	HI (PTVLow)	CI (PTVLow)	PTV count
better	0.00	0.00	0.00	4.00	1.43	4.29	1
worse	0.00	0.00	0.00	4.00	0.00	4.29	0
similar	100.00	100.00	100.00	92.00	98.57	91.43	-

Table 4: Structure-wise plan quality comparison between RapidPlan and our FDP. Each structure is categorized as “better”, “worse”, or “similar”, indicating whether the plan generated by FDP shows improved, reduced, or comparable quality relative to RapidPlan. The “similar” thresholds are 1 Gray for OARs; 0.015 for PTV indices (HI and CI as defined in Sec. 2.1). Percentages (%) are reported for each category per structure, while the **bold** column presents the column counts.




Figure 3: Demonstration of user preference with sliding bars. The left/right of sliding bars represents better/worse quality for the specific structure respectively. The Preference 1 (P1) has the preferences of OAR sparing over PTV homogeneity, and Preference 2 (P2) has the opposite. The preferences have been captured in the dose prediction, as shown in the DVHs in the right panel.

OAR	SpinalCord05	Larynx-PTV	Lips	Mandible-PTV	OCavity-PTV	ParotidCon-PTV	ParotidIps-PTV	Esophagus
P1 MeanDose	14.59	14.41	13.83	26.25	28.29	4.84	22.03	4.89
P2 MeanDose	16.33	15.47	13.44	27.25	28.04	5.24	23.4	5.55
P1 MaxDose	34.2	62.2	35.3	71.3	69.2	12.5	68.4	16.0
P2 MaxDose	38.3	63.3	36.6	71.5	69.7	15.5	68.8	18.7

OAR	CochleaL	CochleaR	Shoulders	SubmandR-PTV	PosteriorNeck	PharConst-PTV	BrainStem03	Trachea
P1 MeanDose	6.28	1.92	4.13	10.01	26.04	19.82	4.64	9.8
P2 MeanDose	6.33	1.93	4.09	11.51	26.3	20.41	4.83	11.21
P1 MaxDose	9.2	1.9	61.4	27.5	73.9	68.8	21.1	44.2
P2 MaxDose	9.6	1.9	60.6	27.1	73.8	69.4	20.7	47.5

Table 5: The achieved mean dose from Eclipse with different flavors based on Figure 3.

Sliding Bars to Adapt the Preference. As demonstrated in Figure 3, we include two user preferences (P1: focusing more on OAR sparing, P2: focusing PTV over OAR). our FDP can have dose prediction responding to the preferences. The planning system achieved plans based on different flavors have been shown in Table 5. More demonstration details can be found in Appendix 5 and 6.

Ablation of Stage I. In the context of user preference encoding, we find that without pre-trained Stage I, the reconstruction error is slightly larger, as shown in Table 6 (MAE represents the mean absolute error masked by 5 Gy isodose lines). More importantly, the predicted dose is less realistic as boundary artifacts of PTVs (black arrow) and OARs (white arrow) are shown as in Figure 3.

Methods	MAE ($\downarrow$)
w/o. S1-pretrain	2.63
w. S1-pretrain	2.56

Table 6: test set MAEs of with-out vs. with Stage I pretrain.




Figure 4: Example demonstrates the value of Stage I pre-training.

4 Discussion

Conclusion. We propose a novel two-stage training framework for flexible dose prediction, representing a pioneering approach to integrate user preferences (e.g., OAR sparing, PTV homogeneity) with sliders. Unlike prior work that primarily emphasizes dose prediction accuracy, we validate the clinical utility of our AI-generated plans within a widely adopted treatment planning system. Compared to RapidPlan, our model demonstrates superior generalization in plan quality. We hope this work encourages further research and accelerates AI application for routine radiotherapy practice.

Limitation and Future Work. Our current evaluation focuses on head-and-neck cancer site, the most challenging region for RT planning due to its complex anatomy and stringent dose constraints. Extending our approach to other treatment sites is expected to be straightforward and is part of our future work. Second, the training paradigms of RapidPlan (e.g., conventional methods with smaller datasets) and our model (e.g., deep learning with larger-scale data) differ fundamentally, and clinical quality comparison is inherent complex and subjective. More rigorous comparisons are planned for future investigation in diverse clinical scenarios.

Disclaimer. The information in this paper is based on research results that are not commercially available. Future commercial availability cannot be guaranteed.

