

Stefan Gerlach*, Christoph Fürweger, Theresa Hofmann and Alexander Schlaefer

Multicriterial CNN based beam generation for robotic radiosurgery of the prostate

<https://doi.org/10.1515/cdbme-2020-0030>

Abstract: Although robotic radiosurgery offers a flexible arrangement of treatment beams, generating treatment plans is computationally challenging and a time consuming process for the planner. Furthermore, different clinical goals have to be considered during planning and generally different sets of beams correspond to different clinical goals. Typically, candidate beams sampled from a randomized heuristic form the basis for treatment planning. We propose a new approach to generate candidate beams based on deep learning using radiological features as well as the desired constraints. We demonstrate that candidate beams generated for specific clinical goals can improve treatment plan quality. Furthermore, we compare two approaches to include information about constraints in the prediction. Our results show that CNN generated beams can improve treatment plan quality for different clinical goals, increasing coverage from 91.2 to 96.8% for 3,000 candidate beams on average. When including the clinical goal in the training, coverage is improved by 1.1% points.

Keywords: machine learning; robotic radiosurgery; treatment planning.

Problem

The key idea of radiation therapy is treating lesions with overlapping X-ray beams from multiple directions. One system used in clinical practice is the CyberKnife system [1] where a linear accelerator is mounted on a robotic arm. This flexible

beam placement can allow for high coverage of the target while sparing healthy tissue. However, generating a treatment plan which satisfies all clinical goals can be computationally challenging and a time consuming process for the planner. Typically, a subset of candidate beams sampled from a randomized heuristic is selected and weighted by optimization with respect to constraints, e.g., on the minimum and maximum dose in the volumes of interest (VOIs). The weights correspond to the activation time of the beams. Generating a clinically acceptable treatment plan is an iterative process in which the planner adjusts the constraints and reoptimizes the beam weights. Here, multiple clinical goals, e.g., minimizing the treatment time or maximizing the dose to the target are addressed sequentially. Therefore, fast optimization is desirable.

Deep learning based methods have been studied for multiple tasks in medical imaging and radiation therapy including segmentation [2] and classification [3]. Other knowledge based methods are employed to optimize beam related parameters in intensity modulated radiation therapy (IMRT) [4, 5] or to optimize beam orientations, positions, shapes, and weights directly (direct aperture optimization). The latter either requires solving a computationally demanding mixed integer problem [6] or they combine the dose in the target, dose constraints, and apertures in the objective function [7], not allowing setting hard constraints on the doses of critical organ structures. Furthermore, the effect of clinical goals on the generation of candidate beams has not been studied sufficiently for robotic radiosurgery. However, it has been observed that different beam parameters influence different clinical goals [8]. E.g., beams that cover a larger volume are more common in treatment plans focusing on the clinical goal of reducing the total activation time of all beams.

In this paper, we investigate the multicriterial aspect of treatment planning on the CNN based candidate beam generation. We extend an earlier approach [9] and present different setups to train CNNs using radiological features for predicting each beam's influence on the dose with various clinical goals. We use these predictions to select new candidate beams, improving plan quality while using fewer candidate beams. We train and evaluate the CNNs on different subsets of 50 patients previously treated for prostate tumor.

***Corresponding author: Stefan Gerlach**, Institute of Medical Technology, Hamburg University of Technology, Hamburg, Germany, e-mail: stefan.gerlach@tuhh.de

Christoph Fürweger, Europäisches Cyberknife Zentrum München-Großhadern, Munich, Germany; Department of Stereotaxy and Functional Neurosurgery, University of Cologne, Faculty of Medicine and University Hospital Cologne, Cologne, Germany

Theresa Hofmann, Europäisches Cyberknife Zentrum München-Großhadern, Munich, Germany

Alexander Schlaefer, Institute of Medical Technology, Hamburg University of Technology, Hamburg, Germany

Methods

Data set

Generating treatment plans typically involves several distinct steps. First, the delineated VOIs, i.e., the planning target volume (PTV) and organs at risk (OARs), are discretized into voxels. Then, the dose coefficients, i.e., the dose each beam delivers to each voxel for a given beam weight, are calculated for randomly sampled candidate beams. Finally, the inverse planning problem is solved with respect to constraints on dose and total beam weight or beam monitor units resulting in a weighted subset of beams.

Our CNNs are trained on treatment plans generated by optimizing the coverage using a set of 6,000 candidate beams. The coverage represents one clinical goal and describes the fraction of the PTV that receives at least the prescribed dose. We use our in-house planning software to maximize the coverage indirectly by minimizing the missing dose of voxels receiving less than the prescribed dose [10].

