

Feasibility and analysis of CNN-based candidate beam generation for robotic radiosurgery

Stefan Gerlach^{a)}

Institute of Medical Technology, Hamburg University of Technology, Hamburg 21073, Germany

Christoph Fürweger

Europäisches Cyberknife Zentrum München-Großhadern, Munich 81377, Germany

Department of Stereotaxy and Functional Neurosurgery, University of Cologne, Faculty of Medicine and University Hospital Cologne, Cologne 50937, Germany

Theresa Hofmann

Europäisches Cyberknife Zentrum München-Großhadern, Munich 81377, Germany

Alexander Schlaefer

Institute of Medical Technology, Hamburg University of Technology, Hamburg 21073, Germany

(Received 27 January 2020; revised 24 May 2020; accepted for publication 7 June 2020; published 8 July 2020)

Purpose: Robotic radiosurgery offers the flexibility of a robotic arm to enable high conformity to the target and a steep dose gradient. However, treatment planning becomes a computationally challenging task as the search space for potential beam directions for dose delivery is arbitrarily large. We propose an approach based on deep learning to improve the search for treatment beams.

Methods: In clinical practice, a set of candidate beams generated by a randomized heuristic forms the basis for treatment planning. We use a convolutional neural network to identify promising candidate beams. Using radiological features of the patient, we predict the influence of a candidate beam on the delivered dose individually and let this prediction guide the selection of candidate beams. Features are represented as projections of the organ structures which are relevant during planning. Solutions to the inverse planning problem are generated for random and CNN-predicted candidate beams.

Results: The coverage increases from 95.35% to 97.67% for 6000 heuristically and CNN-generated candidate beams, respectively. Conversely, a similar coverage can be achieved for treatment plans with half the number of candidate beams. This results in a patient-dependent reduced averaged computation time of 20.28%–45.69%. The number of active treatment beams can be reduced by 11.35% on average, which reduces treatment time. Constraining the maximum number of candidate beams per beam node can further improve the average coverage by 0.75 percentage points for 6000 candidate beams.

Conclusions: We show that deep learning based on radiological features can substantially improve treatment plan quality, reduce computation runtime, and treatment time compared to the heuristic approach used in clinics. © 2020 The Authors. *Medical Physics* published by Wiley Periodicals LLC on behalf of American Association of Physicists in Medicine. [<https://doi.org/10.1002/mp.14331>]

Key words: beam optimization, machine learning, radiation therapy, robotic radiosurgery, treatment planning

1. INTRODUCTION

An important advancement in radiotherapy is the development of robotic radiosurgery. The key idea of this method is to utilize the kinematic flexibility of a robot to deliver dose from multiple positions around the patient. Furthermore, the beams of radiation can be nonisocentric to achieve high conformity and allow for a steep dose gradient around the target. However, the optimal arrangement of beams becomes a challenging task due to the arbitrarily large search space of potential beam directions. In clinical practice, a beam generation heuristic based on a randomized selection of candidate beams has proven to be robust. Since there is no known analytical solution to the planning problem, the beams to be delivered during treatment are typically chosen from the set of candidate beams by solving an

optimization problem. We model this optimization problem as a linear program and employ the simplex algorithm to solve for Pareto efficient solutions.

However, planning based on candidate beams is necessarily incomplete. While larger candidate beam sets generally result in better performance, the computational effort increases disproportionately with the number of candidate beams and can get prohibitive. Additionally, while improvements to clinical plan quality may diminish with increasing numbers of candidate beams, computation time is of interest, too. Furthermore, there are similarities in the patient geometry between different patients for prevalent types of cancer. Therefore, we propose a convolutional neural network (CNN)-based approach that is trained on conventionally computed treatment plans to generate candidate beams for new patients.

A CNN represents a trainable model that learns the relation between input and output data without requiring much domain knowledge. The network essentially consists of multiple convolutional layers followed by one or multiple fully connected layers prior to the output. Each layer has trainable weights that are learned by minimizing an error function for pairs of input and output data. Simply put, the convolutional layers extract features from the input images that are then evaluated by the fully connected layers to make a prediction.

Convolutional neural networks have been applied to different image processing tasks in the medical domain like classification^{1,2} or segmentation³ and have shown significant improvement over conventional methods.⁴ Recently, methods for predicting dose distributions from three-dimensional volumes or two-dimensional (2D) slices of computed tomography (CT) scans using fully convolutional neural networks^{5,6} have been proposed. These distributions can then be used for inverse planning and are a promising step toward automatic treatment plan generation.

Approaches to optimize beam orientations, positions, shapes, and weights directly (direct aperture optimization) have been proposed. Those involve either solving a mixed integer problem⁷ or combining the optimization of the dose in the target, dose constraints, and apertures in the objective function^{8,9}. The former means a significant increase in computational effort while the latter sacrifices the ability to set hard constraints on the doses of critical organ structures.

Furthermore, knowledge-based methods were also applied to optimization of beam-related parameters in intensity-modulated radiation therapy (IMRT) like statistical and machine learning methods^{10–12} or case- and atlas-based methods^{13–15}.

However, the problem of robotic radiosurgery is considerably more complex since more than 100 directions are typically considered in comparison to five to nine beam angles that are used for IMRT. Note that identifying apertures in IMRT corresponds to finding a spacial arrangement and shape for the treatment beams in robotic radiosurgery.

We propose an approach to represent the patient geometry as projections of radiological features of the organ structures on 2D planes in node eye view. This approach is motivated by an earlier method based on case-based reasoning¹⁶. With these features, the CNN is trained to predict each beam's influence on the dose distribution. These predictions then determine the probability for the beam to be included in the candidate beam set. Finally, the solution of the inverse planning problem selects and weights the treatment beams. We evaluate possible improvements on treatment plan quality and computation time for more than 30 patients previously treated for prostate cancer and discuss the difference in candidate beam sets generated by this approach to randomized candidate beam sets generated by the conventional approach.

2. MATERIALS AND METHODS

2.A. Treatment planning

For robotic radiosurgery, for example, with the CyberKnife, the beam source is mounted on a 6 degrees-of-freedom

robotic arm.¹⁷ To avoid collisions with the patient and medical equipment, discrete positions are defined from where beams can be delivered. These beam nodes lie roughly on the surface of a cylinder around the patient with a radius of 900–950 mm in the case of prostate cancer therapy. Beams can be oriented arbitrarily, starting at the beam nodes. We consider nonisocentric beams with circular Iris¹⁸ collimators. A treatment plan then defines a weighted subset of candidate beams for treatment. Here, each beam's weight corresponds to the activation time of the respective beam.

