



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

# Subgroup-Specific Risk-Controlled Dose Estimation in Radiotherapy

Paul Fischer<sup>\*1,2</sup>, Hannah Willms<sup>\*1</sup>, Moritz Schneider<sup>3</sup>, Daniela Thorwarth<sup>1,3</sup>, Michael Muehlebach<sup>4</sup>, and Christian F. Baumgartner<sup>1,2</sup>

<sup>1</sup> Cluster of Excellence – ML for Science, University of Tübingen, Germany

<sup>2</sup> Faculty of Health Sciences and Medicine, University of Lucerne, Switzerland

<sup>3</sup> Section for Biomedical Physics, Department of Radiation Oncology, University of Tübingen, Germany

<sup>4</sup> Max Planck Institute for Intelligent Systems, Tübingen, Germany

**Abstract.** Cancer remains a leading cause of death, highlighting the importance of effective radiotherapy (RT). Magnetic resonance-guided linear accelerators (MR-Linacs) enable imaging during RT, allowing for inter-fraction, and perhaps even intra-fraction, adjustments of treatment plans. However, achieving this requires fast and accurate dose calculations. While Monte Carlo simulations offer accuracy, they are computationally intensive. Deep learning frameworks show promise, yet lack uncertainty quantification crucial for high-risk applications like RT. Risk-controlling prediction sets (RCPS) offer model-agnostic uncertainty quantification with mathematical guarantees. However, we show that naive application of RCPS may lead to only certain subgroups such as the image background being risk-controlled. In this work, we extend RCPS to provide prediction intervals with coverage guarantees for multiple subgroups with unknown subgroup membership at test time. We evaluate our algorithm on real clinical planning volumes from five different anatomical regions and show that our novel subgroup RCPS (SG-RCPS) algorithm leads to prediction intervals that jointly control the risk for multiple subgroups. In particular, our method controls the risk of the crucial voxels along the radiation beam significantly better than conventional RCPS.

## 1 Introduction

Cancer remains one of the leading causes of death for people under the age of 70. Radiotherapy (RT) has been proven to be a critical treatment modality for various tumour entities. Recent advancements in medical imaging have led to the development of the magnetic resonance-guided linear accelerator (MR-Linac), which integrates MR imaging with RT [8]. Aligning the computed tomography (CT) planing volume to an MR scan acquired at the beginning of each treatment fraction, and re-evaluating the treatment plan based on the transformed CT volume allows to better adjust the plan to the patient's inter-fraction changes. This

---

\* Contributed equally


holds the potential for more precise tumor targeting, enhancing treatment outcomes. However, in order to successfully realize these benefits, fast and accurate dose deposition calculations are needed. Current algorithms are based on Monte Carlo simulations, which provide highly accurate results ([4]). However, these calculations are computationally intensive, taking minutes to hours. Therefore recent research has focused on optimizing treatment planning efficiency [17]. Deep learning (DL) frameworks have shown encouraging results in dose estimation for RT, with fast calculation time and accurate predictions [13,15,16,18]. However, despite RT’s high-risk nature, prior approaches have not adequately addressed risk assessment within DL-based dose estimation.

A natural approach for assessing the risk associated with a prediction is quantifying the prediction uncertainty. Recent years have seen the development of various methods for uncertainty quantification in medical imaging. Examples include deep ensembles [14], Monte Carlo dropout [11], or approaches based on variational autoencoders [3,6,12]. A major limitation of those techniques is that they do not provide any guarantees about the usefulness or correctness of the uncertainty estimates. Recently, risk-controlling prediction sets [2] (RCPS) has gained popularity as a simple, model-agnostic strategy to adapt heuristic notions of uncertainty into uncertainty measures with guarantees. RCPS allows to construct a set of predictions with a guarantee that the correct solution is inside this set with a user-defined probability. Such prediction sets can indicate poor model performance through excessively large intervals, revealing that the models may not be acceptable for certain high-risk applications.

