



Published in final edited form as:

Med Phys. 2020 March ; 47(3): 837–849. doi:10.1002/mp.13955.

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

Dan Nguyen, Rafe McBeth, Azar Sadeghnejad Barkousaraie, Gyanendra Bohara, Chenyang Shen, Xun Jia, Steve Jiang

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

Abstract

Purpose: We propose a novel domain specific loss, which is a differentiable loss function based on the dose volume histogram, and combine it with an adversarial loss for the training of deep neural networks. In this study, we trained a neural network for generating Pareto optimal dose distributions, and evaluate the effects of the domain specific loss on the model performance.

Methods: In this study, 3 loss functions—mean squared error (MSE) loss, dose volume histogram (DVH) loss, and adversarial (ADV) loss—were used to train and compare 4 instances of the neural network model: 1) MSE, 2) MSE+ADV, 3) MSE+DVH, and 4) MSE+DVH+ADV. The data for 70 prostate patients, including the planning target volume (PTV), and the organs-at-risk (OAR) were acquired as $96 \times 96 \times 24$ dimension arrays at 5 mm^3 voxel size. The dose influence arrays were calculated for 70 prostate patients, using a 7 equidistant coplanar beam setup. Using a scalarized multicriteria optimization for intensity modulated radiation therapy, 1200 Pareto surface plans per patient were generated by pseudo-randomizing the PTV and OAR tradeoff weights. With 70 patients, the total number of plans generated was 84,000 plans. We divided the data into 54 training, 6 validation, and 10 testing patients. Each model was trained for a total of 100,000 iterations, with a batch size of 2. All models used the Adam optimizer, with a learning rate of 1×10^{-3} .

Results: Training for 100,000 iterations took 1.5 days (MSE), 3.5 days (MSE+ADV), 2.3 days (MSE+DVH), 3.8 days (MSE+DVH+ADV). After training, the prediction time of each model is 0.052 seconds. Quantitatively, the MSE+DVH+ADV model had the lowest prediction error of 0.038 (conformation), 0.026 (homogeneity), 0.298 (R50), 1.65% (D95), 2.14% (D98), 2.43% (D99). The MSE model had the worst prediction error of 0.134 (conformation), 0.041 (homogeneity), 0.520 (R50), 3.91% (D95), 4.33% (D98), 4.60% (D99). For both the mean dose PTV error and the max dose PTV, Body, Bladder and rectum error, the MSE+DVH+ADV outperformed all other models. Regardless of model, all predictions have an average mean and max dose error less than 2.8% and 4.2%, respectively.

Conclusion: The MSE+DVH+ADV model performed the best in these categories, illustrating the importance of both human and learned domain knowledge. Expert human domain specific knowledge can be the largest driver in the performance improvement, and adversarial learning can be used to further capture nuanced attributes in the data. The real-time prediction capabilities allow for a physician to quickly navigate the tradeoff space for a patient, and produce a dose distribution as a tangible endpoint for the dosimetrist to use for planning. This is expected to considerably reduce the treatment planning time, allowing for clinicians to focus their efforts on the difficult and demanding cases.

I. Introduction

External beam radiation therapy is one of the major treatments available to cancer patients, with major modalities available including intensity modulated radiation therapy (IMRT)¹⁻⁷ and volume modulated arc therapy (VMAT)⁸⁻¹⁵. IMRT and VMAT have revolutionized the treatment planning over the past decades, drastically improving the treatment plan quality and patient outcome. However, many tedious and time consuming aspects still exist within the clinical treatment planning workflow. In particular, there are two aspects: 1) The dosimetrist must tediously and iteratively tune the treatment planning hyperparameters of the fluence map optimization in order to arrive at a planner-acceptable dose, and 2) many feedback loops between the physician and dosimetrist occur for the physician to provide his comments and judgement on the plan quality, until a physician-acceptable dose is finally produced. For a patient, this process can continually repeat for many hours to many days, depending on the complexity of the plan.

Much work over the years has been focused on reducing the treatment complexity by simplifying certain aspects in the planning workflow, such as feasibility seeking¹⁶, multicriteria optimization for tradeoff navigation on the Pareto surface¹⁷⁻¹⁹, and other algorithms for performance improvements²⁰⁻²⁶. While effective, these methods still require a large amount of intelligent input from the dosimetrist and physician, such as in weight tuning and deciding appropriate dose-volume constraints and tradeoffs.

To address this, the field developed machine learning models to predict relevant dosimetric endpoints, which can be categorized into 1 of 3 categories: 1) predicting single dose constraint points directly, 2) predicting dose volume histograms (DVH), and 3) predicting the 3D dose distribution of the plan. With a 3D dose distribution, one can fully reconstruct the DVH, and with the DVH, the dose constraints can then be calculated. Many studies focused on either predicting dose constraints or the dose volume histogram, eventually forming the backbone of knowledge-based planning (KBP)²⁷⁻⁴². KBP used machine learning techniques and models to predict clinically acceptable dosimetric criteria, utilizing a large pool of historical patient plans and information to draw its knowledge from. Before the era of deep neural networks, KBP's efficacy was heavily reliant on not only the patient data size and diversity, but also on the careful selection of features extracted from the data to be used in the model^{32-39,42,43}. This limited the model to predict small dimensional data, such as the DVH or specific dosimetrist criteria.

With the advancements in deep learning, particularly in computer vision^{44–46} and convolutional neural networks⁴⁷, many studies have investigated clinical dose distribution prediction using deep learning on several sites such as for prostate IMRT^{48,49}, prostate VMAT³², lung IMRT⁵⁰, head-and-neck IMRT^{51–54}, head-and-neck VMAT⁵⁵. In addition to clinical dose prediction, deep learning models are capable of accurately generating Pareto optimal dose distributions, navigating the various tradeoffs between planning target volume (PTV) dose coverage and organs-at-risk (OAR) dose sparing⁵⁶. Most of these methods utilize a simple loss function for training the neural network—the mean squared error (MSE) loss. MSE loss is a generalized, domain-agnostic loss function that can be applied to many problems in many domains. Its large flexibility also means that it is incapable of driving its performance in a domain-specific manner.

Mahmood and Babier et al.^{52–54} investigated the use of adversarial learning for dose prediction. Since the development of generative adversarial networks (GAN) by Goodfellow⁵⁷, adversarial loss has been popularized in the deep learning community for many applications. While used heavily for generative models, such as GANs, the adversarial loss can be applied to almost any neural network training. The adversarial loss's emerging success in deep learning application is largely due to the discriminator capability to calculate its own feature maps during the training process. In essence, the discriminator is learning its own domain knowledge of the problem. However, an adversarial framework is not without its issues. The user has little control over what kinds of features the discriminator may be learning during the training process. It is possible for the discriminator to learn the correct answer for the wrong reason. In addition, careful balancing of the learning between the two networks is essential for preventing catastrophic failure. These may affect the overall performance of the prediction framework.

In 2018, Muralidhar et al.⁵⁸ proposed a domain adapted loss into their neural network training, in order to address deep learning in cases of limited and poor-quality data, which is a problem commonly found within the medical field. They found that, by including domain-explicit constraints, the domain adapted network model had drastically improved performance over its domain-agnostic counterpart, especially in the limited, poor-quality data situation. We realize the importance of including domain specific losses into the radiation therapy problem of dose prediction. We propose the addition of a differentiable loss function based on the dose volume histogram (DVH), one of the most important and commonly used metrics in radiation oncology, into the training of deep neural networks for volumetric dose distribution prediction. In this study, we will train a neural network for generating Pareto optimal dose distributions, and evaluate the effects of MSE loss, DVH loss, and adversarial loss on the network's performance. Pareto optimal plans are the solutions to a multicriteria problem with various tradeoffs. In particular, the tradeoff lies with the dose coverage of the tumor and the dose sparing of the various critical structures. The benefit of such a model is two-fold. First, the physician can interact with the model to immediately view a dose distribution, and then adjust some parameters to push the dose towards their desired tradeoff in real time. This also allows for the physician to quickly comprehend the kinds of the tradeoffs that are feasible for the patient. Second, the treatment planner, upon receiving the physician's desired dose distribution, can quickly generate a

fully deliverable plan that matches this dose distribution, saving time in tuning the optimization hyperparameters and discussing with the physician.