Computed tomography image data with contoured volumes of interest (VOIs), namely the planning target volume (PTV) and organs at risk (OARs), are the basis for treatment planning. Additional shell structures around the PTV are defined to control the dose gradient in normal tissue. The goal of treatment planning then is to find a configuration of beams that deliver the prescribed dose to the PTV while constraining the maximal dose to PTV, OARs, and shells.

This optimization problem is often solved by linear programming¹⁹ which has the advantage over heuristic optimization, for example, simulated annealing or gradient descent-based methods, of being optimal with respect to the set of candidate beams. The optimization is separated into distinct steps. First, a large set of candidate beams is randomly sampled and dose coefficients are determined. These coefficients represent the dose delivered by a beam at a discrete location proportional to the beam's weight. Second, the inverse optimization problem is formulated and solved with respect to clinical goals and dose constraints. Finally, from the solution, the subset of weighted beams is identified as part of the treatment plan. Generally, the quality of the treatment plan increases with increasing number of candidate beams as there are more options for treatment beams. However, computation time also increases and can become prohibitive. Furthermore, the number of actual weighted treatment beams is typically much smaller than the number of candidate beams. This motivates the use of fewer promising candidate beams to achieve comparable treatment plan quality.

2.B. Dataset

Our dataset for training and cross-validation consists of ten patient cases with prostate cancer that were previously treated with the CyberKnife with VOI sizes that represent a wide range of patients (Table I). We additionally evaluate our trained CNN on a separate cohort of 27 patients also treated with the CyberKnife. For each case, the prostate is defined as the PTV and bladder and rectum as OARs. For treatment planning we adopt a five-fraction protocol with a prescribed dose of 36.25 Gy and set hard constraints on the dose of PTV and OARs to 40.25 and 36 Gy, respectively. To achieve a steep dose gradient, we introduce two shells enclosing the PTV at 3 and 9 mm distance. We also constrain the total and per beam monitor units to 40 000 and 300 MU, respectively. To generate training data for the CNNs, we compute 30 treatment plans with 6000 random candidate beams each for every patient. Candidate beams originate at a random beam node

and are oriented toward a random point in the projection of the PTV.

The diameters of the beams' collimators are chosen from 10, 15, 20, 30, or 40 mm at 800-mm distance with equal probability. To make the treatment plans comparable between patients, the constraints on the dose for the shell structures are tuned to achieve roughly 95 % coverage of the PTV. The coverage is a dose criterion which describes the fraction of the PTV which receives at least the prescribed dose. It is commonly used to evaluate the plan quality of a treatment plan and is optimized when generating treatment plans in our setup. We maximize the coverage indirectly by minimizing the underdosage of the PTV's fraction that receives less than the prescribed dose by formulating the inverse planning problem as a linear program using our in-house planning software.¹⁹

2.C. Feature generation

One possible approach to feature generation could be a representation of the complete planning problem. The corresponding label would then be the entire set of weighted beams. However, this would result in a highly complex description of the planning problem which would make learning meaningful patterns hard and require a large amount of training data to generalize sufficiently well. Thus, we setup our CNN to predict the influence of a single beam on the delivered dose, independent of other beams.

A common method to describe the dosimetric influence of a beam is the projection of the VOIs from the beam node on a plane perpendicular to the line from beam node to beam target, that is, the beam's eye view.²⁰ A similar method has been applied in a case-based approach,¹⁶ where, among other features, similarity of beam nodes is established by projections of VOIs on a plane perpendicular to the line from beam node to PTV centroid. We describe various properties of a beam by projections onto this plane and create a separate gray-scale image for each property. Each image has a pixel size of 1 mm × 1 mm and a size of 150 × 150 pixels, which covers the PTV and most parts of the OARs. We concatenate the gray-scale images in the channel dimension and use the resulting image tensor as input for CNN. The feature creation is shown in Fig. 1 and explained in detail in the following.

A first image represents the beam relative to the target. We intersect the central beam with the projection plane and compute the effective radius at the source-plane distance. We further refer to the resulting mask of the area covered by the beam as *beam feature*.

TABLE I. Volumes of interest sizes of patients in the dataset in cm^3 .

	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10
PTV	67	93	82	81	109	70	115	86	87	52
rectum	50	84	44	34	66	90	56	53	41	114
bladder	76	95	90	119	74	187	57	37	110	259

Additional images represent the VOIs as viewed from the beam node. We refer to them as *VOI features* and study three different approaches. First, we compute separate projections of each VOI on the plane, effectively masking the outline of the VOI. This leads to input data of size $[W, H, D] = [150, 150, 1 + n_v] = [150, 150, 4]$ corresponding to one *beam feature* and one *VOI feature* per VOI when considering $n_v = 3$ VOIs.

Second, we include the distance along the beam in the projection. Creating two images per VOI that encode the minimum and maximum distance in the projection, respectively, we represent the three-dimensional VOI shape. Therefore, we generate two VOI features per VOI and have an input of size $[W, H, D] = [150, 150, 1 + 2n_v] = [150, 150, 7]$.

Third, we consider that the effect of the beam does depend on the tissue density. We, therefore, replace the minimum and maximum distance projections by the respective minimum and maximum radiological depth. Note that this results in a deformed three-dimensional representation of the VOI that accounts for the effective dose coefficients. The input size is again $[W, H, D] = [150, 150, 1 + 2n_v] = [150, 150, 7]$.

To assess the relevance of the convolutional layers of the CNN in our approach, we compare the CNN to fully connected neural networks (FCNN). We compare flattened images and hand-crafted features as input. Flattened images are the linearized entries of the described image tensors while the hand-crafted features describe the overlapping area of VOI projections and beam projection, that is, the overlap of beam and VOIs, the overlap between each VOI, and additionally the overlap of the area where a VOI is in front of another and the overlap of the area where a VOI is behind of another, totaling in 36 features. These features are similar to those used in literature.¹²

2.D. Model architecture and training

We adapt DenseNet-121,²¹ which is a state of the art CNN architecture for natural image classification. The core idea of DenseNet is to reuse features to reduce the number of trainable parameters and thereby improve the network's computational efficiency to allow for a deeper network. This is facilitated by the organization of multiple convolutional layers in densely connected blocks where inside a block the output of a layer is concatenated with the input of all following layers (Fig. 2). Therefore, the feature tensor of size $[B, W, H, D_l]$ of the dense block grows in D_l with the l th convolutional block (CB) by the growth rate g , which we set to $g = 32$. Here, B is the batch size, H and W are the spacial dimensions, and D_l is the depth of the concatenated feature map tensor. Each CB consists of a 1×1 convolution with output feature map depth $D' = 4 \cdot g$ and a 3×3 convolution with $D' = g$. Here, the first convolution acts as a bottleneck layer because it limits the feature map depth to $4 \cdot g$ for the second convolution and forces the CNN to represent the data more efficiently.