To generate reference treatment plans, we adopt a 5-fraction protocol [11] with a prescribed dose of 36.25 Gy for the prostate and hard upper dose constraints on the PTV (prostate) and OARs (bladder, rectum) of 40.25 and 36 Gy, respectively. Total and per beam monitor units are constrained to 40,000 and 300 MU, respectively. We also introduce artificial shell structures at 3 and 9 mm distance around the prostate to control dose in normal tissue. We tune the shell constraints to achieve roughly 95% coverage for every patient.

To generate treatment plans with different clinical goals, we vary the maximum constraints on either PTV, total monitor units, or shells and tune the remaining constraints to again achieve roughly 95% coverage. Note that treatment plans are still optimized by maximizing the coverage.

Feature generation

We train the CNNs to predict the influence on the dose for each beam independently. As features for a beam we concatenate gray-scale images of projections described in the following on a plane perpendicular to the line from beam origin to PTV centroid. We construct the first image as the intersection of the beam’s central axis, extended by the effective radius of the beam at the source-plane distance. This represents the volume influenced by the beam. Further images represent the VOIs in relation to the beam as shown in Figure 1. Each image is constructed as the projection of the VOIs. To also encode the volume in the projection, as well as the tissue density, we create two images per VOI encoding the minimum and maximum radiological depth, calculated from the computed tomography scan of the patient.

To include information about the clinical goal, we input the normalized constraints as an additional feature. We normalize each constraint by the respective constraint from the treatment plan of the same patient with the reference clinical goal. Since shell constraints of treatment plans for all clinical goals are constructed such that they have roughly 95% coverage, the constraints on the shells are implicitly defined and are not included as a feature.

CNN model setup

Figure 2 shows how we adapt the DenseNet-121 to predict the beam weight normalized by the maximum allowed weight per beam.

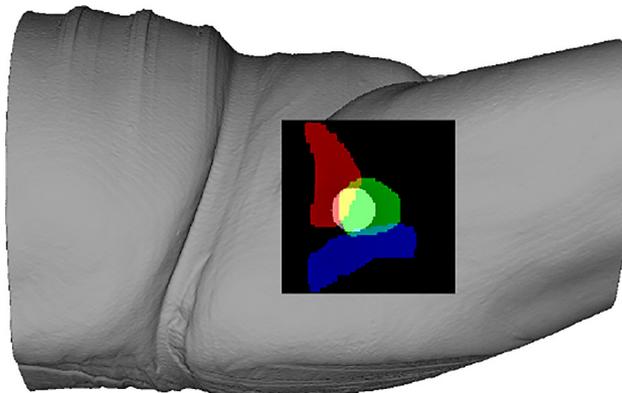


Figure 1: Partial image feature for one beam relative to the patient. Minimum radiological depth of PTV, rectum, and bladder in the channels green, blue, and red respectively. The beam is superimposed in white. Maximum radiological depth features are not shown here.

Furthermore, we study two different ways to combine image features and constraint features. First, we train individual models for every set of constraints and combine the predictions with the respective normalized constraints in a separate fully-connected layer. Second, we concatenate the constraint features with the feature vector obtained from the convolutional layers and train the model end-to-end.

We use the data of nine patients for hyper-parameter optimization with 3-fold cross-validation (CV), where each fold contains three patients. Here, we found a learning rate of $l_r = 10^{-3}$ halved every five epochs and a batch size of $b=32$ to be optimal. In each CV iteration, we train the models for 15 epochs using the Adam optimizer. Data of remaining 41 patients is used for evaluation using 3-fold CV with 13 or 14 patients in each fold.

Since there is an imbalance of $i \approx 19$ times more unweighted than weighted beams we consider unweighted beams with a probability of $p=1/i$ during training to improve training time and convergence. Furthermore, since the treatment plans used for training slightly differ in coverage from the targeted ($c_{des} = 95\%$), we weight each training example by $w = \frac{c_{des}}{c_p}$ for unweighted beams and the inverse for weighted beams with actual treatment plan coverage c_p . This increases the training weight for unweighted beams in low coverage treatment plans ($c_p < 95\%$) and also increases the training weight for weighted beams in high coverage treatment plans ($c_p > 95\%$).

For inference, we consider the predicted normalized beam weight as the probability of accepting the beam into the set of candidate beams. Therefore, we generate random beams and evaluate them using the CNNs until the set of candidate beams contains the desired number of candidate beams. Then, treatment plan generation continues regularly by optimizing beam weights.