While RCPS has already seen successful adoption in medical image analysis (e.g. [1]), a remaining limitation is that it can only provide guarantees on a global level. There are many situations where we are interested in obtaining guarantees for different subgroups of our population, or different image regions. For instance, in RT, we would like the method to be calibrated along the beam as well as the background. If there is an imbalance between different subgroups (e.g. more background voxels) naive application of RCPS will focus mostly on the majority group (e.g. background) and fail to meet the guarantees for the minority subgroup (e.g. the beam). Calculating RCPS for the different subgroups separately is only a solution if the subgroup is known at test time. However, if this is not the case, as in our RT example, it is not possible to determine the correct RCPS model to use for prediction.

In this paper, we address this problem by proposing a novel calibration algorithm for RCPS that takes into account subgroups and can provide subgroup as well as global guarantees. Our contributions are:

1. The first application of uncertainty quantification in neural network-based dose estimation for RT.
2. A novel algorithm that yields mathematical guarantees for uncertainty intervals for subgroups in the dataset.
3. A quantitative and qualitative evaluation of the algorithm on RT dose prediction on a real-world multi-organ dataset.



**Fig. 1. Overview.** We use the DeepDose network [13] to convert a personalized RT plan defined by input CT scan, beam center distance, radiological depth, source distance map, and beam shape ( $X_1$  to  $X_5$ ) to a voxel-wise dose prediction  $\hat{f}(x)$ . Extending DeepDose by our novel subgroup risk-controlled prediction sets algorithm (SG-RCPS) allows to obtain a calibrated upper and lower bound for the dose ( $\hat{u}(X)$  &  $\hat{l}(X)$ ), as well as the voxel-wise size of the interval ( $\hat{u}(X) - \hat{l}(X)$ ) which serves as final uncertainty measure.

## 2 Methods

The RCPS framework [1,2] ensures that a set-valued predictor  $\mathcal{T}$  maintains a *risk* below a user-specified level  $\alpha$  with a user-defined probability of  $1 - \delta$ . In regression problems such as ours, the prediction set is often a prediction interval characterised by a lower and an upper bound value. The risk  $\mathcal{R}(\mathcal{T}) = \mathbb{E}[L(Y, \mathcal{T}(X))]$  is defined through a loss function  $L$  tailored to the application, which encodes a notion of consequence if the desired property is not fulfilled. In this work, we demonstrate the application of RCPS to DL-based dose prediction and extend the method to ensure risk guarantees for multiple subgroups.

In the following, we will first describe our dose estimation framework (Sec. 2.1). We will then show how heuristic prediction intervals can be obtained using quantile regression (Sec. 2.2). Next, we will discuss how to define the concept of risk, and how the heuristic prediction intervals can be adjusted to control the risk with the desired levels (Sec. 2.3). Lastly, we will describe our novel subgroup RCPS algorithm which allows controlling the risk for multiple subgroups without knowledge of the subgroup membership at test time (Sec. 2.4).

### 2.1 Dose Estimation using DeepDose

In order to construct a voxel-wise dose predictor  $\hat{f}$  we build on the previously proposed DeepDose network [13,18], which is in turn derived from a 3D UNet architecture [5]. DeepDose takes the beam shape, center beam line distance, source distance, CT image and radiological depth as input  $X \in \mathbb{R}^{5 \times W \times H \times D}$  and outputs a predicted dose  $\hat{f}(X)_i \in \mathbb{R}$  for each voxel  $i$  (see Fig.1). The network is trained with a highly accurate Monte Carlo simulation as ground truth  $Y_i$  for each voxel. For notational clarity, we will omit the index  $i$  in the following.

## 2.2 Heuristic Dose Prediction Intervals using Quantile Regression

To extend the DeepDose network by the ability to provide prediction intervals, we adopt a voxel-wise quantile regression approach [1]. Specifically, we add two additional output channels  $\tilde{l}(X)$  and  $\tilde{u}(X)$  to estimate the voxel-wise upper and lower bound, respectively. Similar to [1], we train the two additional network heads using pinball losses, which allow to estimate a specific quantile. For a general feature  $x$ , label  $y$ , and quantile  $\beta$ , the pinball loss is given by

$$\mathcal{L}_\beta(\hat{q}_\beta(x), y) = (y - \hat{q}_\beta(x))\beta \mathbb{1}_{\{y > \hat{q}_\beta(x)\}} + (\hat{q}_\beta(x) - y)(1 - \beta) \mathbb{1}_{\{y \leq \hat{q}_\beta(x)\}}, \quad (1)$$

