

PATIENT-SPECIFIC FINETUNING OF DEEP LEARNING MODELS FOR ADAPTIVE RADIOTHERAPY IN PROSTATE CT

Mohamed S. Elmahdy^{†*}, Tanuj Ahuja^{†◦*}, U. A. van der Heide^{‡*}, and Marius Staring^{†‡}

[†]Department of Radiology, Leiden University Medical Center, Leiden, The Netherlands.

[◦] Computer Science and Engineering, Guru Gobind Singh Indraprastha University, Delhi, India.

[‡] Department of Radiation Oncology, Leiden University Medical Center, Leiden, the Netherlands.

^{*} Department of Radiation Oncology, the Netherlands Cancer Institute, Amsterdam, the Netherlands.

ABSTRACT

Contouring of the target volume and Organs-At-Risk (OARs) is a crucial step in radiotherapy treatment planning. In an adaptive radiotherapy setting, updated contours need to be generated based on daily imaging. In this work, we leverage personalized anatomical knowledge accumulated over the treatment sessions, to improve the segmentation accuracy of a pre-trained Convolution Neural Network (CNN), for a specific patient. We investigate a transfer learning approach, fine-tuning the baseline CNN model to a specific patient, based on imaging acquired in earlier treatment fractions. The baseline CNN model is trained on a prostate CT dataset from one hospital of 379 patients. This model is then fine-tuned and tested on an independent dataset of another hospital of 18 patients, each having 7 to 10 daily CT scans. For the prostate, seminal vesicles, bladder and rectum, the model fine-tuned on each specific patient achieved a Mean Surface Distance (MSD) of 1.64 ± 0.43 mm, 2.38 ± 2.76 mm, 2.30 ± 0.96 mm, and 1.24 ± 0.89 mm, respectively, which was significantly better than the baseline model. The proposed personalized model adaptation is therefore very promising for clinical implementation in the context of adaptive radiotherapy of prostate cancer.

Index Terms— transfer learning, segmentation, prostate cancer, adaptive radiotherapy, organs at risk.

1. INTRODUCTION

Prostate cancer is the second leading cause of cancer death in American men, behind lung cancer. About 1 man in 41 will die of prostate cancer. The American Cancer Society estimates about 174,650 new cases of prostate cancer and about 31,620 deaths from prostate cancer in the United States for 2019 only [1]. Radiation Therapy (RT) is one of the treatment options for prostate cancer that uses a highly localized dose distribution to kill cancer cells. RT dose is usually fragmented over 4 to 8 weeks resulting in 20 to 40 daily fractions

[2]. Since the dose is delivered in several sessions, variations in the size and shape of the target area and Organs-At-Risk (OARs) is bound to take place. Often these changes are due to organ deformation such as variations in the rectum and bladder filling [3]. Continuing RT using the initial planning CT scan and the corresponding dose distribution despite these changes may lead to under-dosing of the target area or overdosing the OARs.

Adaptive radiation therapy aims to adapt the treatment dose distribution to the daily anatomy, in order to achieve safe dose escalation [4] or smaller margins. This adaptation can be done by re-imaging, re-contouring, and re-planning at every session. However, manual re-contouring takes a lot of time and consequently new variations in the size and the shape of the OARs could be introduced meanwhile. Therefore, there is a need to re-contour the CT images accurately and in a time-efficient manner, i.e. (semi-)automatically.

Literature of prostate segmentation can be broadly categorized into two approaches, non-learning and learning approaches. For non-learning approaches, Qiao *et al.* used the open source `elastix` software to apply deformable image registration and subsequently propagate the contours from the planning to the daily scans [5]. For learning-based approaches, Elmahdy *et al.* proposed a hybrid learning and iterative approach, where they used a CNN network to segment the bladder and explicitly feed it to the registration model as prior knowledge on the underlying anatomy [6]. A substantial improvement was observed compared to the results reported in [5]. Recently, deep learning and specifically convolutional neural networks (CNNs) are being used to automatically segment OARs. Deep learning has the power to extract information from the data rather than depending on hand-crafted features. Milletari *et al.* developed an end-to-end CNN network to segment 3D medical images and reported a Dice Similarity Coefficient (DSC) of 0.869 ± 0.033 on the PROMISE 2012 challenge dataset [7] (MR data). Tong *et al.* proposed the use of a shape representation model, to learn highly representative shape characteristics of OARs and help the final segmentation of a Fully Convolution

*denotes equal contribution

Neural Network (FCNN) [8].