Results and discussion

We evaluate the coverage of treatment plans with clinical goals of different homogeneity of dose in the PTV by varying PTV dose constraints, different monitor units resulting in different treatment time, and different dose conformity to the PTV by varying shell dose constraints. Figure 3 shows

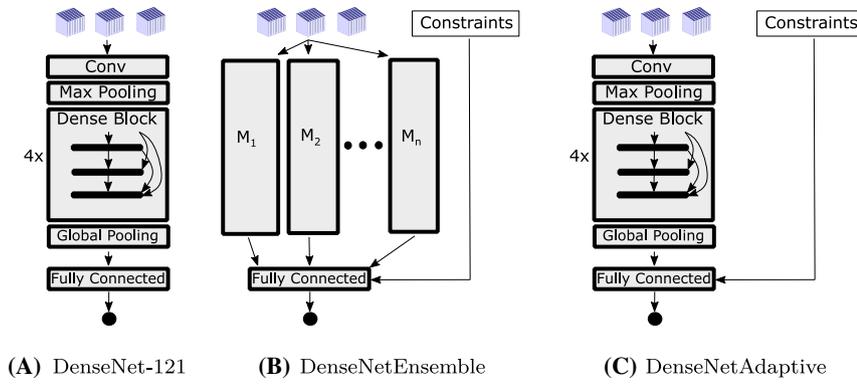


Figure 2: The three CNN architectures we employ. M_1 to M_n refer to DenseNet-121 trained on different clinical goals.

the coverage after optimization of 3,000 randomized and CNN generated candidate beams. Average coverage can be improved from 91.2 to 96.8% over all evaluated clinical goals. Here, the CNN was trained on the same constraints it was evaluated on. Note that constraints are tuned to achieve 95% coverage with 6,000 randomized beams. Figure 4 shows a comparison of the different CNN architectures. DenseNetEnsemble and DenseNetAdaptive are trained on treatment plans with all constraints, DenseNetRef is

trained on reference constraints, and DenseNetFit is the same as in Figure 3C. The latter two use DenseNet-121. Coverage can be improved from 94.9 to 96.0% for DenseNetFit and to 95.2% for DenseNetAdaptive and decreases for DenseNetEnsemble to 94.8% on average. Differences to DenseNetRef are statistically significant with respect to the Wilcoxon rank sum test (significance level $\alpha=0.05$) for DenseNetFit ($p=1.1 \cdot 10^{-8}$) and not significant for DenseNetEnsemble ($p=0.18$) and DenseNetAdaptive ($p=0.26$). This

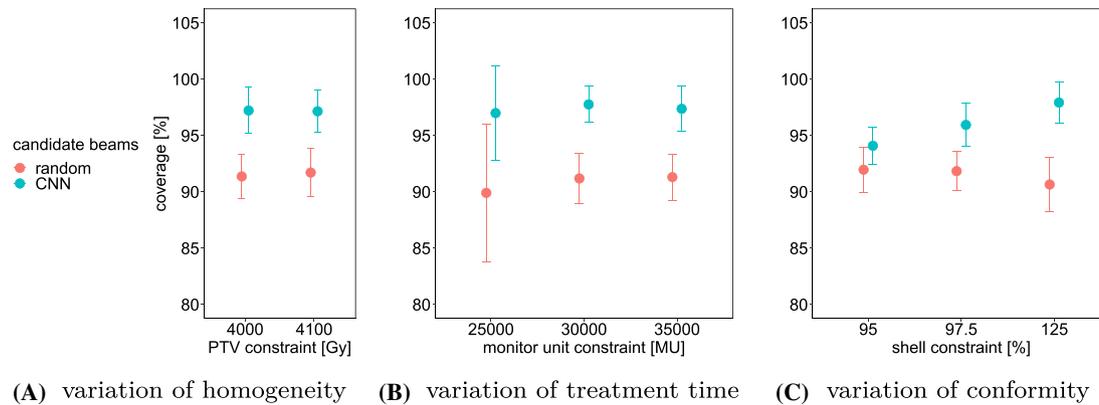


Figure 3: Coverage mean and standard deviation for 3,000 random and CNN generated candidate beams and different clinical goals. CNNs are trained on the target constraints but different patients using DenseNet-121. Shell constraint relative to reference.

