Addressing the Gap in Oncologist Training: Generating Clinically Realistic, Suboptimal Dose Distributions through Geometry Aware Convolutions

Skylar S. Gay[1,2], Mary P. Gronberg[3], Raymond Mumme[1], Christine Chung[1], Meena Khan[1], Chelsea Pinnix[4], Sanjay Shete[1], Brent Parker[1,5], Tucker J. Netherton[1], Carlos E. Cardenas[5], and Laurence E. Court[1]

[1]Department of Radiation Physics, The University of Texas MD Anderson Cancer Center, Houston, USA
[2]The University of Texas MD Anderson Cancer Center UTHealth Houston Graduate School of Biomedical Sciences, Houston, USA
[3]Department of Radiation Oncology, UT Southwestern Medical Center, Dallas, USA
[4]Department of Radiation Oncology, The University of Texas MD Anderson Cancer Center, Houston, USA
[5]Division of Physics and Engineering, The University of Texas Medical Branch at Galveston, Galveston, USA
[5]Department of Radiation Oncology, The University of Alabama at Birmingham, Birmingham, USA

Abstract Radiation oncologists are responsible for determining the clinical acceptability of radiotherapy treatment plans. Their training typically revolves around routine clinical practice, observing the approval of patient plans. This approach depends on a sufficiently rich patient population (which is difficult for smaller clinics and less common cancer types) and is example-starved as the trainee typically only sees a single plan example for each patient. To address this, we have developed techniques to rapidly create clinically realistic plans that are suboptimal in a controllable way. These plans can then be used to train oncologists to understand what dose distributions are possible, and how to spot plans that could be improved. Here we first use a deep learning-based dose prediction to create a high-quality treatment plan (trained using only high-quality plans). We then increase dose by controllable amounts around selectable organs-at-risk. This approach is shown to generate clinically realistic, but suboptimal plans, as scored by experienced dosimetrists.

1 Introduction

Radiotherapy plan quality and radiation oncologist experience directly correlate to patient outcome and quality of life[1–4]. Therefore, it is vital that the quality of the final approved plan be high – that is, the prescription to the target be met while minimizing dose delivered to healthy tissue. Unfortunately, the literature reveals that suboptimal plans are approved frequently[5–8], even by experienced radiation oncologists. A contributing factor is the training environment during residency in which plan review is taught. Residents report the absence of examples across diverse treatment sites; the substantial clinical time pressures and high-stakes environment which make effective education difficult; and discontinuities introduced between identifying suboptimalities and making requests for edits and receiving the updated plan.

It is important to address these challenges by providing a wide range of example plans that need review and improvement in an environment that is low-stakes and without the pressure to have plans approved quickly for treatment. Simply curating lower-quality plans would be infeasible, as clinics typically have limited numbers available and inadvertent example memorization by the trainees would become common. Instead, a method of generating clinically realistic but suboptimal plans is needed. To date, however, most research has focused on generating high-quality plans[9–18]. Therefore, the current study develops a tool to create suboptimal dose distributions. These will be clinically realistic and suitable for integration into the training curriculum, providing the ability for real-time education in a low-stakes environment with virtually limitless examples.

2 Materials and Methods

In this study, suboptimal plans are modeled as potentially ideal dose distributions to which error has been added. This approach allows starting with consistent, high-quality dose distributions, and then creating areas of suboptimality around select OARs.

High-quality plans are required as the standard and what residents should be learning to achieve during plan review training. In this study, volume-modulated arc therapy (VMAT) dose distributions were predicted using the 3-Dimensional Dense Dilated U-Net model[10]. Two sets of models were used, one for head and neck (HN) radiotherapy[11] and the other for gynecological (GYN) radiotherapy[12]. Each set of models was previously trained using three-fold cross-validation on clinical data. This training data was carefully selected to be high-quality based upon both DVH metrics as well as upon physician review and consisted of planning computed tomography (CT) images, the clinical dose map from the treatment planning system, planning target volume (PTV) contours of varying prescriptions, and organ at risk (OAR) contours.