To the best of our knowledge, no deep learning based approach has been used to adapt a CNN model to a particular patient at a certain time point. In this study, we propose to adapt a baseline CNN model as the patient goes through their RT treatment. Thus, instead of depending on a static deep learning model, we accumulate the knowledge over successive sessions for the same patient. This accumulated knowledge is then used to encourage the model into predicting a more personalized segmentation. Moreover, we study the performance of the fine-tuned network when more and more imaging becomes available over the course of treatment.

2. MATERIALS AND METHODS

2.1. Dataset

This study includes two different datasets from two different institutes and scanners, for patients who underwent intensity-modulated RT for prostate cancer. First, a dataset from Leiden University Medical Center (LUMC), Netherlands, has a total of 379 patients with one CT scan each. The scans were acquired using a Toshiba scanner, having 68 to 240 slices with a voxel size of approximately $1.0 \times 1.0 \times 3.0$ mm. The second dataset is from Haukeland University Hospital, Norway, and includes 18 patients with 8-11 CT scans each corresponding to multiple fractions. These scans were acquired using a GE scanner, having 90 to 180 slices with a voxel size of approximately $0.9 \times 0.9 \times 3.0$ mm. The target structures (prostate and seminal vesicles) as well as OARs (bladder and rectum) were manually delineated by oncologists. For the LUMC data informed consent was waived by the local Medical Ethical Committee, while for the Haukeland data informed consent was given by all included patients.

2.2. Baseline Segmentation CNN Model

In this study, for the baseline CNN model we adopted the network introduced in Elmahdy *et al.* [6], which has a straightforward architecture, but increased the number of the output labels from 2 to 5 for multi-organ segmentation. Similar to the standard U-net, it has an encoder and decoder path with four resolutions. Each encoder block consists of $3 \times 3 \times 3$ convolutions with a stride of two in each dimension, followed by a rectified linear unit (ReLU). In the synthesis path, each block of the decoder consists of $3 \times 3 \times 3$ convolutions with a stride of one in each dimension, followed by a ReLU and an upsampling of $2 \times 2 \times 2$. For a better localization, high resolution features from the contracting path are combined with the upsampled output, using skip connections. This aids to recover fine-grained details that were lost during the compression phase. The upsampling path is then followed by one Fully Connected (FC) layer, similar to the original paper, or three FC layers. These models are denoted by **base^a** and **base^b**, respectively. A softmax layer finalizes the models.

2.3. Patient-specific Daily Model Adaptation

In order to adapt the model to the patient-specific anatomy, we update the model with the anatomical knowledge from previous treatment sessions for the same patient. We thus personalize the pre-trained baseline CNN and utilize the available imaging of the patient. Since the new data (daily scans) as well as the task (segmentation) are the same as the base models, we re-use the base model weights as initialization to the model adaptation, as in a transfer learning approach. The weights in the FC layers are then fine-tuned using the imaging from previous fractions, while all other weights remain fixed. For each patient, at the first fraction we fine-tune the base model based on the planning scan. At subsequent fractions we start with the previous fine-tuned model and continue fine-tuning based on the scan and segmentation pair of the previous fraction. Here we consider that the segmentation can be manually corrected to clinical quality before the start of the current fraction. To be precise, let M_0 be the base model and M_j the fine-tuned model at fraction j , Φ the model adaptation process, (I_j, S_j) the image and segmentation pair at fraction j . Model adaptation is then performed as follows:

$$M_j = \Phi(M_{j-1}; I_{j-1}, S_{j-1}). \quad (1)$$

The segmentation prediction of the current fraction is then given by $M_j(I_j)$. The base models are trained on the LUMC dataset. The networks are adapted to the Haukeland dataset, which has multiple follow-up scans per patient, mimicking an adaptive RT scenario.

3. EXPERIMENTS AND RESULTS

3.1. Evaluation Measures and Implementation

The Dice Similarity Coefficient (DSC), Mean Surface Distance (MSD), and 95% Hausdorff Distance (HD) are used to evaluate the error between the ground truth delineations and the segmentations predicted by the models. These are evaluated for the prostate, seminal vesicles, bladder, and rectum. A Wilcoxon signed rank test at $p = 0.05$ is used to assess statistical significance. Evaluation results are shown for the Haukeland data only, as only this dataset has follow-up scans.

We used the Tensorflow library for the implementation of the 3D CNNs. Out of the 379 LUMC patients, 70% (259 cases) were used for training the baseline models and the remaining 30% (111 cases) were used for validation. A total of 1000 patches of size $128 \times 128 \times 128$ were extracted from each CT volume, sampling the classes with a uniform distribution to handle class imbalance. All scans were resampled to a fixed voxel size of $1 \times 1 \times 2$ mm in order to handle variations in voxel size. The Dice Similarity Coefficient (DSC) is deployed as a cost function and the network was trained using the state-of-the-art Rectified Adam (RAdam) optimizer introduced in [9], with a fixed learning rate of 10^{-4} . The network

Table 1. DSC value of the target volumes and OARs.