References

- [1] V. M. Affairs. Bilateral head&neck 70/63/56gy (hn-sib-bpi) [rapidplan]. <https://medicalaffairs.varian.com/hn-sib-bpi-rapidplan-vmat2>, 2024. Accessed: 2024-10-19. 4, 14
- [2] V. M. Affairs. Rapidplan knowledge-based planning. <https://medicalaffairs.varian.com/halcyon-case-studies>, 2024. Accessed: 2024-09-19. 14
- [3] S. Arberet, F. C. Ghesu, R. Gao, M. Kraus, J. Sackett, E. Kuusela, and A. Kamen. Beam’s eye view to fluence maps 3d network for ultra fast vmat radiotherapy planning. *Medical Physics*, 2025. 1
- [4] A. Babier, R. Mahmood, B. Zhang, V. G. L. Alves, A. M. Barragán, B. Barragán-Montero, J. Beaudry, C. E. Cardenas, Y. Chang, Z. Chen, J. Chun, K. Diaz, H. David Eraso, E. Faustmann, S. Gaj, S. Gay, M. Gronberg, B. Guo, J. He, G. Heilemann, S. Hira, Y. Huang, F. Ji, D. Jiang, J. Carlo, J. Giraldo, H. Lee, J. Lian, S. Liu, K.-C. Liu, J. Jos´, J. Marrugo, K. Miki, K. Nakamura, T. Netherton, D. Nguyen, H. Nourzadeh, A. F. I. Osman, Z. Peng, J. Darío, Q. Muñoz, M. Muñoz, C. Rams, D. J. Rhee, J. David Rodriguez, H. Shan, J. V. Siebers, M. H. Soomro, K. Sun, A. A. Andrés, U. Hoyos, C. Valderrama, R. Verbeek, E. Wang, S. Willems, Q. Wu, X. Xu, S. Yang, L. Yuan, S. Zhu, L. Zimmermann, K. L. Moore, T. G. Purdie, A. L. McNiven, and T. C. Y. Chan. OpenKBP-Opt: An international and reproducible evaluation of 76 knowledge-based planning pipelines. *Physics in Medicine & Biology*, 2022. 1, 2
- [5] A. Babier, B. Zhang, R. Mahmood, K. L. Moore, T. G. Purdie, A. L. McNiven, and T. C. Y. Chan. OpenKBP: The open-access knowledge-based planning grand challenge. *Medical Physics*, 2020. 1
- [6] M. J. Cardoso, W. Li, R. Brown, N. Ma, E. Kerfoot, Y. Wang, B. Murrey, A. Myronenko, C. Zhao, D. Yang, et al. Monai: An open-source framework for deep learning in healthcare. *arXiv preprint arXiv:2211.02701*, 2022. 3
- [7] R. E. Drzymala, R. Mohan, L. Brewster, J. Chu, M. Goitein, W. Harms, B. S. ’ And, and M. Urie. DOSE-VOLUME HISTOGRAMS. *Inr. J. Radrarion Oncology Biol. Phys*, 1991. 1, 4
- [8] J. Fan, J. Wang, Z. Chen, C. Hu, Z. Zhang, and W. Hu. Automatic treatment planning based on three-dimensional dose distribution predicted from deep learning technique. *Medical physics*, 2019. 1
- [9] R. Gao, M. Diallo, H. Liu, A. Magliari, J. Sackett, W. Verbakel, S. Meyers, R. Mcbeth, M. Zarepisheh, S. Arberet, et al. Automating high quality rt planning at scale. *arXiv preprint arXiv:2501.11803*, 2025. 1, 2
- [10] R. Gao, F.-C. Ghesu, S. Arberet, S. Basiri, E. Kuusela, M. Kraus, D. Comaniciu, and A. Kamen. Multi-agent reinforcement learning meets leaf sequencing in radiotherapy. In *International Conference on Machine Learning*, 2024. 1, 2
- [11] R. Gao, B. Lou, Z. Xu, D. Comaniciu, and A. Kamen. Flexible-cm gan: Towards precise 3d dose prediction in radiotherapy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 1, 3
- [12] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 2014. 2
- [13] M. P. Gronberg, B. M. Beadle, A. S. Garden, H. Skinner, S. Gay, T. Netherton, W. Cao, C. E. Cardenas, C. Chung, D. T. Fuentes, C. D. Fuller, R. M. Howell, A. Jhingran, T. Y. Lim, B. Marquez, R. Mumme, A. M. Olanrewaju, C. B. Peterson, I. Vazquez, T. J. Whitaker, Z. Wooten, M. Yang, and L. E. Court. Deep Learning–Based Dose Prediction for Automated, Individualized Quality Assurance of Head and Neck Radiation Therapy Plans. *Practical Radiation Oncology*, 2023. 1