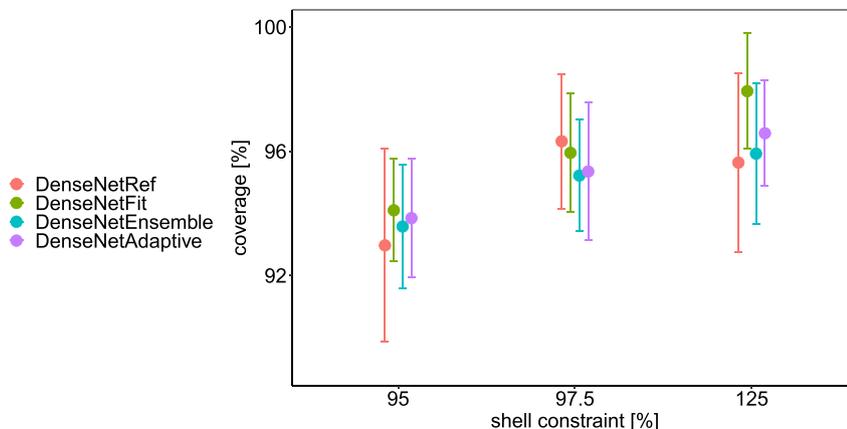


Figure 4: Coverage mean and standard deviation for 3,000 CNN generated beams. DenseNetRef and DenseNetFit use the DenseNet-121 architecture, trained on the reference constraints and target constraints, respectively. Shell constraint relative to reference.

suggests that the information drawn from the constraint features is not yet optimally exploited by the CNNs that fuse information. However, including training data from more constraint sets could improve results further.

Conclusion

We have shown that improving coverage with CNN generated candidate beams is feasible for different clinical goals. Including the clinical goal in the training can further improve treatment plan quality by a small amount. However, the results for CNNs trained on different constraints suggests that further improvements in architecture can be made.

Research funding: This work was partially funded by Deutsche Forschungsgemeinschaft (grant SCHL 1844/3-1).

Author contributions: All authors have accepted responsibility for the entire content of this manuscript and approved its submission.

Competing interests: Authors state no conflict of interest.

Informed consent: Informed consent has been obtained from all individuals included in this study.

Ethical approval: This article is based on anonymized treatment planning data and does not contain any studies with human participants or animals performed by the authors.

References

1. Kilby W, Michael N, Dooley JR, Maurer CR, Sayeh S. 2 – A technical overview of the CyberKnife system. In: Abedin-Nasab M.H., editor Handbook of robotic and image-guided surgery. Amsterdam, Netherlands: Elsevier; 2020:15–38 pp.
2. Cha KH, Hadjiiski LM, Samala RK, Chan HP, Cohan RH, Caoili EM, et al. Bladder cancer segmentation in CT for treatment response assessment: application of deep-learning convolution neural network - a pilot study. *Tomography* 2016;2:421–9.
3. Kajikawa T, Kadoya N, Ito K, Takayama Y, Chiba T, Tomori S, et al. Automated prediction of dosimetric eligibility of patients with prostate cancer undergoing intensity-modulated radiation therapy using a convolutional neural network. *Radiol Phys Technol* 2018;11:320–7.
4. Huang Y, Yue H, Wang M, Li S, Zhang J, Liu Z, et al. Fully automated searching for the optimal VMAT jaw settings based on Eclipse Scripting Application Programming Interface (ESAPI) and RapidPlan knowledge-based planning. *J Appl Phys* 2018;19:177–82.
5. Yuan L, Wu QJ, Yin F, Li Y, Sheng Y, Kelsey CR, et al. Standardized beam bouquets for lung IMRT planning. *Phys Med Biol* 2015;60:1831–43.
6. Lee EK, Fox T, Crocker I. Simultaneous beam geometry and intensity map optimization in intensity-modulated radiation therapy. *Int J Radiat Oncol Biol Phys* 2006;64:301–20.
7. MacFarlane M, Hoover DA, Wong E, Goldman P, Battista JJ, Chen JZ. A fast inverse direct aperture optimization algorithm for intensity-modulated radiation therapy. *Med Phys* 2019;46:1127–39.
8. Schlaefer A, Jungmann O, Schweikard A, Kilby W. Objective specific beam generation for image guided robotic radiosurgery. *Int J Comput Assist Radiol Surg (Print)* 2007;2:58–60.
9. Gerlach S, Fürweger C, Hofmann T, Schlaefer A. Feasibility and analysis of CNN based candidate beam generation for robotic radiosurgery. *Med Phys* 2020. Accepted Author Manuscript. <https://doi.org/10.1002/mp.14331>.
10. Schlaefer A, Schweikard A. Stepwise multi-criteria optimization for robotic radiosurgery. *Med Phys* 2008;35:2094–103.
11. King CR, Freeman D, Kaplan I, Fuller D, Bolzicco G, Collins S, et al. Stereotactic body radiotherapy for localized prostate cancer: pooled analysis from a multi-institutional consortium of prospective phase ii trials. *Radiother Oncol* 2013;109:217–21.