Reusing the output of earlier layers as input allows for smaller D' in each convolution and thus reduces the number

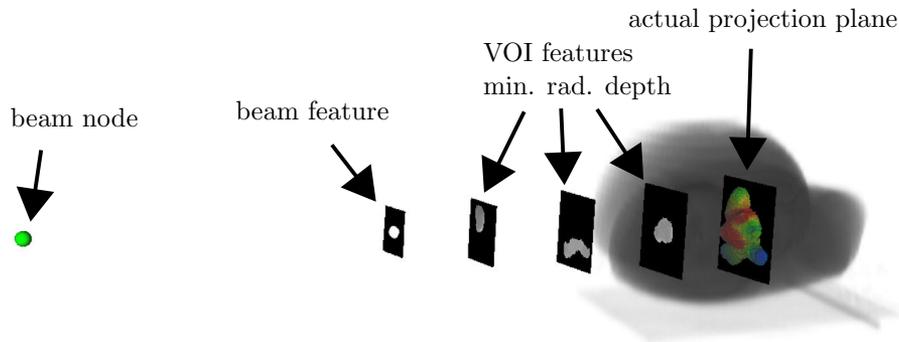


FIG. 1. Feature generation for one beam. Displayed are the generation of the beam feature and minimum radiological depth volumes of interest features schematically. Maximum radiological depth features are omitted here. Note that the actually used features are projections on the same plane and have the same size. [Color figure can be viewed at wileyonlinelibrary.com]

of parameters for similar performance of the network. It also improves the convergence speed during training because the error gradient is connected directly from the output of the dense block to each layer.

Down-sampling, or pooling, layers are an essential part for many CNNs. DenseNet uses average pooling and max-pooling layers, which compute the average or maximum value for an input patch of the feature map, to decrease the spacial size of feature maps. However, this technique is not feasible within a dense block since constant spacial size is required when concatenating the feature maps in the way described above. To still be able to incorporate down-sampling in DenseNet, dense blocks are connected via transition layers. These transition layers essentially consist of a 1×1 convolution with $D' = D/2$ followed by a 2×2 average pooling layer with stride 2. Thereby, H , W , and D are halved in each transition. A summary of the DenseNet structure is given in Fig. 3.

We use transfer learning to initialize the weights of the CNN with weights trained from the ImageNet dataset,²² which consists of natural RGB images classified into 1000 classes. This method has been shown to improve results in the medical setting especially with a limited dataset.²³ We adapt the pretrained DenseNet by replacing the fully connected layer to have a single output and linear activation for regression. Furthermore, we copy the weights of the first convolutional layer to the remaining channels since our images have four or seven channels, depending on the type of VOI features used. We then fine-tune the weights using the Adam optimizer²⁴ and a learning rate of 10^{-3} . The learning rate is determined by hyperparameter optimization using grid search

where the search is guided by the validation loss. Evaluation of the resulting CNN-generated beams is skipped to reduce evaluation effort.

The label for each beam is derived from the associated beam weight that is assigned by the corresponding treatment plan in the training set. It is normalized by the maximum allowed weight per beam. Since after treatment planning there is a class imbalance of multiple unweighted beams per weighted beam, $w_b \approx 19$, we have to address this during training. Although the dose constraints are tuned to achieve roughly 95% coverage in the PTV, different sets of candidate beams result in different actual coverage. Since PTV coverage is an important quality metric, we weight each training example by

$$w_i = \begin{cases} \frac{c_{des}}{c_p}, & \text{for } w_{mu,i} = 0 \\ w_b \cdot \frac{c_p}{c_{des}}, & \text{for } w_{mu,i} > 0 \end{cases}$$

to include the variation in coverage in the training process. Here, $c_{des} = 95\%$ is the desired coverage, c_p is the coverage of the treatment plan that includes beam i , and $w_{mu,i}$ is the number of monitor units assigned to the beam i . The key idea, besides addressing the weight imbalance, is to increase the training weight w_i of beams that contribute to the dose delivery, if the plan quality is high, and to also increase w_i of noncontributing beams if the plan quality is low. The former ensures that truly good beams have a high weight during training. The latter reflects that beams, which were not even part of a treatment plan with poor plan quality, are truly irrelevant for the planning problem and therefore also have a high weight during training. Note that in each set the best beams have been selected by optimization.

We compare DenseNet to a FCNN to assess the relevance of the convolutional layers. For a fair comparison, we remove the convolutional layers and extend the fully connected layer to have a comparable number of parameters as the original CNN.

2.E. Inference and implementation

We use Keras [<https://github.com/fchollet/keras>], which is a high-level API for tensorflow²⁵ to train our CNNs. To

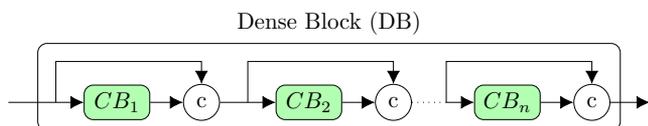


FIG. 2. Dense Block with n Convolutional Blocks (CB), each of which applies a 3×3 convolution with output feature map depth $D' = g = 32$ following a 1×1 convolution with $D' = 4 \cdot g$ that serves as a bottleneck. Batch normalization and ReLU activation is performed before each convolution. [Color figure can be viewed at wileyonlinelibrary.com]

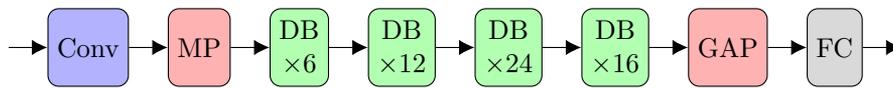


FIG. 3. Summary of the complete DenseNet. *Conv* is a convolution with kernel size of 7×7 , output feature map depth $D' = 64$, and stride 2 followed by batch normalization and ReLU activation. *MP* is 3×3 max-pooling with stride 2. $DB \times c$ is a dense block with c convolutional blocks. *GAP* is 3×3 global average pooling with stride 2. *FC* is a fully connected layer. [Color figure can be viewed at wileyonlinelibrary.com]

generate new beams, we integrate the inference in our in-house Java planning system. A new candidate beam is CNN-generated by first computing the beam and VOI features for a beam with random position, orientation, and collimator. The prediction of the CNN for that beam is then used as the probability to accept the beam into the set of candidate beams. An overview of the complete workflow is given in Fig. 4. Note that the VOI features are computed only once for every beam node and can then be reused for multiple beams starting at that beam node.