- [14] M. P. Gronberg, A. Jhingran, T. J. Netherton, S. S. Gay, C. E. Cardenas, C. Chung, D. Fuentes, C. D. Fuller, R. M. Howell, M. Khan, T. Y. Lim, B. Marquez, A. M. Olanrewaju, C. B. Peterson, I. Vazquez, T. J. Whitaker, Z. Wooten, M. Yang, and L. E. Court. Deep learning–based dose prediction to improve the plan quality of volumetric modulated arc therapy for gynecologic cancers. *Medical Physics*, 2023. 1
- [15] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 2
- [16] W. T. Hrinivich, M. Bhattacharya, L. Mekki, T. McNutt, X. Jia, H. Li, D. Y. Song, and J. Lee. Clinical vmat machine parameter optimization for localized prostate cancer using deep reinforcement learning. *Medical physics*, 2024. 1
- [17] X. Huang and S. Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE international conference on computer vision*, pages 1501–1510, 2017. 3
- [18] E. Huynh, A. Hosny, C. Guthier, D. S. Bitterman, S. F. Petit, D. A. Haas-Kogan, B. Kann, H. J. Aerts, and R. H. Mak. Artificial intelligence in radiation oncology. *Nature Reviews Clinical Oncology*, 2020. 1
- [19] X. Kui, F. Liu, M. Yang, H. Wang, C. Liu, D. Huang, Q. Li, L. Chen, and B. Zou. A review of dose prediction methods for tumor radiation therapy. *Meta-Radiology*, 2024. 1
- [20] A. Magliari, R. Clark, L. Rosa, and S. Beriwal. Hn-sib-bpi: A single click, sub-site specific, dosimetric scorecard tuned rapidplan model created from a foundation model for treating head and neck with bilateral neck. *Medical Dosimetry*, 2024. 3, 4
- [21] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022. 2
- [22] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022. 2
- [23] S. Roy, G. Koehler, C. Ulrich, M. Baumgartner, J. Petersen, F. Isensee, P. F. Jaeger, and K. H. Maier-Hein. Mednext: transformer-driven scaling of convnets for medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2023. 3
- [24] A. Van Den Oord, O. Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 2, 3
- [25] Varian. Rapidplan knowledge-based planning. <https://www.varian.com/products/radiotherapy/treatment-planning/rapidplan-knowledge-based-planning>, 2024. Accessed: 2024-09-19. 1
- [26] T. Wan, A. Wang, B. Ai, B. Wen, C. Mao, C.-W. Xie, D. Chen, F. Yu, H. Zhao, J. Yang, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025. 2
- [27] T. Wang and P. Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International conference on machine learning*, pages 9929–9939. PMLR, 2020. 3
- [28] W. Wang, Y. Sheng, C. Wang, J. Zhang, X. Li, M. Palta, B. Czito, C. G. Willett, Q. Wu, Y. Ge, F. F. Yin, and Q. J. Wu. Fluence Map Prediction Using Deep Learning Models – Direct Plan Generation for Pancreas Stereotactic Body Radiation Therapy. *Frontiers in Artificial Intelligence*, 2020. 1
- [29] Y. Xing, D. Nguyen, W. Lu, M. Yang, and S. Jiang. A feasibility study on deep learning-based radiotherapy dose calculation. *Medical physics*, 2020. 1
- [30] Y. Zhang, C. Li, L. Zhong, Z. Chen, W. Yang, and X. Wang. Dosediff: distance-aware diffusion model for dose prediction in radiotherapy. *IEEE Transactions on Medical Imaging*, 2024. 1

5 User Interface with Sliding Bars (Demo 1)




Figure 5: Screenshot of the Gradio demo interface. A full demo video recording can be found in an [anonymous HuggingFace repo](#).

This demo includes four steps, executed when clicking the corresponding buttons in the upper panel:

- Load yaml file. The yaml file path can be inputted by user. The yaml file includes data parameters (e.g., DICOM data paths, PTV prescribed doses, RT structures mapping) and model parameters (e.g., trained model path, model configurations, save output roots).
- Load trained model. Load the trained model with PyTorch code.
- Load and preprocess DICOM data. prepare the data from DICOM for deep learning input. Currently our research code is pure Python not optimized from speed, this process takes about half minutes. This is one-time running for each patient.
- Inference with sliding bar parameters. The model forward only takes about 30 ms for model forward, and takes about 1.5 seconds to visualize the results. When moving the sliding and click the 'inference' button, the results can be shown almost real-time.