Because generating Pareto optimal plans for the patient requires for the network to learn how to map many dose distributions with tradeoffs from a single anatomy, the neural network must learn to differentiate the potentially small nuances between the different doses that may have substantial clinical consequences. While these small nuances may not be well reflected in a metric such as voxel-wise mean squared error, our own domain metrics can amplify the clinically relevant differences. The usage of an adversarial loss can further aid the neural network in learning important differences that cannot be easily formulated into a loss function. We believe that training a network to generate Pareto optimal dose distribution is well suited for testing the effects of MSE loss, DVH loss, and adversarial loss.

II. Methods

II.1. Patient and Pareto Optimal Plan Data

The data for 70 prostate patients, including the planning target volume (PTV), and the organs-at-risk (OAR)—body, bladder, rectum, left femoral head, and right femoral head—were acquired as $96 \times 96 \times 24$ dimension arrays at 5 mm^3 voxel size. Ring and skin structures were added as tuning structures. The dose influence arrays were calculated for the 70 patients, using a 7 equidistant coplanar beam plan IMRT setup. The beamlet size was 2.5 mm^2 at 100 cm isocenter. Using this dose influence data, we generated IMRT plans that sampled the Pareto surface, representing various tradeoffs between the PTV dose coverage and OAR dose sparing. The multicriteria objective can be written as

$$\underset{x}{\operatorname{argmin}} \left\{ f_{PTV}(x), f_{OAR_1}(x), f_{OAR_2}(x), \dots, f_{OAR_{n_{OAR}}}(x) \right\}$$

$$\text{subject to} \quad x \geq 0, \tag{1}$$

where x is the fluence map intensities to be optimized. The individual objectives, $f_s(x) \forall s \in \{PTV, OAR_u \forall u \in \{1, 2, \dots, n_{OAR}\}\}$, are for the PTV and each of the OARs used in the optimization problem, where n_{OAR} represents the total number of OARs. The index s represents a structure used in the optimization, which is the PTV or one of the OARs. For simplicity, we define S as the set of all structures used in the optimization. In this case, $S = \{PTV, OAR_u \forall u \in \{1, 2, \dots, n_{OAR}\}\}$ and $s \in S$. In radiation therapy, the objective function is formulated with the intention to deliver the prescribed dose to the PTV, while minimizing the dose to each OAR. Because to the physical aspects of radiation in external beam radiation therapy, it is impossible to deliver to the PTV the prescription dose without irradiating normal tissue. In addition, it has been shown that the integral dose to the body is similar regardless of the plan^{59–62}, so, in essence, one can only choose how to best distribute the radiation in the normal tissue. For example, by reducing the dose to one OAR, either the PTV coverage will worsen or more dose will be delivered to another OAR. Therefore, we arrive at a multicriteria objective, where there does not exist a single optimal x^* that would minimize all $f_s(x) \forall s \in PTV, OAR$. In this study, we choose to use the ℓ_2 -norm to represent

the objective, $f_s(x) = \frac{1}{2} \|A_s x - p_s\|_2^2$. In this formulation, A_s is the dose influence matrix for a given structure, and p_s is the desired dose for a given structure, assigned as the prescription dose if s is the PTV, and 0 otherwise. Our beamlet-based dose influence matrix was calculated using the Eclipse AAA dose calculation engine (Varian Medical Systems, Inc.). This allows for us to linearly scalarize the multicriteria optimization problem⁶³, by reformulating it into a single-objective, convex optimization problem

$$\begin{aligned} \underset{x}{\operatorname{argmin}} \quad & \frac{1}{2} \sum_{s \in \mathcal{S}} w_s^2 \|A_s x - p_s\|_2^2 \\ \text{subject to} \quad & x \geq 0. \end{aligned} \tag{2}$$

Scalarizing the problem required the addition of new hyperparameters, w_s , which are the tradeoff weights for each objective function, $f_s(x) \forall s \in \mathcal{S}$. By varying w_s to different values, different Pareto optimal solutions can be generated by solving the optimization problem. Using an in-house GPU-based proximal-class first-order primal-dual algorithm, Chambolle-Pock⁶⁴, we generated many pseudo-random plans, by assigning pseudo-random weights to the organs-at-risk. The weight assignment fell into 1 of 6 categories as shown in Table 1.

For each patient, 100 plans of the single organ spare category (bladder, rectum, left femoral head, right femoral head, shell, skin) were generated for each critical structure, yielding 600 organ sparing plans per patient. To further sample the tradeoff space, another 100 plans of the high, medium, low, and extra low weights category were generated, as well as 200 plans of the controlled weights category. In the high, medium, low, extra low, and controlled weights category, the PTV had a 0.05 probability of being assigned $\operatorname{rand}(0,1)$ instead of 1. The bounds for the controlled weights were selected through trial-and-error such that the final plan generated was likely to fall within clinically relevant bounds, although it is not necessarily acceptable by a physician. In total 1200 plans were generated per patient. With 70 patients, the total number of plans generated was 84,000 plans.

Regarding time for data generation, for each patient, on average it takes 32 minutes to use the Eclipse AAA engine to compute beamlet-based dose influence matrices for a 5 beam IMRT plan. Using our GPU-accelerated optimization algorithm, it takes roughly 2 seconds to generate 1 Pareto optimal plan. These exclude the time it takes to identify and gather the original patient data, as well as preprocessing steps required for converting the data into python-friendly formats. While the optimization of the Pareto optimal plans is fast, the bottleneck is the dose influence matrix calculation for each patient. This is an additional motivation for using neural networks, which do not require dose influence matrices for predicting Pareto optimal dose distributions. This has been shown to yield substantial time savings in generating Pareto optimal plans⁵⁶.

II.2. Loss Functions

In this study, 3 loss functions—mean squared error (MSE) loss, dose volume histogram (DVH) loss, and adversarial (ADV) loss—were used to train and compare 4 instances of the

neural network model. The first model used only the voxel-wise MSE loss. The second model's loss function used the MSE loss in conjunction with the ADV loss. The third model used the MSE loss in conjunction with the DVH loss. The fourth and last model's loss function combined MSE, DVH, and ADV losses all together. Respectively, the study will denote each variation as MSE, MSE+ADV, MSE+DVH, and MSE+DVH+ADV. The following section will describe the ADV and DVH losses in detail.

II.1.1. Adversarial Loss—Our adversarial-style training utilizes a framework similar to that of generative adversarial networks (GAN)⁵⁷, with respect to having another model acting as a discriminator to guide the main network to produce a dose distribution close to the real data. The major benefit to this approach is that the discriminator is calculating its own features and metrics to distinguish the ground truth data and predicted data. Effectively, this is allowing the discriminator to learn its own domain knowledge, and then provide feedback to update the main model. For this study, we utilized the Least Squares GAN (LSGAN)⁶⁵ formulation:

$$\underset{\theta_{N_D}}{\operatorname{argmin}} L_{ADV_D} = \frac{1}{2} \|N_D(y_{true}) - b\|_2^2 + \frac{1}{2} \|N_D(N_G(x)) - a\|_2^2 \quad (3)$$

$$\underset{\theta_{N_G}}{\operatorname{argmin}} L_{ADV_G} = \frac{1}{2} \|N_D(N_G(x)) - c\|_2^2 \quad (4)$$

where θ_{N_D} and θ_{N_G} are the trainable weights parameterizing the discriminator, N_D , and generator, N_G , respectively. L_{ADV_D} and L_{ADV_G} are the loss functions to be minimized with respect to θ_{N_D} and θ_{N_G} . The variable x represents the input into the generator, N_G , which, because of L_{ADV_G} , tries to create data that has a similar distribution to that of y_{true} , our target data. The discriminator tries to distinguish y_{true} from the data created from the generator. As per suggestion by the LSGAN publication⁶⁵, to minimize the Pearson χ^2 divergence, we set $a = -1$, $b = 1$, and $c = 0$.