However, the proposed method evaluates a beam without considering other beams already in the set of candidate beams. This can lead to a large number of similar beams for beam nodes with disproportionately many candidate beams and thus reduce the options for the optimizer. Note that in CyberKnife planning, the number of beams per beam node corresponds to the spacial distribution. Therefore, we study the effect of constraining the maximum allowed proportion of candidate beams per beam node, that is, the maximum number of beams starting on one node relative to the number of all beams.

2.F. Experimental setup

We train each CNN for only 15 epochs due to the long training time of about 3.3 days on a Nvidia GeForce GTX 1080 Ti. We split the complete data set into two sets. The data of patients 3, 6, 7, and 10 are used for hyperparameter optimization using fourfold cross-validation. The data of the remaining six patients are used for testing where we fold the data six times and train on five patients while testing on the remaining one, respectively. Due to the heuristic nature of the beam generation methods, we

compute each treatment plan with ten beam sets to increase statistical significance.

To validate dosimetric accuracy, a CNN-generated plan was randomly selected and delivered to an Easy Cube solid water phantom (Euromechanics, Schwarzenbruck, Germany). Absolute point dose was measured by an inserted PinPoint ionization chamber (PTW, Freiburg, Germany). A dose plane including the target region was recorded on radiochromic film (Gafchromic EBT3, Ashland) inserted in coronal orientation. Measured and planned dose distributions were compared in FilmQA software (3cognition, Inc.) by Gamma analysis.

3. RESULTS

Figure 5 shows the resulting coverage from treatment plan optimization with random candidate beams and neural network-generated beams using different features for various numbers of candidate beams. We evaluated our CNN approach with cross-validation [Fig. 5(a)] and the two best performing features on a separate patient cohort of 27 patients [Fig. 5(b)]. Clearly, in both cases the resulting coverage with CNN-generated improves over randomized candidate beams for all features and numbers of candidate beams studied with the same size of the optimization problem. The differences are statistically significant with respect to the Wilcoxon rank sum test and $P < 10^{-17}$ for the null hypothesis of identical distributions. Note that only half the number of candidate beams is needed for comparable coverage. Also note that the difference in coverage improvements of CNN-generated beams over randomized beams between the results in Figs. 5(a) and 5(b) are insignificant for more than 1500 beams ($p > 0.07$).

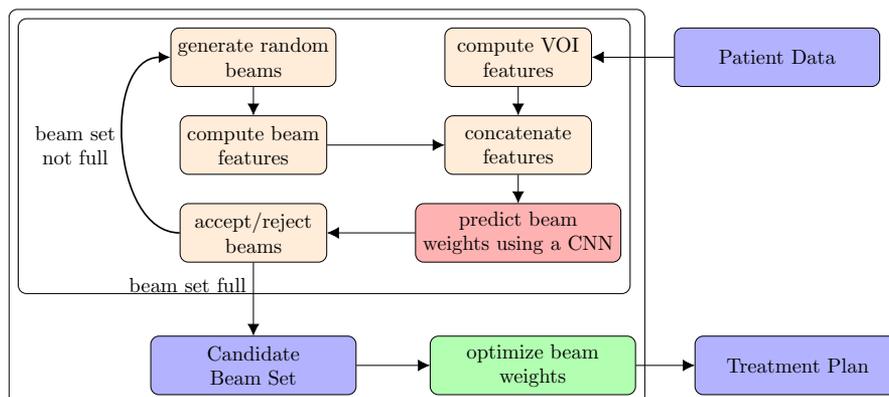


FIG. 4. The complete approach to convolutional neural network-based beam generation and treatment plan computation. [Color figure can be viewed at wileyonlinelibrary.com]

Figure 5(c) compares the coverage for CNNs and FC neural networks. While overlap features show little improvement over random beam selection, training on the flattened images of projections comes closer to CNN-generated beams. However, the benefit of convolutions is still clearly visible.

As Fig. 6 illustrates, there are differences in coverage improvements between different patients. From the patients studied, patient 1 has one of the highest and patient 8 one of the lowest absolute improvement in coverage.

Furthermore, we evaluated the runtime of the randomized and the CNN-based beam generation approach. Treatment plans were computed with an Intel Xeon E3-1200 v3 and a Nvidia GeForce GTX 1080 Ti. Figure 7(a) compares the runtime with the tested numbers of candidate beams. To quantify the reduction in runtime, we compare the runtime of treatment plans between 90% and 96% coverage. The relative reduction was computed by

$$r_{mean} = 1 - \text{mean}\left(\frac{r_{cnn}}{c_{cnn}}\right) \cdot \text{mean}\left(\frac{c_r}{r_r}\right)$$

where r_{cnn} and r_r are the runtimes of treatment plan computation with CNN-generated beams or randomized beams, respectively, and c_{cnn} and c_r are the coverage with CNN-generated beams or randomized beams, respectively. Following this equation, the average runtime for all patients is reduced by $r_{mean} = 34.21\%$. The extremes for runtime reduction occurs for patient 1 with $r_{mean} = 45.69\%$ and for patient 8 with $r_{mean} = 20.28\%$. Figure 7(b) shows the number of weighted beams similarly. Here, the average reduction for high coverage is 11.35%, 16.19%, and 5.15% for all patients, patient 1, and 2, respectively.