As shown in Figure 5, the **upper panel** displays the buttons for the four main steps. The **left panel** presents the sliding bar interface along with the running logs. The **middle section on the right** shows the predicted dose overlaid on the CT image, while the **bottom right panel** compares the dose-volume histograms (DVHs) of the current prediction (after adjusting the sliding bars) versus the previous prediction (before sliding bars adjustment).

6 Deliverable Plan in Eclipse Derived from Our FDP (Demo 2)




Figure 6: Screenshot of deliverable plan in Eclipse derived from our flexible dose proposer (FDP). The upper right shows the BEV of the patients with RT structures and Multi-Leaf Collimator overlaid. Other panels show the achieved dose overlaid with CT. A full demo video recording can be downloaded from an [anonymous HuggingFace repo](#).

7 Model Structures of Our Flexible Dose Proposer

Stage I model structure parameters, following the style of MONAI¹:

Table 7: Stage I network structure parameters.

Parameter	Value
spatial_dims	3
in_channels	1
out_channels	1
channels	[32, 256, 512, 512]
downsample_parameters	[[2, 4, 1, 1], [2, 4, 1, 1], [1, 3, 1, 1], [1, 3, 1, 1]]
upsample_parameters	[[1, 3, 1, 1, 0], [1, 3, 1, 1, 0], [2, 4, 1, 1, 0], [2, 4, 1, 1, 0]]
num_res_channels	256
num_res_layers	1
num_embeddings	512
embedding_dim	4

Stage II image encoder structure parameters, following the style of MedNext²:

Table 8: Stage I network structure parameters.

Parameter	Value
in_channels	7
n_channels	24
exp_r	[2, 3, 4, 4, 4, 4, 4, 3, 2]
kernel_size	3
deep_supervision	False
do_res	True
do_res_up_down	True
block_counts	[2, 3, 4, 4, 8, 8, 8, 4, 3]

¹<https://github.com/Project-MONAI/MONAI/blob/dev/monai/networks/nets/vqvae.py>

²https://github.com/MIC-DKFZ/MedNeXt/blob/main/nnunet_mednext/network_architecture/mednextv1/MedNextV1.py

8 Illustration and Metrics about DVH Differences




Figure 7: The cartoon illustration about intra- and inter- patient differences.

A cartoon illustration is shown in Figure 7 with DVHs. The solid line represents the achieved plan, and dash line represents the target (estimated) plan. Each subfigure plot the DVHs for one organ per patient. The metrics of intra- and inter- patient differences are defined below:

$$\mu_{\text{tra}} = \frac{1}{N} \frac{1}{L} \sum_{j=1}^N \sum_{i=1}^L D_i^j, \sigma_{\text{tra}} = \frac{1}{N} \sum_{j=1}^N \sqrt{\frac{1}{L} \sum_{i=1}^L (D_i^j - \mu_{\text{tra}}^j)^2} \quad (4)$$

$$\mu_{\text{ter}} = \frac{1}{N} \sum_{j=1}^N D_m^j, \sigma_{\text{ter}} = \sqrt{\frac{1}{N} \sum_{j=1}^N (D_m^j - \mu_{\text{ter}})^2} \quad (5)$$

where the L is the length of DVH dimensions, and N is the number of testing patients.

The assumption is that DVH estimations are more accurate when the achieved DVHs are well-aligned following complex optimization using a comprehensive objective list within the Eclipse treatment planning system.

Since the mean of the differences ($D_{\text{intra}}, D_{\text{inter}}$) may be offset by consistent margins with some rule-based methods, so we pay less attention to the mean values in this study and shade the comparison in the tables. Constrastively, large standard deviations are indicative of poor estimation consistency and reliability.

9 Results when Planned with IMRT

The tables presented in the main text reflect results from VMAT-based planning. Below, we provide the corresponding outcomes using IMRT. The consistency across both approaches highlights the robustness of our model.