	#Itr	Prostate ^{†‡}	SV ^{†‡}	Rectum ^{†‡}	Bladder ^{†‡}
		$\mu \pm \sigma$	$\mu \pm \sigma$	$\mu \pm \sigma$	$\mu \pm \sigma$
base ^a		0.76 ± 0.06	0.55 ± 0.19	0.79 ± 0.08	0.86 ± 0.14
proposed ^a	500	0.78 ± 0.06	0.61 ± 0.17	0.79 ± 0.07	0.89 ± 0.09
proposed ^a	2000	0.80 ± 0.05	0.63 ± 0.17	0.80 ± 0.07	0.90 ± 0.07
proposed ^a	5000	0.82 ± 0.05	0.65 ± 0.16	0.80 ± 0.07	0.91 ± 0.07
base ^b		0.77 ± 0.07	0.46 ± 0.20	0.78 ± 0.08	0.87 ± 0.08
proposed ^b	500	0.80 ± 0.05	0.50 ± 0.18	0.81 ± 0.06	0.91 ± 0.07
proposed ^b	2000	0.83 ± 0.04	0.56 ± 0.15	0.83 ± 0.05	0.91 ± 0.06
proposed ^b	5000	0.83 ± 0.04	0.55 ± 0.17	0.83 ± 0.05	0.91 ± 0.06

was trained for 1,000,000 iterations with a batch size of 4 using an NVIDIA Titan Xp GPU with 12 GB of memory. For model adaptation, a total of 2000 patches were extracted from a single CT volume with the same patch size, cost function, and learning rate as the base model. For investigating the effect of the number of adaptation iterations on the network performance, we varied this parameter between 500, 2000, and 5000 iterations. This adaptation can be performed offline, and took less than an hour per fraction at 2000 iterations.

3.2. Results

Tables 1, 2 and 3 show the quantitative results of all the models, averaged over all the treatment fractions. Adapted models are denoted by proposed^a and proposed^b. Here, † and ‡ represent a statistically significant difference between base^a vs proposed^a, and base^b vs proposed^b at 2000 iterations, respectively. In terms of DSC, MSD, and 95% HD, the larger base^b model is slightly better than base^a, except for the seminal vesicles. This pattern is also visible for the proposed models. All proposed models outperform the baseline models, sometimes by a margin. Furthermore, adapting the network for 5000 iterations seems as good as 2000 iterations as shown in the tables. Predicting a segmentation took ~0.6 seconds.

To investigate the relation between the amount of patient-specific data available for fine-tuning and segmentation performance, we applied the models $M_j, j \in \{0, 5\}$ on scan I_6 from each patient. Boxplots of the results are shown in Figure 1 for the MSD. Segmentation of most structures tend to improve when more data is available, while such a trend is not visible for the rectum. Some example segmentation results are given in Figure 2 for the models base^b and proposed^b at 2000 iterations.

4. DISCUSSION AND CONCLUSION

In this study we investigated the hypothesis of personalizing the automatic contouring by deep learning, of the target organs and OARs for adaptive radiotherapy of prostate cancer. Unlike traditional CNN networks, the proposed model adaptation strategy predicts the segmentation of the daily CT based on the accumulated anatomical knowledge from the previous

Table 2. MSD value of the target volumes and OARs.

	#Itr	Prostate ^{†‡}	SV ^{†‡}	Rectum ^{†‡}	Bladder ^{†‡}
		$\mu \pm \sigma$	$\mu \pm \sigma$	$\mu \pm \sigma$	$\mu \pm \sigma$
base ^a		2.70 ± 0.88	3.78 ± 4.76	2.77 ± 1.23	2.66 ± 9.46
proposed ^a	500	2.38 ± 0.64	2.58 ± 2.67	2.76 ± 1.05	1.62 ± 1.28
proposed ^a	2000	2.03 ± 0.55	2.38 ± 2.76	2.66 ± 1.03	1.48 ± 1.17
proposed ^a	5000	1.78 ± 0.52	2.41 ± 3.17	2.59 ± 1.00	1.39 ± 1.12
base ^b		2.29 ± 0.70	5.08 ± 5.33	2.78 ± 1.17	1.83 ± 0.96
proposed ^b	500	1.95 ± 0.53	5.21 ± 16.71	2.46 ± 1.05	1.40 ± 1.03
proposed ^b	2000	1.64 ± 0.43	3.18 ± 3.34	2.30 ± 0.96	1.24 ± 0.89
proposed ^b	5000	1.61 ± 0.41	4.67 ± 16.96	2.30 ± 0.96	1.25 ± 0.91

Table 3. 95% HD value of the target volumes and OARs.