Additionally, we study the difference in beam distribution between randomized and CNN-generated beam sets. Figure 8 shows the impact of the CNN based approach on the beam collimators. The average collimator size and standard

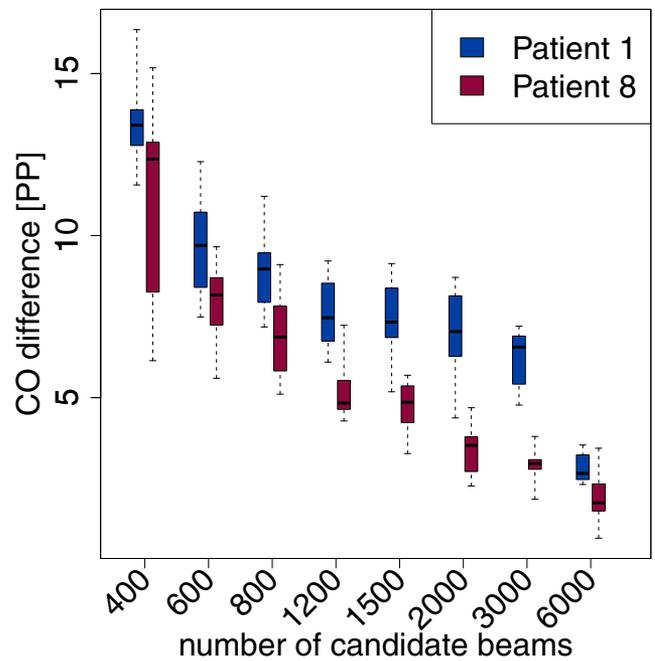


FIG. 6. Boxplot of coverage difference of convolutional neural network-generated beams with radiological depth features to randomized beams. Results for patient 1 and 8 with whiskers and boxes at 0%, 25%, 50%, 75%, and 100% quantiles. PP is percentage points. [Color figure can be viewed at wileyonlinelibrary.com]

deviation in the unweighted case increase from (22.9 ± 10.7) mm for randomized beams to (24.6 ± 10.8) mm for CNN-generated beams and from (25.7 ± 10.7) mm to (25.8 ± 10.9) mm in the weighted case.

Figure 9 shows an example of the beams' target distribution for one beam node. The distribution shows all candidate beams from the beam node with their respective diameter at the projection surface. While randomized beams are

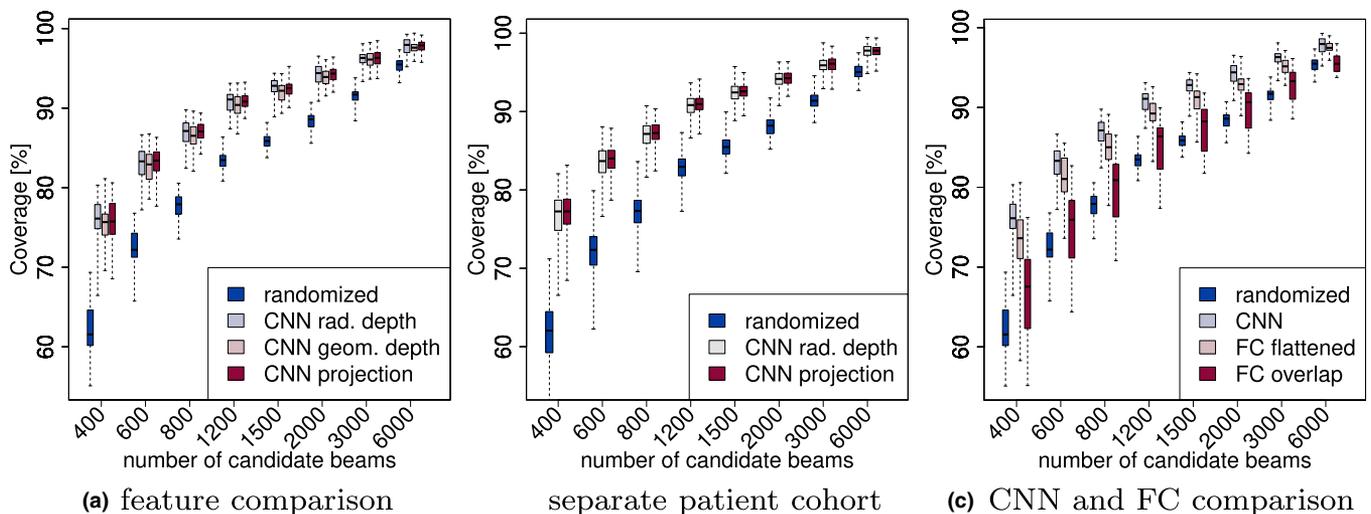


FIG. 5. Boxplot of coverage for randomized beams and convolutional neural network (CNN)-generated beams (a,b) with radiological, geometric, and projection features evaluated by cross-validation (a) and on a separate patientcohort of 27 patients (b). The CNN with radiological features is also compared to fully connected networks (c). Whiskers and boxes are displayed at 0%, 25%, 50%, 75%, and 100% quantiles. [Color figure can be viewed at wileyonlinelibrary.com]

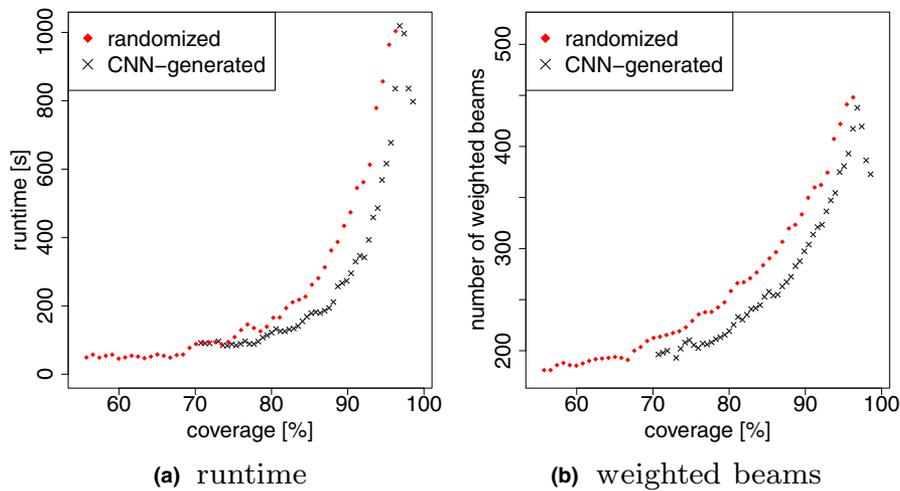


FIG. 7. Runtimes (a) and weighted beams (b) for randomized and convolutional neural network-generated candidate beams with radiological depth features. Each point represents the average over a range of 2% coverage. [Color figure can be viewed at wileyonlinelibrary.com]

generally more evenly distributed, CNN-generated beams concentrate in the upper region for the displayed beam node. This region corresponds to the part of the PTV that is not covered by the bladder. Note that the rectum lies behind the PTV in the view from the shown beam node. This is represented by larger values in the depth features of the rectum compared to the PTV. Figure 9 also illustrates how different the shape and location of projected VOIs can be, even for the same patient.