OAR	SpinalCord 05	Larynx-PTV	Lips	Mandible-PTV	OCavity-PTV	ParotidCon-PTV	ParotidIps-PTV	Esophagus
RP std	3.92	2.90	3.15	2.91	4.21	3.25	3.09	3.85
ours std	1.63	1.71	1.26	1.49	2.15	1.06	1.21	1.63
RP mean	5.62	10.75	5.39	8.08	8.68	6.33	7.85	3.89
ours mean	2.57	4.33	1.78	1.99	3.31	1.92	1.90	1.69

OAR	SubmandL-PTV	Shoulder	SubmandR-PTV	Posterior Neck	PharConst-PTV	Brain Stem 03	Trachea	<i>better count</i>
RP std	3.65	1.59	2.30	4.64	3.16	1.49	3.66	0
ours std	1.29	0.57	1.30	1.25	1.61	1.02	1.72	15
RP mean	6.36	1.33	8.06	10.42	8.02	2.47	5.42	0
ours mean	2.96	0.67	3.53	2.58	2.48	1.41	2.83	15

Table 9: Statistics of intra-patient differences of expected and achieved DVHs ($\downarrow$).

OAR	SpinalCord 05	Larynx-PTV	Lips	Mandible-PTV	OCavity-PTV	ParotidCon-PTV	ParotidIps-PTV	Esophagus
RP std	2.41	4.44	2.76	2.72	5.28	2.20	2.61	0.57
ours std	1.92	2.52	2.25	0.93	3.62	1.08	1.05	0.54
RP mean	0.88	8.80	1.98	3.05	6.07	2.93	3.66	0.40
ours mean	0.79	4.69	1.41	0.86	2.91	0.92	1.10	0.42

OAR	SubmandL-PTV	Shoulder	SubmandR-PTV	Posterior Neck	PharConst-PTV	Brain Stem 03	Trachea	<i>better count</i>
RP std	7.72	0.23	4.99	4.18	6.14	1.41	7.70	0
ours std	3.89	0.23	2.94	2.88	2.27	0.86	3.48	14
RP mean	11.89	0.15	7.28	4.69	9.18	1.38	3.45	0
ours mean	3.46	0.15	3.22	2.40	2.79	0.73	1.64	14

Table 10: Statistics of inter-patient differences of expected and achieved DVHs ($\downarrow$).

OAR	SpinalCord 05	Larynx-PTV	Lips	Mandible-PTV	OCavity-PTV	ParotidCon-PTV	ParotidIps-PTV	Esophagus
better	36.36	17.65	37.66	10.39	50.00	10.39	13.70	21.95
worse	6.49	41.18	3.90	6.49	6.58	29.87	26.03	2.44
similar	57.14	41.18	58.44	83.12	43.42	59.74	60.27	75.61

OAR	SubmandL-PTV	Shoulders	SubmandR-PTV	Posterior Neck	PharConst-PTV	BrainStem 03	Trachea	OAR count
better	38.30	0.00	47.73	1.30	34.29	7.79	36.59	10
worse	23.40	0.00	20.45	10.39	11.43	3.90	0.00	4
similar	38.30	100.00	31.82	88.31	54.29	88.31	63.41	-

PTV	HI (PTVHigh)	CI (PTVHigh)	HI (PTVMid)	CI (PTVMid)	HI (PTVLow)	CI (PTVLow)	PTV count
better	5.20	0.00	13.64	0.00	2.98	1.49	3
worse	1.30	0.00	0.00	0.00	2.98	0.00	0
similar	93.51	100.00	86.36	100.00	94.03	98.51	-

Table 11: Quality per structure of achieved plans of RapidPlan and our FDP. “better”, “worse”, “similar” represent the plan derived from FDP has better, worse, or similar quality respectively. Percentages (%) are reported for each category per structure, while the **bold** column presents the column counts. If the percent of “better” is larger than “worse” of one structure, the “better” count will add 1.

10 Small Validation Sets Reflect RapidPlan Distribution More Than FDP

OAR	SpinalCord05	Larynx-PTV	Lips	Mandible-PTV	OCavity-PTV	ParotidCon-PTV	ParotidIps-PTV	Esophagus
RP std	1.48	1.70	1.47	2.91	2.05	2.33	1.44	2.95
ours std	1.57	1.42	0.99	1.76	1.83	1.31	1.90	1.81
RP mean	1.83	4.32	4.53	5.28	4.61	4.95	3.54	2.58
ours mean	1.82	2.96	4.59	2.62	3.93	2.00	2.83	1.62

OAR	SubmandL-PTV	Shoulders	SubmandR-PTV	Posterior Neck	PharConst-PTV	BrainStem03	Trachea	better count
RP std	3.25	1.25	2.18	1.41	2.48	1.06	4.43	6
ours std	1.58	0.45	2.65	1.74	2.66	1.33	1.91	9
RP mean	5.65	1.21	4.94	1.78	3.21	1.89	3.48	4
ours mean	5.53	0.69	6.75	2.22	5.58	0.83	1.88	11

Table 12: Statistics of intra-patient differences of expected and achieved DVHs (↓).