II.1.2. Dose Volume Histogram Loss—The DVH is one of the most commonly used metrics in radiation oncology for evaluating the quality of a plan, so it is natural to assume that including this metric as part of the loss would be beneficial. However, the calculation of the DVH involves non-differentiable operations, which means any loss based on it cannot provide a gradient to update the neural network. We propose a differential approximation of the DVH, which we define as \overline{DVH} . Given a binary segmentation mask, M_s , for the s^{th} structure, and a volumetric dose distribution, D , the volume at or above a given dose threshold value, d_b can be approximated as

$$v_{s, d_t}(D, M_s) = \frac{\sum_{i, j, k} \operatorname{Sigmoid}\left(\frac{m}{\beta_t}(D(i, j, k) - d_t)\right) M_s(i, j, k)}{\sum_{i, j, k} M_s(i, j, k)} \quad (5)$$

where $Sigmoid(x) = \frac{1}{1 + e^{-x}}$ is the sigmoid function, m controls the steepness of the curve, β_t is the bin width of the histogram, and i, j , and k are the voxel indices for the 3D arrays. The t is an index for the dose threshold values and bin widths. The total number of thresholds is defined as n_t , which is also equivalent to the number of bins in \overline{DVH} . We also constrain the dose to be monotonically increasing with increasing index, $d_1 \leq d_2 \leq \dots \leq d_{n_t}$. Based on this, the \overline{DVH}_s for any structure, s , can then be defined as:

$$\overline{DVH}_s(D, M_s) = (v_{s, d_1}, v_{s, d_2}, \dots, v_{s, d_{n_t}}) \quad (6)$$

The bin centers and widths are then defined as

$$center_{bin} = \left(\frac{d_1 + d_2}{2}, \frac{d_2 + d_3}{2}, \dots, \frac{d_{n_t} + d_{n_t+1}}{2} \right) \quad (7)$$

$$width_{bin} = (\beta_1, \beta_2, \dots, \beta_{n_t}) = (d_2 - d_1, d_3 - d_2, \dots, d_{n_t+1} - d_{n_t}) \quad (8)$$

To illustrate Equation 3, we calculated the DVH and the approximated DVH, of varying steepness values of $m = \{1, 2, 4, 8\}$, of a PTV and OAR or an example prostate patient. As demonstrated by Figure 1, when the steepness of the curve $m \rightarrow \infty$, then $\overline{DVH} \rightarrow DVH$.

Because \overline{DVH} is computed using sigmoid, the gradient,

$$\frac{\partial \overline{DVH}(D, M)}{\partial D} = \left(\frac{\partial v_{s, d_0}}{\partial D}, \frac{\partial v_{s, d_1}}{\partial D}, \dots, \frac{\partial v_{s, d_n}}{\partial D} \right),$$

can be computed, allowing for a loss function utilizing \overline{DVH} to be used to update the neural network weights. We can then define a mean squared loss of the DVH as

$$L_{DVH}(D_{true}, D_{pred}, M) = \frac{1}{n_s} \frac{1}{n_t} \sum_s \left\| \overline{DVH}_s(D_{true}, M_s) - \overline{DVH}_s(D_{pred}, M_s) \right\|_2^2 \quad (9)$$

where D_{true} and D_{pred} are the ground truth and predicted doses, respectively. While a gradient of L_{DVH} exists, it is possible that the gradient space is ill-behaved and would be not suitable for use. We studied the properties of this approximation using a simple toy example. Letting $D = (1, 2)$, The exact DVH and approximate DVH with varying values of $m = \{1, 2, 4, 8\}$ can be calculated, shown in Figure 2.

It can be observed that in this toy example in Figure 2 the approximation is smoother and has larger error with smaller m , which agrees with Figure 1. To investigate the gradient properties of the loss using the approximate DVH, we calculated

$$\|DVH([1, 2], M) - \overline{DVH}([i, j], M)\|_2^2 \quad \forall i, j \in (0, 3) \text{ for } M = [1, 1].$$

Figure 3 shows the squared error value of the difference between DVH for the data (1,2) and the \overline{DVH} for the data $(i, j) \forall i, j \in (0, 3)$. There are multiple local minima. For our case it is when $(i, j) = (1, 2)$ or $(2, 1)$, since either will produce the same DVH. For higher m , the local

minimas become more defined, with steeper gradients surrounding them, an undesirable quality for optimization and learning problems. While a lower steepness, m , may not approximate the DVH as well, the loss function involving the \overline{DVH} maintains the same local minima, and provides a smoother, and more well-behaved gradient than its sharper counterparts. For the remainder of this study, we chose to use \overline{DVH} with $m = 1$.

II.3. Model Architecture

In this study the dose prediction model that was utilized was a U-net style architecture⁶⁶, and the discriminator model was a classification style architecture^{67,68}.

Specifically, the models were adjusted to match the data shape. The architectures shown in Figure 4 depict the models used in the study. The U-net takes as input a 3 channel tensor that consists of, 1) the PTV mask with the value w_{PTV} assigned as its non-zero value, 2) the OAR masks that include all the OARs respectively assigned their w_{OAR_s} , and 3) the body mask as a binary. The U-net then performs multiple calculations, with maxpooling operations to reduce the resolution for more global feature calculation, and then upsampling operations to eventually restore the image resolution back to the original. Concatenation operations are used to merge the local features calculated in the first half of the U-net with the global features calculated at the bottom and later half of the U-net.

The discriminator architecture is of an image classification style network. The goal of the discriminator is to learn how to distinguish the optimized dose distributions versus the U-net predicted dose distribution. Similar to conditional generative adversarial network framework⁶⁹, the discriminator will additionally take the same input as the U-net. In total, the input data has 4 channels—3 channels of the U-net's input and 1 channel of either the optimized or predicted dose distribution. As shown in Figure 4, the discriminator goes through a process of convolutions and strided convolutions to calculate new feature maps and to reduce the image resolution, respectively. It is important to note that the strided convolution is used to reduce one or more of the image dimensions by half, but differ in which dimensions are being reduced in order to eventually reduce the image to $4 \times 4 \times 4$ pixels. For example, the first strided convolution is applied to the first 2 image dimensions, reducing the image from $92 \times 92 \times 20$ to $46 \times 46 \times 20$, but the last strided convolution is reducing the 3rd image dimension. The specific details can be seen in the image sizes specified in Figure 4.

In addition, Group Normalization⁷⁰ was used in place of Batch Normalization⁷¹, which has been shown to allow for the models to effectively train on small batch sizes. All activations in the hidden layer are rectified linear units (ReLU) activations. Final activations for both the U-net and discriminator are linear activations.

II.4. Training and Evaluation

We first notate the mean squared error loss, dose volume histogram loss, and U-net's adversarial loss as $L_{MSE}(y_t, y_p)$, $L_{DVH}(y_t, y_p, M)$, and $L_{ADV_G}(y_t, y_p, x)$, where x is the input into the U-net model, y_t is the ground truth optimized dose distribution, y_p is the predicted

dose distribution, and M contains the binary segmentation masks. The total objective for training the U-net is then defined as

$$L_{Total} = L_{MSE}(y_t, y_p) + \lambda_{DVH}L_{DVH}(y_t, y_p, M) + \lambda_{ADV}L_{ADV_G}(y_p, x), \quad (10)$$

and the objective for training the discriminator is simply $L_{ADV_D}(y, x)$, where y can either be y_t or y_p for a given training sample. For each study—MSE, MSE+ADV, MSE+DVH, and MSE+DVH+ADV—the weightings, λ_{DVH} and λ_{ADV} , used are shown in Table 2. These were chosen by evaluating the order of magnitude of the values that each loss function exhibits for a converged model. From previous dose prediction studies and results^{48,55}, we can estimate that the $L_{MSE} \sim 10^{-4}$ and $L_{DVH} \sim 10^{-3}$ for a converged model. Since we are using least squares GAN framework, we estimate the loss L_{ADV_G} ranges from 10^{-1} to 10^0 . Our choice of λ_{DVH} and λ_{ADV} , shown in Table 2, is to have each component of the loss to be within a similar order of magnitude for when the model is converged.