To further analyze the CNN-generated set of candidate beams, we study the distribution of beams among beam nodes (Fig. 10). While the randomized candidate beams are evenly distributed, CNN-generated candidate beams are biased toward certain beam nodes and more closely resemble the distribution of weighted beams. The difference of distributions for weighted beams between randomized and CNN-generated beams is small.

The difference in coverage when constraining the maximum number of candidate beams per beam node is shown in Table II. Note that the difference to the unconstrained method is significant ($P < 0.004$) for every entry in the table with respect to the Wilcoxon rank sum test.

Expectedly, dosimetric verification of the exemplary CNN-generated plan yielded clinically acceptable results: The difference between measured and calculated point dose was less than 0.1%, which is well within the uncertainty of the dose measurement, and 2D Gamma pass rate (global, 3%/1 mm) was 98.0%.

4. DISCUSSION

Considering Fig. 5, the results show a significant improvement of coverage with CNN-generated candidate beams over randomized candidate beams. With our approach, a fairly small number of patients in the training set is already enough for the CNN to generalize very well, which we have illustrated on more than 30 independent cases covering a wide range of organ structure sizes. The differences in coverage improvement between patient cohorts are only significant for small numbers of beams due to the larger variance for fewer beams. Furthermore, improvement varies between patients. For example, patient 8 shows smaller improvement than other studied patients. For this patient the prostate size is average while the bladder is smaller than that of any patient in the

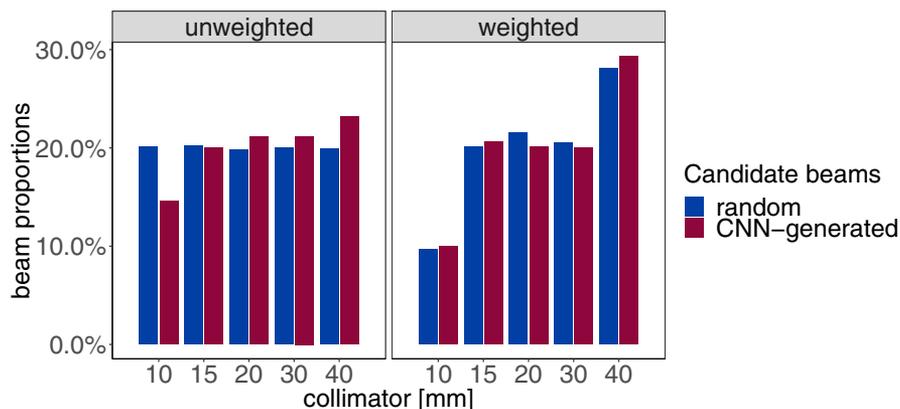


FIG. 8. Histogram of proportion of collimator frequencies for randomized and convolutional neural network-generated beams, respectively, before and after treatment plan optimization [Color figure can be viewed at wileyonlinelibrary.com]

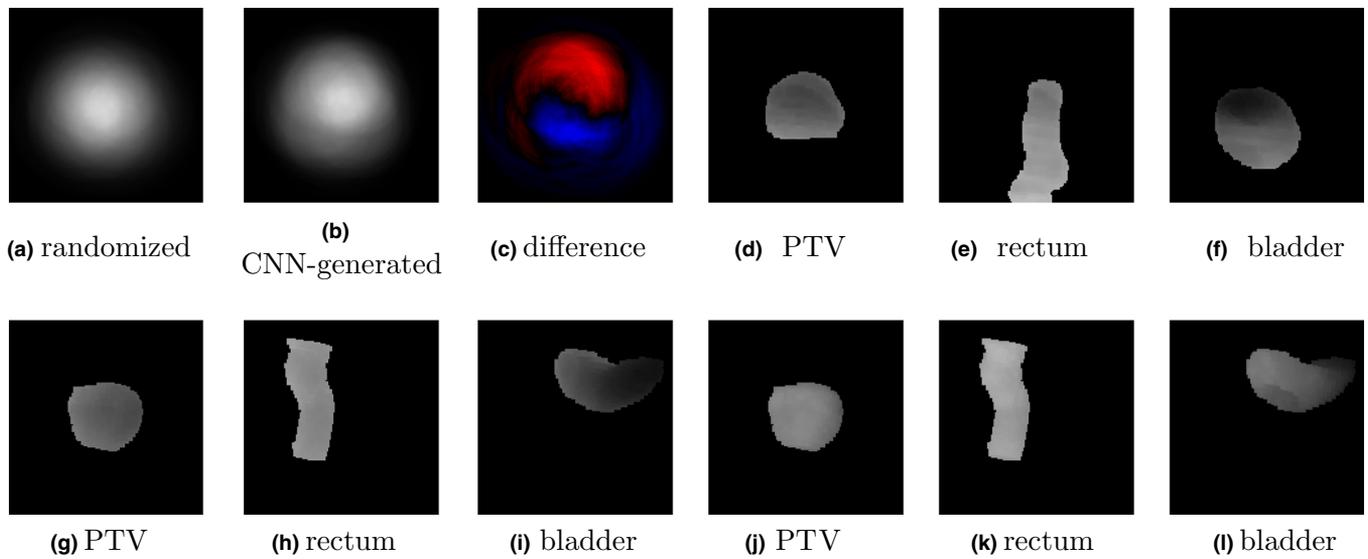


FIG. 9. Distribution of candidate beams for one beam node for randomized beams (a) and convolutional neural network (CNN)-generated beams (b), normalized to the number of beams starting at the node and the difference between distributions (c). Red colors indicate a higher density in the CNN-generated beam distribution, blue colors indicate the opposite. The minimal radiological depth of the volumes of interest (d-f) from the same beam node and the minimal (g-i) and maximal (j-l) radiological depth of a different beam node for the same patient is shown with brighter pixels meaning larger distance. [Color figure can be viewed at wileyonlinelibrary.com]