where  $\hat{q}_\beta(x)$  is the corresponding quantile estimator and  $\mathbb{1}$  denotes the indicator function. We use  $\hat{f}(X) - \tilde{l}(X)$  and  $\hat{f}(X) + \tilde{u}(X)$  for the respective lower and upper quantile estimators. Our overall training objective  $\mathcal{L}$  is comprised of losses for the upper and lower quantiles as well as the standard MSE loss for the point prediction

$$\mathcal{L} = \mathcal{L}_{\alpha/2}(\tilde{l}(X), Y) + \mathcal{L}_{1-\alpha/2}(\tilde{u}(X), Y) + \text{MSE}(Y, \hat{f}(X)), \quad (2)$$

where each loss is only applied to the corresponding head. This gives our framework the ability to not only output a per-voxel dose prediction  $\hat{f}(X)$ , but also a heuristic prediction interval

$$\mathcal{T}(X) = [\hat{f}(X) - \tilde{l}(X), \hat{f}(X) + \tilde{u}(X)]. \quad (3)$$

## 2.3 RCPS for Radiotherapy Dose Estimation

We now show how the RCPS framework [1,2] can be used to obtain dose prediction intervals that are guaranteed to keep a *risk* below a user-specified level. First, we define the risk of  $\mathcal{T}$  as the predicted interval *not* containing the ground truth dose  $Y$

$$\mathcal{R}(\mathcal{T}) = \mathbb{E} [\mathbb{1}_{\{Y \notin \mathcal{T}(X)\}}] = \Pr(Y \notin \mathcal{T}(X)). \quad (4)$$

We then define new lower and upper bounds by scaling them with a non-negative factor  $\hat{\lambda}$

$$\hat{l}(X) = \hat{\lambda}\tilde{l}(X) \quad \text{and} \quad \hat{u}(X) = \hat{\lambda}\tilde{u}(X). \quad (5)$$

RCPS provides a strategy to choose  $\hat{\lambda}$  based on a calibration dataset such that  $\mathcal{R}(\mathcal{T}) \leq \alpha$  with a probability of at least  $1 - \delta$  on future test data under the assumption of exchangeability of the test and calibration sets. We use  $\alpha = \delta = 0.1$  for all experiments, meaning that with a probability of at least 90%, a minimum of 90% of the ground truth dose depositions should be contained in the predicted dose interval.

Since the calibration data is only a random sample of the data distribution,  $\hat{\lambda}$  cannot be chosen by simply minimising the risk on the calibration set. Rather a point-wise *upper confidence bound (UCB)*  $\mathcal{R}^+ : \Pr(\mathcal{R}(\lambda)) \leq \mathcal{R}^+$  must be obtained. This UCB accounts for the calibration sample size and the desired

probability of the guarantee holding on future data. Following [1] we use the Hoeffding bound [9] to define  $\mathcal{R}^+$  as

$$R^+(\lambda) = \frac{1}{nWHD} \sum_{k=1}^n \#\{\mathcal{T}_\lambda(X_k) \notin Y_k\} + \sqrt{\frac{1}{2n} \log \frac{1}{\delta}}. \quad (6)$$

In the original RCPS approach, Bates et al. [2] proposed a greedy optimization algorithm for obtaining the smallest possible prediction interval still fulfilling the desired guarantees. It requires initializing  $\hat{\lambda}$  with a very large value and reducing it until  $\mathcal{R}^+$  falls under the desired risk level, that is,

$$\hat{\lambda} = \min \left\{ \lambda : \hat{R}^+(\lambda) \leq \alpha \right\}. \quad (7)$$

This procedure produces a dose interval predictor

$$\mathcal{T}_{\hat{\lambda}}(X) = [\hat{f}(X) - \hat{\lambda}\tilde{l}(X), \hat{f}(X) + \hat{\lambda}\tilde{u}(X)], \quad (8)$$

for the predicted dose that satisfies the desired risk-properties *on average* for all voxels of the test data distribution. However, it does not offer conditional guarantees for individual data subgroups. For instance, as we will show in Sec. 3 a naive application of RCPS results in the voxels along the radiation beam violating the desired guarantees.