We divided the 70 prostate patients into 54 training, 6 validation, and 10 testing patients, yielding 64,800 training, 7,200 validation, and 12,000 testing plans. For the training that involved adversarial loss, the U-net and discriminator would alternate every 100 iterations, to allow for some stability in the training and loss. The discriminator was trained to take as input the same inputs as the u-net as well as a dose distribution, either from the real training data or from the U-net’s prediction. With a 0.5 probability, the discriminator would receive either real training data or data predicted from the U-net. Each U-net model was trained for a total of 100,000 iterations, using a batch size of 2. All models used the Adam optimizer, with a learning rate of 1×10^{-3} . All training was performed on an NVIDIA 1080 Ti GPU with 11 GB RAM. After training, the model with the lowest total validation loss was used to assess the test data.

All dose statistics will also be reported relative to the relative prescription dose (i.e. the errors are reported as a percent of the prescription dose). As clinical evaluation criteria PTV coverage (D_{98} , D_{99}), PTV max dose, homogeneity $\left(\frac{D_2 - D_{98}}{D_{50}}\right)$, van’t Riet conformation number⁷² $\left(\frac{(V_{PTV} \cap V_{100\%Iso})^2}{V_{PTV} \times V_{100\%Iso}}\right)$, the dose spillage $R_{50} \left(\frac{V_{50\%Iso}}{V_{PTV}}\right)$, and the structure max and mean doses (D_{max} and D_{mean}) were evaluated. D_{max} is defined as the dose to 2% of the structure volume, as recommended by the ICRU report no 83⁷³.

III. Results

For each model, training for 100,000 iterations took 1.5 days (MSE), 3.5 days (MSE+ADV), 2.3 days (MSE+DVH), 3.8 days (MSE+DVH+ADV). After training, the prediction time of each U-net is the same at 0.052 seconds, since all 4 U-net models in the study are identical in architecture.

Figure 5 shows the validation losses for each model. The top row shows the total validation loss, while the bottom row shows just the mean squared error loss. Overall the loss curve had flattened by the end of the 100,000 iterations, indicating that each model converged. The

final instances of the models chosen for evaluation were the models that performed the best on their respective total validation loss. Each model has achieved similar MSE losses, with our chosen models having their MSE validation losses at 2.46×10^{-4} (MSE), 2.40×10^{-4} (MSE+ADV), 2.26×10^{-4} (MSE+DVH), 2.5×10^{-4} (MSE+DVH+ADV).

Figure 6 shows a comparison of the predictions between each of the 4 models in the study on 1 example patient and Pareto optimal plan. The “optimized” dose is the ground truth Pareto optimal dose that was obtained by solving the optimization problem outlined in Equations 1 and 2. The avoidance map is a sum of the critical structures, including a ring and skin tuning structure, with their assigned tradeoff weights. The 4 models each take in the top row of Figure 6 as their input, and then predict what the Pareto optimal plan should look like. Visually, with the same input, each model produces a strikingly similar dose distribution to the optimized case. The MSE model visually did slightly better in sparing the dose in the normal tissue region posterior of the rectum.

The DVHs of the dose predictions are more revealing to the dose prediction errors in a clinically relevant manner, shown in Figure 7. The two models involving DVH loss (red and green) have less error in predicting the dose in the PTV, Body, and Bladder, with respect to its DVH, and visually similar predictions for the remaining OARs. Overall, including the domain specific DVH based loss has overall improved the model’s dose prediction in regards to the structure’s DVH.

Figure 8 shows the errors for several clinical metrics calculated from the predicted dose distributions, as compared to that of the optimized dose. The MSE model had the largest prediction error of 0.134 (conformation), 0.041 (homogeneity), 0.520 (R50), 3.91% (D95), 4.33% (D98), 4.60% (D99). The additional adversarial and DVH losses further improved the prediction error, with the MSE+DVH+ADV model having the lowest prediction error of 0.038 (conformation), 0.026 (homogeneity), 0.298 (R50), 1.65% (D95), 2.14% (D98), 2.43% (D99). The other two models, MSE+ADV and MSE+DVH, had errors that were between the other two, with the MSE+DVH model’s having less error than MSE+ADV. In terms of these dosimetric criteria, including the DVH loss has the best performance, even more than just including adversarial loss. Figure 8 is the prime example of how expert human domain knowledge can be used to greatly improve the model performance towards the domain relevant criteria. Since conformation, homogeneity, dose spillage, and dose coverage all use particular DVH values in its calculation, they all improved greatly from usage of L_{DVH} . Adversarial training for automatic learning of domain knowledge can further augment the performance by further capturing details that we did not specifically quantify in the loss. It does performs worse on its own, compared to the domain knowledge loss, due to the fact that the adversarial learning model does not truly know what is important to the user. However, when combined with our own domain knowledge, the performance can be maximized.

For both the mean dose PTV error and the max dose PTV, Body, Bladder and rectum error, the same improving trend can be observed in the order of the MSE model, MSE+ADV model, MSE+DVH model, and MSE+DVH+ADV model shown in Figures 9 and 10. However, there is not a clear trend in the mean dose for the OARs, due to the fact that MSE

is already designed for reducing average errors, making it competitive for the mean dose error performance. There also lacks a trend for the max dose point for the femoral heads, which are further away from the PTV and are in the lower dose region that has higher variability in the dose distribution. All predictions have very low average mean and max dose errors of less than 2.8% and 4.2%, respectively.

Due to the large number of test plans we have found that, for conformity, homogeneity, and dose coverage (D95, D98, and D99), the MSE+DVH+ADV model had a statistically significant lower error than the other predictive models, with the largest p-value=0.007. For mean and max doses to the OARs, only 2 comparisons against the MSE+DVH+ADV model were found to be not significantly different, which was the mean dose error to the bladder against the MSE model (p-value=0.894), and the max dose error to the left femoral head against the MSE+DVH model (p-value=0.409). All other mean and max dose comparisons against the MSE+DVH+ADV had found statistically significant differences, with the largest p-value=0.042.

IV. Discussion

To our knowledge, this is the first usage of a domain specific loss function, the DVH loss, for volumetric dose distribution prediction in radiation therapy. We compare the performance of deep neural networks trained using various loss combinations, including MSE loss, MSE+ADV loss, MSE+DVH loss, and MSE+DVH+ADV. Inclusion of the DVH loss had improved the model's prediction in almost every aspect, except for mean dose to the OARs and the max dose the femurs. The DVH loss does not directly represent mean or max dose, and thus is not directly minimizing these aspects. In addition, MSE loss is inherently designed to minimize average error, thus it is not surprising that MSE loss alone is competitive for driving the organ mean dose error down, since the additional DVH and ADV losses may have the model focus on aspects other than mean error. Regardless of the model, all predictions have an average mean and max dose error less than 2.8% and 4.2%, respectively, of the prescription dose for every structure of interest.

To be specific, the performance of our model improved with respect to our domain relevant metrics, because our domain knowledge-based losses are designed to reduce the error in very specific areas of the model's prediction, while focusing less on minimizing error for irrelevant metrics. In addition, having multiple losses can have a regularization effect, which can improve model generalizability and overall performance on unseen data. However, this does not guarantee that the model's performance will improve in all aspects, as indicated by the competitive organ mean dose error with the MSE model, since mean dose error is already heavily related to mean squared error.

Overall, the MSE+DVH+ADV performed the best in most of the categories, particularly the conformity, heterogeneity, high dose spillage, and planning target volume (PTV) dose coverage. This illustrates the importance of both human and learned domain knowledge. Expert human domain specific knowledge can greatly improve the performance of the model, tailoring the prediction towards domain relevant aspects. However, by having to explicitly formulate this domain knowledge into an equation, it is difficult to capture the

nuanced aspects of the problem. Using adversarial learning can then be used to further augment the model’s performance, since the discriminator network can pick out the subtle aspects that the domain specific formulation may have missed.

Due to the non-convexity of both the DVH and ADV losses, as well as the inherent non-convex nature of neural networks, the MSE loss was utilized in every variant of the study, acting as the initial driving force and guide for the model to reasonably converge before the DVH and/or ADV losses began to take effect on the model’s prediction. MSE loss still has many desirable properties from an optimization perspective. It is convex and has an extremely well behaved gradient. In addition the properties of the squared ℓ_2 -norm, where $\ell_p(x) = \sqrt[p]{\sum_i |x_i|^p}$, is one of the most understood and utilized functions in optimization⁷⁵. It is not surprising that the previous studies achieved the state-of-the-art performance for dose prediction utilizing only MSE loss.