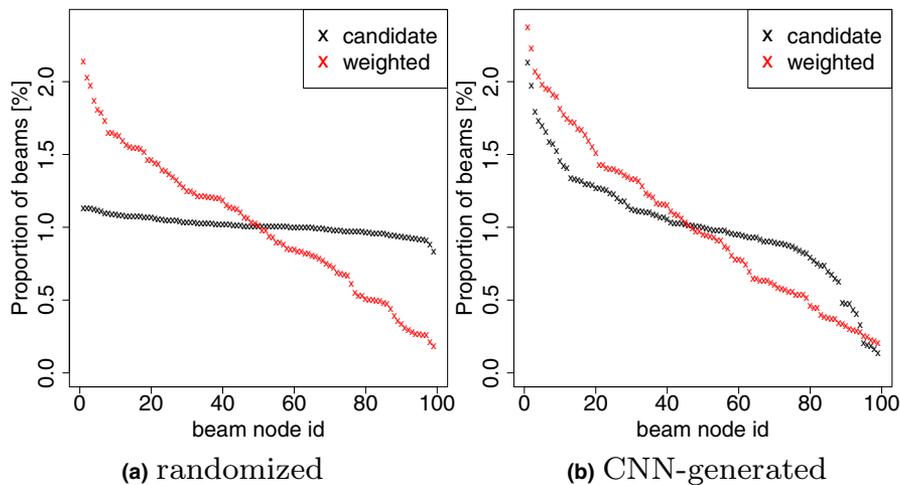


FIG. 10. Proportion of randomized (a) and convolutional neural network-generated (b) candidate and weighted beams per beam node. Beam nodes are sorted by proportion of beams. [Color figure can be viewed at wileyonlinelibrary.com]

training set. Compared to the rectum, the bladder usually covers a larger area of the prostate from the majority of beam nodes which can explain the larger impact of its size on the predictive capability of the CNN. Furthermore, a small bladder is less restrictive on the dose distribution such that the selection of candidate beams is less challenging and a set of randomized beams cannot be improved much.

Note that with 6000 CNN-generated beams, the optimization function attains a value of zero for some patients, which corresponds to 100% coverage. Thus, there would have been options to further improve the plan quality in these cases, for example, by setting stricter dose constraints. Furthermore, note that we have trained the CNN on a single GPU for relatively few epochs. Training on multiple GPUs makes training on larger datasets for more epochs feasible.

The differences in coverage between the different features are small compared to the difference to randomized beams. This is due to the setup for radiation therapy of the pelvis. Relevant tissue structure in that region is homogeneous compared to other structures, for example, in the lung where bones may influence the dose delivery more severely. Thus, the radiological depth for targets in the pelvis can be approximated well by the geometrical depth. Therefore, the learned information for both features is similar. When we compare the projection features to the other studied features, no statistically significant difference can be observed. However, the average coverage is higher with radiological depth features but lower for geometric depth features. This suggests that the CNN prediction for beam generation improves with more radiological information about the VOIs but it does not fail

TABLE II. Difference in coverage for convolutional neural network-generated treatment plans with constrained maximum number of candidate beams per beam node compared to the unconstrained approach in percentage points.

Candidate beams	Upper limit for proportion of candidate beams per beam node [%]				
	1.2	1.25	1.3	1.35	1.4
400	2.11 ± 2.55	1.96 ± 2.43	1.73 ± 2.87	1.57 ± 2.45	1.82 ± 2.35
600	1.66 ± 1.98	1.37 ± 1.92	1.25 ± 2.04	1.24 ± 2.08	1.66 ± 1.93
800	1.53 ± 1.54	1.22 ± 1.64	0.99 ± 1.84	1.21 ± 1.82	1.47 ± 1.60
1200	1.64 ± 1.23	1.34 ± 1.51	1.47 ± 1.35	1.56 ± 1.48	1.66 ± 1.28
1500	1.73 ± 1.18	1.50 ± 1.23	1.60 ± 1.34	1.51 ± 1.28	1.62 ± 1.26
2000	1.72 ± 0.99	1.46 ± 0.97	1.65 ± 1.23	1.73 ± 1.04	1.54 ± 0.93
3000	1.68 ± 0.67	1.81 ± 0.68	1.55 ± 0.73	1.67 ± 0.69	1.73 ± 0.74
6000	0.59 ± 0.50	0.69 ± 0.48	0.66 ± 0.55	0.75 ± 0.49	0.57 ± 0.51

Positive values represent higher coverage for the constrained approach with largest difference for the respective number of candidate beams printed in bold. Additionally, the standard deviation is shown.

critically without it. Furthermore, the simplicity of the projection features improves the coverage over geometric depth features. Also, the general shape of the considered VOIs is similar for prostate radiation therapy between patients. Therefore, the volume can approximately be inferred by the projection alone. Furthermore, the shape of the prostate is usually regular and close to a sphere. For radiation therapy of other types of cancer, for example, lung carcinoma, spinal metastases, or complex skull base meningioma, this can be more challenging.

While there have been approaches to solve dense linear optimization problems in part on a GPU^{26,27}, they require inherently more sequential processing. As CNNs, on the other hand, are structured in a way that makes them highly parallelizable, they are well suited to be executed on a GPU. Therefore, the speedup of inference on a GPU decreases the computation time below that of the conventional approach as Fig. 7(a) shows. The CNN-based beam generation adds computational effort that grows linearly with increasing number of beams. The runtime of a linear optimization problem, on the other hand, grows disproportionally with increasing number of candidate beams. Thus, the advantage of the CNN-based approach is more significant for more candidate beams. However, this only holds while the number of candidate beams is still small enough that not virtually every orientation is considered by the set of beams. The results show that this point is not yet reached for 6000 beams since improvements in coverage over randomized beams are still significant.

Due to the large search space of possible candidate beams, previous work also considered heuristic methods like simulated annealing or genetic algorithms^{28,29} to search for promising beam configurations for IMRT. These methods, however, would drastically increase the computational effort in our case because the linear program would have to be solved for every evaluated set of candidate beams.

An additional advantage of optimizing with fewer candidate beams is a lower number of weighted beams [Fig. 7(b)]. The number of weighted beams is a factor in the treatment time of the patient as the radiation source has to be adjusted

for each beam. Thus, fewer beams generally result in a shorter treatment time. Note that we did not employ beam reduction methods to maintain comparability. The decrease in runtime and the number of treatment beams for coverage close to 100% is due to the optimization function reaching the minimum possible value. Therefore, the optimizer requires fewer iterations and weighted beams to reach the optimum.

The distribution of collimator sizes of CNN-generated candidate beams is closer to that of weighted beams than randomized beams are. The distribution of collimator sizes is closely linked to the clinical goal during treatment plan optimization. For example, limiting the MU would increase the use of larger collimators to still deliver the necessary dose. Thus, the learned bias is most useful for generating treatment plans with similar clinical goals.

