



Published in final edited form as:

Med Phys. 2020 September ; 47(9): 3898–3912. doi:10.1002/mp.14374.

Using Deep Learning to Predict Beam-Tunable Pareto Optimal Dose Distribution for Intensity Modulated Radiation Therapy

Gyanendra Bohara, Azar Sadeghnejad Barkousaraie, Steve Jiang, Dan Nguyen

Medical Artificial Intelligence and Automation (MAIA) Laboratory, Department of Radiation Oncology, UT Southwestern Medical Center, Dallas TX, USA

Abstract

Purpose: Many researchers have developed deep learning models for predicting clinical dose distributions and Pareto optimal dose distributions. Models for predicting Pareto optimal dose distributions have generated optimal plans in real time using anatomical structures and static beam orientations. However, Pareto optimal dose prediction for Intensity Modulated Radiation Therapy (IMRT) prostate planning with variable beam numbers and orientations has not yet been investigated. We propose to develop a deep learning model that can predict Pareto optimal dose distributions by using any given set of beam angles, along with patient anatomy, as input to train the deep neural networks. We implement and compare two deep learning networks that predict with two different beam configuration modalities.

Methods: We generated Pareto optimal plans for 70 patients with prostate cancer. We used fluence map optimization to generate 500 IMRT plans that sampled the Pareto surface for each patient, for a total of 35,000 plans. We studied and compared two different models, Model I and Model II. Although they both used the same anatomical structures—including the planning target volume (PTV), organs at risk (OARs), and body—these models were designed with two different methods for representing beam angles. Model I directly uses beam angles as a second input to the network as a binary vector. Model II converts the beam angles into beam doses that are conformal to the PTV. We divided the 70 patients into 54 training, 6 validation, and 10 testing patients, thus yielding 27,000 training, 3,000 validation, and 5,000 testing plans. Mean square loss (MSE) was taken as the loss function. We used the Adam optimizer with a default learning rate of 0.01 to optimize the network's performance. We evaluated the models' performance by comparing their predicted dose distributions with the ground truth (Pareto optimal) dose distribution, in terms of DVH plots and evaluation metrics such as PTV D_{98} , D_{95} , D_{50} , D_2 , D_{\max} , D_{mean} , Paddick Conformation Number, R50 and Homogeneity index.

Results: Our deep learning models predicted voxel-level dose distributions that precisely matched the ground truth dose distributions. The DVHs generated also precisely matched the ground truth. Evaluation metrics such as PTV statistics, dose conformity, dose spillage (R50) and homogeneity index also confirmed the accuracy of PTV curves on the DVH. Quantitatively, Model I's prediction error of 0.043 (confirmation), 0.043 (homogeneity), 0.327 (R50), 2.80% (D_{95}), 3.90% (D_{98}), 0.6% (D_{50}), 1.10% (D_2) was lower than that of Model II, which obtained 0.076 (confirmation), 0.058 (homogeneity), 0.626 (R50), 7.10% (D_{95}), 6.50% (D_{98}), 8.40% (D_{50}),

6.30% (D2). Model I also outperformed Model II in terms of the mean dose error and the max dose error on the PTV, bladder, rectum, left femoral head, and right femoral head.

Conclusions: Treatment planners who use our models will be able to use deep learning to control the tradeoffs between the PTV and OAR weights, as well as the beam number and configurations in real time. Our dose prediction methods provide a stepping stone to building automatic IMRT treatment planning.

1. Introduction

Today, it is estimated that about two-thirds of all patients with cancer receive Radiation Therapy as a unique treatment or in combination with more complex treatment procedures.¹ One of the remarkable achievements in External Beam Radiation Therapy (EBRT) is the development of Intensity Modulated Radiation Therapy (IMRT),²⁻⁷ which uses variable beam intensities to treat cancer. IMRT allows the delivery of less dose to the organs at risk (OARs) and more dose to the planning target volume (PTV) than 3D conformal radiation therapy,⁸⁻¹¹ but its treatment planning process is more difficult and time consuming. IMRT treatment planning consists of two iterative processes: first, the planner uses dose-volume constraints and other hyper-parameters to obtain an optimal plan to deliver as much of the prescribed dose to the PTV as possible while minimizing the dose to critical structures. The planner has to iteratively and tediously tune the hyper-parameters in a trial-and-error fashion. Second, the physician reviews the plan and provides further comments and feedback to get the outcome that achieves the best tradeoffs between PTV and OARs.^{12,13} These two processes loop until the final plan is approved. It can take from multiple hours to a week to generate an acceptable plan, depending on the treatment site and its complexity.

Several studies have tried to improve the treatment planning process by using mathematical optimization algorithms to account for various aspects. Multicriteria optimization¹³⁻¹⁵ focuses on generating multiple plans with tradeoffs between the PTV and OARs on the Pareto surface, which allows the clinician to then choose the plan with their desired tradeoffs. Beam orientation optimization¹⁶⁻²¹ focuses on finding a suitable set of beam directions that improves upon manually selected or protocol-based beam orientations. Direct aperture optimization, also called machine parameter optimization,²²⁻²⁷ focuses on generating deliverable, high quality plans by determining the optimal aperture shapes and their intensities. There are many commercial software packages available based on mathematical optimization algorithms, such as Eclipse™ comprehensive treatment planning (Varian Medical systems, Palo Alto, CA, USA), Pinnacle treatment planning (Philips Radiation Oncology, Fitchburg, WI, USA),²⁸ and RayPlan treatment planning system (RaySearch Laboratories, Stockholm, Sweden). Still, these systems require manual, tedious tuning of hyper-parameters, such as structure weights, beam geometries, appropriate dose-volume constraints, and tradeoffs between the PTV and OARs. Also, treatment plans generated on these systems differ from planner to planner and from physician to physician based on their work experience and preferences. The customized tedious process, planning variations based on personal experiences and preferences, and the need for strong domain knowledge expertise could lead to suboptimal plans that compromise patient care.²⁹⁻³¹

A new set of methods, called knowledge-based planning (KBP),^{32–37} has been developed to address the shortcomings of mathematical optimization algorithms and improve the quality and efficiency of treatment planning by learning a database of carefully designed past clinical plans. KBP uses machine learning algorithms and is a powerful tool for guiding treatment planners and physicians to achieve high quality plans. RapidPlan™ is an example of KBP that was developed by Varian Medical Systems. This system estimates the dose volume histogram (DVH) for the new plan by using patient-specific geometry. Many researchers have reported on RapidPlan's performance and compared it with that of conventional treatment planning, and they have found that, in its current state, RapidPlan is much faster and can generate clinically acceptable plans with higher quality than conventional treatment planning for about half of the cases.^{38–43} However, it is not fully automated yet, and for the remaining treatment cases, manual tuning is still necessary to make acceptable plans. In addition, KBP relies heavily on small datasets because datasets have not been integrated between different institutions, so caution should be taken when applying these methods to patients whose geometry falls outside the plan library.³⁹ Also, before the deep learning era, KBP methods used more traditional machine learning algorithms, and they were limited to predicting DVH or particular dosimetric criteria from user-defined features.⁴⁴

Deep learning has advanced many areas such as image recognition, speech recognition, natural language translation towards automation⁴⁵ and has addressed the shortcomings of traditional machine learning by learning its own features from data without the need for human intervention. Likewise, deep learning has the potential to automate the IMRT treatment planning process by removing its dependence on handcrafted features. The development of the fully convolutional network (FCN)⁴⁶ allowed for pixel-wise prediction using supervised learning, which opened the door for voxel-wise dose prediction and generation of DVH curves in treatment planning. Recently, many researchers have developed different deep learning models for predicting clinical dose distributions for IMRT and Volumetric Modulated Arc Therapy (VMAT) modalities on different treatment sites such as lung, prostate, and head-and-neck.^{47–53} However, all of these models used static beam orientations for their study, thus limiting their uses in the treatment planning workflow to a subset of common treatment plans based on protocol. One approach that uses varying beam angles to predict the clinical dose for lung IMRT patients has been developed recently.⁵⁴

Clinical dose prediction models are often limited to a single dose predicted per patient. This contrasts with Pareto optimal dose prediction models, which can generate multiple plans that have differing tradeoffs between the different critical structures. Previous studies found that deep learning models that use anatomical structures and static beam orientations to predict Pareto optimal dose distributions could generate multiple optimal plans with differing tradeoffs in real time.^{44,55} However, Pareto optimal dose predictions for IMRT prostate plans with variable beam numbers and orientations have not yet been studied. In this paper, we present an approach that uses deep learning networks to predict Pareto optimal dose distributions for prostate IMRT plans that involve anatomical structures and varying beam numbers and orientations. Specifically, our contribution to the current literature is the addition of the ability to tune the beam orientations in a deep learning-based, Pareto optimal dose prediction model. Such a model would allow for a treatment planner to quickly explore

the beam orientation space, and select a beam arrangement that can even be outside of the typical clinical protocol. We implement and compare two deep learning networks that predict with two different beam configuration modalities: Model I, which uses the direct input of the beam angles in the network as a binary vector, and Model II, which uses the conformal beam dose that corresponds to the beam angles used in Model I. Model II serves here as a state-of-the-art model for comparison; this model is similar to the AB model introduced by Ana et al.,⁵⁴ where beam setup information was represented by the cumulative dose distribution for all the beams in the plan computed by using the fluence-convolution broad beam (FCBB)⁵⁶ dose calculation method. In our case, we used a simple algorithm to generate a beam conformal to the PTV structure (see section 2.2).

This work will provide treatment planners with the advantage of using deep learning to control the tradeoffs between the PTV and OAR weights, as well as the beam number and configurations, in real time.

2. Materials and Methods

For this study, we generated Pareto optimal plans for 70 patients with prostate cancer. We used fluence map optimization to generate 500 IMRT plans that sampled the Pareto surface for each patient, for a total of 35,000 plans. More details about generating the Pareto optimal dose distribution dataset are presented in section 2.1. The deep learning models used for predicting Pareto optimal dose distributions are described in section 2.3. We studied and compared two different models. Although they both used the same anatomical structures—which included the planning target volume (PTV), organs at risk (OARs), and body—these models were designed with two different methods for representing the beam angles. For Model I, we directly used beam angles as a binary vector for the second input to the network. For Model II, we converted the beam angles into beam doses that were conformal to the PTV. More details about generating the conformal beam doses are provided in section 2.2. We divided the 70 patients into 54 training, 6 validation, and 10 testing patients, yielding 27,000 training, 3,000 validation, and 5,000 testing plans. Detailed explanations of the model training, validation, and testing are presented in section 2.4.