OAR	SpinalCord05	Larynx-PTV	Lips	Mandible-PTV	OCavity-PTV	ParotidCon-PTV	ParotidIps-PTV	Esophagus
RP std	0.12	0.29	1.37	1.16	4.49	1.11	0.66	0.39
ours std	0.12	0.41	0.67	0.85	1.87	0.63	0.68	0.39
RP mean	0.38	4.23	4.93	2.80	9.90	2.61	2.59	0.38
ours mean	0.37	0.72	4.09	1.11	2.69	0.98	1.00	0.36

OAR	SubmandL-PTV	Shoulders	SubmandR-PTV	PosteriorNeck	PharConst-PTV	BrainStem03	Trachea	better count
RP std	0.14	0.08	3.14	0.13	2.14	0.00	0.53	5
ours std	0.57	0.10	3.94	0.10	0.97	0.00	0.36	7
RP mean	8.29	0.15	4.20	0.21	5.33	1.05	1.18	3
ours mean	3.44	0.15	6.06	0.24	1.61	1.24	0.70	11

Table 13: Statistics of inter-patient differences of expected and achieved DVHs (↓).

OAR	SpinalCord 05	Larynx-PTV	Lips	Mandible-PTV	OCavity-PTV	ParotidCon-PTV	ParotidIps-PTV	Esophagus
better	25.00	0.00	50.00	75.00	75.00	25.00	50.00	50.00
worse	0.00	50.00	0.00	25.00	0.00	0.00	0.00	0.00
similar	75.00	50.00	50.00	0.00	25.00	75.00	50.00	50.00

OAR	SubmandL-PTV	Shoulders	SubmandR-PTV	Posterior Neck	PharConst-PTV	BrainStem03	Trachea	OAR count
better	50.00	0.00	66.67	0.00	100.00	0.00	75.00	10
worse	50.00	0.00	0.00	25.00	0.00	0.00	0.00	2
similar	0.00	100.00	33.33	75.00	0.00	100.00	25.00	-

PTV	high homo	high conf	mid homo	mid conf	low homo	low conf	PTV count
better	0.00	0.00	0.00	0.00	0.00	0.00	0
worse	0.00	0.00	0.00	0.00	0.00	0.00	0
similar	100.0	100.0	100.0	100.0	100.0	100.0	-

Table 14: Quality per structure of achieved plans of RapidPlan and our FDP. “better”, “worse”, “similar” represent the plan derived from FDP has better, worse, or similar quality respectively. Percentages (%) are reported for each category per structure, while the **bold** column presents the column counts. If the percent of “better” is larger than “worse” of one structure, the “better” count will add 1.

11 Translate 3D Dose Prediction to Optimization Objectives

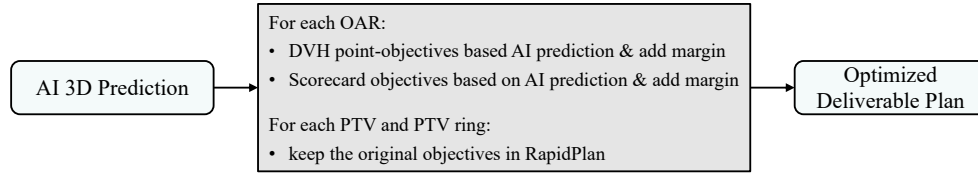


Figure 8: Illustration of translating 3D dose prediction to optimization objectives in Eclipse TPS.

Currently, Eclipse TPS does not support the optimization with 3D dose distribution as the direct objective. Therefore, we use a method to convert the 3D dose prediction into a set of dose-volume and mean dose objectives that can be utilized within Eclipse TPS for plan optimization. As shown in Figure 8, for each organ at risk (OAR) we sample the predicted dose-volume histogram (DVH) at specific volume percentages to create corresponding dose-volume objectives after adding some rule-based margins. Additionally, based on metrics from a scorecard [1], we further adding mean dose and dose-volume objectives to better align with clinical goals.

For PTVs and PTV rings, we follow the settings of the RapidPlan [2], where the objectives are not from a model estimation but from dose prescriptions.