

# Automated Treatment Planning in Radiation Therapy using Generative Adversarial Networks

**Rafid Mahmood**

*Department of Mechanical and Industrial Engineering  
University of Toronto, Toronto, ON, Canada*

RMAHMOOD@MIE.UTORONTO.CA

**Aaron Babier**

*Department of Mechanical and Industrial Engineering  
University of Toronto, Toronto, ON, Canada*

ABABIER@MIE.UTORONTO.CA

**Andrea McNiven**

*Radiation Medicine Program  
Princess Margaret Cancer Centre, Toronto, ON, Canada*

ANDREA.MCNIVEN@RMP.UHN.CA

**Adam Diamant**

*Schulich School of Business  
York University, Toronto, ON, Canada*

ADIAMANT@SCHULICH.YORKU.CA

**Timothy C. Y. Chan**

*Department of Mechanical and Industrial Engineering  
University of Toronto, Toronto, ON, Canada*

TCYCHAN@MIE.UTORONTO.CA

## Abstract

Knowledge-based planning (KBP) is an automated approach to radiation therapy treatment planning that involves predicting desirable treatment plans before they are then corrected to deliverable ones. We propose a generative adversarial network (GAN) approach for predicting desirable 3D dose distributions that eschews the previous paradigms of site-specific feature engineering and predicting low-dimensional representations of the plan. Experiments on a dataset of oropharyngeal cancer patients show that our approach significantly outperforms previous methods on several clinical satisfaction criteria and similarity metrics.

## 1. Introduction

Radiation therapy (RT) is one of the primary methods for treating cancer and is recommended for over 50% of all cancer patients (Delaney et al., 2005). In RT, a linear accelerator (linac) outputs high-energy x-ray beams from multiple angles around a patient to deliver a prescribed dose of radiation to a tumor while minimizing dose to the healthy tissue. An RT treatment plan is the result of a complex design process involving multiple medical professionals and several software systems. This includes specialized optimization software that determines the beam characteristics (e.g., aperture shapes for each beam angle, dose delivered from each aperture) required to deliver the final dose distribution. The optimization model takes as input a set of computed tomography (CT) images of the patient, various dosimetric objectives and constraints, and other parameters that guide the optimization process. The model outputs a treatment plan that is subsequently evaluated by an oncol-

ogist. The oncologist usually proposes modifications to the plan, which then requires the treatment planner to re-solve the optimization model using updated parameters. The total process is labor intensive, time-consuming, and costly, as the back-and-forth between the planner and oncologist is often repeated multiple times until the plan is finally approved.

The significant manual effort associated with the current treatment planning paradigm, along with the fact that RT plans are generally quite similar for patients with similar geometries, has motivated researchers to investigate how automation can be used in the planning process (Sharpe et al., 2014). A key enabler of automation is known as knowledge-based planning (KBP), which leverages historically delivered treatments to generate new plans for similar patients. Figure 1 depicts the two main components of a KBP-driven automated planning system: (i) a machine learning model that uses CT-derived patient geometric features to predict a clinically acceptable three-dimensional dose distribution (Appenzoller et al., 2012; Yang et al., 2013; Shiraishi et al., 2015; Younge et al., 2018); and (ii) an optimization model that converts the prediction into a “deliverable” plan (McIntosh and Purdie, 2017; Wu et al., 2017; Babier et al., 2018b). The second step is needed to ensure the treatment plan produced by the machine learning model satisfies the physical delivery constraints imposed by the linac.

A major drawback of most existing KBP prediction methods is their reliance on low-dimensional hand-tailored features derived from patient geometry to predict new dose distributions. In contrast, we propose a new paradigm for generating KBP predictions that automatically learns to predict a 3D dose distribution directly from a CT image. More specifically, we recast the dose prediction problem as an image colorization problem, which we solve using a generative adversarial network (GAN) (Goodfellow et al., 2014). GANs, which have produced impressive results in other image colorization applications (Isola et al., 2017; Zhu et al., 2017), involve a pair of neural networks: a generator that performs a task and a discriminator that evaluates how well the task is performed. In our application, the generator serves as a treatment planner that designs a treatment, while the discriminator plays the role of the oncologist who critiques the generated dose distribution by comparing it to the real treatment plan. Both neural networks train simultaneously on historical data, effectively replicating and aggregating the combined knowledge gained during the iterative manual process used to design clinically acceptable treatments.

In this paper, we develop a novel automated treatment planning pipeline for oropharyngeal cancer that uses a GAN to predict 3D dose distributions. In contrast to previous machine learning methods, our approach does not require the pre-specification of an extensive set of feature variables for prediction. Instead, our model learns what features are important to produce clinically acceptable treatment plans. We apply our KBP methodology to a dataset consisting of 26,279 CT images from 217 patients with oropharyngeal cancer that have undergone radiation therapy. Approximately 60% of these images are used to train the GAN, which is used to predict high quality dose distributions for the remaining out-of-sample patients. These predictions are used as input into an optimization model to produce deliverable plans. We compare our approach to several other techniques, including three feature-based machine learning models and a standard convolutional neural network (CNN). We demonstrate that our approach outperforms all other models in achieving several clinically relevant criteria and in matching the clinical (benchmark) plans.




Figure 1: Overview of KBP-driven automated treatment planning pipeline.

**Technical Significance** We demonstrate the first use of GANs for generating radiation treatment plans in cancer. We recast KBP prediction as an image colorization problem for which GANs are known to perform well. Moreover, we provide the first full pipeline comparison between different KBP prediction methods by optimizing the predicted dose distribution and comparing the final result to deliverable plans. We find that, in this setting, our GAN approach outperforms all other methods, including the latest in machine learning-based KBP approaches, in meeting clinical criteria.

**Clinical Relevance** Oropharyngeal cancer is one of the most difficult cancers to plan a treatment for, and as a result, generating deliverable treatment plans is particularly time consuming (Das et al., 2009). Our GAN approach automates the planning approach producing, on average, plans that are superior to clinical ones in several key metrics. Our site-independent method suggests similar performance for simpler sites, such as prostate and stomach cancers, while showing that high-quality oropharynx treatment plans can be automatically generated.

## 2. Related work

### 2.1. Knowledge-based planning

Many different approaches have been tested for the machine learning component of a KBP-driven automated planning pipeline (cf. Figure 1). Query-based methods identify previously treated patients who are sufficiently similar to the new patient, and use the historically achieved dose metrics as predictions for the new patient (Wu et al., 2009, 2011). Another common approach uses principal component analysis (PCA), in conjunction with linear regression, to predict dose metrics for new patients (Zhu et al., 2011; Yuan et al., 2012). However, these well-established techniques only predict two-dimensional dose metrics. Recent research has shown that 3D dose distribution predictions can also be generated using random forest or neural network-based models (Shiraishi and Moore, 2016; McIntosh et al., 2017; Nguyen et al., 2017). Nevertheless, for many of these approaches to work effectively, significant effort must be spent in feature engineering, i.e., introducing features specific to the cancer site. Furthermore, some of these approaches compare the predicted dose distributions, rather than deliverable plans post-optimization, to the clinical plans.

For the optimization phase of KBP, there are two main approaches for turning predictions into treatments: dose mimicking (Pettersson et al., 2016) and inverse optimization (Chan et al., 2014). The dose mimicking model minimizes the  $L_2$  loss between the predicted dose distribution and one that satisfies all physical constraints. Alternatively, inverse optimization (IO) is a methodology that estimates parameters of an optimization problem from its observed solutions (Ahuja and Orlin, 2001). In the RT context, IO finds parameters, e.g., objective function weights, that allow a deliverable treatment plan to re-create the predicted dose distribution as closely as possible (Chan et al., 2014). A key

advantage of inverse optimization is that it better replicates the trade-offs implicit in clinical treatment plans (Chan and Lee, 2018).

## 2.2. Generative adversarial networks

GANs are a well-studied class of deep learning algorithms used in *generative* modeling, i.e., in the creation of new data (Goodfellow et al., 2014). Although initially used to artificially generate 2D images, and later 3D models (Wu et al., 2016), their success has garnered increasing interest for healthcare applications. GANs have been used for medical drug discovery (Kadurin et al., 2017), generating artificial patient records (Choi et al., 2017; Esteban et al., 2017), the detection of brain lesions (Alex et al., 2017), and image augmentation for improved liver lesion classification (Frid-Adar et al., 2018).

A GAN consists of two neural networks, a generator and a discriminator, working in tandem. The generator  $G(\cdot)$  takes an initial random input  $\mathbf{z} \sim p_{\mathbf{z}}$  and attempts to generate an artificial data sample  $\mathbf{x} = G(\mathbf{z})$  (i.e., the 3D dose distribution). The discriminator  $D(\cdot)$  is a classifier that takes generated and real data samples, and tries to identify which is which, i.e.,  $D(\mathbf{x}) \in [0, 1]$  where 1 suggests the generated sample is satisfactory. The interaction between the networks can be formalized mathematically as a minimax game. If  $\mathbf{x} \sim p_{\text{data}}$  is the probability distribution over the real data samples, then the game is defined as

$$\min_G \max_D \left\{ V(G, D) = \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}} [\log(1 - D(G(\mathbf{z})))] \right\}.$$

GANs have been proven effective in style transfer problems, where the generator input  $\mathbf{z}$  is a data sample corresponding to one style (or characteristic) and the output  $\mathbf{x}$  is a mapping to a different style (Isola et al., 2017; Zhu et al., 2017). For example, style transfer can be used to transform grayscale images to colored photos (Sangkloy et al., 2017), in facial recognition for surveillance-based law enforcement (Wang et al., 2017), and in 3D reconstruction of damaged artifacts (Hermezo and Sipiran, 2017). Here, the generator  $G(\mathbf{z})$  learns the mapping between styles that generates samples resembling the ground truth. Since key structures in the output may be entangled with noise from the generator, the desired output is often achieved by modifying the original minimax game with a penalty term on large deviations between the real and generated samples:

$$\min_G \max_D \left\{ V(G, D) + \lambda \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}, \mathbf{z} \sim p_{\mathbf{z}}} [\|\mathbf{x} - G(\mathbf{z})\|_1] \right\}, \quad (1)$$

where  $\lambda$  is a regularizer that balances the trade-off between learning style and the real data.

## 3. Methods

We used contoured CT images and clinically acceptable dose distributions from the treatment plans of past oropharyngeal cancer patients to train a style transfer GAN. We then passed out-of-sample predicted dose distributions through an IO pipeline (Babier et al., 2018b) to generate the final treatment plans. For baseline comparisons, we also implemented several methods from the literature using the complete pipeline. Figure 2 shows a high-level overview of this automated planning pipeline.

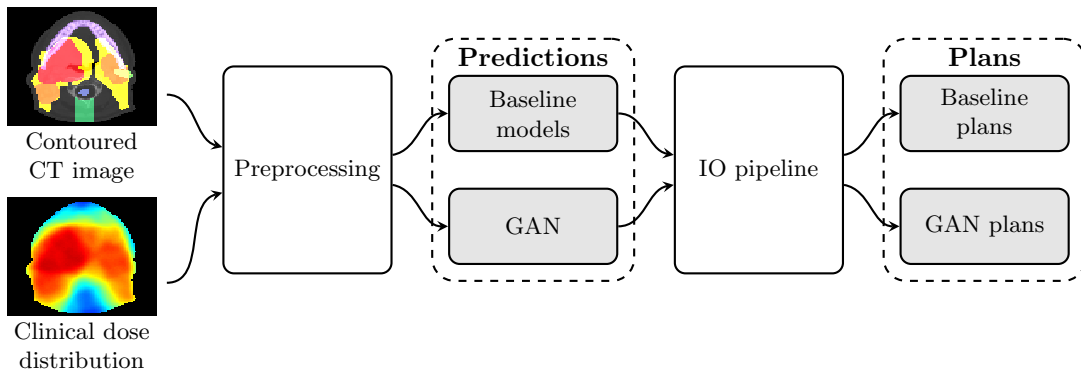


Figure 2: An schematic of our KBP-based automated planning pipeline.

### 3.1. Data

We obtained treatment plans from 217 oropharyngeal cancer patients treated at a single institution with 6 MV, step-and-shoot, intensity-modulated radiation therapy machine. All plans were for a prescription of 70 Gy, 63 Gy, and 56 Gy in 35 fractions to the gross disease, intermediate risk, and elective target volumes, respectively.

For each patient, we identified a set of targets and healthy organs-at-risk (OARs). Targets were denoted as planning target volumes (PTVs) along with the oncologist-prescribed dose (e.g., PTV70 corresponds the target with the highest dose prescription). OARs included the brainstem, spinal cord, right and left parotids, larynx, esophagus, and mandible. Every voxel (a 3D pixel of size  $4\text{ mm} \times 4\text{ mm} \times 2\text{ mm}$ ) of a CT image was classified by their clinically drawn contours. All voxels were assigned a structure-specific color, and in cases where the voxel was classified as both target and OAR, we reverted to target. All unclassified tissue was left as the original CT image grayscale.

### 3.2. GAN model

We first divided each 3D CT image into 2D slices of  $128 \times 128$  pixels. The generator used a single CT image slice to predict the dose distribution along that same plane without considering the vertical relationship between different slices. This process was repeated for every slice until a full 3D dose distribution was produced. Our training set consisted of all 2D slices from the 3D CT images for 130 patients, totaling 15,657 images. The CT images from the remaining 87 patients were used for out-of-sample evaluation.

Our GAN learning model was built on the `pix2pix` style transfer architecture of [Isola et al. \(2017\)](#). We used a U-net generator that passed a 2D contoured CT image slice through consecutive convolution layers, a bottleneck layer, and then through several deconvolution layers. The U-net also employed skip connections, i.e., the output of each convolution layer was concatenated to the input of a corresponding deconvolution layer. This allowed the generator to easily pass “high dimensional” information (e.g., structural outlines) between the inputted CT image slice and the outputted dose slice. The discriminator passed a 2D slice of the dose distribution along several consecutive convolution layers, outputting a single scalar value. In the training phase, the discriminator received one real and one generated dose distribution before backpropagation. We disconnected the discriminator after training,

at which point the generator only received a contoured CT slice. We refer the reader to Appendix A for additional details regarding the network architectures.

We used the loss function given by (1) with  $\lambda = 90$ , and trained using Adam (Kingma and Ba, 2014), with learning rate 0.0002 and  $\beta_1 = 0.5$  and  $\beta_2 = 0.999$  for 25 epochs. We used the default Adam settings from Isola et al. (2017), as they were proven to be good for a variety of different style transfer problems. While we swept through various values for  $\lambda$  and the number of epochs, we found these default settings to be sufficient, with minimal subsequent improvement. We found it useful to stop training when the loss functions were roughly equal; if the loss from the  $l_1$  penalty fell too low, the GAN began to simply memorize the dataset. The code for all experiments, along with the parameter settings is provided at <http://github.com/rafidrm/gancer>.

### 3.3. Plan generation

Predicted dose distributions were inputted into an IO pipeline to generate optimized plans. The IO model determined the weights of a parametric “forward” optimization model given a predicted dose distribution. The objective of the forward model was to minimize the sum of 65 objective functions: seven per OAR and three per target. Terms for the OARs included the mean dose, max dose, and the percentile (0.25, 0.50, 0.75, 0.90, and 0.975) above the maximum predicted dose to the OAR. Similarly, terms for the target included the maximum dose, average dose below prescription, and average dose above prescription. The complexity of the KBP-generated treatment plan was constrained to match the clinical treatment (Craft et al., 2007) where complexity represents a (convex) surrogate measure for the physical deliverability of a plan. We note that in reality, there are additional constraints in the IO pipeline that we omit for tractability. Thus, our notion of a deliverable plan does not include all physical constraints. Physical parameters for the optimization model were derived from A Computational Environment for Radiotherapy Research (Deasy et al., 2003). To replicate the clinical plans, all KBP-generated plans were delivered from nine equidistant coplanar beams at angles  $0^\circ$ ,  $40^\circ$ ,  $\dots$ ,  $320^\circ$ . We used Gurobi 7.5 to solve the inverse and forward optimization problems associated with the IO pipeline. Additional details of the IO model can be found in Babier et al. (2018a).

### 3.4. Baseline approaches

We compared our GAN approach to generating predicted dose distributions with several state-of-the-art techniques. We briefly describe the baseline approaches here.

- **Bagging query (BQ):** A look-up method identifies patients with similar geometries who have undergone radiation therapy and outputs their doses as predictions. This approach predicts dose volume histograms (DVHs), i.e., 2D summaries of the 3D dose delivered to specific targets and OARs (e.g., Wu et al. (2009); Babier et al. (2018a)).
- **Generalized PCA (gPCA):** A method combining PCA with linear regression using patient geometry features. Similar to BQ, this method also predicts DVHs (e.g., Yuan et al. (2012); Babier et al. (2018a)).

- **Random forest (RF):** Predicts dose to each voxel (3D dose prediction) using ten customized features based on patient geometry (inspired by [McIntosh et al. \(2017\)](#)). Additional details can be found in Appendix B.
- **U-net (CNN):** Predicts dose to each voxel in 2D slices from a CT image using a U-net convolution neural network architecture (e.g., [Nguyen et al. \(2017\)](#)).

All baseline predictions were fed into the same IO pipeline as the GAN approach to ensure a fair comparison between deliverable plans.

## 4. Results

### 4.1. Sample generated dose distributions

We observed that the style transfer function mapping the CT image to the predicted dose distribution appeared easy to learn. This is because the GAN generated dose distributions had the hallmarks of a deliverable plan, like the sharp dose gradients that are generated by individual beams. However, there were subtle deliverability characteristics that the GAN could not always identify. The optimization step enforced these physical deliverability constraints to correct for these idiosyncracies. This result can be observed in Figure 3, where five sample slices of a clinical, predicted, and optimized plan are presented.

### 4.2. Clinical criteria satisfaction

We measured plan quality by evaluating how frequently they satisfied the standard clinical criteria for oropharyngeal cancer treatment plans; see Table 1. Clinicians commonly use criteria satisfaction as a metric to evaluate plan quality and approve a treatment plan after it satisfies a sufficient number of the criteria. Thus, each criterion (one per OAR and target) was measured on a pass-fail basis depending on whether the mean dose  $\mathcal{D}_{mean}$ , maximum dose  $\mathcal{D}_{max}$ , or the dose to 99% of the volume of that structure  $\mathcal{D}_{99}$ , was above or below a given threshold. To facilitate the comparisons, we scaled the GAN and baseline treatment plans so that their PTV  $\mathcal{D}_{99}$  was equal to the PTV  $\mathcal{D}_{99}$  of the corresponding clinical plan.

Table 2 presents the percentage of the GAN and baseline treatment plans that satisfied the clinical criteria. We note that clinically acceptable plans typically cannot satisfy all criteria simultaneously because of the proximity of the targets to the OARs and the complexity of the head-and-neck site in general. We observed that the BQ and gPCA plans tended to satisfy PTV criteria more frequently, which suggested that they may recommend delivering a higher dose to the target relative to the clinical plan. However, they failed to achieve mean and maximum dose criteria to the OARs (note: there are more than triple the number of OAR criteria as PTV criteria once all plans are normalized to  $\mathcal{D}_{99}$  of the PTV70). On the other hand, the RF plans appeared to satisfy fewer clinical criteria associated with the target as compared to the clinical plans. The CNN plans achieved the closest level of performance to the clinical plans. However, the GAN plans had the best overall performance among all approaches. They offered a balanced trade-off between the OARs and targets, and even outperformed the clinical plans on clinical criteria satisfaction.

The previous results focused on pass-fail performance with respect to the clinical criteria. We also examined the magnitude of passing or failing via head-to-head comparisons of




Figure 3: Sample of slices from a test patient. From top to bottom: contoured CT image (generator input), clinical plan (ground truth), GAN prediction, and GAN plan (post optimization).

the GAN/baseline plans to the clinical plans, and between the GAN and CNN plans (see Figure 4). The x-axis in each figure is the difference in Gray (Gy) between the KBP and the clinical plans (KBP minus clinical) for the criterion on the corresponding y-axis. We found that for each criterion, the majority of GAN plans outperformed their clinical counterparts by several Gy (Figure 4(e)). This is a significant result given that the clinical plans were heavily optimized and delivered to actual patients. The BQ, gPCA, and RF plans displayed substantial variability in performance when compared to the clinical plan. Consistent with Table 2, performance of the CNN plans were closest to the GAN plans although, as shown in Figure 4(f), the GAN plans maintained a small, yet consistent, advantage.

Finally, we compared the KBP plans against the clinical plans using the gamma passing rate (GPR) metric. GPR measures the similarity between two dose distributions on a voxel-by-voxel basis, computing for each voxel, a pass-fail test. We considered the standard choice of GPR, i.e., a 3%/3 mm tolerance (Low et al., 1998), which roughly means a voxel in the evaluated dose distribution (KBP) “passes” if there is at least one voxel in the reference dose distribution (clinical) within 3 mm that receives a dose that is within  $\pm 3\%$  of the reference dose. Table 3 summarizes the average GPR achieved over all KBP-generated plans. A score of 1.0 means that every voxel has passed the criteria; in other words, the two dose distributions were considered identical (within the tolerance). Overall, we observed



Structure	Criteria
Brainstem	$\mathcal{D}_{max} \leq 54$ Gy
Spinal Cord	$\mathcal{D}_{max} \leq 48$ Gy
Right Parotid	$\mathcal{D}_{mean} \leq 26$ Gy
Left Parotid	$\mathcal{D}_{mean} \leq 26$ Gy
Larynx	$\mathcal{D}_{mean} \leq 45$ Gy
Esophagus	$\mathcal{D}_{mean} \leq 45$ Gy
Mandible	$\mathcal{D}_{max} \leq 73.5$ Gy
PTV56	$\mathcal{D}_{99} \geq 53.2$ Gy
PTV63	$\mathcal{D}_{99} \geq 59.9$ Gy
PTV70	$\mathcal{D}_{99} \geq 66.5$ Gy

Table 1: Clinical criteria used to evaluate all plans.  $\mathcal{D}_{mean}$  refers to the mean dose,  $\mathcal{D}_{max}$  the maximum dose, and  $\mathcal{D}_{99}$  dose to 99% of the structure.

	BQ	gPCA	RF	CNN	GAN	Clinical
OAR criteria	61.6%	65.8%	71.5%	72.5%	72.8%	72.0%
PTV criteria	83.5%	85.7%	68.0%	76.3%	81.3%	76.8%
All criteria	67.6%	71.2%	70.7%	73.6%	75.2%	73.3%

Table 2: Frequency of clinical criteria satisfaction.

that the GAN plans generated dose distributions that most closely resembled the clinical dose distributions, followed by the CNN, and then the gPCA plans. Notably, the GAN dose distributions best resembled the clinical dose distribution around the target, which is of primary importance. The GAN plans performed less well on the OARs, but this result was expected given the results from Table 2, which indicated that the GAN plans achieved more OAR clinical criteria than the clinical plan (i.e., the GAN was able to deliver a lower dose to the OARs as compared to the clinical dose distribution).

	BQ	gPCA	RF	CNN	GAN
All OARs	0.548	0.584	0.535	0.566	0.549
All PTVs	0.533	0.728	0.503	0.741	0.761
All Structures	0.536	0.669	0.518	0.670	0.675

Table 3: Average GPR for each population of KBP plans compared to clinical plans.

## 5. Discussion and Future Work

In this paper, we proposed the first GAN-based KBP method to generate radiation therapy treatment plans. We trained our complete pipeline on 130 patients, tested on 87 out-of-sample patients diagnosed with oropharyngeal cancer, and compared our technique with

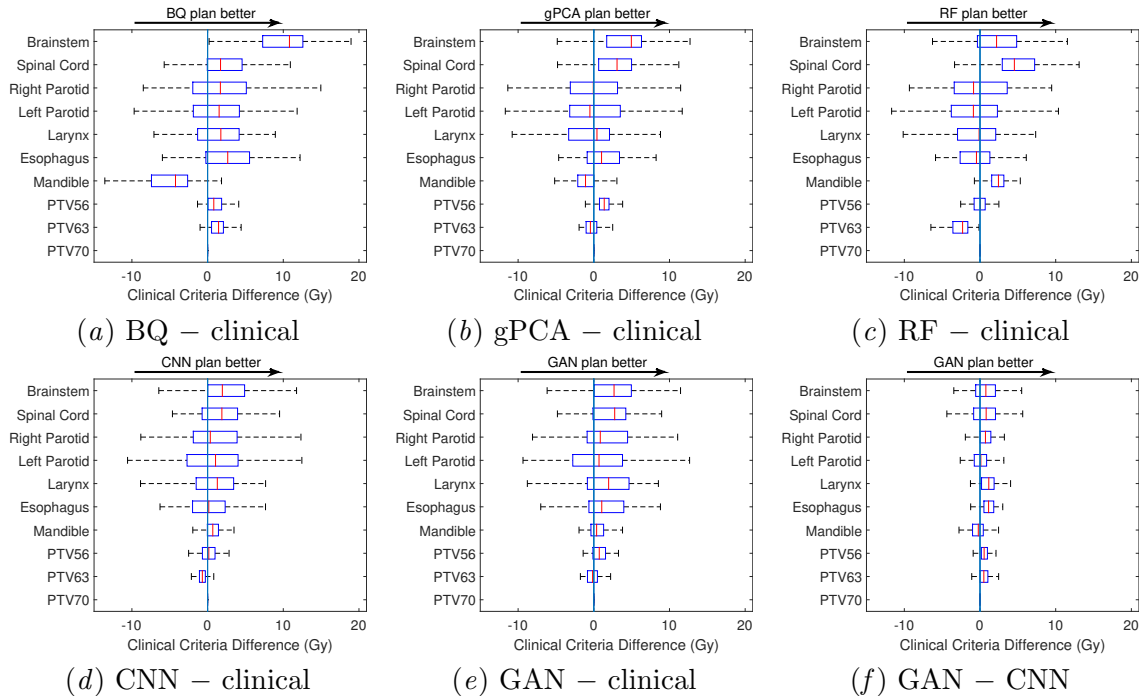


Figure 4: Head-to-head comparisons: (a)–(e) the plans from each KBP-generated model versus their clinical counterparts where positive difference implies the KBP-generated plans were better; (f) the plans from the GAN versus the CNN. Upper and lower boundaries of each box represent the 75th and 25th percentiles respectively, and the vertical line in the box depicts the median. Whiskers extend to 1.5 times the interquartile range. The line across each plot provides a reference for zero difference.

several state-of-the-art planning methods including a query-based approach, a PCA-based method, a random forest, and a CNN. All methods were evaluated on standard clinical criteria for plan evaluation (i.e., OARs sparing and target coverage), showing that the GAN plans outperformed all baseline KBP methods. We also demonstrated that the GAN plans outperformed the clinical plans by satisfying additional criteria on OAR dose sparing and target dose coverage. Finally, we used the gamma passing rate, a standard metric in the radiation therapy literature, to evaluate the similarity of the full 3D dose distribution between the KBP and clinical plans demonstrating that the GAN plans were the most similar to clinical plans on average. Note that the performance of automated planning methods should be measured based on their ability to re-create clinical quality plans with minimal manual effort. Of course, if the auto-generated plans manage to improve upon the clinical plans, that would be even better.

Our approach eschews the classical paradigm of predicting low-dimensional representations, or engineering features, by training a generic neural network to learn desirable dose distributions. Specifically, the GAN recasts KBP prediction as an image colorization problem. Moreover, the GAN is trained by mimicking the iterative process between the

treatment planner and oncologist; the generator network acts as the treatment planner by designing dose distributions while the discriminator acts as the oncologist by determining whether the plans are good or bad. The implication is that selecting the appropriate neural network architecture may be sufficient when creating an automated KBP pipeline that generates deliverable plans. Further, our approach does not add site-specific feature variables which suggests that the good performance we observe may not be limited to patients with oropharyngeal cancer. Finally, since the GAN plans improve upon the clinical plans, it may be useful to analyze the results to generate useful insights for practitioners.

We envision two interesting directions for future work. First, we plan to explore how GANs can develop treatment plans for different cancer sites. By adding site labels, we expect that a GAN can learn from the augmented training set of different cancer sites to better develop plans for specific sites. Second, we hope to automate the preprocessing stage by using uncontoured CT images. As neural networks show increasing promise for automated image segmentation (i.e., tumor and healthy organ identification), we hope to leverage this work to improve our treatment plan prediction model.

## Acknowledgments

This study was approved by the institutional research ethics board. Support for this research was provided by the Natural Sciences and Engineering Research Council of Canada.

## References

- R. K. Ahuja and J. B. Orlin. Inverse optimization. *Operations Research*, 49(5):771–783, 2001.
- V. Alex, M. S. KP, S. S. Chennamsetty, and G. Krishnamurthi. Generative adversarial networks for brain lesion detection. In *Medical Imaging 2017: Image Processing*, volume 10133, page 101330G. International Society for Optics and Photonics, 2017.
- L. M. Appenzoller, J. M. Michalski, W. L. Thorstad, S. Mutic, and K. L. Moore. Predicting dose-volume histograms for organs-at-risk in imrt planning. *Medical physics*, 39(12):7446–7461, 2012.
- A. Babier, J. J. Boutilier, A. L. McNiven, and T. C. Y. Chan. Knowledge-based automated planning for oropharyngeal cancer. *Med Phys*, 45(7):2875–2883, Jul 2018a.
- A. Babier, J. J. Boutilier, M. B Sharpe, A. L. McNiven, and T. C. Y. Chan. Inverse optimization of objective function weights for treatment planning using clinical dose-volume histograms. *Phys Med Biol*, 63(10):105004, May 2018b.
- T. C. Y. Chan and T. Lee. Trade-off preservation in inverse multi-objective convex optimization. *accepted to European Journal of Operations Research*, 2018.
- T. C. Y. Chan, T. Craig, T. Lee, and M. B. Sharpe. Generalized inverse multiobjective optimization with application to cancer therapy. *Oper. Res.*, 62(3):680–95, 2014.

- E. Choi, S. Biswal, B. Malin, J. Duke, W. F. Stewart, and J. Sun. Generating multi-label discrete electronic health records using generative adversarial networks. *arXiv preprint arXiv:1703.06490*, 2017.
- D. Craft, P. Suss, and T. Bortfeld. The tradeoff between treatment plan quality and required number of monitor units in intensity-modulated radiotherapy. *Int. J. Radiat. Oncol. Biol. Phys.*, 67(5):1596–605, 2007.
- I. J. Das, V. Moskvin, and P. A. Johnstone. Analysis of treatment planning time among systems and planners for intensity-modulated radiation therapy. *J Am Coll Radiol*, 6(7):514–7, Jul 2009. doi: 10.1016/j.jacr.2008.12.013.
- J. O. Deasy, A. I. Blanco, and V. H. Clark. CERR: a computational environment for radiotherapy research. *Med. Phys.*, 30(5):979–85, 2003.
- G. Delaney, S. Jacob, C. Featherstone, and M. Barton. The role of radiotherapy in cancer treatment. *Cancer*, 104(6):1129–1137, 2005.
- C. Esteban, S. L. Hyland, and G. Rätsch. Real-valued (medical) time series generation with recurrent conditional gans. *arXiv preprint arXiv:1706.02633*, 2017.
- M. Frid-Adar, I. Diamant, E. Klang, M. Amitai, J. Goldberger, and H. Greenspan. Gan-based synthetic medical image augmentation for increased cnn performance in liver lesion classification. *arXiv preprint arXiv:1803.01229*, 2018.
- I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- R. Hermoza and I. Sipiran. 3d reconstruction of incomplete archaeological objects using a generative adversary network. *arXiv preprint arXiv:1711.06363*, 2017.
- P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. *arXiv preprint*, 2017.
- A. Kadurin, S. Nikolenko, K. Khrabrov, A. Aliper, and A. Zhavoronkov. druGAN: An Advanced Generative Adversarial Autoencoder Model for de Novo Generation of New Molecules with Desired Molecular Properties in Silico. *Molecular Pharmaceutics*, 14(9):3098–3104, 2017.
- D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- D. A. Low, W. B. Harms, S. Mutic, and J. A. Purdy. A technique for the quantitative evaluation of dose distributions. *Medical physics*, 25(5):656–661, 1998.
- C. McIntosh and T. G. Purdie. Voxel-based dose prediction with multi-patient atlas selection for automated radiotherapy treatment planning. *Phys Med Biol*, 62(2):415–431, Jan 2017. doi: 10.1088/1361-6560/62/2/415.

- C. McIntosh, M. Welch, A. McNiven, D. A. Jaffray, and T. G. Purdie. Fully automated treatment planning for head and neck radiotherapy using a voxel-based dose prediction and dose mimicking method. *Phys. Med. Biol.*, 62(15):5926–5944, 2017.
- D. Nguyen, T. Long, X. Jia, W. Lu, X. Gu, Z. Iqbal, and S. Jiang. Dose prediction with u-net: A feasibility study for predicting dose distributions from contours using deep learning on prostate imrt patients. *arXiv preprint arXiv:1709.09233*, 2017.
- K. Petersson, P. Nilsson, P. Engström, T. Knöös, and C. Ceberg. Evaluation of dual-arc vmat radiotherapy treatment plans automatically generated via dose mimicking. *Acta Oncologica*, 55(4):523–525, 2016.
- P. Sangkloy, J. Lu, C. Fang, F. Yu, and J. Hays. Scribbler: Controlling deep image synthesis with sketch and color. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2, 2017.
- M. B. Sharpe, K. L. Moore, and C. G. Orton. Within the next ten years treatment planning will become fully automated without the need for human intervention. *Medical physics*, 41(12), 2014.
- S. Shiraishi and K. L. Moore. Knowledge-based prediction of three-dimensional dose distributions for external beam radiotherapy. *Med. Phys.*, 43(1):378, 2016.
- S. Shiraishi, J. Tan, L. A. Olsen, and K. L. Moore. Knowledge-based prediction of plan quality metrics in intracranial stereotactic radiosurgery. *Med. Phys.*, 42(2):908, 2015.
- N. Wang, W. Zha, J. Li, and X. Gao. Back projection: an effective postprocessing method for gan-based face sketch synthesis. *Pattern Recognition Letters*, 2017.
- B. Wu, F. Ricchetti, G. Sanguineti, M. Kazhdan, P. Simari, M. Chuang, R. Taylor, R. Jacques, and T. McNutt. Patient geometry-driven information retrieval for IMRT treatment plan quality control. *Med. Phys.*, 36(12):5497–505, 2009.
- B. Wu, F. Ricchetti, G. Sanguineti, M. Kazhdan, P. Simari, R. Jacques, R. Taylor, and T. McNutt. Data-driven approach to generating achievable dose-volume histogram objectives in intensity-modulated radiotherapy planning. *Int. J. Radiat. Oncol. Biol. Phys.*, 79(4):1241–7, 2011.
- B. Wu, M. Kusters, M. Kunze-Busch, T. Dijkema, T. McNutt, G. Sanguineti, K. Bzdusek, A. Dritschilo, and D. Pang. Cross-institutional knowledge-based planning (KBP) implementation and its performance comparison to auto-planning engine (APE). *Radiother. Oncol.*, 123(1):57–62, 2017.
- J. Wu, C. Zhang, T. Xue, B. Freeman, and J. Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In *Advances in Neural Information Processing Systems*, pages 82–90, 2016.
- T. Yang, E. C. Ford, B. Wu, M. Pinkawa, B. van Triest, P. Campbell, D. Y. Song, and T. R. McNutt. An overlap-volume-histogram based method for rectal dose prediction

- and automated treatment planning in the external beam prostate radiotherapy following hydrogel injection. *Med. Phys.*, 40(1):011709, 2013.
- K. C. Younge, R. B. Marsh, D. Owen, H. Geng, Y. Xiao, D. E. Spratt, J. Foy, K. Suresh, Q. J. Wu, F. Yin, S. Ryu, and M. M. Matuszak. Improving quality and consistency in nrg oncology radiation therapy oncology group 0631 for spine radiosurgery via knowledge-based planning. *Int J Radiat Oncol Biol Phys*, 100(4):1067–1074, Mar 2018. doi: 10.1016/j.ijrobp.2017.12.276.
- L. Yuan, Y. Ge, W. R. Lee, F. F. Yin, J. P. Kirkpatrick, and Q. J. Wu. Quantitative analysis of the factors which affect the interpatient organ-at-risk dose sparing variation in IMRT plans. *Med. Phys.*, 39(11):6868–78, 2012.
- J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. *arXiv preprint arXiv:1703.10593*, 2017.
- X. Zhu, Y. Ge, T. Li, D. Thongphiew, F. Yin, and Q. J. Wu. A planning quality evaluation tool for prostate adaptive IMRT based on machine learning. *Med. Phys.*, 38(2):719–26, 2011.