As the resultant dose distributions will be used for training, it is important that error be added to the high-quality dose predictions in a controllable manner that results in dose distributions free from obvious manipulation and appearing to have originated in the clinic. Dose changes should be localized to the OARs in which suboptimality is desired without substantial excess change elsewhere, but also without introducing visual errors such as introducing unrealistic dose gradients or other artifacts around the regions. To accomplish this, a convolutional technique was developed which adjusts changes based upon geometry –

that is, large changes in dose are only made near the desired OARs.

To start, a geometric feature λ_{OAR} was derived for each group of OARs. It describes proximity to the region in which suboptimality is desired through a combination of the Euclidean distances outside the OAR boundary and a directionally insensitive beam's eye view (BEV) angle. The geometric features were normalized to the range [0, 1], where 0 indicates least physical and angular distance from the OAR and 1 indicates greatest. Then a sigmoid-like function was applied with parameters controlling the full-width half-max (FWHM) and steepness of the sigmoid curve β, to provide additional control over the geometric space in which the dose was increased.

Equation 1

$$f(\lambda, FWHM, \beta) = \frac{\left(\lambda^{-\frac{\log(2)}{\log(FWHM)}}\right)^{\beta}}{\left(\lambda^{-\frac{\log(2)}{\log(FWHM)}}\right)^{\beta} + \left(1 - \lambda^{-\frac{\log(2)}{\log(FWHM)}}\right)^{\beta}}$$

As the target region contains the highest dose, a similar geometric feature λ_{PTV} was derived and normalized to the range [0, 0.98]. This feature indicated proximity for the PTV for which to make very small perturbations to the dose. This allowed changes to be made to the dose delivered to target, but not unrealistically large relative to the dose increase around each OAR. The final geometric feature was the 3-dimensional element-wise minimum of the OAR feature and the target feature.

Equation 23

$$\lambda = min\left(\lambda_{OAR,ijk}, \lambda_{PTV,ijk}\ \forall i,j,k\ in\ (Z,X,Y)\right)$$

Next, a kernel was designed to take this geometric information into account. There are two parameters for this kernel, geometric feature λ which varies at every step in the convolution, and the constant AUC_m, which is always 1 or greater and determines the maximum area under the curve of the kernel and therefore the resultant dose scaling.

Equation 4

$$g(\lambda, AUC_m) = s(\lambda, AUC_m) \cdot \frac{1}{\sum k} k(\lambda)$$

This may be further divided into two components. First, it is important that the overall kernel area under the curve (AUC) changes with response to geometry, so that greatest dose scaling is applied only near the OARs of interest. The linear function $s(\lambda, AUC_m) = (1 - AUC_m)\lambda + AUC_m$ controls the overall AUC of the kernel, such that it is equal to 1 as λ is largest (far from the OAR) and equal to AUC_m as λ is smallest (closes to the OAR and along the BEV angle).

The second component, the inner kernel k, is self-normalizing to AUC=1 and changes its receptive field based upon geometry. At the limits of the geometric feature, k is the identity kernel to preserve dose as-is ($\lambda = 1$), or the mean kernel to incorporate the most information from neighboring values ($\lambda = 0$).

Equation 5

$$k(\lambda) = \begin{bmatrix} 1-\lambda & 1-\lambda & 1-\lambda \\ 1-\lambda & 1 & 1-\lambda \\ 1-\lambda & 1-\lambda & 1-\lambda \end{bmatrix}$$

Finally, suboptimalities were introduced with greatest relative increase near the selected OARs through convolving the high-quality dose distribution with the geometry aware kernel.

Equation 6

$$D' = D * g(\lambda, AUCm)$$

To test this process, two datasets were used. First, a set of twenty-five HN VMAT radiotherapy cases were curated from our institution. These cases included a planning CT image, the clinical dose map from the treatment planning system, PTV contours of varying prescriptions, and OAR contours required to generate the high-quality dose predictions as described by Gronberg et al.[11] To assess translatability of our approach, an additional fourteen gynecological (GYN) VMAT cases were curated from our institution. These cases included a planning CT image, the clinical dose map from the treatment planning system, a single PTV contour prescribed to 45Gy, and OAR contours required to generate the high-quality dose predictions as described by Gronberg et al.[12]

Five HN OAR groupings were selected in this study (Table 1). The generation parameters listed in Table 1 were selected based upon preliminary experimentation to provide realistic-appearing dose distributions.

Group	*OARs*	*AUC_m*	*FWHM/β*
1	Esophagus Larynx	1.1-1.7	0.15/4
2	Brainstem	1.1-1.7	0.15/4
3	Optic Nerves Lenses	1.1-1.7	0.15/4
4	Parotids	1.1-1.7	0.15/4
5	Cochleae	1.1-1.7	0.15/4

Table 1: Groups of HN OARs for which the dose was increased in this study, with the generation parameters.

Changes to the plan dose was assessed as relative difference in structure mean dose D_{mean} between the edited plans and the predicted high-quality plans. For each OAR in Table 1,

the relative change was calculated as $\Delta D_{mean,OAR} = \frac{D'_{mean,OAR} - D_{mean,OAR}}{D_{mean,OAR}}$. A paired one-sided T-test was applied to assess significance in change to D_{mean} for each OAR, with significance taken at the $p<0.01$ level.

An experienced dosimetrist also reviewed a subset of the dose distributions, all generated with AUC_m=1.4, to assess clinical realism. To avoid bias, the dosimetrist was blinded to plan origin. A random mix of 10 clinical dose distributions and 10 suboptimal distributions was generated. The dosimetrist reviewed each dose distribution individually and answered if they thought the doses were (1) generated by a human with a treatment planning system, (2) generated by a deep learning dose prediction model, or (3) unable to determine or ambiguous.

3 Results

Dose was increased around the selected OARs in each group in a controllable way, as shown in Figure 2. This increase was statistically significant. Increasing kernel AUC_m increased mean dose as shown in Figure 1, where the relatively small error bars indicate strong control of change in dose. Figures 1 and 3 also demonstrate the localization of changes; mean dose is increased most around the OARs of interest in each group. The smaller increase in dose to other OARs are appropriate for clinical realism; while not changed as substantially as for the OARs of interest, they still exhibit realistic dose increases. While Figures 1 & 3 show a subset of groups and OARs for clarity, the general trend remains across all groups.

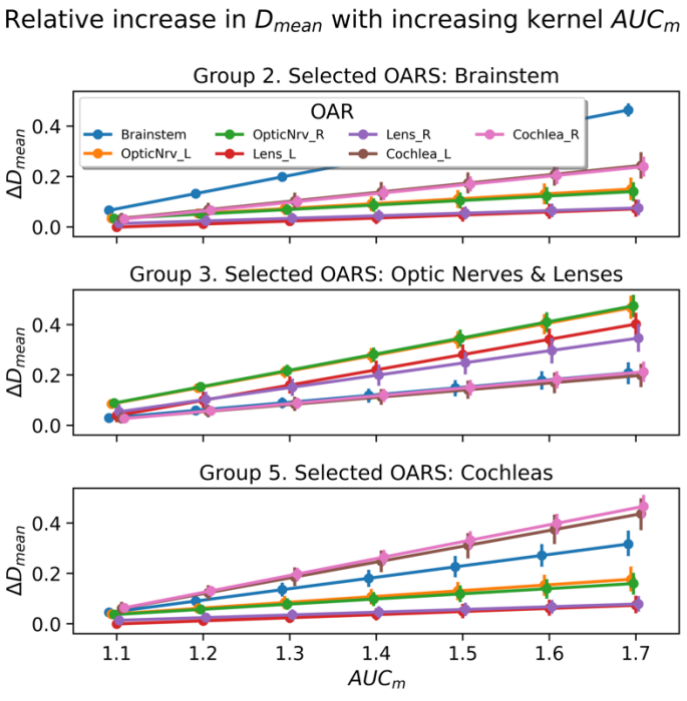

Figure 1: Relative increase in mean dose for selected OAR groups as the maximum kernel AUC increases.

The rate at which the dosimetrist was unable to determine if the dose distribution originated from a clinical treatment plan or a generated prediction was 80% for the suboptimal dose distributions, comparable to 70% for the clinical dose distributions. This indicates that the suboptimal dose distributions do not appear substantially different to plans that might be encountered in the clinic – that is, the generated plans are suboptimal but still clinically realistic.

The GYN dose distributions were edited using the same parameters as in Table 1, but with three groups of OARS: Bladder, Rectum, and the Femoral Heads together. D_{mean} was again increased for each OAR with increasing AUC_m (Table 2). This demonstrates the translatability of our approach across treatment sites.

AUC_m	*Bladder*	*Rectum*	*L. Femur Head*	*R. Femur Head*
1.1	0.02	0.01	0.07	0.07
1.2	0.05	0.03	0.13	0.14
1.3	0.07	0.04	0.20	0.20
1.4	0.09	0.06	0.26	0.27
1.5	0.12	0.07	0.33	0.34
1.6	0.14	0.09	0.39	0.40
1.7	0.16	0.10	0.46	0.47

Table 2: Relative increase in D_{mean} for the GYN cases.

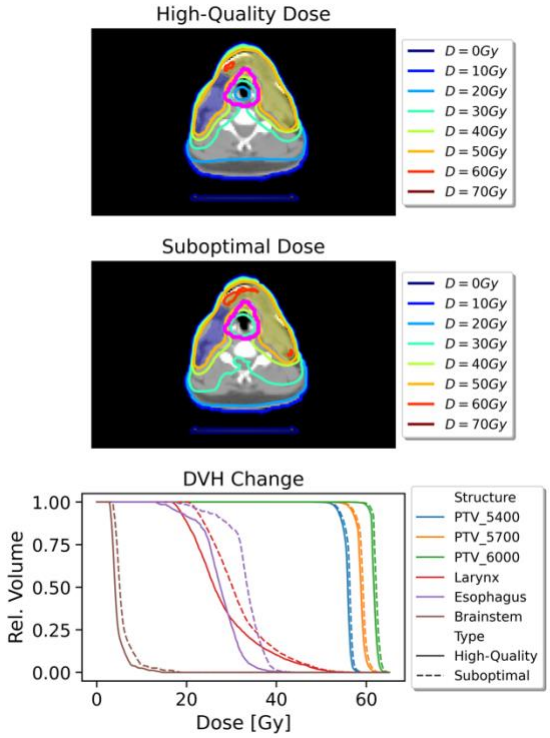

Figure 2: Comparison of the suboptimal dose distribution with the initial high-quality dose distribution. This example came from Group 1 where the larynx (magenta) and esophagus (only seen in the DVH)) were the OARs for which the changes are focused. Blue and yellow regions: low- and mid-risk PTV, respectively.

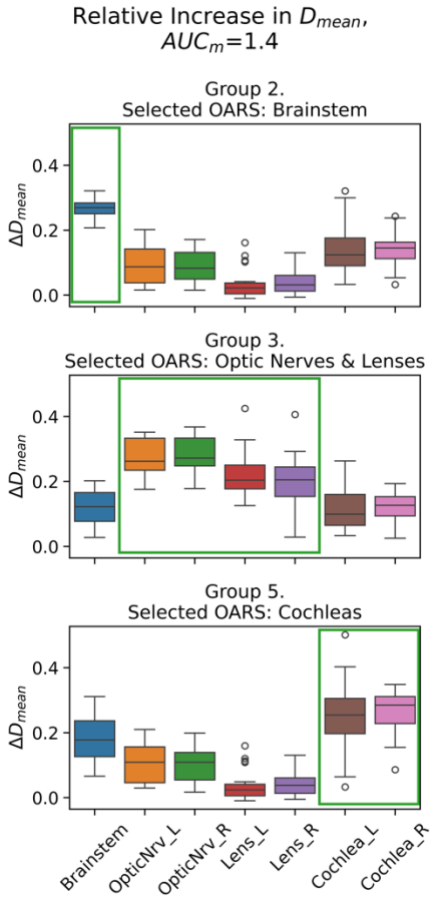

Figure 3: Distribution of relative increase in mean dose for selected OAR groups, with a single kernel AUC_m of 1.4. Green boxes indicate the selected OARs for each group.

4 Discussion

In this study, a tool for generating suboptimal but clinically realistic dose distributions was successfully developed. To our knowledge, this is the first such approach and opens the door to a range of applications for medical education and training.

The next step of this project will be to integrate this tool into the treatment planning system. Using scripting, an interactive model will be designed that presents trainees with different scenarios and asks them to identify where the treatment plan quality can be improved.

5 Conclusion

In this work, we develop a tool for generating artificial dose distributions that are clinically realistic. These dose distributions are suboptimal in predictable and controllable locations.

References

1. Peters, L. J. *et al.* Critical Impact of Radiotherapy Protocol Compliance and Quality in the Treatment of Advanced Head and Neck Cancer: Results From TROG 02.02. *JCO* **28**, 2996–3001 (2010).
2. Chien, C.-R. *et al.* High case volume of radiation oncologists is associated with better survival of nasopharyngeal carcinoma patients treated with radiotherapy: a multifactorial cohort analysis. *Clinical Otolaryngology* **36**, 558–565 (2011).
3. Wuthrick, E. J. *et al.* Institutional Clinical Trial Accrual Volume and Survival of Patients With Head and Neck Cancer. *JCO* **33**, 156–164 (2015).
4. Boero, I. J. *et al.* Importance of Radiation Oncologist Experience Among Patients With Head-and-Neck Cancer Treated With Intensity-Modulated Radiation Therapy. *JCO* **34**, 684–690 (2016).
5. Wu, B. *et al.* Patient geometry-driven information retrieval for IMRT treatment plan quality control. *Medical Physics* **36**, 5497–5505 (2009).
6. Moore, K. L., Brame, R. S., Low, D. A. & Mutic, S. Experience-Based Quality Control of Clinical Intensity-Modulated Radiotherapy Planning. *International Journal of Radiation Oncology*Biology*Physics* **81**, 545–551 (2011).
7. Moore, K. L. *et al.* Quantifying Unnecessary Normal Tissue Complication Risks due to Suboptimal Planning: A Secondary Study of RTOG 0126. *International Journal of Radiation Oncology*Biology*Physics* **92**, 228–235 (2015).
8. Talcott, W. J. *et al.* A Blinded, Prospective Study of Error Detection During Physician Chart Rounds in Radiation Oncology. *Practical Radiation Oncology* **10**, 312–320 (2020).
9. Babier, A. *et al.* OpenKBP: The open-access knowledge-based planning grand challenge and dataset. *Medical Physics* **48**, 5549–5561 (2021).
10. Gronberg, M. P. *et al.* Technical Note: Dose prediction for head and neck radiotherapy using a three-dimensional dense dilated U-net architecture. *Medical Physics* **48**, 5567–5573 (2021).
11. Gronberg, M. P. *et al.* Deep Learning–Based Dose Prediction for Automated, Individualized Quality Assurance of Head and Neck Radiation Therapy Plans. *Practical Radiation Oncology* (2023) doi:10.1016/j.prro.2022.12.003.
12. Gronberg, M. P. *et al.* Deep learning–based dose prediction to improve the plan quality of volumetric modulated arc therapy for gynecologic cancers. *Medical Physics* **50**, 6639–6648 (2023).
13. Barragán-Montero, A. M. *et al.* Three-dimensional dose prediction for lung IMRT patients with deep neural networks: robust learning from heterogeneous beam configurations. *Medical Physics* **46**, 3679–3691 (2019).
14. Chen, X., Men, K., Li, Y., Yi, J. & Dai, J. A feasibility study on an automated method to generate patient-specific dose distributions for radiotherapy using deep learning. *Medical Physics* **46**, 56–64 (2019).
15. Fan, J. *et al.* Automatic treatment planning based on three-dimensional dose distribution predicted from deep learning technique. *Medical Physics* **46**, 370–381 (2019).
16. Kearney, V., Chan, J. W., Haaf, S., Descovich, M. & Solberg, T. D. DoseNet: A volumetric dose prediction algorithm using 3D fully-convolutional neural networks. *Physics in Medicine and Biology* **63**, (2018).
17. Nguyen, D. *et al.* 3D radiotherapy dose prediction on head and neck cancer patients with a hierarchically densely connected U-net deep learning architecture. *Physics in Medicine and Biology* **64**, (2019).
18. Zhang, J. *et al.* Voxel-Level Radiotherapy Dose Prediction Using Densely Connected Network with Dilated Convolutions. in *Artificial Intelligence in Radiation Therapy* (eds. Nguyen, D., Xing, L. & Jiang, S.) 70–77 (Springer International Publishing, Cham, 2019). doi:10.1007/978-3-030-32486-5_9.