#### 2.4 Risk-Controlled Prediction Sets with Multiple Subgroups

If our dataset comprises  $M$  imbalanced subgroups, using Eq. 7 to obtain  $\hat{\lambda}$  provides guarantees for the overall dataset but not for each individual subgroup. This can lead to a systematic miscalibration of uncertainty intervals for under-represented subgroups. If the subgroups are known at test time, this problem can be addressed by calibrating for each subgroup separately, by using a subgroup-specific parameter  $\hat{\lambda}_z$  during test time. In our RT application, we are interested in calibrated uncertainty quantification in the area of the beam as well as the background, which also receives small dose intensities. While we know the absorbed dose during training, we do not have direct access to this information during testing. Hence, naive application of RCPS yields good calibration overall, but not in the critical area of the beam.

Therefore, we propose an extension of the risk-controlling framework for scenarios where the subgroup membership  $Z$  is unknown at test time. Our proposed extension provides the same guarantees as RCPS for each individual subgroup. That is for every subgroup  $Z$  in our dataset, it holds that, at a risk level of  $\alpha$ , the ground truth is included with a probability of at least  $1 - \delta$ .

In order to achieve the desired guarantees we reformulate the risk conditioned on the subgroups  $Z$  as follows

$$\mathcal{R}(\mathcal{T}) = \Pr(Y \notin \mathcal{T}(X)) = \mathbb{E}_Z[\mathbb{E}_{X|Y|Z}[\mathbb{1}_{\{Y \notin \mathcal{T}(X)\}}]]. \quad (9)$$

This leads to a novel subgroup risk-controlled predictions set (SG-RCPS) procedure which is summarised in Algorithm 1. Similar to the original RCPS algorithm, we start off with a large  $\hat{\lambda}$  that satisfies the risk for each subgroup. We then iteratively reduce the interval size until the first confidence bound of a subgroup no longer satisfies the criterion<sup>1</sup>. A proof that Algorithm 1 leads to the desired subgroup guarantees is presented in Appendix A.

---

**Algorithm 1:** Pseudocode for SG-RSPC
 

---

**Input** : Calibration sets  $(X_k, Y_k)_z$ ,  $k = 1, \dots, n_z$  where  $z = 1, \dots, M$  indicates the subgroup; in our case we have  $M = 3$  for foreground, background and the combined image; risk level  $\alpha$ ; error rate  $\delta$ ; predictor  $\hat{f}$ ; heuristic lower and upper interval predictions  $\hat{l}$  and  $\hat{u}$ ; initial max value  $\lambda_{\max}$ ; step size  $d\lambda > 0$

**Output:** Optimal interval scaling  $\hat{\lambda}$

```

 $\lambda \leftarrow \lambda_{\max}$ 
for  $z \leftarrow 1$  to  $M$  do
  |  $UCB_z \leftarrow 1$ 
end
while  $UCB_1 \leq \alpha \ \& \dots \ \& \ UCB_M \leq \alpha$  do
  |  $\lambda \leftarrow \lambda - d\lambda$ 
  | for  $z \leftarrow 1$  to  $M$  do
  | | for  $k \leftarrow 0$  to  $n_z$  do
  | | |  $L_{k,z} \leftarrow \#\{T_\lambda(X_{k,z}) \notin Y_{k,z}\} / WHD$ 
  | | end
  | |  $UCB_z \leftarrow \frac{1}{n_z} \sum_{k=1}^{n_z} L_{k,z} + \sqrt{\frac{1}{2n_z} \log \frac{1}{\delta}}$ 
  | end
end
 $\hat{\lambda} \leftarrow \lambda + d\lambda$ 