2.1 Pareto Optimal Plans

The Pareto optimal solutions for 70 patients with prostate cancer were generated by minimizing the objective function defined from Equations 1–3. These resulting dose distributions for training are calculated after the fluence map optimization step, but before any machine parameter sequencing is performed. The parameters involved were anatomical structures and 10 different sets of 1 to 10 random beam angles. Anatomical structures included the planning target volume (PTV) and the organs at risk (OARs): body, bladder, rectum, left femoral head, and right femoral head. Shell and skin structures were also included in the plan as tuning structures. Pareto optimal solutions represent the various tradeoffs between tumor coverage and normal tissue sparing. This is associated with a multicriteria objective, which can be written as

$$\begin{aligned}
 & \underset{x_\theta}{\text{minimize}} \quad \{f_{PTV}(x_\theta), f_{OAR_1}(x_\theta), f_{OAR_2}(x_\theta), \dots, f_{OAR_n}(x_\theta)\} \\
 & \text{subject to} \quad x_\theta \geq 0, \quad \theta \in A
 \end{aligned} \tag{1}$$

where f_s is the objective function and $A = [\theta_1, \dots, \theta_{10}]$ is the collection of all 10 sets of beam angles.

For example, $\theta_1 = [10]$, $\theta_2 = [24, 38], \dots, \theta_{10} = [20, 26, 30, 38, 46, 6, 56, 98, 64, 120]$ are the randomly generated angles, and x_θ refers to the fluence map intensities to be optimized. Here, we use the ℓ_2 -norm to formulate the objective,

$$f_s(x_\theta) = \frac{1}{2} \|d_{\theta, s} x_\theta - p_s\|_2^2 \tag{2}$$

where $d_{\theta, s}$ is the dose influence matrix for the θ_{th} beam and the s_{th} structure. The dose influence matrices were determined using 1 to 10 random coplanar beams where the beamlet size was 2.5 mm^2 at a 100 cm isocenter. p_s is the desired dose for a given structure, assigned as the prescription dose if s is the PTV, and otherwise 0. The dose influence calculation was performed using the Analytical Anisotropic Algorithm (AAA) provided by the Eclipse treatment planning system, using the built-in application programming interface (Varian Medical systems, Palo Alto, CA, USA). Now, we can reformulate the multicriteria optimization^{14,57,58} as a single-objective, convex optimization problem:

$$\begin{aligned}
 & \underset{x_\theta}{\text{minimize}} \quad \sum_{s \in S} w_s^2 f_s(x_\theta) \\
 & \text{subject to} \quad x_\theta \geq 0, \quad \theta \in A
 \end{aligned} \tag{3}$$

where w_s are the user-defined tradeoff weights for each structure. Different Pareto optimal plans can be generated by varying the w_s to different values. We generated many pseudo-random plans by assigning random weights, as described below, to the organs at risk by using an in-house GPU-based proximal-class first-order primal-dual algorithm, Chambolle-Pock.⁵⁹ While there are some uses of the Chambolle-Pock algorithm used in radiation therapy^{60–64}, any convex solver will theoretically arrive at the same solution as Chambolle-Pock for solving the optimization problem.

The weight generation structure fell into one of three categories, as shown in Table 1.

For each patient, we created 500 plans spanning the low, extra low, and controlled weights categories. These bounds for the controlled weights were chosen through a trial-and-error method so that the plan generated would fall within clinically relevant limits, even though it is not necessarily acceptable by a physician. A total of 35,000 IMRT plans were created, each as $96 \times 96 \times 32$ dimension arrays with a voxel size of 5 mm^3 that sampled the Pareto surface. Table 2 shows the distribution of Pareto optimal plans generated in each weight category per beam set and its assignment as training, validation and testing for the study. Since the table shows the number of plans per beam set, and since there are 10 beam sets (i.e. 1-beam plan, 2-beam plan, ..., 10-beam plan), in total, there are 10 times the number of plans shown in Table 2, which equates to 35,000 plans.

2. Conformal Beam Dose Data

Conformal beam dose describes the high dose volume that is shaped to closely conform to the desired PTV structure. There are numerous ways to make a broad beam conform to the PTV.^{56,65} In this study, we used a simple method to generate a beam that is conformal to the PTV structure. For each beam we first selected a square field of 20×20 beamlets, with dimensions of $2.5\text{mm} \times 2.5\text{mm}$ per beamlet. This the same beamlet data mentioned in section 2.1 that was generated using AAA dose calculation algorithm. We then scale this beam such that its mean dose contribution to the PTV is equal to the prescription dose. Beamlets with an integral dose contribution of less than the threshold of 1% of the prescription dose to the PTV are removed. The remaining beamlets then create a conformal beam around the PTV.

For each patient, 500 plans were generated using all 10 sets of beam angles. Representative images of conformal doses are shown in Figure 1.

2.3. Network Architecture

The network architecture used in this study is depicted in Figure 2. The dose prediction models used a U-Net style architecture.⁶⁶ We used group normalization⁶⁷ instead of batch normalization. This network consists of three major parts: downsampling, bottom, and upsampling. The rectified linear unit (ReLU), group normalization and dropout were applied immediately after every convolution operation in the hidden layers. For the sake of clarity, the following paragraphs will assume these operations are included when “convolution” is mentioned. More details on ReLU, group normalization, and dropout are mentioned later, after the main network architecture description. With the exception of the strided convolution, all other convolutions are zero-padded before the convolution, to maintain the same data shape.

The downsampling part of the U-Net is constructed to contain a five-level hierarchy. We chose five levels with four $2 \times 2 \times 2$ downsampling operations to halve the feature size four times, reducing the data from $96 \times 96 \times 32$ voxels to $6 \times 6 \times 2$ voxels. In each level, two $3 \times 3 \times 3$ convolution operations took place, followed by one downsampling operation. The downsampling operation is a concatenation of two other operations: 1) $2 \times 2 \times 2$ max pooling and 2) $2 \times 2 \times 2$ strided convolution with a $2 \times 2 \times 2$ kernel. In this process, feature channels were doubled in each level.

The bottom part is between the downsampling and upsampling parts of the network. This part takes the last downsampled feature map and performs two $3 \times 3 \times 3$ convolutions. In addition, for one of the models to be evaluated (Section 2.3.1), beam angle information is added as a binary vector that is then processed through a fully connected network. This data is reshaped and concatenated onto the bottom level, prior to the two convolutions.

The upsampling part of the network also consists of five levels, just like the downsampling part. The purpose of this part is to combine the features and spatial information through a sequence of upsampling $3 \times 3 \times 3$ and convolution operations and to concatenate high resolution features from the downsampling part. This part consists of 4 upsampling layers, in addition to the final convolution output layer. In this process, feature channels are reduced

by half, but feature size is doubled after each convolution to maintain symmetry and original feature size.

All activation functions in the hidden layers are rectified linear units (ReLU), but the final activation function is a linear activation function. Group normalization was used in all hidden layers after the convolution and ReLU operations, which normalizes the weights by grouping feature channels of 32, 64, 128, 256, and 512 into 1, 2, 4, 8, and 16 groups, respectively, of 32 channels each (Fig. 2); this allows faster convergence. The dropout scheme, from a previous paper⁴⁷, described in Table 3 was applied after each group normalization.

Models used in this investigation are of two types.

2. 3.1. Model I—The models used in this study are depicted in Figure 2. Model I's architecture is exactly what is shown in Figure 2. Model 1 takes in 3 channels: a PTV channel, a body channel, and an OARs channel. Instead of just binary masks as input, we multiply the masks with their corresponding weights, w_s , from Equation 3, such that a voxel is defined as w_s if a voxel is defined inside a structure and 0 otherwise. The PTV and body channels have just their respective data, while the OARs channel contains the bladder, rectum, femoral heads, and tuning structures information. Mentioned in Section 2.1, the data was resized to 5 mm³ voxels. To maintain uniform data shape for model training, all patient data was filled into a $96 \times 96 \times 32$ array. The body segmentation covers CT slices. The input data for Model 1 ranges from 0 to 1 since the structure weights used for optimization, w_s , were also defined from 0 to 1. As the second input, randomly generated beam angles as a Boolean array of 180 elements—representing angles with 2 degrees separation—are concatenated in the bottom part of the network. In the Boolean array, the current selected beam angles are considered as ones, and all other unselected beam angles are considered as zeros. The Boolean array of 180 elements was the input of a fully connected layer with output of 2304 elements. After that a reshaped operation was applied to change the 2304 length data to (6,6,2, n) where n was 32 (i.e., $2304 = 6 \times 6 \times 2 \times 32$). The reshaped beam angles were matched to the downsampling feature size, allowing for them to be concatenated together along the channels axis, for further processing in the network.

2. 3.2 Model II—Model II's architecture is the same as shown in Figure 2 except without the beam angle binary vector input in the bottom. Model II consists of a single four-channel input: the first three channels are the same as Model I, and the last channel is the conformal dose information for a set of selected beam angles (see section 2.2). The size of each input channels is also $96 \times 96 \times 32$. Analogous to Model I, the anatomical inputs range from 0 to 1. For the additional conformal dose channel, the dose was divided by its maximum dose, to also range from 0 to 1.

2. 4. Model Training, Validation, and Testing

For each model, we divided the 70 patients into 54 training, 6 validation, and 10 testing patients, thus yielding 27,000 training, 3,000 validation, and 5,000 testing plans. We implemented a maximum dropout rate of 0.25 to regularize the network and, thus, avoid overfitting. We applied these dropout rates after group normalization so that the highest

group, 16, got the dropout rate of 0.25 and the lowest group, 1, got 0.125. The dropout rate scheme is presented in Table 3. During training, we used a batch size of 1, which was due to memory constraints.

The mean square loss (MSE)

$$MSE = \frac{1}{n} \sum_v^m (D_{test}^v - D_{pred}^v)^2 \quad (4)$$

was taken as a loss function, where v refers to the voxel index and m represents the total number of voxels. We used the Adam optimizer⁶⁸ with a default learning rate of 0.01 to optimize the network's performance. All training was performed on an NVIDIA Quadro P6000 GPU with 24 GB RAM. The models were trained with an early stopping scheme,⁶⁹ which is a regularization method that prevents overfitting. This scheme stops the network training when the model is no longer improving the validation loss after a set number of iterations, then it saves the model with the lowest validation loss. For our study, we trained the model for an additional 40,000 iterations after finding the best performing model and terminated the training process if no further improving was observed. The validation loss was checked every 100 iterations. The best models with the lowest total validation loss were used to work out the test data after the completion of training.