The final errors were assessed with 12,000 plans from 10 test patients, with varying tradeoff combinations. The total large number of plans with the randomization scheme given in Table 1 gives us confidence that the entire tradeoff space has been reasonably sampled. Theoretically, the space can be fully sampled from just using the “high weights” randomization scheme outlined in Table 1. However, it would take far more sampling points, since most of the plans deriving from this scheme would not be considered even close to clinically acceptable. By including weight randomizations in the “Single organ spare” category, we are able to create a particular single-organ-sparing plan, with varying degrees of sparing through randomization. Furthermore, the remaining randomization schemes allow for us to create general plans that are closer to clinical relevance than the “high weights” scheme. These effectively allow for us to smooth the tradeoff surface between the single-organ-sparing plans, and easily interpolate in between, making the total set of 12,000 plans representative of the different obtainable dose distribution.

The low prediction errors on the test patients signify that the model is capable of reliably generating Pareto optimal dose distributions with high accuracy. In addition, the raw prediction time of the neural network, including data movement to and from the GPU, is at 0.052 seconds. Realistically, with data loading, prediction, DVH calculation, and displaying the dose wash to a user, it takes approximately 0.6 seconds. This is still fast enough for real time interaction with the model to quickly explore the tradeoff space for a patient. The optimization based approach is much slower, first requiring, on average, 32 minutes for dose influence matrix calculation, and then 2 seconds for the optimization of the each Pareto optimal dose. This allows for us to focus this tool towards empowering physicians. Immediately after segmentation of the cancer patient, the physician can immediately begin to generate a dose distribution with realistic and patient-specific tradeoffs between the PTV and various OARs. Not only does this give the physician a sense of the available and achievable tradeoffs, the resulting dose can then be given to a dosimetrist as a tangible and physician-preferred endpoint, alongside the other typical planning information provided by the physician. Usage of such a model is expected to drastically reduce the treatment planning time by reducing the number of feedback loops between the physician and dosimetrist, as well as how much the dosimetrist has to iterate through tuning

hyperparameters in the fluence map optimization. The clinical relevance regarding the predictive improvement of the model—using MSE+DVH+ADV losses versus using only MSE—on the dosimetric constraints still require assessment through a clinical study to be properly quantified.

The addition of the adversarial loss increases the training time the most for training, since the discriminator has to be trained concurrently. The additional DVH loss does slow down the training as well, but has a much smaller effect than the adversarial loss. While the training times were wildly different, the final trained neural networks all yield the same exact prediction time, due to the fact that they have identical network architectures. The network that took the longest training time, MSE+DVH+ADV, took just under 4 days to train, which is still a very reasonable training time to prepare a model.

When training the multiple models, the weights λ_{DVH} and λ_{ADV} were chosen to be 0.1 and 0.001, respectively, in order to keep the losses at a similar order of magnitude until convergence. In general, this technique of assigning the human and learned domain knowledge weightings can be performed similarly. First the general loss model—for example, our MSE model—can be run first until convergence. The predictions of the general model can be used to assess its current loss value and the human domain knowledge metric to obtain orders of magnitude in the error. The adversarial model may have different orders of magnitude in its loss depending on the exact loss function used, but this may be estimated by bound evaluation of when the discriminator is able to perfectly distinguish the real data vs the predicted data and when it is unable to. The human and learned domain knowledge weightings can then be assigned by orders of magnitude, but some fine-tuning may be necessary depending on the problem to solve. From the clinical perspective, by setting the weights λ_{DVH} and λ_{ADV} such that the losses operate in a similar order of magnitude, we are setting equal importance to the overall general dose distribution (e.g. MSE), domain relevant metrics (e.g. DVH), and letting the network itself decide what is important (e.g. ADV). However, it may be necessary to fine tune the weightings, putting even more emphasis on a particular aspect for tackling a particular problem.

Since the invention and adoption of intensity modulated radiation therapy, optimization has become the backbone of radiation therapy systems. The user will place their ideal dose constraints in the system, and an inverse optimization process occurs to solve for the best treatment parameters to realize the dose distribution. While today's cost functions used commercially may be different than what is used in this study, the core concepts remain. We have our dose constraints, which is simplified in this study to have the PTV treated to prescription dose, and to minimize the dose to any organs-at risk. We also have structure importance weightings, w_s , where increasing this value for particular structure means to more heavily weight the imposed dose constraint. Typically, this means to further improve the dose coverage or homogeneity for the target, or to further reduce the dose for a particular critical structure. This study can be further extended in a future investigation where the optimization problem is replaced with a more complex formulation.

While this study was primarily focused on the evaluation of the DVH, ADV, and MSE losses, the final trained models do have their limitations. While these models are capable of

generating dose distributions on the Pareto surface, it is currently limited to prostate cancer patients with 7 beam IMRT. In addition, the predicted dose distributions are not guaranteed to be deliverable, hence the current need for heavier dosimetrist involvement in the treatment planning. As a future study, we plan to broaden our Pareto optimal dose prediction work by improving our models to predict on different cancer sites, to handle a different number and orientation of beams for IMRT, and to work on the VMAT modality. In addition we plan to investigate the development of a fully automated treatment planning pipeline, starting with the implementation of a robust dose mimicking optimization engine, as the threshold-driven optimization for reference-based auto-planning (TORA) algorithm²⁵, which can be capable of generating a deliverable plan given a dose distribution or constraints. We expect such a pipeline would radically reduce the entire treatment planning time, especially for simple cases, allowing for the physician and dosimetrist to focus their efforts on more challenging patients.

V. Conclusion

In this study, we proposed a novel domain specific loss function, the dose volume histogram (DVH) loss, and evaluated its efficacy alongside two other losses, mean squared error (MSE) loss and adversarial (ADV) loss. We trained and evaluated four instances of the models with varying loss combinations, which included 1) MSE, 2) MSE+ADV, 3) MSE+DVH, 4) MSE+DVH+ADV. We found that the models that included the domain specific DVH loss outperformed the models without the DVH loss in most of the categories, particularly on the evaluations of conformity, heterogeneity, high dose spillage, and planning target volume (PTV) dose coverage. The MSE+DVH+ADV model performed the best in these categories, illustrating the importance of both human and learned domain knowledge. Expert human domain specific knowledge can be the largest driver in the performance improvement, but it is difficult to capture nuanced aspects of the problem in an explicit formulation. Adversarial learning can be used to further capture these particular subtle attributes as part of the loss. The prediction of Pareto optimal doses can be performed in real-time, allowing for a physician to quickly navigate the tradeoff space for a patient, and produce a dose distribution as a tangible endpoint for the dosimetrist to use for planning. Eventually we plan to develop a fully automated treatment planning system. This is expected to considerably reduce the treatment planning time, while improving the treatment planning quality, allowing for clinicians to focus their efforts on the difficult and demanding cases.

VI. Acknowledgements

This study was supported by the National Institutes of Health (NIH) R01CA237269 and the Cancer Prevention & Research Institute of Texas (CPRIT) IIRA RP150485.