When studying the distribution of beams to nodes, we can observe that the CNN-generated beams follow the distribution of weighted beams more closely (Fig. 10). However, our approach does not take already accepted beams in the candidate set into account when evaluating new beams. This can lead to many similar candidate beams starting at the same beam node which do not add further options for the optimizer. Therefore, we constrain the maximum number of beams starting at a single node. The results displayed in Table II show an increased coverage over the results presented in Fig. 5(a) for every constraint studied. However, the optimal constraint cannot be inferred explicitly by the results. This suggests a nontrivial relationship and motivates the inclusion of information from all nodes for the evaluation of beams.

The presented training was based on a specific treatment scenario and therefore we cannot expect that the CNN would generalize to different tumor sites or treatment protocols. However, the approach demonstrates that a fairly small set of patient data can be used to derive a beam selection method that improves over the conventional approach. Moreover, clinical planning often follows guidelines or protocols. Given that computing the projections is generic, it should be possible to employ the same approach for other targets.

5. CONCLUSION

Treatment plan generation for radiosurgery is a computationally challenging task. We show that a CNN can be applied to predict the influence of candidate beams on the delivered dose. Using CNN-predicted candidate beams significantly improved coverage, decreased plan computation time, and shortened treatment delivery. Constraining the maximum number of beams per node increases plan quality, which indicates a potential to further improve CNN predictions by considering the complete beam set during beam generation.

ACKNOWLEDGMENTS

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

CONFLICT OF INTEREST

The authors declare that they have no conflict of interest.

ETHICAL APPROVAL

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

INFORMED CONSENT

For this type of study informed consent is not required.

^{a)}Author to whom correspondence should be addressed. Electronic mail: stefan.gerlach@tuhh.de.

REFERENCES

- Gessert N, Lutz M, Heyder M, et al. Automatic plaque detection in IVOCT pullbacks using convolutional neural networks. *IEEE Trans Med Imaging*. 2019;38:426–434.
- Kajikawa T, Kadoya N, Ito K, 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;11:320–327.
- Ibragimov B, Xing L. Segmentation of organs-at-risks in head and neck CT images using convolutional neural networks. *Med Phys*. 2017;44:547–557.
- Litjens G, Kooi T, Bejnordi BE, et al. A survey on deep learning in medical image analysis. *Med Image Anal*. 2017;42:60–88.
- Kearney V, Chan JW, Haaf S, Descovich M, Solberg TD. DoseNet: a volumetric dose prediction algorithm using 3D fully-convolutional neural networks. *Phys Med Biol*. 2018;63:235022.
- Nguyen D, Long T, Jia X, et al. A feasibility study for predicting optimal radiation therapy dose distributions of prostate cancer patients from patient anatomy using deep learning. *Sci Rep*. 2019;9:1076–1075.
- 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;64:301–320.
- 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;46:1127–1139.
- Men C, Jia X, Jiang SB. GPU-based ultra-fast direct aperture optimization for online adaptive radiation therapy. *Phys Med Biol*. 2010;55:4309–4319.
- Huang Y, Yue H, Wang M, 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 Clin Med Phys*. 2018;19:177–182.
- Yuan L, Wu QJ, Yin F, et al. Standardized beam bouquets for lung IMRT planning. *Phys Med Biol*. 2015;60:1831.
- Amit G, Purdie TG, Levinshtein A, et al. Automatic learning-based beam angle selection for thoracic IMRT. *Med Phys*. 2015;42:1992–2005.
- Schreibmann E, Fox T. Prior-knowledge treatment planning for volumetric arc therapy using feature-based database mining. *J Appl Clin Med Phys*. 2014;15:19–27.
- Petrovic S, Khussainova G, Jagannathan R. Knowledge-light adaptation approaches in case-based reasoning for radiotherapy treatment planning. *Artif Intell Med*. 2016;68:17–28.
- Zhang X, Li X, Quan EM, Pan X, Li Y. A methodology for automatic intensity-modulated radiation treatment planning for lung cancer. *Phys Med Biol*. 2011;56:3873.
- Schlaefler A, Dieterich S. Feasibility of case-based beam generation for robotic radiosurgery. *Artif Intell Med*. 2011;52:67–75.
- Kilby W, Dooley JR, Kuduvali G, Sayeh S, Maurer CR. The CyberKnife robotic radiosurgery system in 2010. *Technol Cancer Res Treat*. 2010;9:433–452.
- Echner GG, Kilby W, Lee M, et al. The design, physical properties and clinical utility of an iris collimator for robotic radiosurgery. *Phys Med Biol*. 2009;54:5359–5380.
- Schlaefler A, Schweikard A. Stepwise multi-criteria optimization for robotic radiosurgery. *Med Phys*. 2008;35:2094–2103.
- Kalet IJ, Austin-Seymour MM. The use of medical images in planning and delivery of radiation therapy. *J Am Med Inform Assoc*. 1997;4:327–339.
- Huang G, Liu Z, van der Maaten L, Weinberger KQ. Densely Connected Convolutional Networks. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*; 2017.
- Krizhevsky A, Sutskever I, Hinton GE. ImageNet classification with deep convolutional neural networks. *Neural Inf Process Syst*. 2012;25:1097–1105.
- Shin H-C, Roth HR, Gao M, et al. Summers, Deep convolutional neural networks for computer-aided detection: CNN architectures, dataset characteristics and transfer learning. *IEEE Trans Med Imaging*. 2016;35:1285–1298.
- Kingma DP, Ba J. Adam: a method for stochastic optimization. International Conference for Learning Representations; 2014.
- Abadi M, Barham P, Chen JEA. TensorFlow: large-scale machine learning on heterogeneous systems; 2015.
- Spampinato DG, Elstery AC. Linear optimization on modern GPUs. In: *IEEE International Symposium on Parallel Distributed Processing*; 2009: 1–8.
- Jung JH, O’Leary DP. Implementing an interior point method for linear programs on a CPU-GPU system. *Electron Trans Numer Anal*. 2008;28:174–189.
- Bangert M, Ziegenhein P, Oelfke U. Comparison of beam angle selection strategies for intracranial IMRT. *Medical Physics*. 2012;40:011716.
- Hou Q, Wang J, Chen Y, Galvin JM. Beam orientation optimization for IMRT by a hybrid method of the genetic algorithm and the simulated dynamics. *Med Phys*. 2003;30:2360–2367.