We evaluated the models' performance by comparing their predicted dose distributions with the Pareto optimal dose distribution (ground truth) in terms of DVH plots and evaluation metrics, such as PTV D_{98} , D_{95} , D_{50} , D_2 , Paddick Conformation Number,⁷⁰ R50 and Homogeneity index and the structure max and mean doses (D_{max} and D_{mean}). Readers should refer to the literature for more details about these evaluation metrics.^{47,54,71} D_{max} is considered as the dose delivered to 2% of the structure volume, as recommended by the ICRU report.⁷² We also compared the predicted dose distributions with the ground truth dose distribution using dose map differences in the clinically relevant PTV and OARs regions.

In addition, we have shown an example of beam tuning in the treatment plan with 9 fields Protocol Based IMRT setup.

3. Results

The instance of Model I with the lowest validation loss was found at 356,500 iterations, which took about 175 hours to obtain with the early stopping scheme. The instance of Model II with the lowest validation loss was found at 132,300 iterations, which took about 65 hours to obtain. Further iterations after the lowest validation model did not improve the result. After training, the prediction time of each model is less than 1 second. The loss versus iterations evaluated for the training (blue line) and validation sets (red line) from both models are presented in Figure 3. Observing these loss trends gives us an idea of the models' performance. We observed the decreasing trend of losses during training and validation until the best model was achieved. Each model achieved similar MSE losses, with training losses at 1.007×10^{-4} (Model I) and 1.463×10^{-4} (Model II) and validation losses at 1.251×10^{-4} (Model I) and 1.469×10^{-4} (Model II).

Figure 4 shows the colored dose wash distributions in the PTV and in the organs at risk. These distributions are overlaid with the original CT slices. Each row of dose distributions in Figure 4 represents a different treatment plan. Visually, the dose distributions that are predicted from Models I and II are similar to that of the ground truth. Particularly, they are better matched in the high dose region surrounding the PTV, and have less accuracy in the lower dose regions.

Figure 5 shows the dose map differences between the predicted dose distributions and the ground truth dose distributions. This treatment plan is exactly same plan as mentioned in the bottom row of Figure 4. Visually, the color intensity indicates that the dose differences predicted from Model I are better matched with the ground truth than the predicted dose distribution from Model II.

DVHs obtained from dose predictions for the two representative plans are presented in Figure 6. The DVH curves obtained from Model I corresponding to PTV are better matched with the ground truth than the DVH curves obtained from Model II corresponding to PTV, especially in the shoulder of the PTV DVH. The DVH curves corresponding to other critical structures except without body show the fluctuations in the prediction accuracy from both models. Visually, from the test cases in Figure 6, it is not fully clear whether one model outperforms the other in dose prediction accuracy to the OARs.

Dose evaluation metrics are calculated for all 5000 plans from each testing and predicting dataset. Also, these values are the mean values and deviation from the mean values from all the plans containing 1–10 beam orientations. The metrics presented in the Table 4 represent the dose coverage in the PTV, conformity, dose spillage and homogeneity. For the PTV D98, D95, D50 and D2, the highest mean difference we obtained from Model I was less than 4% and the highest mean difference we obtained from Model II was less than 9%, of the prescription dose. Similarly, for other parameters such as Paddick Confirmation number, dose spillage (R50) and PTV homogeneity, predicted mean value differences are less for Model I in comparison to that of Model II. Overall, the predicted mean differences obtained from Model I are less than that of predicted mean differences obtained from Model II.

The absolute mean dose (D_{mean}) and max dose (D_{max}) values reported in Table 5 give us an idea of how the dose distributed over the voxels of PTV and other critical structures. For D_{mean} values, the highest mean difference we obtained from Model I was less than 2% and the highest mean difference we obtained from Model II was less than 6%, of the prescription dose. Similarly, for D_{max} values, the highest mean difference we obtained from Model I was less than 5% and the highest mean difference we obtained from Model II was less than 12%, of the prescription dose. In all cases, the prediction errors obtained from Model I are less than that of Model II.

A paired t-test was used to determine if there is a statically significant difference between the performance of Model I and Model II, with respect to how accurately they predicted the ground truth Pareto optimal dose. The largest p-value that we obtained from a two tailed paired t-test was 6.82×10^{-63} .

Figure 7 was obtained based on the information presented in Table 4. This shows the errors for several clinical metrics evaluated from the predicted dose distributions, as compared to the metrics of the Pareto optimal dose distributions. Model I's prediction error of 0.043 (confirmation), 0.043 (homogeneity), 0.327 (R50), 2.80% (D95), 3.90% (D98), 0.6% (D50), 1.10% (D2) was lower than that of Model II, from which we obtained 0.076 (confirmation), 0.058 (homogeneity), 0.626 (R50), 7.10% (D95), 6.50% (D98), 8.40% (D50), 6.30% (D2). In terms of these dosimetric criteria, Model I performed better than Model II.

Figure 8 shows the errors for mean dose and max dose evaluated from the predicted dose distributions, as compared to the metrics of the Pareto optimal dose distributions. Model I had low prediction errors of average mean dose (D_{mean}) 0.871% (PTV), 0.214% (Body), 1.10% (Bladder), 1.76% (Rectum), 2.03% (Left Femoral Head), and 1.62% (Right Femoral Head), and average max dose (D_{max}) errors of 1.13% (PTV), 2.04% (Body), 1.41% (Bladder), 1.02% (Rectum), 5% (Left Femoral Head), and 3.48% (Right Femoral Head). Model II had relatively high prediction errors of D_{mean} 5.56% (PTV), 0.382% (Body), 1.94% (Bladder), 4.60% (Rectum), 4.90% (Left Femoral Head), and 3.45% (Right Femoral Head), and average D_{max} errors of 6.33% (PTV), 5.62% (Body), 8.98% (Bladder), 7.25% (Rectum), 11.54% (Left Femoral Head), and 6.89% (Right Femoral Head).

As with the dosimetric criteria shown in Figure 7, Model I outperformed Model II for both the mean dose and the max dose errors on PTV, Body, Bladder, Rectum, Left Femoral Head and Right Femoral Head.

The absolute mean dose (D_{mean}) values reported in Table 6 give us an idea of how the dose distributed over the voxels of body and 10% isodose volume. For D_{mean} values, the highest mean differences we obtained from Model I and Model II were less than 1%, of the prescription dose. In all cases, the prediction errors obtained from Model I are less than that of Model II.

Figure 9 shows the Average voxel wise dose errors for the Body and the 10% isodose volume obtained from the predicted dose distributions, as compared to the metrics of the Pareto optimal dose distributions. Model I had low prediction errors of average mean dose (D_{mean}) 0.61% (Body), 0.49% (10% isodose volume). Model II had relatively high prediction errors of D_{mean} 0.9% (Body), 0.8% (10% isodose volume). In terms of these dosimetric criteria as well, Model I performed better than Model II.

Figure 10 shows the errors for mean dose evaluated from the predicted dose distributions, as compared to the Pareto optimal dose distributions corresponding to each beam numbers in plan. It can be seen that that the prediction error is relatively even regardless of the number of beams in the plan. These values corresponding to each beam geometry set are agree with that of average for all 10 sets of beam geometries shown in Figure 8. Model I had low prediction errors of average mean dose (D_{mean}) 0.871% (PTV), 0.214% (Body), 1.10% (Bladder), 1.76% (Rectum), 2.03% (Left Femoral Head), and 1.62% (Right Femoral Head). Model II had relatively high prediction errors of D_{mean} corresponding to each beam numbers in plan, with less than 0.40% (Body), 2.00% (Bladder) which are uniform with average prediction errors for all sets of beam geometries reported in Figure 8. For PTV and the rest

of the OARs, the prediction errors corresponding to each beam numbers in plan are within 6.00% (PTV), 5.00% (Rectum), 6.00% (Left Femoral Head), and 4.00% (Right Femoral Head) with maximum of 1% fluctuation from the average prediction errors for all sets of beam geometries reported in Figure 8.

Also, from the observation of prediction errors corresponding to each beam numbers in plan, Model I outperformed Model II for the mean dose errors on PTV, Body, Bladder, Rectum, Left Femoral Head and Right Femoral Head.

Figure 11 shows the colored dose wash distributions in the PTV and in the organs at risk for Protocol Based Setup IMRT and Tuned IMRT. The protocol based setup shows 9 equidistant beam angles of [0, 40, 80, 120, 160, 200, 240, 280, 320] degrees, which are also shown in Figure 11 with red arrows. By tuning one of the protocol based beam angles from 0 degrees to 30 degrees while keeping other parameters the same, we can find a better plan in near real time, since the model's prediction time is under 1 second. Visually, the dose distributions that are predicted from Models I for Tuned IMRT are similar to that of the Protocol Based Setup IMRT. Particularly, they are better matched in the high dose region surrounding the PTV, and have less matched in the lower dose regions.

DVHs obtained from dose predictions for the protocol based setup IMRT plan and tuned IMRT plan are presented in Figure 12. By tuning the beam angle, we were able reduce OAR doses, particularly in the rectum, bladder, and right femoral head, while maintaining similar PTV dose. In addition, the objective values calculated using Equations 2–3 for protocol based setup IMRT and tuned IMRT are 34.93 and 34.56 respectively, showing that the tuned IMRT was able to better satisfy the objective function. From these test cases, it is fully clear that the tuned IMRT plan is better plan than protocol based setup IMRT plan, which can be obtained by tuning the beam angles in real time.

4. Discussion:

The goal of this study was to predict Pareto optimal dose distributions by using anatomical structures and a varying number of beams (up to 10) and beam angles. To our knowledge, this is the first study to implement a deep learning-based, Pareto optimal dose prediction method with such flexibility of beam configuration. We designed deep learning models with two different types of input to represent the beam angles (Fig. 2). For Model I, we directly represented the beam angles as a binary vector for the second input. The anatomical structures used as the three-channel first input were planning treatment volume (PTV), body, and organs at risk (OARs). For Model II, we represented the beam angles as a conformal beam dose and included it as a fourth channel, with the 3 anatomical structure channels, in the model's single input. We generated conformal beam dose data that corresponded to the beam angles used in Model I (section 2.1).

Each model was trained, validated, and tested on 54, 6, and 10 patients, respectively. These patient data yielded a total of 27,000 training, 3,000 validation, and 5,000 testing plans. We used MSE as the loss function and the Adam optimizer for the model to minimize the loss between the ground truth dose and the predicted dose. The default learning rate of 0.01 for

the optimization resulted in the best model for minimizing the validation loss. To avoid overfitting, we applied a dropout scheme in addition to group normalization, as shown in Table 3. Group normalization is more effective than batch normalization at handling small batch sizes.⁶⁷ Also, since group normalization is independent of batch sizes, using it helps to avoid manually selecting batch sizes for better convergence of the network.