```

---

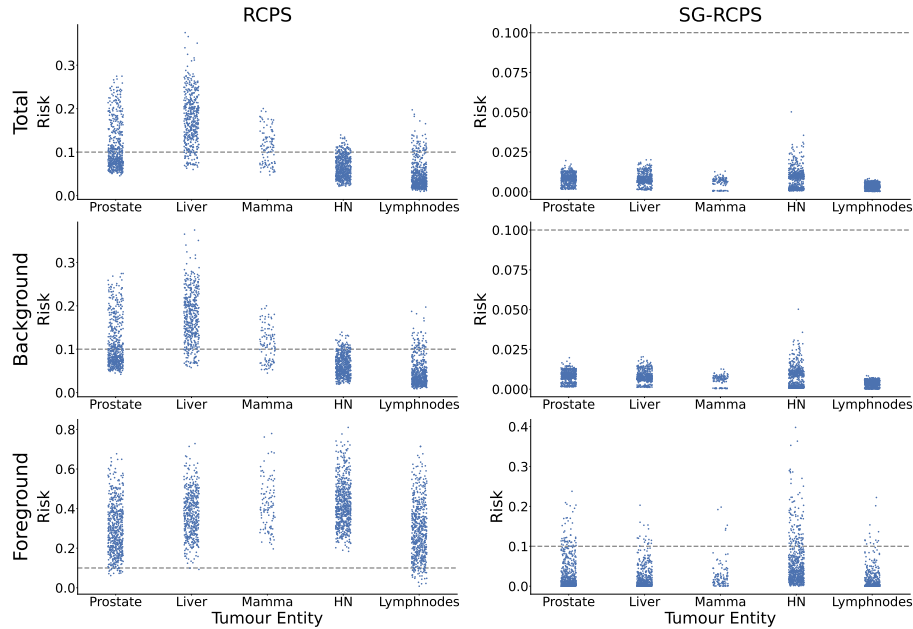
### 3 Experiments and Results

#### 3.1 Dataset

To assess the performance of our model, we trained and tested it on a dataset containing CT data and RT treatment plans of 125 patients obtained from patients at the Department of Radiation Oncology at the University of Tübingen. The study was approved by the institutional review board and all patients gave written informed consent (NCT04172753).

The training dataset comprises four anatomical entities: prostate, liver, breast (mamma), and head and neck (HN). The test and calibration datasets contained data from the same four tumour entities as the training dataset, along with lymph nodes, as an additional out-of-domain (OOD) entity. Testing the neural network on an OOD entity allowed us to assess whether the calibration is able to generalize, which is highly desirable in real-world scenarios. For each patient,

<sup>1</sup> The code is available at <https://github.com/paulkogni/SG-RCPS>



**Fig. 2.** Tumor-specific risks for the original calibration method (left) and our method (right) for the total image (top row), the background radiation (middle row) and foreground radiation (bottom row).

we extracted multi-leaf collimator (MLC) segments, resulting in a total of 6638 segments. For validation, we randomly selected 20 diverse prostate segments from the dataset. For calibration, we randomly selected three random segments from one patient for each entity. More detailed information about the data splits is provided in Appendix B. To train the DeepDose network, each segment was divided into patches of size  $5 \times 32 \times 32 \times 32$ . The ground truth dose estimations were generated using Monte Carlo simulations with the EGSnrc open-source software package [7,10]. For inference, we performed predictions at the patch level and combined them in a sliding window fashion.

### 3.2 Findings

We trained a DeepDose network using training data from all four entities. To evaluate the model’s dose estimation performance we calculated the 3mm/3% gamma pass rate ( $\gamma$ -PR) criterion and observed a  $\gamma$ -PR of 98.9%.

We then compared the two calibration strategies discussed above: RCPS, and our proposed subgroup RCPS (SG-RCPS). We considered three subgroups: the beam foreground and background determined by thresholding the ground truth

**Table 1. Quantitative results.** Empirical risks for RCPS and SG-RCPS averaged over all segments. Controlled risks with  $\alpha \leq 0.1$  for more than  $1 - \delta$  of the cases are highlighted in bold.