	#Itr	Prostate ^{†‡}	SV ^{†‡}	Rectum	Bladder ^{†‡}
		$\mu \pm \sigma$	$\mu \pm \sigma$	$\mu \pm \sigma$	$\mu \pm \sigma$
base ^a		9.1 ± 3.7	12.2 ± 10.7	14.4 ± 7.0	8.7 ± 12.9
proposed ^a	500	7.9 ± 2.8	9.2 ± 8.4	14.7 ± 6.7	7.9 ± 6.7
proposed ^a	2000	6.4 ± 2.1	8.6 ± 9.0	14.3 ± 6.7	7.2 ± 6.5
proposed ^a	5000	5.9 ± 2.4	8.8 ± 10.6	14.0 ± 6.6	6.5 ± 6.1
base ^b		6.9 ± 1.9	15.5 ± 9.7	13.9 ± 9.5	7.0 ± 4.5
proposed ^b	500	6.1 ± 1.7	14.6 ± 20.4	13.5 ± 9.1	6.2 ± 5.8
proposed ^b	2000	5.2 ± 1.4	11.9 ± 9.3	13.5 ± 8.4	5.6 ± 5.3
proposed ^b	5000	5.2 ± 1.4	13.9 ± 20.9	13.3 ± 8.5	5.7 ± 5.6

scans of the same patient. We demonstrate that adapting the model to a specific patient anatomy boosted the performance of the network. Furthermore, increasing the network capacity by adding more fully connected layers was beneficial. Moreover, from Fig. 1 we observe that the rectum delineation did not benefit from fine-tuning, which may be explained by the normal variation in rectum filling over treatment fractions. Such variations in the rectum and seminal vesicles could be further investigated in future work by more sophisticated networks and model generalization methods.

To conclude, we proposed an adaptive training mechanism for personalized automatic contour segmentation for prostate cancer. This adaptation showed potential for improving the prediction of the daily anatomy based on personalized imaging accumulated over fractions. Since the segmentation time is less than a second, the adaptation mechanism is therefore very promising for clinical implementation in the context of adaptive radiotherapy of prostate cancer.

Acknowledgements. This study was financially supported by Varian Medical Systems and ZonMw, grant number 104003012. The Haukeland dataset was provided by oncologist Svein Inge and physicist Liv Bolstad; they are gratefully acknowledged.

5. REFERENCES

- [1] National Cancer Institute. Cancer Stat Facts: Prostate Cancer. <https://seer.cancer.gov/statfacts/html/prost.html>, 2019.




Fig. 1. The effect of model adaptation over sessions in terms of MSD (mm). B is base model, S_i after session i .




Fig. 2. Three example results for the $base^b$ model (top row) and the $proposed^b$ at 2000 itr model (bottom row). The red, yellow, blue and green contours represent the bladder, prostate, rectum, and seminal vesicles, respectively. The solid line indicates manual contouring, while the dotted line the automatic prediction.

[2] Baoshe Zhang, Sung-Woo Lee, et al. Action levels on dose and anatomic variation for adaptive radiation therapy using daily offline plan evaluation: Preliminary results. *Practical radiation oncology*, 9(1):49–54, 2019.

[3] Jan-Jakob Sonke, Marianne Aznar, et al. Adaptive radiotherapy for anatomical changes. In *Seminars in radiation oncology*, volume 29, pages 245–257. Elsevier, 2019.

[4] Di Yan, Frank Vicini, et al. Adaptive radiation therapy. *Physics in Medicine & Biology*, 42(1):123, 1997.

[5] Qiao Yuchuan. *Fast Optimization Methods For Image Registration In Adaptive Radiation Therapy*. PhD dissertation, Leiden University Medical Center, 2017.

[6] Mohamed S. Elmahdy, Thyrsa Jagt, et al. Robust contour propagation using deep learning and image registration for online adaptive proton therapy of prostate cancer. *Medical physics*, 2019.

[7] Fausto Milletari, Nassir Navab, et al. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *Fourth International Conference on 3D Vision (3DV)*, pages 565–571. IEEE, 2016.

[8] Nuo Tong, Shuiping Gou, et al. Fully automatic multi-organ segmentation for head and neck cancer radiotherapy using shape representation model constrained fully convolutional neural networks. *Medical physics*, 45(10):4558–4567, 2018.

[9] Sashank J. Reddi, Satyen Kale, et al. On the convergence of adam and beyond. *arXiv preprint arXiv:1904.09237*, 2019.