Dose color washes in Figure 4 and the dose map differences in Figure 5 show that both models predicted dose distributions within the PTV more accurately than outside regions. This is because the PTV high-dose region is more uniform than other low-dose regions. Overall, our deep learning models (Model I and Model II) predicted voxel-level dose distributions that precisely matched the ground truth dose distributions.

The DVHs generated also precisely matched the ground truth (Fig. 6). Evaluation metrics—such as PTV statistics, dose conformity, dose spillage (R50) and homogeneity index—also confirmed the accuracy of PTV curves on the DVH (Table 4). PTV dose coverage error was within 4% for the prediction from Model I and within 9% for the prediction from Model II (Fig. 7).

Similarly, the predictions of mean,max dose over the PTV and the organs at risk and the mean of voxel wise difference over the body and 10% isodose volume reported in Table 5 and Table 6 respectively indicate the accuracy of both models. The average mean and max dose errors for the prediction from Model I were within 2% and 5%, respectively, for PTV, body, bladder, rectum, left femoral head and right femoral head. Likewise, the average mean and max dose errors for the prediction from Model II were within 6% and 12%, respectively, for PTV, body, bladder, rectum, left femoral head and right femoral head (Fig.8). Also, the mean voxel wise errors for the prediction from Model I and Model II were within 1 % for body and 10% isodose volume (Fig.9). All prediction errors reported in Tables 4, 5 and 6 represent the average errors for all 10 sets of beam geometries.

In addition, the prediction errors of mean dose corresponding to each beam numbers in each plan (1 to 10) over the PTV and the organs at risk, shown in Figure 10, are relatively even regardless of the number of beams the model was using to predict. The average mean dose errors for the prediction from Model I corresponding to each beam numbers in plan agreed with average errors for all 10 sets of beam geometries. The average mean dose errors for the prediction from Model II corresponding to each beam numbers were uniform with Body and Bladder in line with the average errors reported for all 10 sets of beam geometries. For PTV and the rest of the OARs' the prediction errors corresponding to each beam numbers were with maximum of 1% fluctuation from the average errors reported for all 10 sets of beam geometries.

As we know that lower numbers of beam orientations generate low quality plans, and because low-dose region OARs such as femoral heads are further away from the PTV and have higher variability in the dose distribution, prediction errors are higher in these cases. The low prediction errors reported despite such variability in beam angles indicate that both models efficiently predict Pareto optimal dose distributions with high accuracy. However, Model I outperformed Model II in all evaluation criteria mentioned above.

For further verification, we performed t-tests to compare the prediction accuracy between the two models. In all cases, we found that p-values are extremely low (<0.001), which indicates that Model I's performance is statistically significantly superior to that of Model II. Model I outperforming Model II seems counterintuitive at first, as Model II's input of the conformal beam dose seems to be more similar to the final IMRT optimized dose that the model predicts. Although it is possible that including first-order priors or approximations as an input would improve a model's performance, a first-order approximation tends to differ from its exact version only on a local scale. For example, the difference between an accurate dose calculation engine, such as a Monte-Carlo-based engine, and an approximate one is the local scatter contribution from the primary beam. In our case, the difference between our conformal beam dose and the IMRT optimized dose is not local, because changing a beamlet's intensity affects the entire dose distribution along the beamlet's line through the body. This means that the Model II must then learn how to properly add and subtract values from the conformal beam, on a global scale, to transform it to the IMRT dose. This may be equally or more difficult for the neural network than attempting to generate the dose distribution from scratch, as it does with Model I. In addition, Model I is easy to implement and does not need to evaluate the conformal beam dose, which takes extra hours of work. Pareto optimal plans predicted for all 70 patients in this study give physicians the advantage of choosing among different tradeoffs for the critical structures. Physicians can quickly observe multiple Pareto optimal dose predictions in real time and ask the planner for further modifications to obtain the desired tradeoffs. Our method could be the clinical support tool that allows physicians to get the right treatment plan for each patient, and it would give the planner the advantage of making acceptable changes earlier, which would save time in treatment planning.

A potential limitation of this study is the large sampling space from the beam number, beam angles, and the structure weights. There are currently 500 plans per patient, totaling to 35000 plans, and deep learning models do tend to interpolate well between data within its training distribution. However, the true number of needed samples to adequately cover the domain is unknown, and it is possible the model may fail on rare edge cases. Further investigation would be required to determine the number of samples needed in order to prevent such edge failures and to improve the model performance. An alternative could be to limit the sampling to only be in the clinically relevant space, which would drastically reduce the number of required samples.

Modern protocols for prostate IMRT typically ask for 7 or 9 equidistant coplanar beams, while our study uses beam number sampling anywhere from 1 to 10 beam angles, so most of the training samples are different from the IMRT clinical situation, which can raise concerns on whether the model training was compromised. However, as illustrated in Figure 10, we have shown that Model I can predict within roughly 2% mean dose error of the prescription dose for each structure for any beam number setup, which is competitive to any other deep learning-based dose prediction method in literature^{47,48,50–52,54,55,73,74}, most of which use only one type of beam geometry setup for their study. The benefit of training such a beam-flexible model is that, during deployment, the treatment planner may now adjust the beam number in real-time and possibly find a beam configuration the same or better plan quality as before.

Since Model I and Model II both predict well for prostate cancer, we plan to extend these models to other sites under the same IMRT setup. In addition, the Pareto optimal plans are not necessarily deliverable plans, since the machine parameters have yet to be calculated for the predicted dose. Since the additional sequencing steps after optimization may degrade the optimized plan, we plan to examine the extent of plan degradation the dose from applying a sequencing step, as well as investigate adding in the sequencing directly into the Pareto plan optimization as a direct aperture optimization. In addition, we plan to use a threshold-driven optimization engine called TORA,⁷⁵ to create deliverable plans from our current predicted doses in a high quality manner. In addition, we plan to extend this work to the Pareto optimal dose prediction for VMAT plans, given a tunable selection of arc orientations. We hope that this will help us to improve our automated treatment planning system.

While modern protocols for prostate typically ask for a set number of equidistant coplanar beam angles, this can sometimes be varied, where some beams can be dropped or added to tailor the plan to a specific patient. Therefore, we designed to keep the number of beam angles as a flexible parameter, from 1 to 10. In addition, as a future study, we wish to investigate the possible number of beam angles that can be reduced by having mathematically optimized beam angles. This can be achieved and studied by combining our present model with a deep-learning-based beam orientation optimization model.²¹ This beam orientation optimization model can solve for a suitable set of beam angles, given a particular patient anatomy and structure weights. By combining these models, we will develop a framework that can provide plans that are tailored to each patient, in terms of both beam geometries and dosimetric criteria.

5. Conclusion:

We built U-Net-style deep learning models to predict Pareto optimal dose distributions of IMRT prostate plans involving anatomical structures and varying beam angles. We also compared dose predictions between two different beam configuration modalities. We found that both models efficiently predict Pareto optimal dose distributions with high accuracy. However, our deep learning model (Model I) that directly inputs the beam angles into the network as a binary vector was more accurate and robust than the state-of-the-art model (Model II) that inputs a conformal beam dose. Dose predictions from these models take less than a second, which would allow physicians to observe multiple predictions in real time and ask planners for further modifications to achieve the best tradeoffs. From this, planners will be able to make acceptable changes earlier, which will save time in treatment planning. We believe that our method of dose predictions will be a stepping stone to building automatic IMRT treatment planning.

In our future work, we plan to use a threshold-driven optimization engine to generate deliverable plans. By combining our dose prediction model and the beam orientation optimization model, we will develop a unified framework that can provide plans that are tailored to each patient, in terms of both beam geometries and dosimetric criteria. We will further concentrate on building sophisticated dose prediction and beam orientation optimization models that perform well in combination and that can also be independent of treatment sites.

Acknowledgements:

This study was supported by the National Institutes of Health (NIH) R01CA237269. The authors thank Dr. Jonathan Feinberg for editing the manuscript.

7. References:

1. Gianfaldoni S, Gianfaldoni R, Wollina U, Lotti J, Tchernev G, Lotti T. An overview on radiotherapy: from its history to its current applications in dermatology. *Open access Macedonian journal of medical sciences*. 2017;5(4):521. [PubMed: 28785349]
2. Webb S *Intensity-modulated radiation therapy*. CRC Press; 2015.
3. Webb S The physical basis of IMRT and inverse planning. *The British journal of radiology*. 2003;76(910):678–689. [PubMed: 14512327]
4. Nutting C, Dearnaley D, Webb S. Intensity modulated radiation therapy: a clinical review. *The British journal of radiology*. 2000;73(869):459–469. [PubMed: 10884741]
5. Hong T, Ritter M, Tomé WA, Harari P. Intensity-modulated radiation therapy: emerging cancer treatment technology. *British journal of cancer*. 2005;92(10):1819. [PubMed: 15856036]
6. Convery D, Rosenbloom M. The generation of intensity-modulated fields for conformal radiotherapy by dynamic collimation. *Physics in medicine & biology*. 1992;37(6):1359.
7. Bortfeld T IMRT: a review and preview. *Physics in Medicine & Biology*. 2006;51(13):R363. [PubMed: 16790913]
8. Luxton G, Hancock SL, Boyer AL. Dosimetry and radiobiologic model comparison of IMRT and 3D conformal radiotherapy in treatment of carcinoma of the prostate. *International Journal of Radiation Oncology* Biology* Physics*. 2004;59(1):267–284.
9. Kristensen CA, Kjaer-Kristoffersen F, Sapru W, Berthelsen AK, Loft A, Specht L. Nasopharyngeal carcinoma. Treatment planning with IMRT and 3D conformal radiotherapy. *Acta Oncologica*. 2007;46(2):214–220. [PubMed: 17453372]
10. Fenkell L, Kaminsky I, Breen S, Huang S, Van Prooijen M, Ringash J. Dosimetric comparison of IMRT vs. 3D conformal radiotherapy in the treatment of cancer of the cervical esophagus. *Radiotherapy and Oncology*. 2008;89(3):287–291. [PubMed: 18789828]
11. Arbea L, Ramos LI, Martínez-Monge R, Moreno M, Aristu J. Intensity-modulated radiation therapy (IMRT) vs. 3D conformal radiotherapy (3DCRT) in locally advanced rectal cancer (LARC): dosimetric comparison and clinical implications. *Radiation oncology*. 2010;5(1):17. [PubMed: 20187944]
12. Schreiner LJ. On the quality assurance and verification of modern radiation therapy treatment. *Journal of Medical Physics/Association of Medical Physicists of India*. 2011;36(4):189.
13. Craft DL, Hong TS, Shih HA, Bortfeld TR. Improved planning time and plan quality through multicriteria optimization for intensity-modulated radiotherapy. *International Journal of Radiation Oncology* Biology* Physics*. 2012;82(1):e83–e90.
14. Craft DL, Halabi TF, Shih HA, Bortfeld TR. Approximating convex Pareto surfaces in multiobjective radiotherapy planning. *Medical physics*. 2006;33(9):3399–3407. [PubMed: 17022236]
15. Monz M, Küfer K, Bortfeld T, Thieke C. Pareto navigation—algorithmic foundation of interactive multi-criteria IMRT planning. *Physics in Medicine & Biology*. 2008;53(4):985. [PubMed: 18263953]
16. O'Connor D, Yu V, Nguyen D, Ruan D, Sheng K. Fraction-variant beam orientation optimization for non-coplanar IMRT. *Physics in Medicine & Biology*. 2018;63(4):045015. [PubMed: 29351088]
17. Nguyen D, Rwigema J-CM, Victoria YY, et al. Feasibility of extreme dose escalation for glioblastoma multiforme using 4 π radiotherapy. *Radiation Oncology*. 2014;9(1):239. [PubMed: 25377756]
18. Nguyen D, Dong P, Long T, et al. Integral dose investigation of non-coplanar treatment beam geometries in radiotherapy. *Medical physics*. 2014;41(1):011905. [PubMed: 24387513]