VII. References

1. Brahme A Optimization of stationary and moving beam radiation therapy techniques. *Radiotherapy and Oncology*. 1988;12(2):129–140. [PubMed: 3406458]
2. Bortfeld T, Bürkelbach J, Boesecke R, Schlegel W. Methods of image reconstruction from projections applied to conformation radiotherapy. *Physics in Medicine and Biology*. 1990;35(10):1423. [PubMed: 2243845]

3. Bortfeld TR, Kahler DL, Waldron TJ, Boyer AL. X-ray field compensation with multileaf collimators. *International Journal of Radiation Oncology* Biology* Physics*. 1994;28(3):723–730.
4. Webb S Optimisation of conformal radiotherapy dose distribution by simulated annealing. *Physics in Medicine and Biology*. 1989;34(10):1349. [PubMed: 2682694]
5. Convery D, Rosenbloom M. The generation of intensity-modulated fields for conformal radiotherapy by dynamic collimation. *Physics in Medicine and Biology*. 1992;37(6):1359.
6. Xia P, Verhey LJ. Multileaf collimator leaf sequencing algorithm for intensity modulated beams with multiple static segments. *Medical Physics*. 1998;25(8):1424–1434. [PubMed: 9725129]
7. Keller-Reichenbecher M-A, Bortfeld T, Levegrün S, Stein J, Preiser K, Schlegel W. Intensity modulation with the “step and shoot” technique using a commercial MLC: A planning study. *International Journal of Radiation Oncology* Biology* Physics*. 1999;45(5):1315–1324.
8. Yu CX. Intensity-modulated arc therapy with dynamic multileaf collimation: an alternative to tomotherapy. *Physics in Medicine and Biology*. 1995;40(9):1435. [PubMed: 8532757]
9. Otto K Volumetric modulated arc therapy: IMRT in a single gantry arc. *Medical physics*. 2008;35(1):310–317. [PubMed: 18293586]
10. Palma D, Vollans E, James K, et al. Volumetric Modulated Arc Therapy for Delivery of Prostate Radiotherapy: Comparison With Intensity-Modulated Radiotherapy and Three-Dimensional Conformal Radiotherapy. *International Journal of Radiation Oncology* Biology* Physics*. 2008;72(4):996–1001.
11. Shaffer R, Morris WJ, Moiseenko V, et al. Volumetric Modulated Arc Therapy and Conventional Intensity-modulated Radiotherapy for Simultaneous Maximal Intraprostatic Boost: a Planning Comparison Study [published online ahead of print 06/21/2009]. *Clinical Oncology*. 2009;21(5):401–407. [PubMed: 19268554]
12. Shaffer R, Nichol AM, Vollans E, et al. A Comparison of Volumetric Modulated Arc Therapy and Conventional Intensity-Modulated Radiotherapy for Frontal and Temporal High-Grade Gliomas. *International Journal of Radiation Oncology* Biology* Physics*. 2010;76(4):1177–1184.
13. Xing SMCaXWaCTaMWaL. Aperture modulated arc therapy. *Physics in Medicine & Biology*. 2003;48(10):1333. [PubMed: 12812450]
14. Earl M, Shepard D, Naqvi S, Li X, Yu C. Inverse planning for intensity-modulated arc therapy using direct aperture optimization. *Physics in medicine and biology*. 2003;48(8):1075. [PubMed: 12741503]
15. Cao Daliang and Muhammad KNAaJYaFCaDMS. A generalized inverse planning tool for volumetric-modulated arc therapy. *Physics in Medicine & Biology*. 2009;54(21):6725. [PubMed: 19841516]
16. Penfold S, Zalas R, Casiraghi M, Brooke M, Censor Y, Schulte R. Sparsity constrained split feasibility for dose-volume constraints in inverse planning of intensity-modulated photon or proton therapy. *Physics in Medicine & Biology*. 2017;62(9):3599. [PubMed: 28379849]
17. 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]
18. 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.
19. 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]
20. 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]
21. 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]
22. 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]

23. 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]
24. 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]
25. Long T, Chen M, Jiang SB, Lu W. Threshold-driven optimization for reference-based auto-planning. *Physics in medicine and biology*. 2018.
26. Zarepisheh M, Long T, Li N, et al. A DVH-guided IMRT optimization algorithm for automatic treatment planning and adaptive radiotherapy replanning. *Medical Physics*. 2014;41(6Part1):061711–n/a. [PubMed: 24877806]
27. 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. *Medical physics*. 2011;38(2):719–726. [PubMed: 21452709]
28. Appenzoller LM, Michalski JM, Thorstad WL, Mutic S, Moore KL. Predicting dose-volume histograms for organs-at-risk in IMRT planning. *Medical physics*. 2012;39(12):7446–7461. [PubMed: 23231294]
29. Wu B, Pang D, Lei S, et al. Improved robotic stereotactic body radiation therapy plan quality and planning efficacy for organ-confined prostate cancer utilizing overlap-volume histogram-driven planning methodology. *Radiotherapy and Oncology*. 2014;112(2):221–226. [PubMed: 25108808]
30. Shiraishi S, Tan J, Olsen LA, Moore KL. Knowledge-based prediction of plan quality metrics in intracranial stereotactic radiosurgery. *Medical physics*. 2015;42(2):908–917. [PubMed: 25652503]
31. Moore KL, Brame RS, Low DA, Mutic S. Experience-Based Quality Control of Clinical Intensity-Modulated Radiotherapy Planning. *International Journal of Radiation Oncology*Biography*Physics*. 2011;81(2):545–551.
32. Shiraishi S, Moore KL. Knowledge-based prediction of three-dimensional dose distributions for external beam radiotherapy. *Medical physics*. 2016;43(1):378–387. [PubMed: 26745931]
33. Wu B, Ricchetti F, Sanguineti G, et al. Patient geometry-driven information retrieval for IMRT treatment plan quality control. *Medical Physics*. 2009;36(12):5497–5505. [PubMed: 20095262]
34. Wu B, Ricchetti F, Sanguineti G, et al. Data-Driven Approach to Generating Achievable Dose–Volume Histogram Objectives in Intensity-Modulated Radiotherapy Planning. *International Journal of Radiation Oncology*Biography*Physics*. 2011;79(4):1241–1247.
35. Wu B, Pang D, Simari P, Taylor R, Sanguineti G, McNutt T. Using overlap volume histogram and IMRT plan data to guide and automate VMAT planning: A head-and-neck case study. *Medical Physics*. 2013;40(2):021714–n/a. [PubMed: 23387737]
36. Tran A, Woods K, Nguyen D, et al. Predicting liver SBRT eligibility and plan quality for VMAT and 4π plans. *Radiation Oncology*. 2017;12(1):70. [PubMed: 28438215]
37. 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. *Medical Physics*. 2012;39(11):6868–6878. [PubMed: 23127079]
38. Lian J, Yuan L, Ge Y, et al. Modeling the dosimetry of organ-at-risk in head and neck IMRT planning: An intertechnique and interinstitutional study. *Medical Physics*. 2013;40(12):121704–n/a. [PubMed: 24320490]
39. Folkerts MM, Gu X, Lu W, Radke RJ, Jiang SB. SU-G-TeP1–09: Modality-Specific Dose Gradient Modeling for Prostate IMRT Using Spherical Distance Maps of PTV and Isodose Contours. *Medical Physics*. 2016;43(6Part26):3653–3654.
40. Good D, Lo J, Lee WR, Wu QJ, Yin F-F, Das SK. A knowledge-based approach to improving and homogenizing intensity modulated radiation therapy planning quality among treatment centers: an example application to prostate cancer planning. *International Journal of Radiation Oncology*Biography* Physics*. 2013;87(1):176–181.
41. Valdes G, Simone II CB, Chen J, et al. Clinical decision support of radiotherapy treatment planning: a data-driven machine learning strategy for patient-specific dosimetric decision making. *Radiotherapy and Oncology*. 2017;125(3):392–397. [PubMed: 29162279]

42. Yang Y, Ford EC, Wu B, et al. An overlap-volume-histogram based method for rectal dose prediction and automated treatment planning in the external beam prostate radiotherapy following hydrogel injection. *Medical physics*. 2013;40(1):011709. [PubMed: 23298079]
43. Folkerts MM, Long T, Radke RJ, et al. Knowledge-Based Automatic Treatment Planning for Prostate IMRT Using 3-Dimensional Dose Prediction and Threshold-Based Optimization. American Association of Physicists in Medicine. 2017.
44. Krizhevsky A, Sutskever I, Hinton GE. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*. 2012:1097–1105.
45. Girshick R, Donahue J, Darrell T, Malik J. Rich feature hierarchies for accurate object detection and semantic segmentation. *Proceedings of the IEEE conference on computer vision and pattern recognition 2014*:580–587.
46. Simonyan K, Zisserman A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:14091556*. 2014.
47. LeCun Y, Boser B, Denker JS, et al. Backpropagation applied to handwritten zip code recognition. *Neural computation*. 1989;1(4):541–551.
48. 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]
49. 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]
50. 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.
51. 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]
52. Babier A, Mahmood R, McNiven AL, Diamant A, Chan TC. Knowledge-based automated planning with three-dimensional generative adversarial networks. *arXiv preprint arXiv:181209309*. 2018.
53. 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.
54. Babier A, Bouillier JJ, McNiven AL, Chan TC. Knowledge-based automated planning for oropharyngeal cancer. *Medical physics*. 2018;45(7):2875–2883. [PubMed: 29679492]
55. 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]
56. Nguyen D, Barkousaraie AS, Shen C, Jia X, Jiang S. Generating Pareto Optimal Dose Distributions for Radiation Therapy Treatment Planning. *Lecture Notes in Compute Science*. 2019;11769:59–67.
57. Goodfellow I, Pouget-Abadie J, Mirza M, et al. Generative adversarial nets. Paper presented at: *Advances in neural information processing systems 2014*.
58. Muralidhar N, Islam MR, Marwah M, Karpatne A, Ramakrishnan N. Incorporating Prior Domain Knowledge into Deep Neural Networks. Paper presented at: *2018 IEEE International Conference on Big Data (Big Data)2018*.
59. 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–n/a. [PubMed: 24387513]
60. Reese AS, Das SK, Curle C, Marks LB. Integral dose conservation in radiotherapy. *Medical physics*. 2009;36(3):734–740. [PubMed: 19378734]
61. Aoyama H, Westerly DC, Mackie TR, et al. Integral radiation dose to normal structures with conformal external beam radiation. *International Journal of Radiation Oncology* Biology* Physics*. 2006;64(3):962–967.
62. D'souza WD, Rosen II. Nontumor integral dose variation in conventional radiotherapy treatment planning. *Medical physics*. 2003;30(8):2065–2071. [PubMed: 12945972]

63. Jahn J Scalarization in multi objective optimization In: Mathematics of multi objective optimization. Springer; 1985:45–88.
64. 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.
65. Mao X, Li Q, Xie H, Lau RY, Wang Z, Paul Smolley S. Least squares generative adversarial networks. Paper presented at: Proceedings of the IEEE International Conference on Computer Vision2017.
66. Ronneberger O, Fischer P, Brox T. U-net: Convolutional networks for biomedical image segmentation. *International Conference on Medical Image Computing and Computer-Assisted Intervention*. 2015:234–241.
67. LeCun Y, Bottou L, Bengio Y, Haffner P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*. 1998;86(11):2278–2324.
68. LeCun Y, Haffner P, Bottou L, Bengio Y. Object recognition with gradient-based learning In: Shape, contour and grouping in computer vision. Springer; 1999:319–345.
69. Mirza M, Osindero S. Conditional generative adversarial nets. *arXiv preprint arXiv:14111784*. 2014.
70. Wu Y, He K. Group normalization. Paper presented at: Proceedings of the European Conference on Computer Vision (ECCV)2018.
71. Ioffe S, Szegedy C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *International Conference on Machine Learning* 2015:448–456.
72. Van't Riet A, Mak AC, Moerland MA, Elders LH, van der Zee W. A conformation number to quantify the degree of conformality in brachytherapy and external beam irradiation: application to the prostate. *International Journal of Radiation Oncology* Biology* Physics*. 1997;37(3):731–736.
73. Grégoire V, Mackie TR. 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.
74. Savitzky A, Golay MJ. Smoothing and differentiation of data by simplified least squares procedures. *Analytical chemistry*. 1964;36(8):1627–1639.
75. Boyd S, Vandenberghe L. *Convex optimization*. Cambridge university press; 2009.




Figure 1:
 DVH and approximated DVH calculations using varying steepness values of $m = \{1,2,4,8\}$ for an example prostate patient.




Figure 2:
DVH and approximated DVH calculations of toy example for $D = (1,2)$.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript




Figure 3:

Objective value map of the loss function $\|DVH([1, 2], M) - \widehat{DVH}([i, j], M)\|_2^2 \forall i, j \in (0, 3)$ for $M = [1, 1]$. All versions with varying values of m exhibit the same minima.


Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript


U-net Architecture



Legend

- ➔ 3 × 3 × 3 Padded Convolution, Activation, Group Norm, Dropout
- ➔ 2 × 2 × 2 Upsample, 3 × 3 × 3 Padded Convolution, Activation, Group Norm, Dropout
- ➔ Concatenate
- ➔ 2 × 2 × 2 Maxpool

Discriminator Architecture



Legend

- ➔ 3 × 3 × 3 Convolution, Activation, Group Norm, Dropout
- ➔ Full Connection, Activation, Group Norm, Dropout
- ➔ Flatten
- ➔ Strided Convolution, Activation, Group Norm, Dropout

Figure 4: Deep learning models used in the study. The same U-net architecture is utilized in each comparison of MSE, MSE+ADV, MSE+DVH, and MSE+DVH+ADV models. The same discriminator architecture is utilized for training the MSE+ADV and MSE+DVH+ADV models. Black numbers to the left of the feature blocks represent the current data shape. Red numbers above the feature blocks represents the number of features.




Figure 5:
 Top row: Total validation loss (all relevant losses and loss weightings for a specific model are summed are displayed). Bottom row: MSE validation loss (only mean squared error is displayed). Left Column: Raw validation losses at each training iteration. Right Column: Smoothed validation loss using Savitzky–Golay filter⁷⁴.

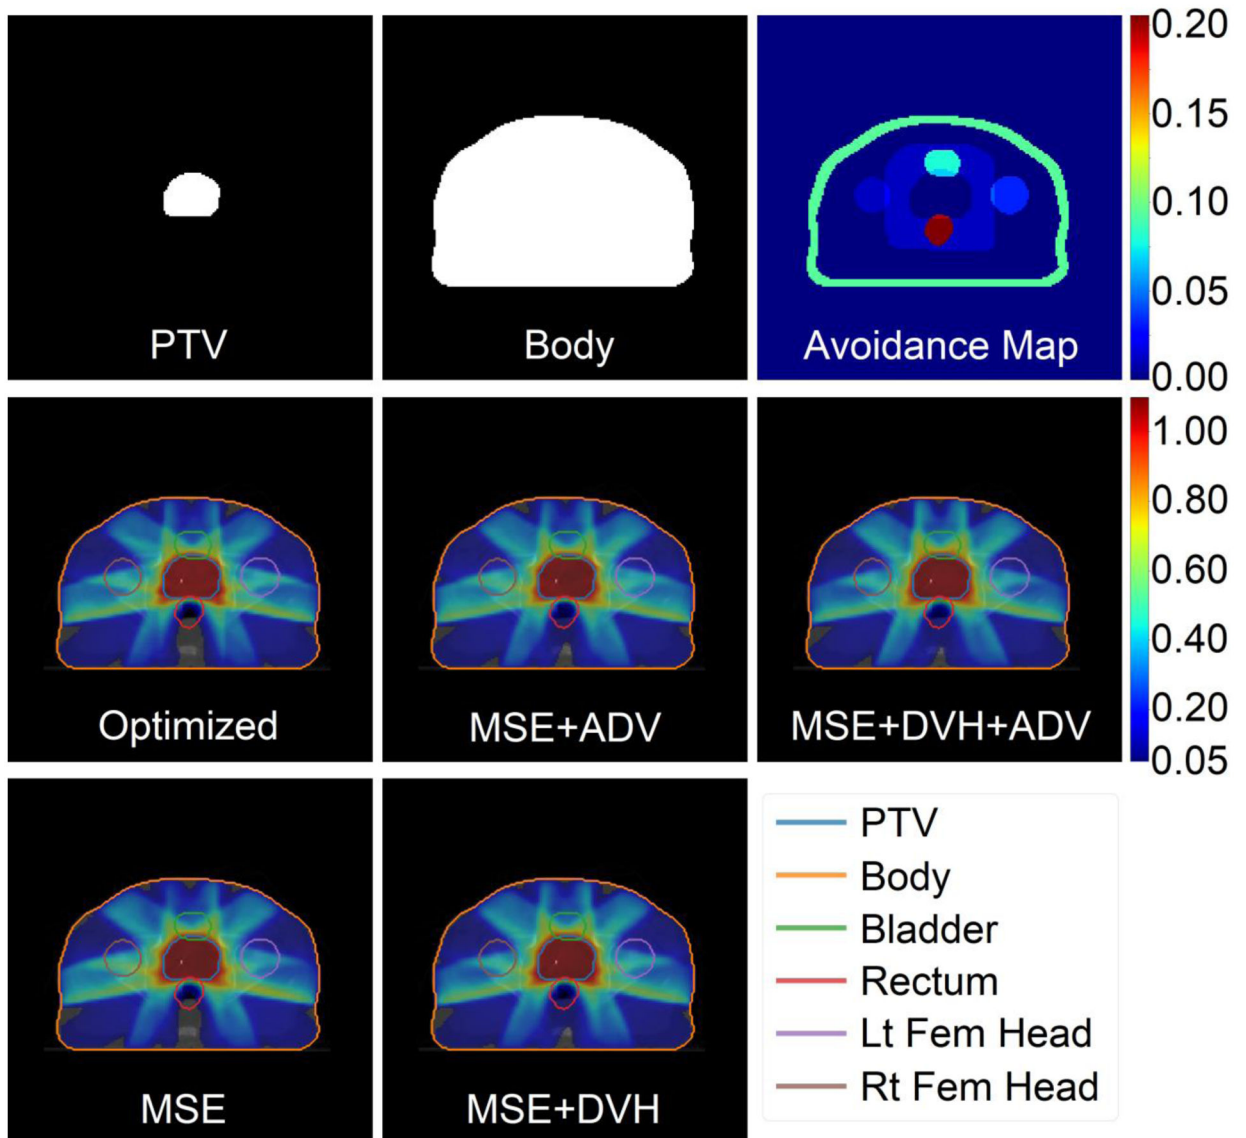


Figure 6: Inputs, optimized dose, and predicted doses for a test patient and a rectum sparing plan. Top row: Inputs of the U-net neural network, which include the PTV assigned to its weight ($w_{PTV}=1$ in this example), a binary mask of the body, and the avoidance map containing the remaining organs-at-risk assigned to their respective tradeoff weight. Bottom two rows: Optimized and predicted dose washes of the Pareto optimal dose. The colorbar shows the doses between 5% and 110% of the prescription dose.

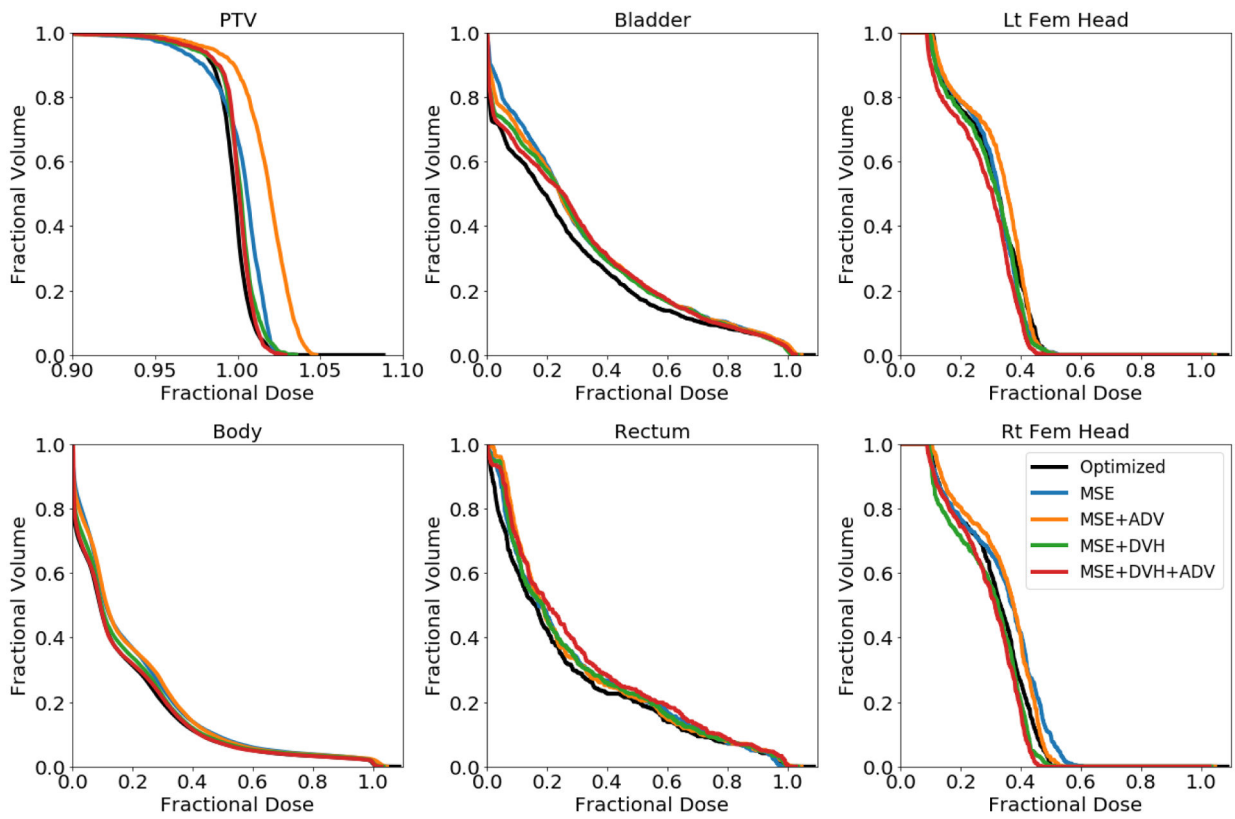


Figure 7: Dose volume histograms (DVH) of optimized dose distribution (black) and predicted dose distributions (various colors) for the same test patient as in Figure 6. Note the x-axis scale for the PTV DVH is different.




Figure 8: Prediction errors for conformation, homogeneity, high dose spillage (R50) and dose coverage on the test data. Error bar represents the 99% confidence interval $\left(\bar{x} \pm 2.576 * \frac{\sigma}{\sqrt{n}}\right)$, taken overall all test patients and plans.




Figure 9: Average error in the mean dose for the PTV and the organs at risk. Error bar represents the 99% confidence interval $\left(\bar{x} \pm 2.576 * \frac{\sigma}{\sqrt{n}}\right)$, taken overall all test patients and plans.




Figure 10: Average error in the max dose for the PTV and the organs at risk. Error bar represents the 99% confidence interval $\left(\bar{x} \pm 2.576 * \frac{\sigma}{\sqrt{n}}\right)$, taken overall all test patients and plans.

Table 1:

Weight assignment categories for the organs at risk. The function $rand(lb, ub)$ represents a uniform random number between a lower bound, lb , and an upper bound, ub . In the high, medium, low, extra low, and controlled weights category, the PTV had a 0.05 probability of being assigned $rand(0,1)$ instead of 1.

Category	Description
Single organ spare	Bladder $w_{bladder} = rand(0,1)$ $w_{OAR bladder} = rand(0,0.1)$
	Rectum $w_{rectum} = rand(0,1)$ $w_{OAR rectum} = rand(0,0.1)$
	Lt Fem Head $w_{lt\ femhead} = rand(0,1)$ $w_{OAR lt\ femhead} = rand(0,0.1)$
	Rt Fem Head $w_{rt\ femhead} = rand(0,1)$ $w_{OAR rt\ femhead} = rand(0,0.1)$
	Shell $w_{shell} = rand(0,1)$ $w_{OAR shell} = rand(0,0.1)$
	Skin $w_{skin} = rand(0,1)$ $w_{OAR skin} = rand(0,0.1)$
High weights	$w_s = rand(0,1) \forall s \in OAR$
Medium weights	$w_s = rand(0,0.5) \forall s \in OAR$
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	$w_{bladder} = rand(0,0.2)$ $w_{rectum} = rand(0,0.2)$ $w_{lt\ femhead} = rand(0,0.1)$ $w_{rt\ femhead} = rand(0,0.1)$ $w_{shell} = rand(0,0.1)$ $w_{skin} = rand(0,0.3)$

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 2:

Choices of λ_{DVH} and λ_{ADV} for the loss function shown in Equation 10.

	λ_{DVH}	λ_{ADV}
MSE	0	0
MSE+ADV	0	0.001
MSE+DVH	0.1	0
MSE+DVH+ADV	0.1	0.001

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript