PAPER • OPEN ACCESS

Probabilistic dose prediction using mixture density networks for automated radiation therapy treatment planning

To cite this article: Viktor Nilsson *et al* 2021 *Phys. Med. Biol.* **66** 055003

View the [article online](#) for updates and enhancements.

END-to-END QA with the
QUASAR™ Multi-Purpose
Body Phantom: TPS and SBRT

WEBINAR

April 08, 2021 - 10 a.m. (EST)

PAPER

OPEN ACCESS

RECEIVED
2 October 2020REVISED
23 December 2020ACCEPTED FOR PUBLICATION
19 January 2021PUBLISHED
12 February 2021

Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Probabilistic dose prediction using mixture density networks for automated radiation therapy treatment planning

Viktor Nilsson^{1,2,5,*} , Hanna Gruselius^{1,5} , Tianfang Zhang^{1,2} , Geert De Kerf³ and Michaël Claessens^{3,4}¹ RaySearch Laboratories, Stockholm, Sweden² Department of Mathematics, KTH Royal Institute of Technology, Stockholm, Sweden³ Iridium Cancer Network, Antwerp, Belgium⁴ Department of Radiation Oncology, Faculty of Medicine and Health Sciences, University of Antwerp, Antwerp, Belgium⁵ These authors contributed equally to this work.

* Author to whom any correspondence should be addressed.

E-mail: vikn@kth.se**Keywords:** mixture density network, dose prediction, dose mimicking, knowledge-based planning, deep learning, radiation therapy treatment planning

Abstract

We demonstrate the application of mixture density networks (MDNs) in the context of automated radiation therapy treatment planning. It is shown that an MDN can produce good predictions of dose distributions as well as reflect uncertain decision making associated with inherently conflicting clinical tradeoffs, in contrast to deterministic methods previously investigated in the literature. A two-component Gaussian MDN is trained on a set of treatment plans for postoperative prostate patients with varying extents to which rectum dose sparing was prioritized over target coverage. Examination on a test set of patients shows that the predicted modes follow their respective ground truths well, both spatially and in terms of their dose–volume histograms. A special dose mimicking method based on the MDN output is used to produce deliverable plans and thereby showcase the usability of voxel-wise predictive densities. Thus, this type of MDN may serve to support clinicians in managing clinical tradeoffs and has the potential to improve the quality of plans produced by an automated treatment planning pipeline.

Introduction

Radiation therapy is a cornerstone treatment used as a single or combined treatment for more than 40 percent of cancer patients (Borras *et al* 2016). Previous and emerging innovations in imaging and treatment techniques have contributed to a more accurate radiation administration to the tumor while simultaneously sparing organs at risk. Despite these improvements, different processes of the radiation therapy workflow are still susceptible to variations. Important sources of uncertainty include heterogeneity in the delineation of regions of interest (ROIs) (Brouwer *et al* 2012, van der Veen *et al* 2019) and differences in preferences associated with conflicting clinical tradeoffs for a particular treatment site (Nelms *et al* 2012). In particular, the radiation therapy treatment planning (RTTP) process requires interaction between dosimetrists, medical physicists and radiation oncologists to find the best treatment plan for each patient. The creation of a deliverable plan typically involves repeatedly solving a large-scale optimization problem with iteratively updated parameters, which makes it a time-consuming and labor-intensive step (Craft *et al* 2012). Moreover, the process is characterized by a user-variability both intra- and inter-institutional due to subjectivity in the management of tradeoffs between planning objectives and institution-specific guidelines (Nelms *et al* 2012). Therefore, efforts have been made to develop methods that can automate the RTTP process to reduce the time required for human intervention and the variation between dosimetrists/physicists to improve overall plan quality and consistency.

In this context, knowledge-based planning (KBP) could remedy the differences coming from subjectivity and has the potential to homogenize the RTTP process (Berry *et al* 2016). Ge and Wu (2019) divide current methods in KBP in two categories: (1) case- and atlas-based methods and statistical modeling and (2) machine

learning (ML) methods. In the former category, a reduced plan quality variation is achieved by seeking similar patients in a cohort of available ones, and then leveraging the knowledge from the previous case to a new patient. In the latter type, a model is trained on a set of previously treated patients, which is then used to predict some dose-related quantity for a new patient that is of assistance when creating a deliverable plan. Examples of such quantities are beam angles, dose–volume histograms (DVHs), voxel-wise dose and mean region doses (Ge and Wu 2019). Voxel-wise dose prediction methods have been created using a variety of ML-based approaches. Nguyen *et al* (2019) validate that a neural network (NN) of U-net type can be applied for voxel-wise dose prediction based on patient contours for prostate intensity-modulated radiation therapy, and Campbell *et al* (2017) use an NN to predict dose for pancreatic stereotactic body radiotherapy (SBRT). Further developments of NNs for KBP have been achieved by Kearney *et al* (2018) and Shiraishi and Moore (2016), where a fully convolutional three-dimensional NN is created for dose prediction in prostate SBRT. Also, generative approaches have been investigated in Mahmood *et al* (2018), Murakami *et al* (1919) and Babier *et al* (2020), where voxel-wise dose is predicted by generative adversarial networks (GANs). Other works have used more classical ML algorithms, e.g., random forest–based methods to predict voxel doses (McIntosh and Purdie 2016, McIntosh *et al* 2017) and the application of flow and shape models for spine SBRT (Liu *et al* 2015).

Common for previous NN-based approaches is that a single dose per voxel is outputted by the ML algorithm—the dose of the predicted optimal treatment plan for the patient. However, for complex cases, inter-planner variations with KBP will remain (Kubo *et al* 2019). Hence, an optimal prediction could consist of the dose distribution that best mimics the behavior of previously seen data, while there is still a possibility to incorporate other preferences by steering the ML algorithm during training. This type of KBP can remedy the variability brought by subjectivity of a planner in the RTTP process. However, since a single dose value is predicted in each voxel, it will not be able to reflect more than one specific protocol and preference in the model. Consequently, if there are potentially conflicting tradeoffs inherent in the treatment site, where one would possibly have different protocols or objectives in place (e.g., complex cases), no such concerns can be reflected in the same ML model.

The present work is an extension of previous methods by incorporating meaningful variations in the RTTP problem in the output of the ML algorithm, in particular demonstrating the variability caused by prioritizing different clinical tradeoffs. This is achieved by using a mixture density network to predict a possibly multimodal probability density function of dose in each voxel. This can add to KBP not only a more homogenized manner of RTTP, but also a way of automatically grasping the viable treatments based on clinically relevant variability, which could be of use as a decision support mechanism in clinical practice. Moreover, an integral part of automated RTTP using ML is the optimization to create a deliverable plan from a predicted dose distribution subject to physical delivery constraints, commonly referred to as dose mimicking (McIntosh *et al* 2017). While dose mimicking can be done in a range of manners, in this paper, we present the idea of incorporating the predictive probability distributions into the objective function as an application of the implied uncertainties. Among other applications, the generation of different deliverable plan alternatives, e.g., corresponding to the mixture density components, may be used in a multicriteria optimization setting (Craft *et al* 2012). It is also possible to use the voxel-wise probability densities to post-process the predicted dose distribution to have the desired DVHs in different regions of interest (McIntosh *et al* 2017), which constitutes a comprehensive way of articulating additional preferences. Hence, a multimodal voxel dose distribution can bring efficiency and improvements to a range of RTTP applications.

Methods and materials

Model setup

In order to reflect a range of possible predictions with an ML model, its output must comprise more than a deterministic output, such as a single dose value per voxel. A more informative output would instead be a probability distribution per voxel, reflecting the range of possible outcomes of dose values and their respective likelihoods given the planning standards learned from the training data. In figure 1, an example of such a probability distribution for a voxel in the target area is depicted.

More formally, let $\{(x_n, d_n)\}_{n=1}^N$ be the training set consisting of independent random variable pairs of contoured images x_n and corresponding dose distributions d_n . Given this and the current patient image $x = (x^v)_v$, as a vector over voxels v , the task is to obtain for each v the probability distribution of the associated dose d^v , i.e.

$$p(d^v | x, \{(x_n, d_n)\}_{n=1}^N). \quad (1)$$

Note that this is the marginal predictive distribution for one particular voxel v . Although one could, in principle, try to model the joint predictive distribution $p(d | x, \{(x_n, d_n)\}_{n=1}^N)$, it is sufficient for the purposes of this study to obtain a collection of marginal distributions over all voxels in the patient volume.

Figure 1. Visualization of a probability distribution of dose versus a deterministic dose for a voxel in the target area.

Mixture density networks

As mentioned in the introduction, deep learning has been applied to the problem of voxel-wise deterministic dose prediction, e.g., using U-nets. A novel idea is to extend the method presented in Nguyen *et al* (2019) around U-net prediction to instead output an inferred predictive distribution $p_\theta(d^v | x)$, where θ denotes the network weights. While there are several approaches to making NNs predict stochastic outputs, such as by sampling from a latent variable space (Goodfellow *et al* 2014, Kingma and Welling 2013), we shall opt for a discriminative approach where the marginal predictive distributions belong to some parametric family and where the NN outputs are the corresponding parameters. The family of Gaussian mixture models has the favorable property of being dense in the space of probability distributions (Nguyen and McLachlan 2019)—in other words, it is a parametric family that can approximate any distribution arbitrarily well. A one-dimensional Gaussian mixture model with C components is a probability distribution with density function

$$\sum_{c=1}^C \alpha_c \mathcal{N}(\cdot | \mu_c, \sigma_c^2), \quad (2)$$

where $\mathcal{N}(\cdot | \mu_c, \sigma_c^2)$ is the density function of the normal distribution with mean μ_c and variance σ_c^2 and where the class probabilities $\{\alpha_c\}_{c=1}^C$ are nonnegative and sum to unity.

Thus, a C -component Gaussian mixture model requires $3C$ parameters (alternatively, $3C - 1$, as one is implicit from the constraint that class probabilities sum to unity), which will be outputted by the NN. Such an NN is known as a mixture density network (Bishop 1994). More precisely, three functions α_θ , μ_θ and σ_θ^2 will be trained in the network, derived from the network weights θ . α_θ maps each input image $x = (x^v)_v$ to a matrix with C rows, where the output $\alpha_\theta(x)_c$ in row c is the vector of the corresponding class probabilities (and similarly for μ_θ and σ_θ^2). The inferred voxel-wise predictive likelihood can thus be written as

$$p_\theta(d^v | x) = \sum_{c=1}^C \alpha_\theta(x)_c^v \mathcal{N}(d^v | \mu_\theta(x)_c^v, \sigma_\theta^2(x)_c^v). \quad (3)$$

This type of model is able to reflect any number of inherent distributional components. Consider a dose inference problem where one knows there is a tradeoff prevalent in the data stemming from two potential treatment protocols T_1, T_2 . An MDN with two components can reflect this behavior by predicting μ_1, μ_2 to match the average behavior of each protocol, and σ_1, σ_2 to reflect the internal uncertainty of each prediction. In figure 2, one can see a visualization of how such a predictive density could be in a voxel where this twofold dose preference is present.

Data

In order to carry out experiments demonstrating the sought variability for a site, a pertinent clinical tradeoff is needed as a basis for the work. Studies have focused on analyzing the radiobiological effect of restricting the dose to the rectum for prostate patients, possibly at the cost of sacrificing target coverage, which is a clinical tradeoff that is central in prostate RTTP (Mavroidis *et al* 2017). In this work, we will use the dose variation in the rectum area to study the ML method's capability to identify and mimic these tradeoffs in its prediction.

The data used in this work originates from Iridium Cancer Network (Antwerp, Belgium) and consists of prostate cancer patients having undergone a prostatectomy prior to radiation therapy. The dataset contains both patients with and without pelvic nodes included in the target volume. All organs were contoured according to the RTOG (Gay *et al* 2012) and ACROP (Salembier *et al* 2018) guidelines. Patients were treated in 35 fractions with a prescribed dose of 56 Gy to the seminal vesicles and eventually also to the pelvic nodes. The prostate bed was treated with a simultaneous integrated boost of 70 Gy. All plans were delivered using two 360° volumetric

Table 1. Vital clinical goals in the Iridium postoperative prostate protocol.

Region	Volume [%]	Dose [Gy]
PTV Prostate	98	≥ 66.5
	0	≤ 74.9
PTV Seminal Vesicles	98	≥ 53.2
Rectum	10	≤ 65
	5	≤ 70
Bladder	10	≤ 70

modulated arc therapy (VMAT) beams.⁶ The clinical goals used during the planning process are summarized in table 1.

In this work, the training data consisted of a pair P_1, P_2 of treatment plans created for each patient, where P_2 has a dose distribution that spares the rectum to a greater extent than P_1 , in general at the cost of sacrificing target coverage. In figure 3, one can see the dose distributions for each such plan in a transverse cross-section for a representative patient. Figure 4 shows the average DVHs for P_1, P_2 in comparison to the original clinical goals in table 1. The training data was generated by the application of an ML-based dose prediction and dose mimicking optimization using RayStation 9A (RaySearch Laboratories, Stockholm, Sweden). The ML algorithm used to predict the doses was trained on approximately 100 clinical plans and is based on the method provided in McIntosh and Purdie (2016). The model used to create the training data was configured to predict one plan P_1 that prioritizes target coverage and another plan P_2 that prioritizes sparing of the rectum area. By applying this planning process for 16 patients from the Iridium dataset, 32 plans were created for the training dataset. The

⁶ This retrospective study did not involve human participants and only incorporated historical, anonymized data.

objective of using this dual input of plans for each patient is that one can easily validate that this twofold dose preference in the rectum is reflected in new MDN predictions. Rather than creating two clinically acceptable plans, P_1 and P_2 just reflect two different planning protocols.

Implementation

The architecture of the neural network used in this work is a U-net (Ronneberger *et al* 2015). U-nets are convolutional neural networks with equal dimension (here, 128×176) in both the and output layers, but possibly different number of channels. For this study, we restrict our attention to predicting dose in each transversal slice separately. As mentioned in the introduction, this type of approach has been used for dose prediction successfully.

The input to the U-net could be a CT image with radiodensities and/or contoured ROIs. As the latter alone has proven to be successful for the task, the network built for this work uses only the ROI contours. The ROIs included are the External (the full body), PTV Prostate, PTV Seminal Vesicles, Rectum and Bladder. Hence, the input layer has five channels. The simplest representation of the input ROIs are their binary encoding in the three-dimensional patient volume, but since the prediction is based on transversal slices, proximity information of ROIs in the craniocaudal direction would be lost. To account for this, a Euclidean distance transform is applied to each ROI before separating the slices. That is, letting d_{Eucl} denote spatial distance, the Euclidean distance transform T_R , with respect to a ROI R , is a map from voxel to scalar, given by

$$T(v) = \min_{v' \in R} d_{\text{Eucl}}(v, v'). \quad (4)$$

In this work, and augmented version $\hat{T}_R$ given by $\hat{T}_R(v) = T_R(v) - T_{R^c}(v)$, R^c denoting the complement of R , is used to provide more local information when $v \in R$. $\hat{T}_R$ is applied in three dimensions before the slices are picked out. In this way, it encodes proximity information from the three-dimensional geometry into the two-dimensional predictors by replacing the binary masks with more informative voxel-wise distance values.

The output of the network is the distributional parameters in each voxel. In this work, attention is restricted to a bimodal Gaussian mixture, as the tradeoff we are focusing on comes from two protocols P_1 , P_2 , i.e., $C = 2$. Thus, the output should have $3C = 6$ channels. U-nets have a contractive phase, followed by an expansive phase. During the contractive phase, the image is successively downsampled using max-pooling. As the image shrinks, the number of channels simultaneously increases, which allows higher-level contextual information to be represented. In the expansive phase, this is propagated up to the original resolution using upsampling. Fine-grained information may be lost in the contractive part of the network. Therefore, there are ‘skip connections’ between the phases where detailed input from each step in the downsampling is propagated to the upsampling part. For an illustration of the complete architecture used, see figure 5.

A loss function that reflects the NN’s performance in this setting is the negative log-likelihood of the data, which measures how probabilistically congruent our data is to the proposed distributions in each voxel. However, since voxels are not equally critical to the plan, we use a loss function where local negative log-likelihoods $-\log p_\theta(d^v | x)$, known from equation (3), have weights w^v reflecting their relative importance. Additionally, to avoid mode collapse—i.e., making sure that the class probabilities do not vanish—we use a regularization penalty

$$b^v(\theta; x) = \max\{(\alpha_0 - \alpha_\theta(x)_c^v)_+\}_{c=1}^C \quad (5)$$

for each voxel, α_0 being some threshold constant. The loss function $L(\theta; x, d)$ associated with the pair (x, d) , which in this case are both slice-specific, can thus be written as

$$L(\theta; x, d) = \sum_v (-w^v \log p(d^v | x) + \lambda b^v(\theta; x)), \quad (6)$$

where λ is some scaling factor. The constants $\{w^v\}_v$, α_0 and λ can be seen as hyperparameters of the model and have been configured by hand. For our computational study, we used

$$w^v = \begin{cases} 0.05 & \text{if } v \text{ is in air,} \\ 5 & \text{if } v \text{ is exclusively in PTV or rectum,} \\ 10 & \text{if } v \text{ is in PTV and rectum,} \\ 1 & \text{otherwise,} \end{cases} \quad \alpha_0 = 0.2, \quad \lambda = 20. \quad (7)$$

Experiments are implemented using TensorFlow 2.2 and TensorFlow Probability 0.10. The model is trained using Adam (Kingma and Ba 2014) with standard parameters. The final model is selected as the one with the lowest validation loss during training and evaluated on an untouched holdout data set. A train/validation/holdout ratio of 50/25/25 is used—specifically, there are 16 patients in the training set, 8 in the validation set, and 9 in the holdout set.

Dose mimicking

To demonstrate the usability of predictive probability distributions such as those outputted by the MDN, we use a specially developed dose mimicking method to create deliverable treatment plans according to the predictions. Since the mixtures comprise two components, two plans may be created for each patient—one where the dose should mimic μ_1 and one for μ_2 . For each v , let F^v be the cumulative distribution function associated with one chosen mixture component of the predictive density $p(d^v | x)$, and let t^v be such that $t^v = 0$ if the dose to voxel v is to be minimized and $t^v = 1$ otherwise. In particular, when v belongs to both the PTV and the rectum, we use $t^v = 1$ if mimicking μ_1 and $t^v = 0$ if mimicking μ_2 . The objective function ψ_{Spat} on the spatial dose distribution is written as sum of cross-entropy losses over all voxels in the patient volume—that is,

$$\psi_{\text{Spat}}(d) = -\sum_v r^v (t^v \log F^v(d^v) + (1 - t^v) \log(1 - F^v(d^v))), \quad (8)$$

where r^v is the voxel volume of v relative to the whole patient volume, satisfying $\sum_v r^v = 1$. To ensure homogeneity in target doses and, in general, that the deliverable dose follows the DVHs of the prediction closely, we add for each ROI R in the model a so called reference-DVH function $\psi_{\text{DVH},R}$ with μ_1 or μ_2 as reference dose. Their definitions are described in detail in Fredriksson (2012).

The optimizations were performed using direct machine parameter optimization with the same beam settings as used in the original data. Let η denote the optimization variables with feasible set $\mathcal{E}$, where the total dose d is determined by some dose deposition mapping $d = d(\eta)$ according to, for example, Unkelbach *et al* (2015). The optimization problem may then be written as

$$\underset{\eta \in \mathcal{E}}{\text{minimize}} \quad w_{\text{Spat}} \psi_{\text{Spat}}(d(\eta)) + \sum_R w_{\text{DVH},R} \psi_{\text{DVH},R}(d(\eta)), \quad (9)$$

w_{Spat} , w_{DVH} being weights. For our computational study, we used $w_{\text{Spat}} = 10^{-4}$ and $w_{\text{DVH},R} = 10^2$ for organs at risk and $w_{\text{DVH},R} = 10^3$ for targets. The problem above was solved using RayStation's native sequential quadratic programming solver, and final doses were calculated using a collapsed cone algorithm.

Results

The MDN trained for 100 epochs, although after around 50 epochs the final model is typically selected due to incipient overfitting. After training the network, its performance is assessed through some qualitative inspections. As it is difficult in general to evaluate fitted predictive distributions, the focus of this section is to provide some key visualizations of the model behavior compared to the ground truth counterpart.

First of all, one would like to see that the two data modes P_1, P_2 have indeed been separated. Thus, a central transversal slice from a patient is firstly studied in figure 6, where the two modes μ_1 and μ_2 are plotted. We also show the predicted standard deviation in each mode. For reference, the ground truth dose for each data mode P_1 and P_2 is plotted.

Examining the modes in figure 6, one sees that μ_2 predicts the dimple present in P_2 , while μ_1 has the higher target coverage characteristic of P_1 . Another way to clearly demonstrate that the modes have separated is to plot the difference $\mu_1 - \mu_2$. In figure 7, one can see that the modes are indeed separated in and around the rectum, while differences in other regions are smaller.

In order to validate the network across the holdout set (as opposed to only the one patient in figures 6, 7), we compare the distribution of dose values in equivalent locations for the patients to the predicted counterparts. The points are drawn along a line in the anterior–posterior direction passing through the center of mass of the rectum in the slice (figures 6, 7) and corresponding slices through the holdout set. Samples are drawn from the two predicted components for each patient and compared to the distribution in P_1 and P_2 using boxplots in figure 8. Five samples are drawn for each patient in order to stabilize the empirical distribution and thus the box positions.

One sees that μ_1 manages to follow P_1 well through the drop-off, although a slight discrepancy can be made out when the offset exceeds 10 mm, while μ_2 follows P_2 well.

Illustrations such as figures 6, 7 and 8 give a detailed assessment of the model in the most critical area of the patient body, where the PTV is the largest and has the most overlap with the rectum. As such, however, they focus on one single representative patient and the dose distribution in a vital slice at a time. A way of examining

the behavior throughout the patient body rather than in just one slice, and also for a larger test set, is by comparing differential or cumulative DVHs. In particular, one would like to see that the DVHs of μ_1 and μ_2 agree with P_1 and P_2 , respectively. Using kernel density estimates, the differential DVHs for the rectum area for a cohort of test patients are shown in figure 9.

In figure 9, it is apparent that the two plans P_1 and P_2 have created two separate and characteristic distributions in the rectum area. Furthermore, it can be seen that the predictive modes μ_1 or μ_2 have successfully recreated this behavior, thus reflecting the clinical tradeoff in the region. The corresponding cumulative DVHs are also examined for ROIs included in the model. In figure 10, the averages of the DVHs of the predictive modes μ_1, μ_2 throughout the holdout set are shown alongside the ground truth counterparts P_1, P_2 , respectively. The predicted modes follow the ground truth well for the risk organs and the prostate, essentially capturing the distinct behaviors of P_1, P_2 . However, the dose to the seminal vesicles is somewhat overestimated, also causing the predicted DVH of the external ROI to deviate slightly. This is not entirely unexpected, considering the fact that the MDN is not trained to necessarily

produce predictive means close to the ground truth in terms of DVH. Apart from the seminal vesicles, the overall plan consistency is kept by the MDN, and not only the variations in the rectum are reflected in the predictions.

Finally, to demonstrate that the proposed dose mimicking method can successfully utilize the predictive distributions outputted by the MDN, we create two deliverable plans for a patient in the holdout set mimicking, respectively, μ_1 and μ_2 . Spatial dose and DVH comparisons are shown in figures 11 and 12. For both cases, one can see that the spatial dose distributions follow the respective predictive modes well, although being slightly less conformal. This is most likely due to the predictions being overly optimistic in the rate by which dose falls off around the target area. One can also see that the DVHs of the mimicked plans followed closely or were in some cases even better than those predicted. While these results show the merits of a probabilistic dose mimicking method, in a clinical setting, one would likely need to further post-process the obtained plans before approval and delivery.

Discussion

In this work, we have demonstrated the application of MDNs for probabilistic dose prediction, which has certain advantages when used in an automated RTTP pipeline. We showed that this type of NN can produce predictive

probability distributions capturing the uncertain nature of decision making associated with conflicting clinical tradeoffs. This was exemplified on a dataset comprising postoperative prostate cancer patients in which two plans had been created for each patient—one prioritizing target coverage and one prioritizing rectum sparing. Visual inspection across the holdout set showed that the predicted modes followed the respective ground truths well, and the predictive uncertainties were assured to be reasonable by examining the distribution of ground truth dose values around different locations. This has the potential to bring together a comprehensive overview of viable plan alternatives for a given patient, which can not only act as a time-saving tool but also facilitate easier decision making during the RTTP process. In particular, as an example of the usability of a probabilistic prediction, we used a specially developed dose mimicking method to create deliverable plans for each of the predicted mixture components.

While most of the previous work has focused on the prediction of a single dose value per voxel, one can, for example, argue that an empirical distribution can be obtained by sampling from generative models, e.g., GANs (Mahmood *et al* 2018, Babier *et al* 2020) or variational autoencoders (Kingma and Welling 2013). However, using such approaches, there is no clear way of obtaining the conditional densities $p(d^v | x)$ we sought to model in this paper. Also, in the atlas-based method by McIntosh and Purdie (2016), a set of dose propositions are predicted for each voxel by the voting of weak learners, which may be seen as a sample from $p(d^v | x)$. The actual density may then be estimated by, e.g., kernel density estimation, but compared to MDNs, this approach requires many more parameters to form the densities, and it is moreover not obvious how one would go about separating the modes as was done for our dose mimicking. In these regards, it is concluded that MDNs present advantages over existing methods.

Examples of future work include studying disjoint sets of patients having different treatment protocols and using more realistic data from actually administered clinical plans. It would also be of interest to investigate data with other types of tradeoffs—for example plans with or without a boost volume inside the PTV or plans with intentional NTCP variations such as for head-and-neck, where saliva glands are spared to different extents in various protocols. Depending on the training data, one could look at other types of MDNs, e.g., an NN with more Gaussian mixture components, should the tradeoffs need more modes to be described. Future work could also focus on implementation parts, such as making the MDN a three-dimensional NN instead of predicting per slice, and various ways to counteract overfitting or stabilizing the training process for this high-dimensional problem that can be prone to mode collapse. Another possible future investigation would be to use the knowledge captured by the MDN in a new form of ML-based multicriteria optimization, where the treatment planner may navigate manually between different generated plan alternatives. This is a natural application of the multimodal predictive distributions in which clinically conflicting tradeoffs may be reflected, leading to a semiautomatic RTTP pipeline where important decision making may be left to be performed by clinicians.

Conclusion

The application of MDNs to dose prediction is a promising method toward the automation of the RTTP process. In this work, we have shown that an MDN can predict dose distributions well and, as opposed to deterministic methods, capture different decision alternatives associated with conflicting clinical tradeoffs. In particular, a two-component Gaussian MDN was trained on a dataset of treatment plans for prostate cancer patients with two distinguished degrees of priority of rectum dose sparing over target coverage. Evaluation on the holdout set showed that the produced probability distributions had predictive means following the respective ground truths

closely in terms of spatial dose and partly in DVH, as well as reasonable predictive variances. Moreover, deliverable plans were created using a specially developed dose mimicking procedure, exemplifying the usability of such probabilistic information. In conclusion, the proposed method may serve to assist clinicians in decision making and improve the quality of treatment plans produced by an automated RTTP pipeline.

Acknowledgments

The authors thank Adnan Hossain and Mats Holmström for their diligent work on the model used for creating the experimental data, and Tatjana Pavlenko, Pierre Nyquist, Fredrik Löfman and Marcus Nordström for valuable input on the manuscript. The work of Michaël Claessens was supported by a grant of the Flemish League against Cancer, Belgium (ref: 000019356).

ORCID iDs

Viktor Nilsson <https://orcid.org/0000-0001-5740-5103>

Hanna Gruselius <https://orcid.org/0000-0002-7944-2893>

Tianfang Zhang <https://orcid.org/0000-0001-6724-2547>

References

- Babier A, Mahmood R, McNiven A L, Diamant A and Chan T C Y 2020 Knowledge-based automated planning with three-dimensional generative adversarial networks *Med. Phys.* **47** 297–306
- Berry S L, Boczkowski A, Ma R, Mechalakos J and Hunt M 2016 Interobserver variability in radiation therapy plan output: results of a single-institution study *Pract. Radiat. Oncol.* **6** 442–9
- Bishop C M 1994 *Mixture Density Networks* (Birmingham: Aston University) <http://publications.aston.ac.uk/id/eprint/373>
- Borras J M, Lievens Y, Barton M, Corral J, Ferlay J, Bray F and Grau C 2016 How many new cancer patients in Europe will require radiotherapy by 2025? An esto-hero analysis *Radiother. Oncol.* **119** 5–11
- Brouwer C L et al 2012 3D variation in delineation of head and neck organs at risk, *Radiat Oncol.* **7** 32
- Campbell W G, Miften M, Olsen L, Stumpf P, Scheffer T, Goodman K A and Jones B L 2017 Neural network dose models for knowledge-based planning in pancreatic sbt *Med. Phys.* **44** 6148–58
- Craft D L, Hong T S, Shih H A and Bortfeld T R 2012 Improved planning time and plan quality through multicriteria optimization for intensity-modulated radiotherapy *Int. J. Radiat. Oncol. Biol. Phys.* **82** e83–90
- Fredriksson A 2012 Automated improvement of radiation therapy treatment plans by optimization under reference dose constraints *Phys. Med. Biol.* **57** 7799–811
- Gay H A et al 2012 Pelvic normal tissue contouring guidelines for radiation therapy: a radiation therapy oncology group consensus panel atlas *Int. J. Radiat. Oncol. Biol. Phys.* **83** e353–362
- Ge Y and Wu Q J 2019 Knowledge-based planning for intensity-modulated radiation therapy: a review of data-driven approaches *Med. Phys.* **46** 2760–75
- Goodfellow I., Pouget-Abadie J., Mirza M., Xu B., Warde-Farley D., Ozair S., Courville A. and Bengio Y. 2014 Generative Adversarial Networks *Adv. Neural Inf. Process. Sys.* **3**
- Kearney V, Chan J W, Haaf S, Descovich M and Solberg T D 2018 Dosenet: a volumetric dose prediction algorithm using 3D fully-convolutional neural networks *Phys. Med. Biol.* **63** 235022
- Kingma D P and Ba J 2014 Adam: a method for stochastic optimization (arXiv:1412.6980)
- Kingma D P and Welling M 2013 Auto-encoding variational Bayes (arXiv:1312.6114)
- Kubo K et al 2019 Inter-planner variation in treatment-plan quality of plans created with a knowledge-based treatment planning system *Phys. Medica.* **67** 132–40
- Liu J, Wu Q J, Kirkpatrick J P, Yin F-F, Yuan L and Ge Y 2015 From active shape model to active optical flow model: a shape-based approach to predicting voxel-level dose distributions in spine SBRT *Phys. Med. Biol.* **60** N83–92
- Mahmood R, Babier A, McNiven A, Diamant A and Chan T C Y 2018 Automated treatment planning in radiation therapy using generative adversarial networks arXiv:1807.06489
- Mavroidis P, Komisopoulos G, Buckley C, Mavroeidi M, Swanson G P, Baltas D, Papanikolaou N and Stathakis S 2017 Radiobiological evaluation of prostate cancer imrt and conformal-RT plans using different treatment protocols *Phys. Medica.* **40** 33–41
- McIntosh C and Purdie T G 2016 Voxel-based dose prediction with multi-patient atlas selection for automated radiotherapy treatment planning *Phys. Med. Biol.* **62** 415–31
- McIntosh C, Welch M, McNiven A, Jaffray D A and Purdie T G 2017 Fully automated treatment planning for head and neck radiotherapy using a voxel-based dose prediction and dose mimicking method *Phys. Med. Biol.* **62** 5926–44
- Murakami Y, Magome T, Matsumoto K, Sato T, Yoshioka Y and Oguchi M 1919 Fully automated dose prediction using generative adversarial networks in prostate cancer patients *PLoS One* **15** e0232697
- Nelms B E, Robinson G, Markham J, Velasco K, Boyd S, Narayan S, Wheeler J and Sobczak M L 2012 Variation in external beam treatment plan quality: an inter-institutional study of planners and planning systems *Pract. Radiat. Oncol.* **2** 296–305
- Nguyen D, Long T, Jia X, Lu W, Gu X, Iqbal Z and Jiang S 2019 A feasibility study for predicting optimal radiation therapy dose distributions of prostate cancer patients from patient anatomy using deep learning *Sci. Rep.* **9** 1079
- Nguyen H D and McLachlan G 2019 Onapproximations via convolution-defined mixture models *Commun. Stat. - Theory Methods* **48** 3945–55
- Ronneberger O, Fischer P and Brox T 2015 U-Net: Convolutional Networks for Biomedical Image Segmentation *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015* (Cham, Switzerland: Springer) 234–41

- Salembier C *et al* 2018 Estro acrop consensus guideline on CT and MRI-based target volume delineation for primary radiation therapy of localized prostate cancer *Radiother. Oncol.* **127** 49–61
- Shiraishi S and Moore K L 2016 Knowledge-based prediction of three-dimensional dose distributions for external beam radiotherapy *Med. Phys.* **43** 378–87
- Unkelbach J *et al* 2015 Optimization approaches to volumetric modulated arc therapy planning *Med. Phys.* **42** 1367–77
- van der Veen J, Gulyban A and Nuyts S 2019 Interobserver variability in delineation of target volumes in head and neck cancer *Radiother. Oncol.* **137** 9–15