Method	Prostate		Liver		Mamma		HN		Lymph.	
	RCPS	SG-RCPS	RCPS	SG-RCPS	RCPS	SG-RCPS	RCPS	SG-RCPS	RCPS	SG-RCPS
Total image	0.415	<b>0.0</b>	0.898	<b>0.0</b>	0.620	<b>0.0</b>	<b>0.089</b>	<b>0.0</b>	<b>0.094</b>	<b>0.0</b>
Backgr. rad.	0.403	<b>0.0</b>	0.896	<b>0.0</b>	0.610	<b>0.0</b>	<b>0.083</b>	<b>0.0</b>	<b>0.094</b>	<b>0.0</b>
Foregr. rad.	0.970	<b>0.084</b>	0.996	<b>0.041</b>	1.0	<b>0.048</b>	1.0	0.176	0.934	<b>0.024</b>

dose (see Fig. 1), and the whole image. We used a target risk level of  $\alpha = 0.1$ , and an error rate of  $\delta = 0.1$ .

Fig. 2 and Tab. 1 show the empirical risk for all segments in the test set grouped by entity. The empirical risk for each segment is defined as proportion of ground truth doses not contained in the predicted interval. Based on our risk settings we expect at most 10% of the segments to fall above the specified risk of 10%.

From Fig. 2 it can be seen that the calibration is dominated by the background class. We found that the normal RCPS algorithm only controlled for the risk in the head & neck, and lymph node entities when considering the total image, and the background only. However, the risk was not controlled to the desired levels in the foreground subgroup (i.e. the beam) for any of the entities. Interestingly, the risks for liver, prostate as well as mamma are not controlled even in the total image. This is likely caused by a mismatch in the proportion of background voxels in the calibration and the test set.


Our proposed SG-RCPS algorithm was able control the risk substantially better for all anatomical areas. There were no dose predictions outside the predicted interval when considering the total image and the background only. We note that because our algorithm estimates a single  $\hat{\lambda}$  that controls the risk for all subgroups jointly, the estimates for the total image and background group were more conservative. When considering the foreground (i.e. the beam) only, we found that all entities except head & neck were risk-controlled with the desired levels. As can be seen in Tab. 1 the empirical risk for the head & neck entity fell slightly short of the desired levels. Notably, the risk for out-of-distribution entity, lymph nodes, was also well controlled. Predicting risk-controlled intervals is particularly important for the foreground, as the the beam is where the most accurate uncertainty estimation is required.

A qualitative example of dose predictions and prediction intervals for RCPS and SG-RCPS is shown in Fig. 3.

## 4 Conclusion

We have proposed subgroup RCPS, an extension of the RCPS algorithm allowing to control risk for multiple subgroups with unknown subgroup membership at test time. We validated our method on a clinical RT dataset comprising





**Fig. 3.** A representative example for a liver tumor visualizing the qualitative differences between the uncertainty intervals for the non-subgroup-specific calibration and our method. The uncertainty intervals provided by our method are significantly wider ( $p < 0.001$ ) than the ones generated by classical RCPS. All values are given in Gray (Gy).

five anatomical entities. Our results demonstrate that in case of imbalances between subgroups our method substantially improves calibration for individual subsets. Specifically, in contrast to regular RCPS, our SG-RCPS approach allows to control the risk for the beam *and* the background thereby increasing safety and trustworthiness in this high-risk application. A potential drawback of our method is that it requires a separate calibration set for each subgroup. Additionally, this method usually yields more conservative prediction intervals. In the future, we will apply this algorithm to datasets that include other under-represented subgroups, such as ethnicity or gender.

**Acknowledgments.** This work was supported by the Excellence Cluster 2064 “Machine Learning — New Perspectives for Science”, project number 390727645 and received funding from the German Research Council under DFG Grant No. ZI 736/2-1. The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting Paul Fischer. The authors acknowledge support by the state of Baden-Württemberg through bwHPC (INST39/963-1 FUGG bwForCluster NEMO) and through the Research and Training Network "AI4MedBW".

**Disclosure of Interests.** The authors have no competing interests to declare.

## References

1. Angelopoulos, A.N., Kohli, A.P., Bates, S., Jordan, M., Malik, J., Alshaabi, T., Upadhyayula, S., Romano, Y.: Image-to-image regression with distribution-free uncertainty quantification and applications in imaging. In: International Conference on Machine Learning. pp. 717–730. PMLR (2022)
2. Bates, S., Angelopoulos, A., Lei, L., Malik, J., Jordan, M.: Distribution-free, risk-controlling prediction sets. *Journal of the ACM (JACM)* **68**(6), 1–34 (2021)

3. Baumgartner, C.F., Tezcan, K.C., Chaitanya, K., Hötter, A.M., Muehlematter, U.J., Schawkat, K., Becker, A.S., Donati, O., Konukoglu, E.: Phiseg: Capturing uncertainty in medical image segmentation. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part II 22. pp. 119–127. Springer (2019)
4. Bol, G., Hissoiny, S., Lagendijk, J., Raaymakers, B.: Fast online monte carlo-based imrt planning for the mri linear accelerator. *Physics in Medicine & Biology* **57**(5), 1375 (2012)
5. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3d u-net: learning dense volumetric segmentation from sparse annotation. In: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2016: 19th International Conference, Athens, Greece, October 17–21, 2016, Proceedings, Part II 19. pp. 424–432. Springer (2016)
6. Fischer, P., Thomas, K., Baumgartner, C.F.: Uncertainty estimation and propagation in accelerated mri reconstruction. In: International Workshop on Uncertainty for Safe Utilization of Machine Learning in Medical Imaging. pp. 84–94. Springer (2023)
7. Friedel, M., Nachbar, M., Mönnich, D., Dohm, O., Thorwarth, D.: Development and validation of a 1.5 t mr-linac full accelerator head and cryostat model for monte carlo dose simulations. *Medical Physics* **46**(11), 5304–5313 (2019)
8. Hall, W., Paulson, E., van der Heide, U., Fuller, C., Raaymakers, B., Lagendijk, J., Li, X., Jaffray, D., Dawson, L., Erickson, B., Verheij, M., Harrington, K., Sahgal, A., Lee, P., Parikh, P., Bassetti, M., Robinson, C., Minsky, B., Choudhury, A., Tersteeg, R., Schultz, C.: Mr linac atlantic consortium and the viewray c2t2 research consortium. the transformation of radiation oncology using real-time magnetic resonance guidance: A review. *European Journal of Cancer* **122**, 42–52 (2019)
9. Hoeffding, W.: Probability inequalities for sums of bounded random variables. The collected works of Wassily Hoeffding pp. 409–426 (1994)
10. Kawrakow, I., Rogers, D., Mainegra-Hing, E., Tessier, F., Townson, R., Walters, B.: Egsnrc toolkit for monte carlo simulation of ionizing radiation transport. doi:10.4224/40001303 (2000)
11. Kendall, A., Badrinarayanan, V., Cipolla, R.: Bayesian segnet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding. arXiv preprint arXiv:1511.02680 (2015)
12. Kohl, S., Romera-Paredes, B., Meyer, C., De Fauw, J., Ledsam, J.R., Maier-Hein, K., Eslami, S., Jimenez Rezende, D., Ronneberger, O.: A probabilistic u-net for segmentation of ambiguous images. *Advances in neural information processing systems* **31** (2018)
13. Kontaxis, C., Bol, G., Lagendijk, J., Raaymakers, B.: Deepdose: towards a fast dose calculation engine for radiation therapy using deep learning. *Physics in Medicine & Biology* **65**(7), 075013 (2020)
14. Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems* **30** (2017)
15. Martinot, S., Bus, N., Vakalopoulou, M., Robert, C., Deutsch, E., Paragios, N.: High-particle simulation of monte-carlo dose distribution with 3d convlstm. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part IV 24. pp. 499–508. Springer (2021)

16. Neishabouri, A., Wahl, N., Mairani, A., Köthe, U., Bangert, M.: Long short-term memory networks for proton dose calculation in highly heterogeneous tissues. *Medical physics* **48**(4), 1893–1908 (2021)
17. Randall, J.W., Rammohan, N., Das, I.J., Yadav, P.: Towards accurate and precise image-guided radiotherapy: clinical applications of the mr-linac. *Journal of Clinical Medicine* **11**(14), 4044 (2022)
18. Tsekas, G., Bol, G., Raaymakers, B., Kontaxis, C.: Deepdose: a robust deep learning-based dose engine for abdominal tumours in a 1.5 t mri radiotherapy system. *Physics in Medicine & Biology* **66**(6), 065017 (2021)