19. Jia X, Men C, Lou Y, Jiang SB. Beam orientation optimization for intensity modulated radiation therapy using adaptive l2, 1-minimization. *Physics in Medicine & Biology*. 2011;56(19):6205. [PubMed: 21891848]
20. Breedveld S, Storchi PR, Voet PW, Heijmen BJ. iCycle: Integrated, multicriterial beam angle, and profile optimization for generation of coplanar and noncoplanar IMRT plans. *Medical physics*. 2012;39(2):951–963. [PubMed: 22320804]
21. Barkousaraie AS, Ogunmolu O, Jiang S, Nguyen D. A Fast Deep Learning Approach for Beam Orientation Optimization for Prostate Cancer IMRT Treatments. arXiv preprint arXiv:190500523. 2019.
22. van Asselen B, Schwarz M, van Vliet-Vroegindeweyj C, Lebesque JV, Mijnheer BJ, Damen EM. Intensity-modulated radiotherapy of breast cancer using direct aperture optimization. *Radiotherapy and oncology*. 2006;79(2):162–169. [PubMed: 16712992]
23. Shepard D, Earl M, Li X, Naqvi S, Yu C. Direct aperture optimization: a turnkey solution for step-and-shoot IMRT. *Medical physics*. 2002;29(6):1007–1018. [PubMed: 12094970]
24. Men C, Romeijn HE, Taşkın ZC, Dempsey JF. An exact approach to direct aperture optimization in IMRT treatment planning. *Physics in Medicine & Biology*. 2007;52(24):7333. [PubMed: 18065842]
25. Cassioli A, Unkelbach J. Aperture shape optimization for IMRT treatment planning. *Physics in Medicine & Biology*. 2012;58(2):301. [PubMed: 23257284]
26. Bedford JL, Webb S. Constrained segment shapes in direct-aperture optimization for step-and-shoot IMRT. *Medical physics*. 2006;33(4):944–958. [PubMed: 16696471]
27. Ahunbay EE, Chen G-P, Thatcher S, et al. Direct aperture optimization-based intensity-modulated radiotherapy for whole breast irradiation. *International Journal of Radiation Oncology* Biology* Physics*. 2007;67(4):1248–1258.
28. Perumal B, Sundaresan HE, Ranganathan V, Ramar N, Anto GJ, Meher SR. Evaluation of plan quality improvements in PlanIQ-guided Autoplanning. *Reports of Practical Oncology & Radiotherapy*. 2019;24(6):533–543. [PubMed: 31641339]
29. Nelms BE, Robinson G, Markham J, et al. Variation in external beam treatment plan quality: an inter-institutional study of planners and planning systems. *Practical radiation oncology*. 2012;2(4):296–305. [PubMed: 24674168]
30. Moore KL, Schmidt R, Moiseenko V, 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*. 2015;92(2):228–235.
31. Fan J, Wang J, Zhang Z, Hu W. Iterative dataset optimization in automated planning: Implementation for breast and rectal cancer radiotherapy. *Medical physics*. 2017;44(6):2515–2531. [PubMed: 28339103]
32. Zhu X, Ge Y, Li T, Thongphiew D, Yin FF, Wu QJ. A planning quality evaluation tool for prostate adaptive IMRT based on machine learning [published online ahead of print 2011/04/02]. *Med Phys*. 2011;38(2):719–726. [PubMed: 21452709]
33. Zhang J, Wu QJ, Xie T, Sheng Y, Yin FF, Ge Y. An Ensemble Approach to Knowledge-Based Intensity-Modulated Radiation Therapy Planning [published online ahead of print 2018/04/05]. *Frontiers in oncology*. 2018;8:57. [PubMed: 29616187]
34. Yuan L, Ge Y, Lee WR, Yin FF, Kirkpatrick JP, Wu QJ. Quantitative analysis of the factors which affect the interpatient organ-at-risk dose sparing variation in IMRT plans [published online ahead of print 2012/11/07]. *Med Phys*. 2012;39(11):6868–6878. [PubMed: 23127079]
35. Nwankwo O, Mekdash H, Sihono DSK, Wenz F, Glatting G. Knowledge-based radiation therapy (KBRT) treatment planning versus planning by experts: validation of a KBRT algorithm for prostate cancer treatment planning. *Radiation Oncology*. 2015;10(1):111. [PubMed: 25957871]
36. Ge Y, Wu QJ. Knowledge-based planning for intensity-modulated radiation therapy: A review of data-driven approaches [published online ahead of print 2019/04/10]. *Med Phys*. 2019;46(6):2760–2775. [PubMed: 30963580]
37. David EW, Marie d. A Call for Knowledge-Based Planning. *AI Magazine*. 2001;22(1).

38. Wu H, Jiang F, Yue H, Li S, Zhang Y. A dosimetric evaluation of knowledge-based VMAT planning with simultaneous integrated boosting for rectal cancer patients [published online ahead of print 2016/12/09]. *J Appl Clin Med Phys*. 2016;17(6):78–85. [PubMed: 27929483]
39. Tol JP, Delaney AR, Dahele M, Slotman BJ, Verbakel WF. Evaluation of a knowledge-based planning solution for head and neck cancer [published online ahead of print 2015/02/15]. *International journal of radiation oncology, biology, physics*. 2015;91(3):612–620.
40. Ma C, Huang F. Assessment of a knowledge-based RapidPlan model for patients with postoperative cervical cancer. *Precision Radiation Oncology*. 2017;1(3):102–107.
41. Kubo K, Monzen H, Ishii K, et al. Dosimetric comparison of RapidPlan and manually optimized plans in volumetric modulated arc therapy for prostate cancer [published online ahead of print 2017/07/15]. *Physica medica : PM : an international journal devoted to the applications of physics to medicine and biology : official journal of the Italian Association of Biomedical Physics (AIFB)*. 2017;44:199–204.
42. Fogliata A, Reggiori G, Stravato A, et al. RapidPlan head and neck model: the objectives and possible clinical benefit [published online ahead of print 2017/04/30]. *Radiation oncology (London, England)*. 2017;12(1):73.
43. Chang ATY, Hung AWM, Cheung FWK, et al. Comparison of Planning Quality and Efficiency Between Conventional and Knowledge-based Algorithms in Nasopharyngeal Cancer Patients Using Intensity Modulated Radiation Therapy [published online ahead of print 2016/06/16]. *International journal of radiation oncology, biology, physics*. 2016;95(3):981–990.
44. Nguyen D, Barkousaraie AS, Shen C, Jia X, Jiang S. Generating Pareto optimal dose distributions for radiation therapy treatment planning. *arXiv preprint arXiv:190604778*. 2019.
45. LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature*. 2015;521(7553):436–444. [PubMed: 26017442]
46. Long J, Shelhamer E, Darrell T. Fully convolutional networks for semantic segmentation. Paper presented at: *Proceedings of the IEEE conference on computer vision and pattern recognition* 2015.
47. Nguyen D, Long T, Jia X, et al. A feasibility study for predicting optimal radiation therapy dose distributions of prostate cancer patients from patient anatomy using deep learning. *Scientific Reports*. 2019;9(1):1076. [PubMed: 30705354]
48. Nguyen D, Jia X, Sher 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 & Biology*. 2019;64(6):065020. [PubMed: 30703760]
49. Mahmood R, Babier A, McNiven A, Diamant A, Chan TC. Automated treatment planning in radiation therapy using generative adversarial networks. *arXiv preprint arXiv:180706489*. 2018.
50. Kearney V, Chan JW, Haaf S, Descovich M, Solberg TD. DoseNet: a volumetric dose prediction algorithm using 3D fully-convolutional neural networks. *Physics in Medicine & Biology*. 2018;63(23):235022. [PubMed: 30511663]
51. Kajikawa T, Kadoya N, Ito K, et al. A convolutional neural network approach for IMRT dose distribution prediction in prostate cancer patients. *Journal of radiation research*. 2019;60(5):685–693. [PubMed: 31322704]
52. 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. *Medical physics*. 2019;46(1):370–381. [PubMed: 30383300]
53. 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*. 2019;46(1):56–64. [PubMed: 30367492]
54. Barragán-Montero AM, Nguyen D, Lu W, et al. Three-Dimensional Dose Prediction for Lung IMRT Patients with Deep Neural Networks: Robust Learning from Heterogeneous Beam Configurations. *Medical physics*. 2019.
55. Nguyen D, McBeth R, Sadeghnejad Barkousaraie A, et al. Incorporating human and learned domain knowledge into training deep neural networks: A differentiable dose volume histogram and adversarial inspired framework for generating Pareto optimal dose distributions in radiation therapy [published online ahead of print 2019/12/11]. *Med Phys*. 2019.

56. Lu W, Chen M. Fluence-convolution broad-beam (FCBB) dose calculation. *Physics in Medicine & Biology*. 2010;55(23):7211. [PubMed: 21081826]
57. Jahn J. Scalarization in multi objective optimization In: *Mathematics of multi objective optimization*. Springer; 1985:45–88.
58. Breedveld S, Craft D, Van Haveren R, Heijmen B. Multi-criteria optimization and decision-making in radiotherapy. *European Journal of Operational Research*. 2019;277(1):1–19.
59. Chambolle A, Pock T. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of mathematical imaging and vision*. 2011;40(1):120–145.
60. Nguyen D, O'Connor D, Ruan D, Sheng K. Deterministic direct aperture optimization using multiphase piecewise constant segmentation. *Medical Physics*. 2017;44(11):5596–5609. [PubMed: 28834556]
61. Nguyen D, Ruan D, O'Connor D, et al. A novel software and conceptual design of the hardware platform for intensity modulated radiation therapy. *Medical Physics*. 2016;43(2):917–929. [PubMed: 26843252]
62. Nguyen D, Thomas D, Cao M, O'Connor D, Lamb J, Sheng K. Computerized triplet beam orientation optimization for MRI-guided Co-60 radiotherapy. *Medical Physics*. 2016;43(10):5667–5675. [PubMed: 27782726]
63. Nguyen D, Lyu Q, Ruan D, O'Connor D, Low DA, Sheng K. A comprehensive formulation for volumetric modulated arc therapy planning. *Medical Physics*. 2016;43(7):4263–4272. [PubMed: 27370141]
64. Nguyen D, O'Connor D, Yu VY, et al. Dose domain regularization of MLC leaf patterns for highly complex IMRT plans. *Medical Physics*. 2015;42(4):1858–1870. [PubMed: 25832076]
65. Lu W. A non-voxel-based broad-beam (NVBB) framework for IMRT treatment planning. *Physics in Medicine & Biology*. 2010;55(23):7175. [PubMed: 21081819]
66. Ronneberger O, Fischer P, Brox T. U-net: Convolutional networks for biomedical image segmentation. Paper presented at: *International Conference on Medical image computing and computer-assisted intervention* 2015.
67. Wu Y, He K. Group normalization. Paper presented at: *Proceedings of the European Conference on Computer Vision (ECCV)* 2018.
68. Kingma D, Ba J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*. 2014.
69. Yao Y, Rosasco L, Caponnetto A. On early stopping in gradient descent learning. *Constructive Approximation*. 2007;26(2):289–315.
70. Paddick I. A simple scoring ratio to index the conformity of radiosurgical treatment plans. *J Neurosurg*. 2000;93(Suppl 3):219–222. [PubMed: 11143252]
71. Greenspan H, Van Ginneken B, Summers RM. Guest editorial deep learning in medical imaging: Overview and future promise of an exciting new technique. *IEEE Transactions on Medical Imaging*. 2016;35(5):1153–1159.
72. Grégoire V, Mackie T. State of the art on dose prescription, reporting and recording in intensity-modulated radiation therapy (ICRU report No. 83). *Cancer/Radiothérapie*. 2011;15(6–7):555–559.
73. Nguyen D, Barkousaraie AS, Shen C, Jia X, Jiang S. Generating Pareto Optimal Dose Distributions for Radiation Therapy Treatment Planning. *Lecture Notes in Computer Science*. 2019;11769:59–67.
74. Babier A, Boutilier JJ, McNiven AL, Chan TC. Knowledge-based automated planning for oropharyngeal cancer. *Medical physics*. 2018;45(7):2875–2883. [PubMed: 29679492]
75. Long T, Chen M, Jiang S, Lu W. Threshold-driven optimization for reference-based auto-planning. *Physics in Medicine & Biology*. 2018;63(4):04NT01.




Figure 1:
Conformal dose corresponding to different beam angles (1–10).

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

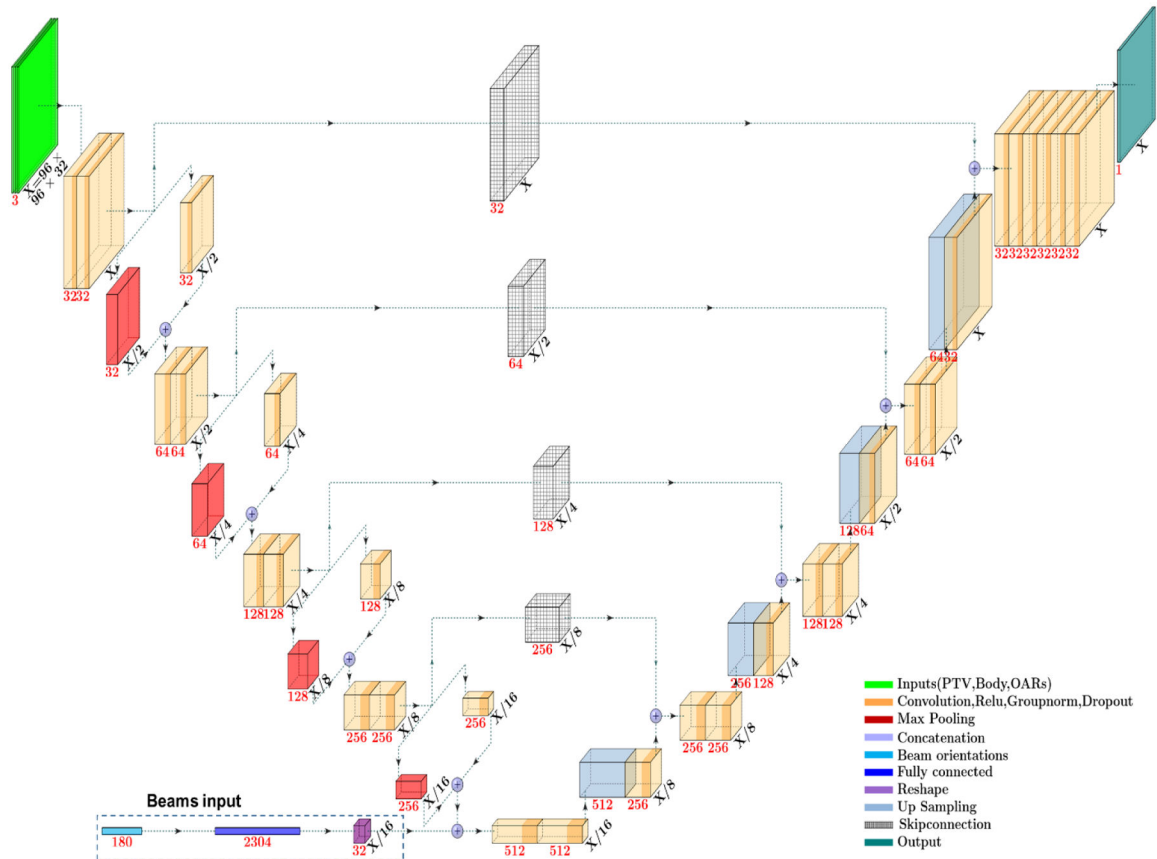


Figure 2:
Deep learning models used in the study.

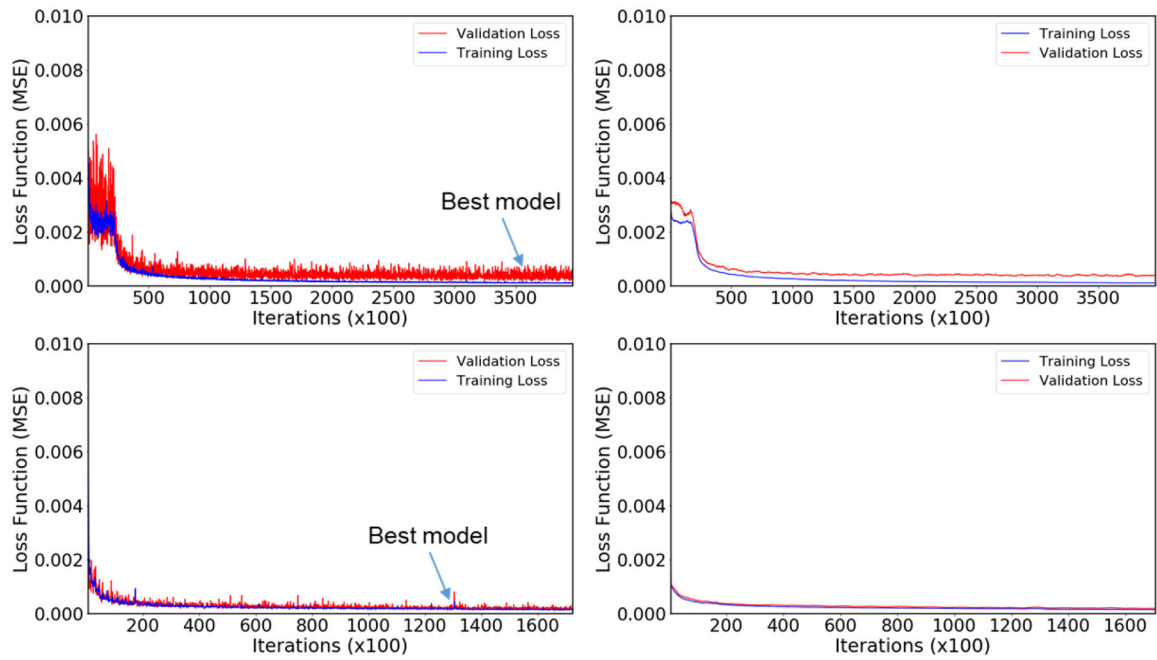


Figure 3: Training vs. Validation loss as a function of iterations for both models. Top row plots are for Model I, and bottom row plots are for Model II. Left column plots represent the actual training and validation loss, and the right column plots represent the smooth training and validation losses obtained by using the moving average method.




Figure 4. Ground truth dose distribution vs. dose distribution predicted by Models I and II. The figures in the first column represents the ground truth dose distributions for two different treatment plans. The second column and third column distributions were predicted by Model I and Model II, respectively.




Figure 5.

Dose map differences between the Model I and Model II predicted dose distributions and the ground truth dose distribution. These differences are obtained taking the differences between Model I and Ground truth, and Model II and Ground truth as shown in second row of Figure 4.

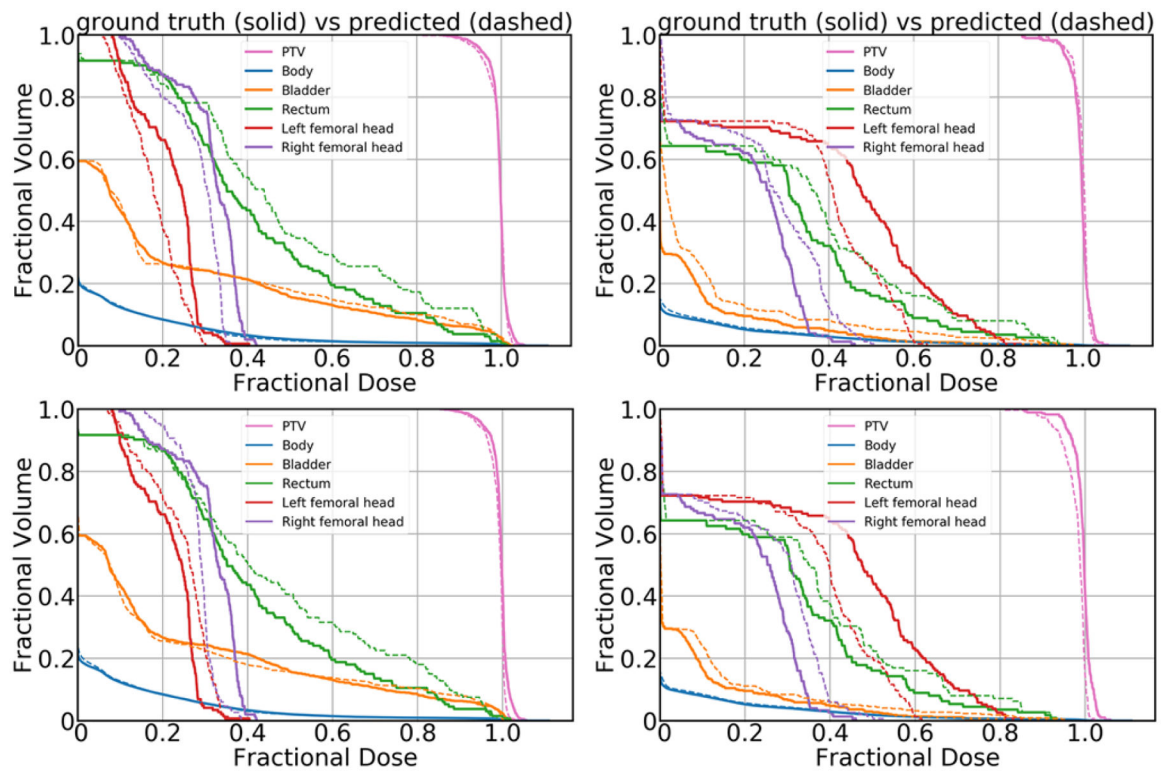


Figure 6.

DVH plots obtained from Model I (top row) and Model II (bottom row). The top row plots were obtained from Model I, and the bottom row plots were obtained from Model II. The solid lines correspond to the ground truth dose, and the dashed lines correspond to the predicted dose.




Figure 7: Prediction errors obtained from Models I and II for conformation, high dose spillage (R50), homogeneity, and PTV dose coverage on the test data. Error bar represents the 99% confidence interval ($\bar{x} \pm 2.576 * \frac{\sigma}{\sqrt{n}}$), where \bar{x} and σ are mean and standard deviation, respectively.

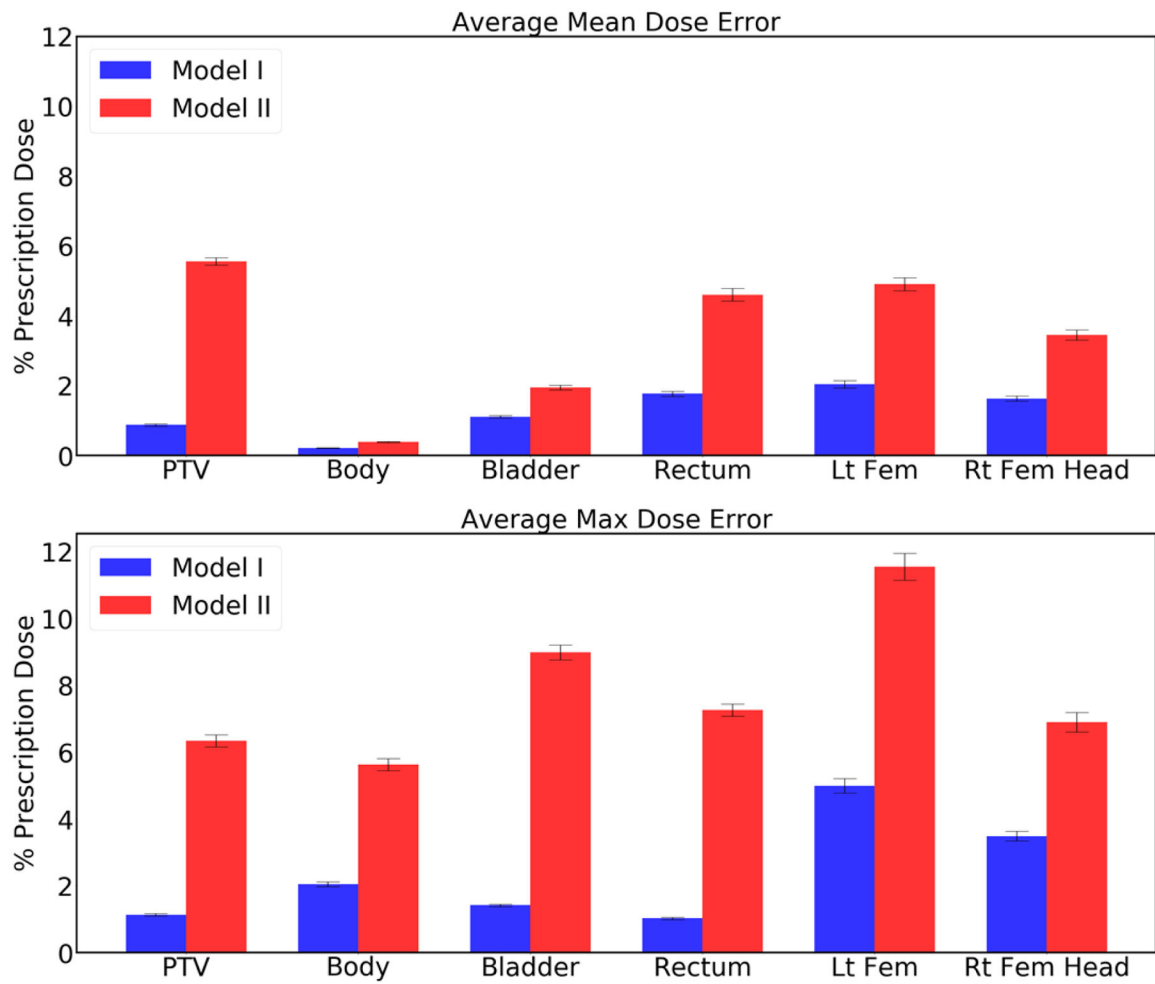


Figure 8: Average error in the mean dose (top plot) and the max dose (bottom plot) for the PTV and the organs at risk. Error bar represents the 99% confidence interval ($\bar{x} \pm 2.576 * \frac{\sigma}{\sqrt{n}}$), where \bar{x} and σ are mean and standard deviation, respectively.




Figure 9:

Average voxel wise dose error for the Body and the 10% isodose volume. Error bar represents the 99% confidence interval $(\bar{x} \pm 2.576 * \frac{\sigma}{\sqrt{n}})$, where \bar{x} and σ are mean and standard deviation, respectively.

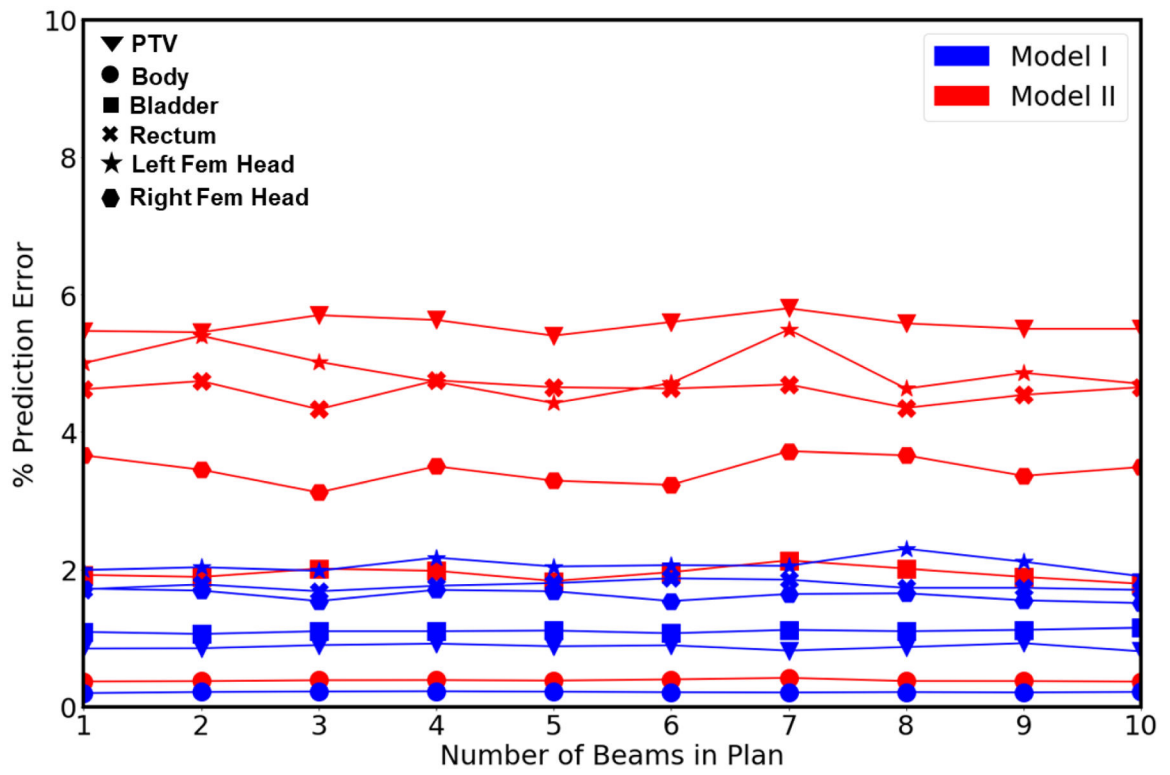


Figure 10: Average error in the mean dose for the PTV and the organs at risk obtained from Model I and II corresponding to number of beams in plan geometry (1–10).

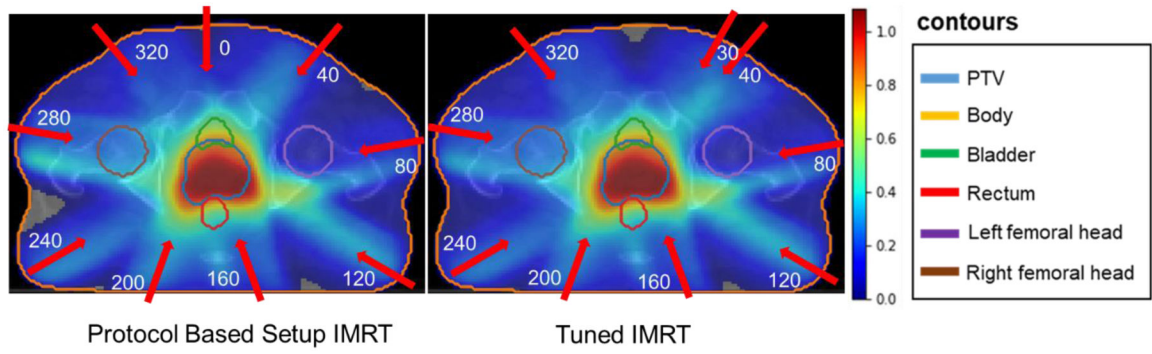


Figure11:
Protocol Based Setup IMRT and Tuned IMRT dose distribution predicted by Models I.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

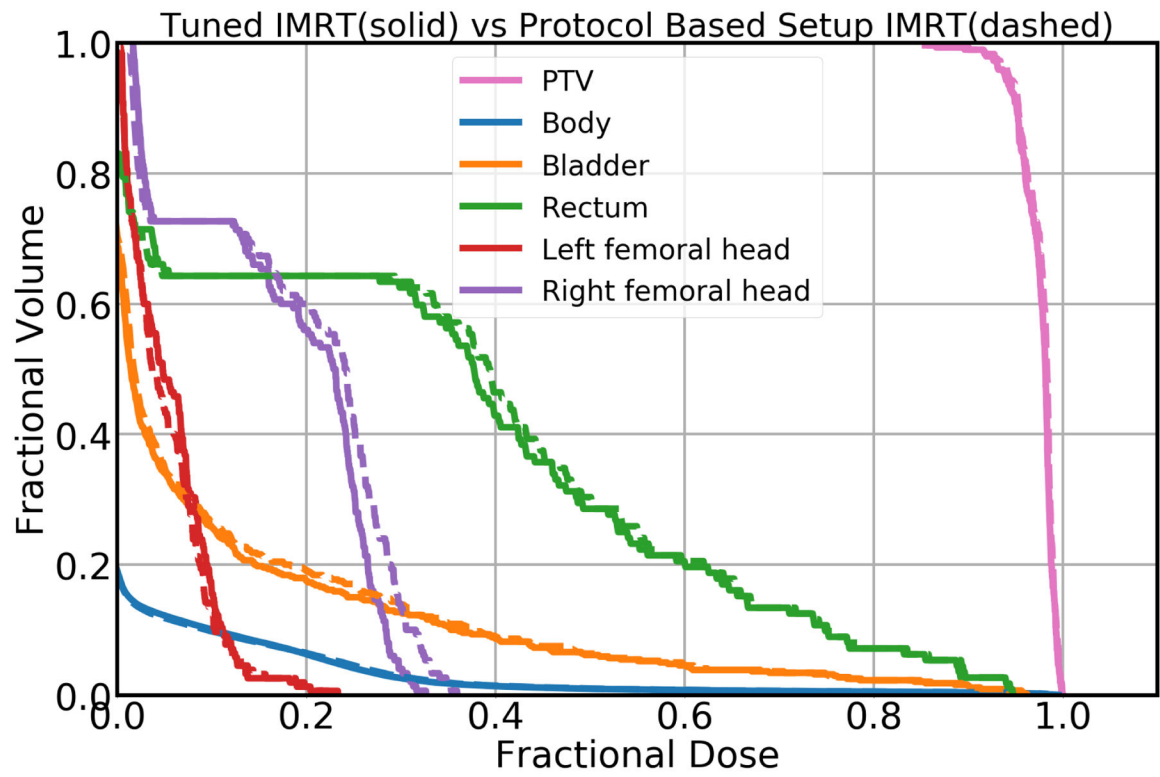


Figure 12: DVH plot obtained from Model I. The solid lines correspond to the Tuned IMRT, and the dashed lines correspond to the Protocol Based Setup IMRT.

Table 1:

Weight generation categories for the organs at risk. The function $rand(L_B, U_B)$ creates a uniform random number between a lower bound (L_B) and an upper bound (U_B). In all categories, the PTV weights were assigned 1.

Category	Description
Low weights	$w_s = rand(0,0.1) \forall s \in OAR$
Extra low weights	$w_s = rand(0,0.05) \forall s \in OAR$
Controlled weights	$w_{bladder} = rand(0,0.2)$ $w_{rectum} = rand(0,0.2)$ $w_{lt\ fem\ head} = rand(0,0.1)$ $w_{rt\ fem\ head} = rand(0,0.1)$ $w_{shell} = rand(0,0.1)$ $w_{skin} = rand(0,0.3)$

Table 2:

Distribution of Pareto plans. 40% of plans were assigned with low weights, 20% of plans were assigned extra low weights, and 40% of plans were assigned with controlled weights. The total number of plans in each category is 10 times the value shown, since there are a total of 10 beam sets (i.e. 1-beam plan, 2-beam plan, ..., 10-beam plan).

Weights	Training Plans per beam set	Validation Plans per beam set	Testing Plans per beam set
Low	1080	120	200
Extra Low	540	60	100
Controlled	1080	120	200

Table 3:

Dropout rate scheme used in the networks.

U-net Hierarchy Level	Groups	Dropout Rate
1	1	0.125
2	2	0.148
3	4	0.176
4	8	0.210
5	16	0.250

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 4:

Means and standard deviations (Mean \pm SD) for clinical DVH metrics of ground truth (Pareto optimal) dose, predicted dose, absolute difference between the predicted dose and the ground truth, conformation, high dose spillage (R50) and homogeneity obtained from Model I and Model II.

		Model I	Model II	Model I	Model II
	Pareto optimal dose	Predicted dose	Predicted dose	 Predicted – Ground truth 	 Predicted – Ground truth
	Mean \pm SD	Mean \pm SD	Mean \pm SD	Mean \pm SD	Mean \pm SD
PTV D98	0.75 \pm 0.08	0.75 \pm 0.09	0.69 \pm 0.09	0.04 \pm 0.03	0.07 \pm 0.05
PTV D95	0.82 \pm 0.06	0.83 \pm 0.06	0.76 \pm 0.06	0.03 \pm 0.03	0.07 \pm 0.05
PTV D50	0.99 \pm 0.01	0.99 \pm 0.01	0.90 \pm 0.03	0.01 \pm 0.01	0.08 \pm 0.02
PTV D2	1.02 \pm 0.04	1.03 \pm 0.04	0.96 \pm 0.02	0.01 \pm 0.01	0.06 \pm 0.05
Paddick Confirmation number	0.63 \pm 0.25	0.66 \pm 0.26	0.59 \pm 0.21	0.04 \pm 0.04	0.08 \pm 0.07
R50	5.2 \pm 1.3	5.0 \pm 1.4	5.4 \pm 1.2	0.33 \pm 0.23	0.63 \pm 0.55
PTV Homogeneity	0.29 \pm 0.11	0.29 \pm 0.12	0.30 \pm 0.10	0.04 \pm 0.04	0.06 \pm 0.05

Table 5:

Mean and standard deviation (Mean \pm SD) of maximum and mean values of the Pareto optimal dose distribution, the predicted dose distribution, and the absolute difference between the predicted dose distribution and the ground truth received on the PTV and other critical structures.

		Model I		Model II		
Pareto optimal dose		Predicted dose	Predicted dose	Predicted – Ground truth	Predicted – Ground truth	
Mean \pm SD		Mean \pm SD	Mean \pm SD	Mean \pm SD	Mean \pm SD	
D_{\max}	PTV	1.02 \pm 0.04	1.03 \pm 0.04	0.96 \pm 0.02	0.01 \pm 0.01	0.06 \pm 0.05
	Body	0.54 \pm 0.01	0.53 \pm 0.15	0.52 \pm 0.12	0.02 \pm 0.02	0.06 \pm 0.05
	Bladder	0.96 \pm 0.08	0.97 \pm 0.08	0.88 \pm 0.06	0.01 \pm 0.01	0.09 \pm 0.06
	Rectum	1.00 \pm 0.05	1.00 \pm 0.05	0.93 \pm 0.04	0.01 \pm 0.01	0.07 \pm 0.05
	Left fem	0.47 \pm 0.30	0.46 \pm 0.30	0.47 \pm 0.25	0.05 \pm 0.06	0.12 \pm 0.11
	Right fem	0.35 \pm 0.28	0.33 \pm 0.27	0.32 \pm 0.23	0.03 \pm 0.04	0.07 \pm 0.08
D_{mean}	PTV	0.87 \pm 0.08	0.87 \pm 0.08	0.82 \pm 0.07	0.01 \pm 0.01	0.06 \pm 0.03
	Body	0.03 \pm 0.01	0.03 \pm 0.01	0.03 \pm 0.01	0.00 \pm 0.00	0.00 \pm 0.00
	Bladder	0.22 \pm 0.11	0.22 \pm 0.10	0.20 \pm 0.10	0.01 \pm 0.01	0.02 \pm 0.02
	Rectum	0.51 \pm 0.14	0.51 \pm 0.14	0.48 \pm 0.12	0.02 \pm 0.02	0.05 \pm 0.04
	Left fem	0.18 \pm 0.15	0.18 \pm 0.15	0.19 \pm 0.14	0.02 \pm 0.03	0.05 \pm 0.05
	Right fem	0.16 \pm 0.16	0.16 \pm 0.15	0.16 \pm 0.15	0.02 \pm 0.02	0.03 \pm 0.04

Table 6:

Mean and standard deviation (Mean \pm SD) of the Body and 10 % isodose volume of the absolute voxel based dose differences between predicted and ground truth dose distribution.

Voxel Wise Dose Difference	Model I	Model II
D-mean	 Predicted – Ground truth (Mean\pm std)	 Predicted – Ground truth (Mean\pm std)
Body	0.0061 \pm 0.0018	0.0090 \pm 0.0029
10% isodose volume	0.0049 \pm 0.0011	0.0088 \pm 0.0022

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript