

PAPER • OPEN ACCESS

A reinforcement learning application of a guided Monte Carlo Tree Search algorithm for beam orientation selection in radiation therapy

To cite this article: Azar Sadeghnejad-Barkousaraie *et al* 2021 *Mach. Learn.: Sci. Technol.* **2** 035013

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

You may also like

- [Localised Wannier orbital basis for the Mott insulators \$\text{GaV}_4\text{S}_8\$ and \$\text{GaTa}_4\text{Se}_8\$](#)
A. Camjayi, R. Weht and M. J. Rozenberg
- [Self-Healing Effect of Hybrid Sol-Gel Coatings Based on GPTMS, TEOS, \$\text{SiO}_2\$ Nanoparticles and \$\text{Ce}\(\text{NO}_3\)_3\$ Applied on Aluminum Alloy 7075-T6](#)
Urša Tiringar, Alicia Durán, Yolanda Castro et al.
- [Learning from troubleshooting activities when contrasting erroneous examples with worked examples in the physics classroom](#)
Rafi' Safadi and Ranin Ababsy



PAPER

OPEN ACCESS

RECEIVED
28 October 2020REVISED
14 January 2021ACCEPTED FOR PUBLICATION
10 February 2021PUBLISHED
13 May 2021

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

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



A reinforcement learning application of a guided Monte Carlo Tree Search algorithm for beam orientation selection in radiation therapy

Azar Sadeghnejad-Barkousaraie^{id}, Gyanendra Bohara, Steve Jiang and Dan Nguyen^{* id}

Medical Artificial Intelligence and Automation (MAIA) Laboratory, Department of Radiation Oncology, UT Southwestern Medical Center, Dallas, TX, United States of America

* Author to whom any correspondence should be addressed.

E-mail: Dan.Nguyen@UTSouthwestern.edu**Keywords:** beam orientation optimization, radiation therapy, intensity modulated radiation therapy (IMRT), Monte Carlo Tree Search, prostate cancer

Abstract

Current beam orientation optimization algorithms for radiotherapy, such as column generation (CG), are typically heuristic or greedy in nature because of the size of the combinatorial problem, which leads to suboptimal solutions. We propose a reinforcement learning strategy using a Monte Carlo Tree Search (MCTS) that can find a better beam orientation set in less time than CG. We utilize a reinforcement learning structure involving a supervised learning network to guide the MCTS and to explore the decision space of beam orientation selection problems. We previously trained a deep neural network (DNN) that takes in the patient anatomy, organ weights, and current beams, then approximates beam fitness values to indicate the next best beam to add. Here, we use this DNN to probabilistically guide the traversal of the branches of the Monte Carlo decision tree to add a new beam to the plan. To assess the feasibility of the algorithm, we used a test set of 13 prostate cancer patients, distinct from the 57 patients originally used to train and validate the DNN, to solve five-beam plans. To show the strength of the guided MCTS (GTS) compared to other search methods, we also provided the performances of Guided Search, Uniform Tree Search and Random Search algorithms. On average, GTS outperformed all the other methods. It found a better solution than CG in 237 s on average, compared to 360 s for CG, and outperformed all other methods in finding a solution with a lower objective function value in less than 1000 s. Using our GTS method, we could maintain planning target volume (PTV) coverage within 1% error similar to CG, while reducing the organ-at-risk mean dose for body, rectum, left and right femoral heads; the mean dose to bladder was 1% higher with GTS than with CG.

1. Introduction

Radiation therapy is one of the main modalities to cure cancer and is used in over half of all cancer treatments, either as a standalone therapy or in conjunction with another modality, such as surgery or chemotherapy [1]. In intensity-modulated radiation therapy (IMRT), the patient's body is irradiated from fixed beam locations around the body, and the radiation field is modulated at each beam position by multi-leaf collimators [2]. In IMRT, the optimal choice of beam orientations has a direct impact on the treatment plan quality, which influences the final treatment outcome [3] and, ultimately, patients' quality of life. Current clinical practice is to select the beam orientations either using protocol or manually by the treatment planner [3, 4]. Beam orientation optimization (BOO) methods solve for a suitable set of beam angles by minimizing an objective function to a local minimum [5, 6]. BOO has been studied extensively in radiation therapy procedures, for both coplanar [3, 5–29] and non-coplanar [3, 5, 9, 10, 23, 26, 29–37] IMRT, as well as for intensity-modulated proton therapy [38–43] over the past three decades. However, BOO methods have not been widely adopted because of their high computational cost and complexity, as they

involve a large-scale (NP) (nondeterministic polynomial time) hard combinatorial problem [34, 44]. Despite extensive research, there are still no clinically practical beam orientation selection algorithms because the procedure is computationally expensive and time intensive, and the final solutions are usually suboptimal, so BOO remains a challenge to treatment planning.

To measure the quality of a BOO solution, the dose influence array of each potential beam orientation must be calculated. The dose influence array for one beam associates all the individual beamlets in the fluence map with the voxels of the patient's body. This calculation is time-consuming, and it requires a large amount of memory for optimization. To manage the limited capacity of computational resources, the treatment planning process, after defining the objective function, may be divided into two major steps: (a) finding a suitable set of beam orientations, and (b) solving the fluence map optimization (FMO) [6] problem for the selected beams. However, these two steps are not independent of each other: the quality of the BOO solution can be evaluated only after the FMO problem has been solved, and the FMO can be defined only after the BOO has been solved. Because of the non-convexity and large scale of the problem, researchers have considered dynamic programming methods that break the problem into a sequence of smaller problems. Column generation (CG) is a successful algorithm specifically developed to solve complex problems such as BOO. Romeijn *et al* [13] initially applied CG to radiotherapy to solve a direct aperture optimization problem. Dong *et al* [45] then proposed a greedy algorithm based on CG, which iteratively adds beam orientations until the desired number of beams has been reached. Rwigema *et al* [46] used CG to find a set of 30 beam orientations for use in 4π treatment planning of stereotactic body radiation therapy for patients with recurrent, locally advanced, or metastatic head-and-neck cancers, to show the superiority of the 4π treatment plans over those created by volumetric modulated arc therapy. Nguyen *et al* [47] used CG to solve the triplet beam orientation selection problem specific to MRI-guided Co-60 radiotherapy. Yu *et al* [32] used an in-house CG algorithm to solve an integrated problem of beam orientation and FMO.

However, CG is a greedy algorithm that cannot guarantee optimal results and typically yields a suboptimal solution. In addition, CG still takes as much as 560 s to suggest a five-beam plan for prostate IMRT [48]. The aim of this work is to find a method that can explore a larger area of the BOO decision space to find higher quality solutions than those achieved by CG in a shorter amount of time. The method we propose here starts with a deep neural network (DNN) that has been trained using CG as a supervisor. This network can mimic the behavior of CG by directly learning CG's fitness evaluations of beam orientations via supervised learning (SL). Our previous work presents the efficiency of this supervised network, which can propose a set of beam angles that is better than that proposed by CG in less than 2 s [48]. Given a set of already selected beams, this network will predict the fitness value of each beam, which is how much the beam will improve the objective function when added in the next iteration.

In this study, we extend our previous work and combine this trained SL network with a reinforcement learning method called the Monte Carlo Tree Search (MCTS). We use these fitness values from the SL network as a guide to efficiently navigate the action space of the reinforcement learning tree. Specifically, this method provides the probability of selecting a beam in the search space of the tree at each iteration, so that the beam with the better likelihood to improve the objective function has the higher chance of being selected at each step. To evaluate our proposed method, we compared its performance against the state-of-the-art CG. We developed three additional combinations of guided and Random Search (RandS) Tree approaches for comparison.

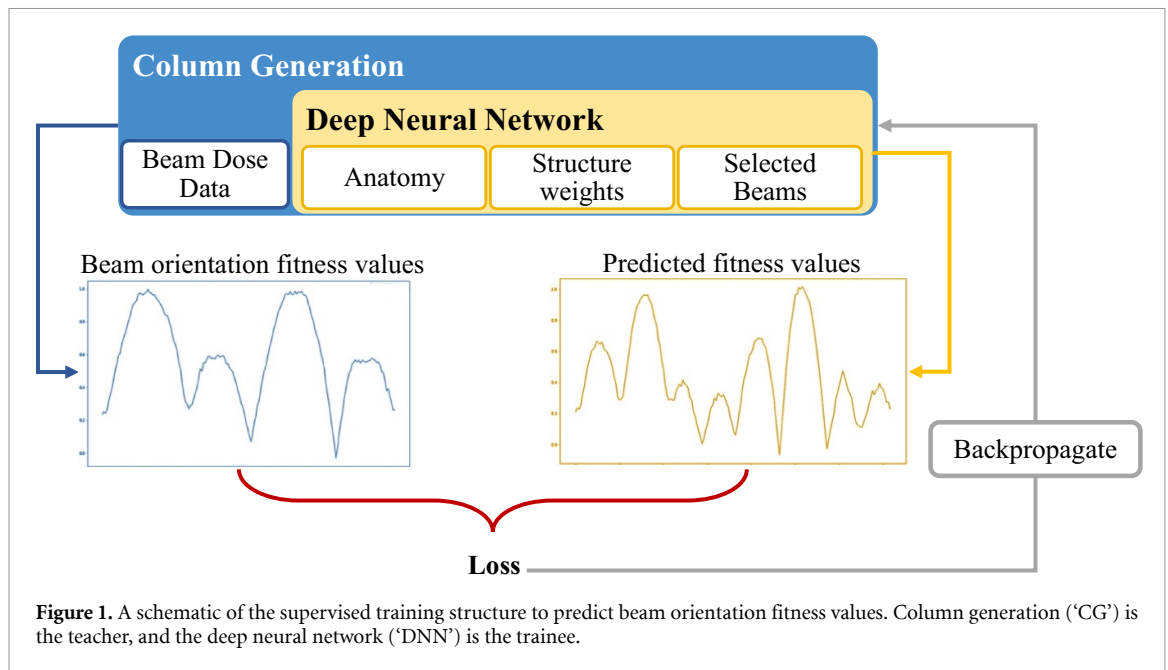
2. Theoretical and experimental methods

The proposed method has a reinforcement learning structure involving an SL network that guides the MCTS to explore the beam orientation selection decision space.

This method, which we named guided MCTS (GTS), consists of two main phases: (a) supervised training of a DNN to predict the probability distribution for adding the next beam, based on patient anatomy, and (b) using this network for a GTS method to explore a larger decision space more efficiently to find better solutions than CG. For the first phase, we use the CG implementation for BOO problems, where CG iteratively solves a sequence of FMO problems [6] using a graphics processing unit (GPU)-based Chambolle–Pock algorithm [49]; the results of the CG method are used to train a supervised neural network. For the second phase, which is the main focus of this work, we present an MCTS algorithm that uses the trained DNN. Each of these phases is presented in the following sections.

2.1. Supervised training of the DNN

We developed a DNN model that learns from CG how to find fitness values for each beam based on the anatomical features of a patient and a set of structure weights for the planning target volume (PTV) and



organs-at-risk (OARs). The greedy CG algorithm starts with an empty set of selected beams, then calculates the fitness value of each beam based on the optimality condition of the objective function shown in equation (1):

$$\min_x \frac{1}{2} \sum_{s \in S} w_s^2 D_s x - p^2 \quad s.t. \ x \geq 0 \quad (1)$$

where w_s is the weight for structure s , which is pseudo-randomly generated between zero and one during the training process to generate many different scenarios. The value p is the prescription dose for each structure, which is assigned 1 for the PTV and 0 for OARs. At each iteration of CG, fitness values are calculated based on Karush–Kuhn–Tucker conditions [50, 51] of a master problem, and they represent how much each beam would improve the objective function value. The beam with the highest fitness value is added to the selected beam set S . Then, FMO is performed for the selected beams, which affects the fitness value calculations for the next iteration. The process is repeated until the desired number of beams has been selected. The supervised DNN learns to mimic this behavior through training, as shown in figure 1. Once trained, this DNN can efficiently provide a suitable set of beam angles in less than 2 s, as opposed to the 560 s that CG needs to solve the same problem. The details of the DNN structure and its training process are described in our previous work [48].

For this study, the patients' anatomical features included the contoured structures (OARs) on images from patients with prostate cancer and the treatment planning weights assigned to each structure.

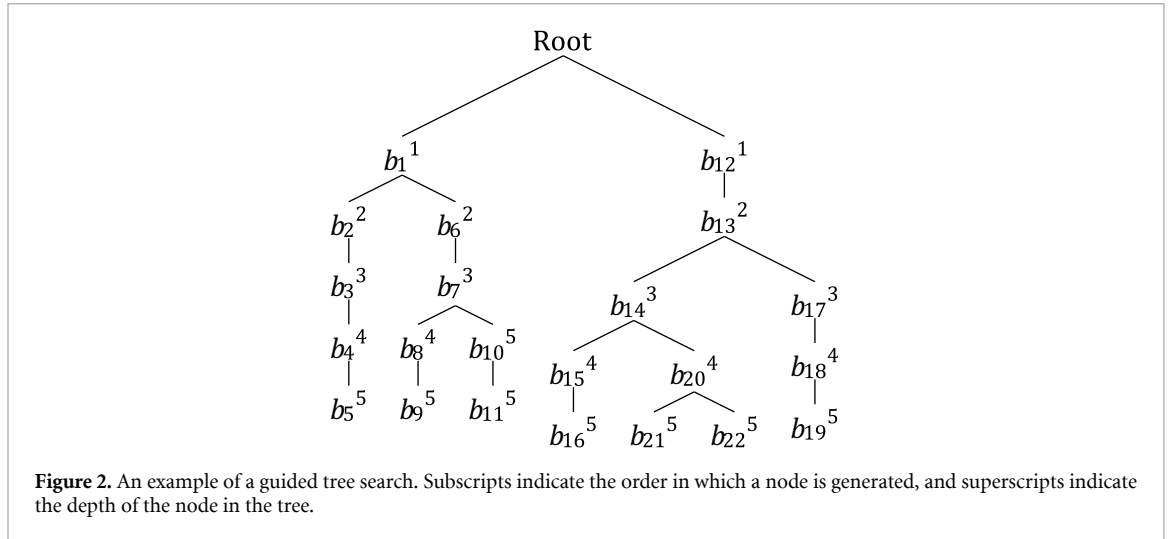
We used the images of 70 prostate cancer patients for this research, each with six contours: PTV, body, bladder, rectum, and left and right femoral heads. We randomly selected 50 of these 70 patients to train the network and seven patients to validate it. We used the images from the remaining 13 patients to test and apply the MCTS method.

2.2. Monte Carlo Tree Search

The pre-trained DNN probabilistically guides the traversal of the branches on the Monte Carlo decision tree to add a new beam to the plan. Each branch of the tree starts from the root as an empty set of beams and continues until it reaches the terminal state. After each complete plan (selection of five beams, in our case) has been explored, the FMO problem is solved; based on its solution, the probability distribution to select the next beam will be updated, using the reward function, in the backpropagation stage. Then, starting from the root, the exploration of the next plan will begin until the stopping criterion is met. Figure 2 shows an example of a tree search that has discovered seven plans so far.

2.2.1. Basics of MCTS

The MCTS uses a decision tree to explore the decision space by randomly sampling from it [52]. The search process of the MCTS consists of four steps: (a) node selection, (b) expansion, (c) simulation, and (d)



backpropagation of the simulation result. To explain these processes in detail, we need to define some properties first:

State of the problem: includes patient's anatomical features and a set of selected beam orientations (B). At the beginning of the planning, this set has no members, and it is updated throughout the solution procedure.

Actions: the selection of the next beam orientation to be added to set B , given the state of the problem.

Solution or terminal state: the state of the problem in which the number of selected beam orientations (size of B) is the same as a predefined number of beams (N) chosen by the user. At this point, a feasible solution for the problem is generated.

Decision tree: the solution space of a set of discrete numbers—in this work, the beam orientations—especially with iterative structures, can be defined as a tree, where each node and branch represent the selection of one number and a subset of available numbers, respectively.

Node (Y): selection of one potential beam orientation.

Root (O): a node with an empty set of beam orientations. Every solution starts from the root.

Path: a unique connected sequence of nodes in a decision tree.

Branch (Q): a path originating from the root node. Each branch represents the iterative structure of the proposed method. The length of a branch is the number of nodes in that branch. In this work, a solution is a branch with size $N + 1$. There is only one branch from each root to any node in a tree.

Leaf: the last node of a branch. There is no exploration of the tree after a leaf is discovered.

Internal node: any node in a tree other than the root and the leaves.

The **selection** process in the proposed method is guided by a pre-trained DNN, as described in section 2.1. This DNN is used to probabilistically guide the traversal of the branches on the Monte Carlo decision tree to add a new beam to the plan. At each node—starting from the root node—the DNN is called to predict an array of fitness values for each beam orientation (P). An element of this array $P[i]$ represents the likelihood of selecting the i th beam orientation. For example, if the number of potential beam orientations is 180, with 2° separation, Y would be an array of size 180, and $P[2]$ is the likelihood of selecting the beam orientation in the 2nd position of the potential beam orientations, $P[2] = 4^\circ$. The **expansion** process happens after the selection process at internal nodes to further explore the tree and create child nodes. The traversal approach in the proposed method is depth first, which means that the branch of a node that is visited or created for the first time continues expanding until there are $N + 1$ nodes in the branch. In this case, the selection and expansion processes overlap, because only one child node is created or visited at a time, though a node can be visited multiple times and several children can be generated from one node. The leaf node is the exception, as it does not have any children. Nodes in a branch must be unique, which means that a branch of each external node (Q) can be expanded only to nodes that are not already in the branch. In fact, the BOO problem can be defined as a dynamic programming problem with the formula

$$G_k^S = S \cup \{k\} \cup G_{n^*}^{S \cup \{k\}} \mid n^* = \arg \max_{n > k} V_{G_n^{S \cup \{k\}}} \quad (2)$$

where S is a set of indices for previously selected beams, k is the index of the currently selected beam, and $G_{n^*}^{S \cup \{k\}}$ is a subset of beams to be selected that has the highest reward value. Each **simulation** iteratively

Algorithm 1. Select N beam orientations from M candidate beams.

```

1. Initialization
2. set selected beam as  $B \leftarrow \emptyset$  an empty set, best cost value as infinity ( $V^* \leftarrow \infty$ ), and best selected beam as  $B^* \leftarrow \emptyset$ 
3. create a root node object ($OS) with the following properties:
4. name( $O_{id} \leftarrow Root$ ), Probability distribution( $O_p \leftarrow \emptyset$ ), number of visits( $O_z \leftarrow 0$ ), beam index( $O_b$ )
5. reward( $O_R \leftarrow 0$ ), cost( $O_v \leftarrow \infty$ ), depth( $O_D \leftarrow 0$ ), parent( $O_{parent} \leftarrow \emptyset$ ), children( $O_{children} \leftarrow \emptyset$ )
6. assign root node to current node( $Y^\# \leftarrow O$ )
7. given the set  $B$  as input to DNN, predict an array of fitness values and assign it to root node  $Y^\#_p \leftarrow Pred(DNN, B)$ 
8. set  $stop \leftarrow False$ 
9. End Initialization
10. while  $stop$  is  $False$  do:
11.     choose the next beam index ( $b$ ) using the probability distribution of the current node  $Y^\#_p$ 
12.     create string name as  $ID \leftarrow Y^\#_{id} \# str(b)$   $\triangleright \#$ : string concatenation
13.     update selected beam set  $B \leftarrow B \cup b$ 
14.     if  $ID \notin Y^\#_{children_{id}}$  (current node does not have a child named  $ID$ ) then:
15.         create a new node  $Y$  with  $Y_{id} \leftarrow ID$ 
16.          $Y_{parent} \leftarrow Y^\#, Y_z \leftarrow 1$ 
17.         predict an array of fitness values  $F(DNN, B)$ 
18.         assign predicted values to new node ( $Y_p \leftarrow Pred(DNN, B)$ )
19.          $Y_D \leftarrow Y^\#_D + 1$ 
20.         set beam index( $Y_b \leftarrow b$ )
21.         add  $Y$  as a new child  $Y^\#_{children} \leftarrow Y^\#_{children} \cup Y$ 
22.         update current node  $Y^\# \leftarrow Y$ 
23.     else
24.         update current node ( $Y^\# \leftarrow X | \{X \in Y^\#_{children} \ \& \ X_{id} = ID\}$ )
25.         update visit parameter of current node, ( $Y^\#_z \leftarrow Y^\#_z + 1$ )
26.     end if
27.     if  $|B| = N$  or  $Y^\#_D = N$  then:
28.         Reward Calculations
29.         solve FMO given set  $B$  and save as  $Y^\#_v \leftarrow Fmo(B)$ 
30.         if  $V^* > Y^\#_v$  then:
31.             set  $V^* \leftarrow Y^\#_v, B^* \leftarrow B, Y^\#_R \leftarrow 1$ 
32.         else:
33.              $Y^\#_R \leftarrow (V^* - Y^\#_v) / Y^\#_v + 0.15$ 
34.         end if
35.         while  $Y^\#_{id} \neq Root$  do:
36.              $Y^\# \leftarrow Y^\#_{parent}$ 
37.              $Y^\#_R \leftarrow \sum_{X \in Y^\#_{children}} X_R / |Y^\#_{children}|$ 
38.             for  $X \in Y^\#_{children}$  do:
39.                  $Y^\#_D[X_b] \leftarrow Y^\#_D[X_b] / X_z + \sqrt{\ln X_z / Y^\#_z}$ 
40.                 if  $Y^\#_v > X_v$  then  $Y^\#_v \leftarrow X_v$  end if
41.             end for
42.         end while
43.         End Reward Calculations
44.         if stopping criteria is met then:
45.              $Stop \leftarrow True$ 
46.         else:
47.              $B \leftarrow \emptyset, Y^\# \leftarrow O$ 
48.         end if
49.     end if
50. end while
51. Return  $B^*$  and  $V^*$ 

```

selects a predefined number of beams (N); in this work, $N = 5$. After each complete plan has been explored, the FMO problem is solved and used for the **backpropagation** step, which updates the probability distribution for beam selection.

2.2.2. Main algorithm

The details of the GTS algorithm in the form of a pseudo-code are provided in algorithm 1. Several properties of each node in the proposed tree are updated after the exploration of final states. To simplify the algorithm, these properties are addressed as variables, defined as follows:

Cost (Y_v): after a leaf is discovered, an FMO problem associated with the beams of that branch is solved, where the value of the FMO cost function is the value associated with its corresponding leaf. The cost value of all other nodes (other than leaves) in a tree is the average cost of its sub-branches. For example, in figure 2, the cost value of node b_1^1 is the average cost of nodes b_2^2 and b_6^2 .

Algorithm 2. The Guided Search algorithm (GuidS).

Initialize B as an empty array, best cost value as infinity ($V^* \leftarrow \infty$), and best selected beam as $B^* \leftarrow \emptyset$

1. set current number of beam orientations in B as 0, $N_B \leftarrow 0$
2. set $stop \leftarrow False$
3. **while** $stop = False$ **do**:
4. **Select B using DNN**
5. **while** $N_B < N$ **do**:
6. predict an array of fitness values $P = F(DNN, B)$
7. Select next node (b) with the probability of $P(b)$
8. Update B : $B = B \cup \{b\}$ and $N_B = N_B + 1$
9. **End while**
10. **return** $\{B\}$
11. **End Select B using DNN**
12. solve FMO given set B and save as $V \leftarrow Fmo(B)$
13. **if** $V < V^*$ **then**:
14. $V^* \leftarrow V$ and $B^* \leftarrow B$
15. **end if**
16. **if** stopping criteria is met **then**:
17. $stop \leftarrow True$
18. **else**:
19. $B \leftarrow \emptyset, N_B \leftarrow 0$
20. **end if**
21. **end while**
22. **return** B^* and V^*

Probability distribution (Y_P): an array of size 180 (the number of potential beam orientations), where the i th element of this array represents the chance of improving the current cost value if the tree branches out by selecting the i th beam. After a node is discovered in the tree, this distribution is populated by using a DNN. After the first discovery of a node, Y_P is updated based on the reward values.

Reward (Y_R): a function of the node's cost values and the best cost value ever discovered in the search process. The reward values are updated after each cost calculation and are calculated and updated by the reward calculation procedure defined in line 28 of algorithm 1.

Depth (Y_D): the number of beam orientations selected after node Y is discovered.

Name (Y_{id}): a unique string value as id for each node. This value is the path from the root to node Y .

Beam Set (Y_B): the set of beams selected for a branch starting from the root and ending in node Y .

Parent (Y_{parent}): the node immediately preceding node Y in a branch from the root to Y . With the exception of the root node, all nodes in a tree have one parent.

Children ($Y_{children}$): the node(s) of the sub-branches that immediately proceed from node Y . With the exception of leaves, all nodes in a tree have at least one child.

2.3. Algorithms for performance comparison

We designed four frameworks to show the efficiency of the proposed GTS method by comparison:

Guided Tree Search (GTS): as presented in algorithm 1. This framework used a pre-trained policy network to guide a Monte Carlo decision tree.

Guided Search (GuidS): this framework used the pre-trained network to search the decision space by iteratively choosing one beam based on the probabilities predicted by the policy network. Unlike GTS, the search policy is not updated as the search progresses. This process is detailed in algorithm 2.

Random Sampling Tree Search (RTS): this algorithm is a simple MCTS method that starts with a uniform distribution to randomly select beam orientations, then updates the search policy as the tree search progresses. Note that all of the tree operations used in GTS are also used in this algorithm, but the tree search is not guided by a pre-trained policy network. This method is proposed to show the impact of using a DNN to guide the decision tree.

Random search (RandS): this method searches the decision space with uniformly random probability until the stopping criterion is met. It randomly selects five beam orientations and solves the corresponding FMO problem. The search policy is not updated. Its procedure is close to algorithm 2, but the 'Select B using DNN' procedure is replaced by randomly selecting five unique beams.

2.4. Data

We used images from 70 patients with prostate cancer at 5 mm³ voxel size. The exact shape of each patient varied, but ranged from 80 to 168 rows, 124 to 261 columns, and 110 to 219 slices. Each image has six contours: PTV, body, bladder, rectum, left femoral head, and right femoral head. The segmentation masks for each structure were contoured by physicians and dosimetrists on the CT images during the clinical workflow. Additionally, the skin and shell¹ tuning structures were added during the FMO process to control high dose spillage and conformity in the body. The dose influence arrays were calculated for 70 patients, which included 180 candidate coplanar beams equidistantly spaced 2 degrees apart, and with a beamlet size of 2.5 mm² at 100 cm isocenter. These dose influence data were then used to generate the various plans with varying beams through CG and FMO optimization. CG was implemented with a GPU-based Chambolle–Pock algorithm [49], a first-order primal-dual proximal-class algorithm, to create scenarios. The patients were divided randomly into two exclusive sets: (a) a model development set consisting of 57 patients, 50 for training and seven for validation, and (b) a test data set consisting of 13 patients. The total number of scenarios created were, 6270 unique training and validation scenarios (ten five-beam plans for each of the 57 patients = 10 × 5 × 57, and then 12 randomly generated beam sets from zero input to four initial beam input scenarios = 12 × 5 × 57). To generate multiple scenarios for each patient, other than randomly selecting the input beam set with sizes from zero to four, the weights of the structures were also generated pseudo randomly, as described in the next paragraph) and 650 test scenarios (ten five-beam plans for each of the 13 test patients = 10 × 5 × 13). To feed into the model, the data were zero-padded and interpolated into a 64 × 64 × 64 shape. The aspect ratio of original data was maintained by the zero padding. We also used 90° rotation of images to augment the training data four times. The DNN trained over 400 epochs, each with 2500 steps and a batch size of one.

We evaluated the performances of four methods: GTS, GuidS, RTS, and RandS, explained in subsection 2.3. Two of these methods, GTS and GuidS, use the pre-trained DNN as a guidance network. We started with the images of 70 patients with prostate cancer, and we used the images of 57 of them to train and validate the DNN; therefore, we could not use them for testing in this project, to keep the proposed methods completely independent from the dataset used for training the DNN. We used the images of the 13 patients that the DNN had never seen before as the test set. Multiple scenarios can be generated for each patient, based on the weights assigned to patient's structures to plan their treatments. We semi-randomly generated ten sets of weights for each patient. In total, we had 130 test plans across 13 test patients for comparison. All the tests in this paper were performed on a computer with an Intel Core I7 processor at 3.6 GHz, 64 GB memory, and an NVIDIA GeForce GTX 1080 Ti GPU with 11 GB video memory.

The structure weight selection scheme is outlined as follows:

- In 50% of the cases, we used a uniform distribution in the range of 0–0.1 to generate a weight for each OAR separately.
- In 10% of the cases, we used a uniform distribution in the smaller range of 0–0.05 to select weights for OARs separately.
- Finally, in 40% of the cases, we used specific ranges for each OAR. Bladder: [0,0.2], rectum: [0,0.2], right femoral head: [0,0.1], left femoral head: [0,0.1], shell: [0,0.1] and skin: [0,0.3].

The weights range from 0 to 1. We found that this weighting scheme is likely to produce a dose distribution that falls into clinically meaningful bounds, but the dose itself may not be approved by the physician for that patient.

Finally, considering only the test scenarios, we compared the FMO solutions of beam sets generated by CG and by the four tree search methods according to the following metrics:

PTV D₉₈, PTV D₉₉: the dose that 98% and 99%, respectively, of the PTV received.

PTV D_{max}: the maximum dose received by the PTV; the value of D₂ is considered for this metric.

PTV Homogeneity: $(PTVD_2 - PTVD_{98}) / PTV D_{50}$, where PTV D₂ and D₅₀ are the doses received by 2% and 50%, respectively, of the PTV.

Paddick Conformity Index (CI_{Paddick}) [53, 54]: $(V_{PTV} \cap V_{100\%Iso})^2 / V_{PTV} \times V_{100\%Iso}$, where V_{PTV} is the volume of the PTV, and V_{100%Iso} is the volume of the isodose region that received 100% of the dose.

High Dose Spillage (R₅₀): $V_{50\%Iso} / V_{PTV}$, where V_{50%Iso} is the volume of the isodose region that received 50% of the dose.

¹ The shell is a tuning structure around the PTV. The shell thickness should be equal to the radius of the PTV (assuming that the PTV is a sphere). The radius of the PTV is calculated by the following equation: $radius = \left(\frac{|V_s| \times 3/4}{\pi} \right)^{1/3}$, where |V_s| is the number of voxels in the PTV.

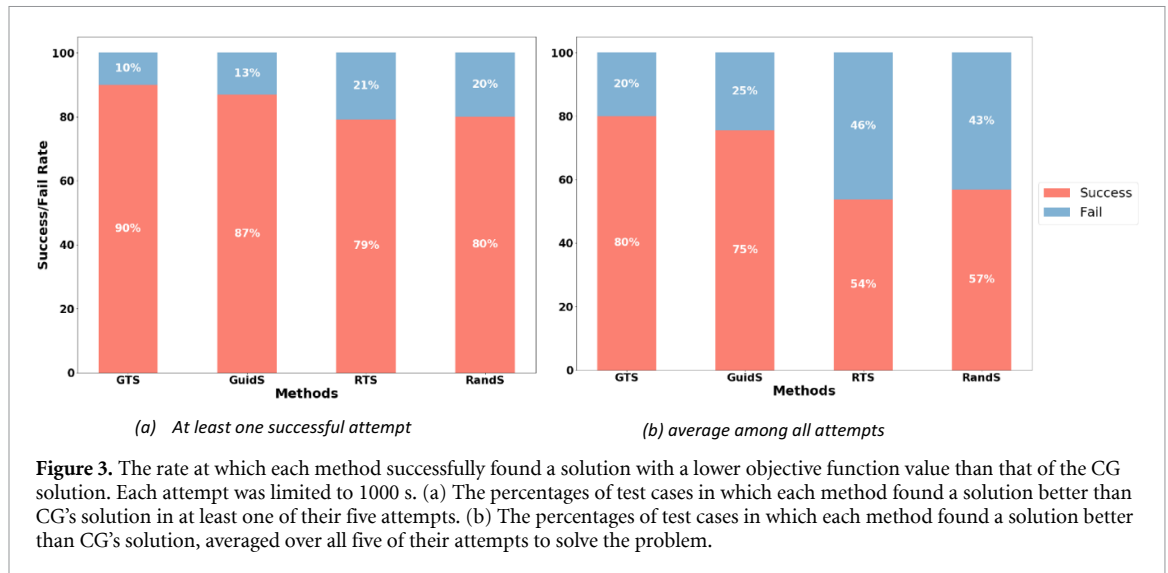


Figure 3. The rate at which each method successfully found a solution with a lower objective function value than that of the CG solution. Each attempt was limited to 1000 s. (a) The percentages of test cases in which each method found a solution better than CG's solution in at least one of their five attempts. (b) The percentages of test cases in which each method found a solution better than CG's solution, averaged over all five of their attempts to solve the problem.

To test and compare the results of the four mentioned methods, for each attempt to solve a test scenario, each method was given 1000 s to search the solution space. Whenever a method found a solution better than CG's, the solution and its corresponding time stamp and the number of total solutions visited by this method were saved. We used these values to analyze the performance of each method. The best solution found in each attempt to solve the problem was used as the final solution of that attempt and for calculating PTV metrics. We used the average objective function value of the final solutions in five attempts to compare the performances of the four methods and the CG solution.

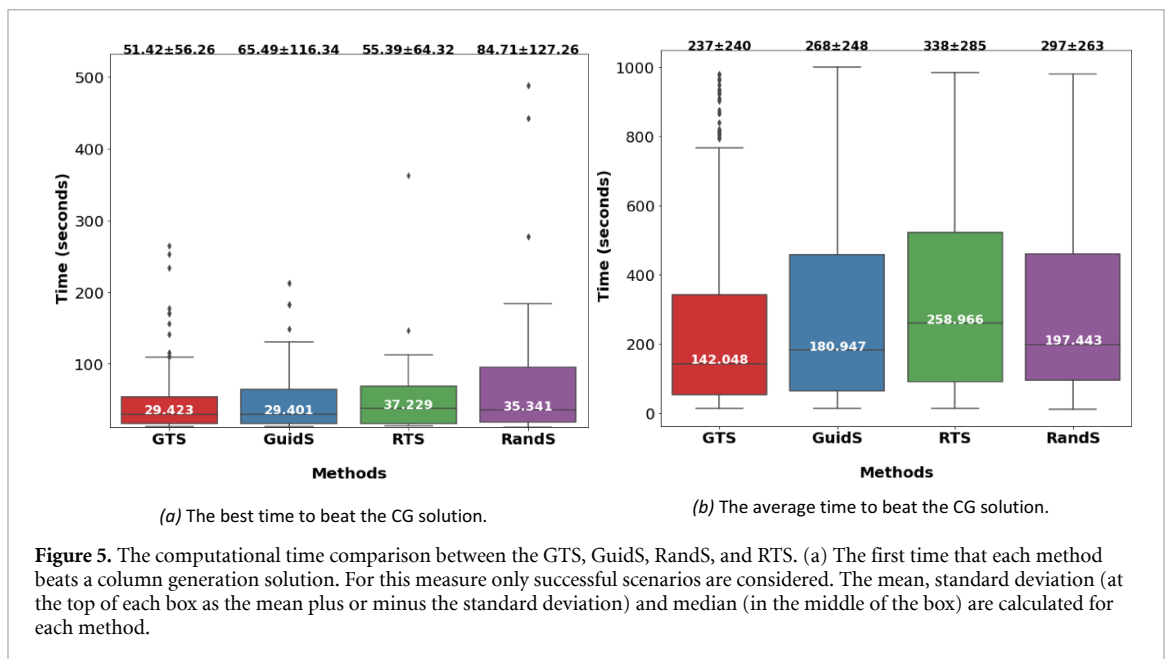
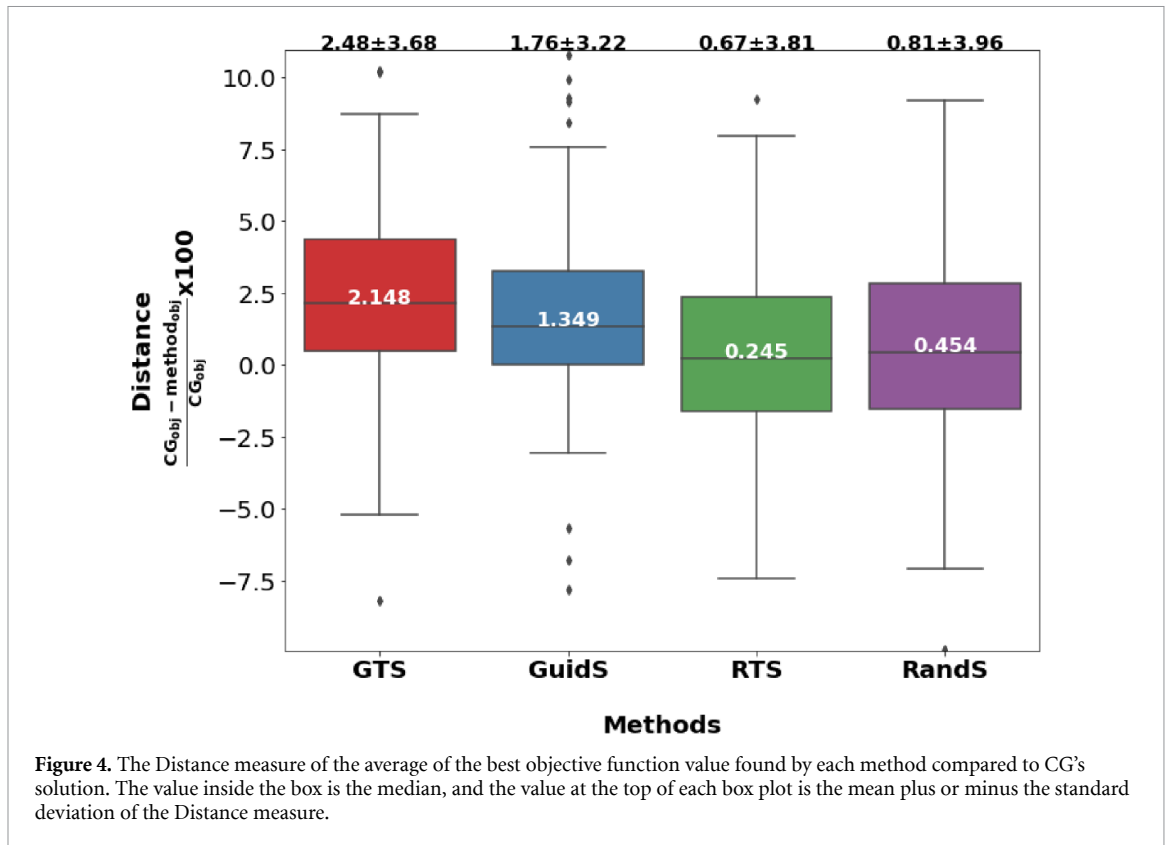
3. Results

First, we compared the efficiency of the four GTS, GuidS, RTS, and RandS methods. Although the main purpose of these methods is to find a solution better than CG, there were some cases where none of these methods beat the CG solution: either the CG solution was very close to optimal, or there were several local optima with a wide search space, which are difficult to explore efficiently, especially for the RTS and RandS methods. The percentages of the total number of attempts required for each method to successfully find a solution better than CG in at least one of five attempts and averaged over all attempts to solve the problem are presented in figures 3(a) and (b), respectively. Note that for this test, the stopping criterion was 1000 s of computational time. As we expected, GTS and GuidS, which use the pre-trained DNN, performed better than the other two methods. However, there were still cases where they could not find a better solution than CG. The maximum number of scenarios in which all four methods successfully found a solution with an objective function value better than CG's solution was 102 out of a total of 130 test cases (78.46%).

The domain of the objective function value varied for different test-case scenarios, so we introduced the **Distance** measure to normalize the objective values of the four methods for further comparison with the CG solutions. The 'Distance' measure is the difference between the objective function value of each method and that of CG, divided by the CG objective value ($\text{Distance} = \frac{\text{CG}_{\text{obj}} - \text{method}_{\text{obj}}}{\text{CG}_{\text{obj}}} \times 100$).

If a method finds a solution better than CG, the Distance measure will be positive; if a method does not find a solution better than CG, this value will be negative. This means that the method with the largest Distance measure found the solutions with the best quality—with objective values smaller than those of CG—more efficiently than CG, as the time was limited to 1000 s. Figure 4 shows the box plot of Distance measures for the GTS, GuidS, RTS, and RandS methods. As shown in this figure, the average Distance measure for the GTS method was 2.48, which was higher than for the other three methods: 1.76 for GuidS, 0.67 for RTS, and 0.81 for RandS.

For the computational time evaluation and comparison measures, we only considered the 102 test cases where all four methods found a better solution than CG within the 1000 seconds time limit. The boxplot of the best time needed to beat the CG solution is presented in figure 5(a). This represents the first time that a method found a solution better than CG's solution within 1000 seconds across all five attempts. The boxplot of the average time to beat CG across all attempts to solve each test case is provided in figure 5(b). To measure the average time to beat CG, we considered all test cases and all attempts. The fastest method according to this measure was GTS, whose average time to beat CG was 51 seconds in the best case and



237 seconds across all cases. The second fastest method was RTS with an average time of 55 seconds for the best case, while GuidS had an average time of 65 seconds for the best case, and RandS had the worst performance of 84 seconds. However, when the average times of all attempts were considered, GuidS beat CG in 268 seconds, RandS in 297 seconds, and RTS had the worst performance of 338 seconds on average.

Figure 6 shows the search rate for each method; the search rate is calculated by dividing the total number of nodes in the search tree by its computational time. The number presented in this figure represents how many nodes are visited per minute for each method. As shown GTS has the highest rate of visiting nodes. We believe the one by one selection of beams and storing the prediction data throughout the tree can create a more efficient method to explore the tree, unlike GuidS, which needs to use model prediction at each node and may cause the occupation of extra time and space. The next method with the highest rate is RandS.

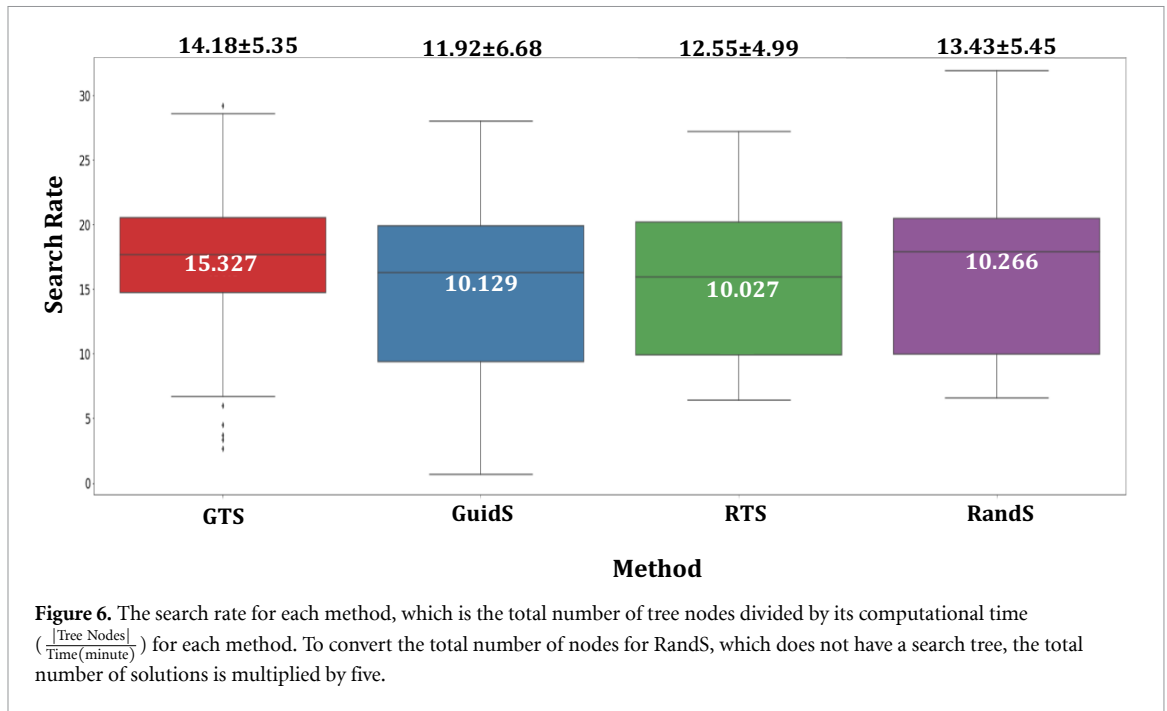


Figure 6. The search rate for each method, which is the total number of tree nodes divided by its computational time $(\frac{\text{Tree Nodes}}{\text{Time (minute)}})$ for each method. To convert the total number of nodes for RandS, which does not have a search tree, the total number of solutions is multiplied by five.

Table 1. A one-tailed paired sample *t*-test to test the average objective function value and Distance measure for every pair of CG, GTS, GuidS, RTS, and RandS methods, with 99% confidence intervals. All highlighted values have *p*-values greater than 0.01.

Tested methods	Objective value		Distance ($\frac{CG-obj}{CG}$)	
	<i>t</i> -statistic	<i>p</i> -value	<i>t</i> -statistic	<i>p</i> -value
CG vs GTS	6.267	2.54×10^{-9}	7.683	1.72×10^{-12}
CG vs GuidS	4.940	1.19×10^{-6}	6.393	1.37×10^{-9}
CG vs RTS	0.885	1.89×10^{-1}	2.014	2.30×10^{-2}
CG vs RandS	1.670	4.87×10^{-2}	2.340	1.04×10^{-2}
RandS vs RTS	-1.496	6.85×10^{-2}	1.096	1.38×10^{-1}
RandS vs GTS	6.826	1.53×10^{-10}	-11.843	1.18×10^{-22}
RandS vs GuidS	3.873	8.51×10^{-5}	-4.839	1.83×10^{-6}
RTS vs GTS	8.245	8.18×10^{-14}	-15.271	5.25×10^{-31}
RTS vs GuidS	5.257	2.95×10^{-7}	-5.832	2.08×10^{-8}
GuidS vs GTS	2.412	8.64×10^{-3}	-3.945	6.52×10^{-5}

RandS can explore more nodes per minute, but this method does not save any history of previous tree traversal, so the computational resources are largely available for this method. However, this method is not guided and therefore not efficient, even by exploring more nodes on the tree, finding better solutions is not guaranteed. Finally, RTS and GuidS have the least rate of visiting nodes, where latter has the lowest rate among all; GuidS is the most computationally demanding method, due to its frequent use of the DNN for each node to be visited. GuidS does not save the history of a node, and therefore needs to use the loaded DNN to predict the next fitness values of each node, while RTS does not need to load and use the DNN, but saves a history of the tree search traversal.

To study the statistical significance of the differences between the performance of GTS and the other four methods (GuidS, RTS, RandS, and CG), we used a one-tailed paired sample *t*-test to compare the objective function values of each pair of methods, as well as the Distance measure. The null hypothesis is that the average objective function and Distance measures of all methods are the same. If we show the null hypothesis as $GTS = GuidS = RandS = RTS = CG = 0$, the alternative hypothesis can be described as $GTS < GuidS < RandS < RTS < CG$ for the objective value parameters and $GTS > GuidS > RandS > RTS > CG$ for the Distance measure. Ten paired sample *t*-tests were performed for objective function values and Distance measures. These statistics are presented in table 1. The distributions of objective values and Distance measures are provided in appendix section. In table 1, the values highlighted in red show that all pairings of CG, RTS, and RandS have *p*-values greater than 0.01 and thus, are not significantly different. In contrast, the average Distance and objective value measures of GuidS and GTS differ significantly from the other methods.

Table 2. The mean \pm standard deviation for PTV Statistics, Paddick Conformity Index (CI_{Paddick}), and high dose spillage (R_{50}) of the GTS, GuidS, RTS, RandS, and CG methods.

Method	PTV D_{98}	PTV D_{99}	PTV D_{max}	PTV		
				homogeneity	CI_{Paddick}	R_{50}
GTS	0.977 \pm 0.011	0.960 \pm 0.019	1.070 \pm 0.040	0.089 \pm 0.045	0.874 \pm 0.061	4.569 \pm 0.994
GuidS	0.976 \pm 0.012	0.960 \pm 0.019	1.071 \pm 0.040	0.089 \pm 0.046	0.874 \pm 0.068	4.487 \pm 0.948
RTS	0.977 \pm 0.011	0.959 \pm 0.020	1.070 \pm 0.039	0.088 \pm 0.045	0.863 \pm 0.085	4.714 \pm 1.312
RandS	0.977 \pm 0.012	0.960 \pm 0.019	1.071 \pm 0.040	0.089 \pm 0.046	0.867 \pm 0.070	4.673 \pm 1.022
CG	0.977 \pm 0.011	0.961 \pm 0.020	1.072 \pm 0.041	0.090 \pm 0.046	0.884 \pm 0.059	4.478 \pm 0.963

Table 3. The average and maximum fractional dose received by each structure in plans generated by the GTS, GuidS, RTS, RandS, and CG methods, where the prescription dose is set to 1. Underlined values highlight the minimum values in each row.

	Structures	Methods				
		GTS	GuidS	RTS	RandS	CG
Mean Dose	PTV	1.039 \pm 0.025	1.039 \pm 0.025	1.038 \pm 0.024	1.039 \pm 0.025	1.040 \pm 0.025
	Body	0.037 \pm 0.012	0.037 \pm 0.012	0.037 \pm 0.012	0.037 \pm 0.012	0.038 \pm 0.013
	Bladder	0.207 \pm 0.125	0.207 \pm 0.122	0.207 \pm 0.126	0.206 \pm 0.122	0.204 \pm 0.116
	Rectum	0.317 \pm 0.109	0.321 \pm 0.111	0.319 \pm 0.110	0.322 \pm 0.115	0.334 \pm 0.116
	L-femoral	0.213 \pm 0.105	0.201 \pm 0.103	0.217 \pm 0.115	0.212 \pm 0.111	0.222 \pm 0.112
	R-femoral	0.214 \pm 0.101	0.221 \pm 0.110	0.227 \pm 0.124	0.224 \pm 0.124	0.217 \pm 0.109
Max Dose	PTV	1.113 \pm 0.055	1.113 \pm 0.055	1.113 \pm 0.055	1.114 \pm 0.057	1.116 \pm 0.058
	Body	1.190 \pm 0.131	1.195 \pm 0.148	1.200 \pm 0.147	1.199 \pm 0.143	1.173 \pm 0.130
	Bladder	1.094 \pm 0.046	1.094 \pm 0.048	1.096 \pm 0.050	1.094 \pm 0.045	1.095 \pm 0.048
	Rectum	1.072 \pm 0.045	1.071 \pm 0.044	1.073 \pm 0.045	1.074 \pm 0.048	1.071 \pm 0.040
	L-femoral	0.609 \pm 0.193	0.596 \pm 0.209	0.619 \pm 0.215	0.609 \pm 0.220	0.613 \pm 0.193
	R-femoral	0.639 \pm 0.242	0.650 \pm 0.249	0.625 \pm 0.236	0.628 \pm 0.244	0.617 \pm 0.245

These results show that GTS outperforms all the other methods significantly and GuidS performs second best, as we expected.

The PTV statistics, Paddick Conformity Index (CI_{Paddick}) and dose spillage (R_{50}) of plans generated by CG, GTS, GuidS, RTS, and RandS are presented in table 2. Note that PTV D_2 is used to measure PTV D_{max} , as recommended by the International Commission on Radiation Units and Measurements (ICRU) Report 83 [55]. The plans generated by all the methods have very similar PTV coverage. CG plans have the highest CI_{Paddick} , followed by the GTS and GuidS plans. The CG and GuidS plans have the lowest dose spillage values, followed by GTS.

The average and maximum doses received by each structure are provided in table 3. These values reflect the fractional dose in plans generated by each method with the assumption that the prescription dose is one; e.g. if the prescription dose is 70 Gy, the average dose of 0.207 in the table means 14.47 Gy (0.207×70) in the prescribed plan. The minimum values in each row are shown in bold for easier interpretation. On average, plans generated by GTS have lower mean doses to OARs than the other methods, while plans generated by CG have the lowest maximum doses to OARs. GTS plans spare the rectum and right femoral head better than the other methods. Although the average fractional dose to the bladder by GTS plans (0.207) is higher than by CG plans (0.204), GTS plans have a lower maximum fractional dose to bladder (1.094) than CG plans (1.095). GuidS plans have the lowest average fractional dose to the left femoral head (0.201), which is considerably lower than for RandS plans (0.212), which performed second best. As an example, figure 7 shows the dose volumes and dose washes of plans generated by GTS and CG for one test case scenario.

4. Discussion

In this research, we propose an efficient BOO framework that can find a better solution than CG in a shorter amount of time by utilizing a reinforcement learning structure involving an SL network, DNN, to guide the MCTS to explore the beam orientation selection decision space.

Although CG is a powerful optimization tool, it is a greedy approach that is computationally expensive and time consuming, and it may also get stuck in a local optimum. This is particularly true for highly non-linear optimization problems with many local optima, such as BOO. In this work, we tried to improve upon CG via four different approaches: (a) GTS, (b) GuidS, (c) RTS, and (d) RandS. Although the quality of solutions obtained using RandS, RTS and CG did not differ significantly, both RandS and RTS, which have

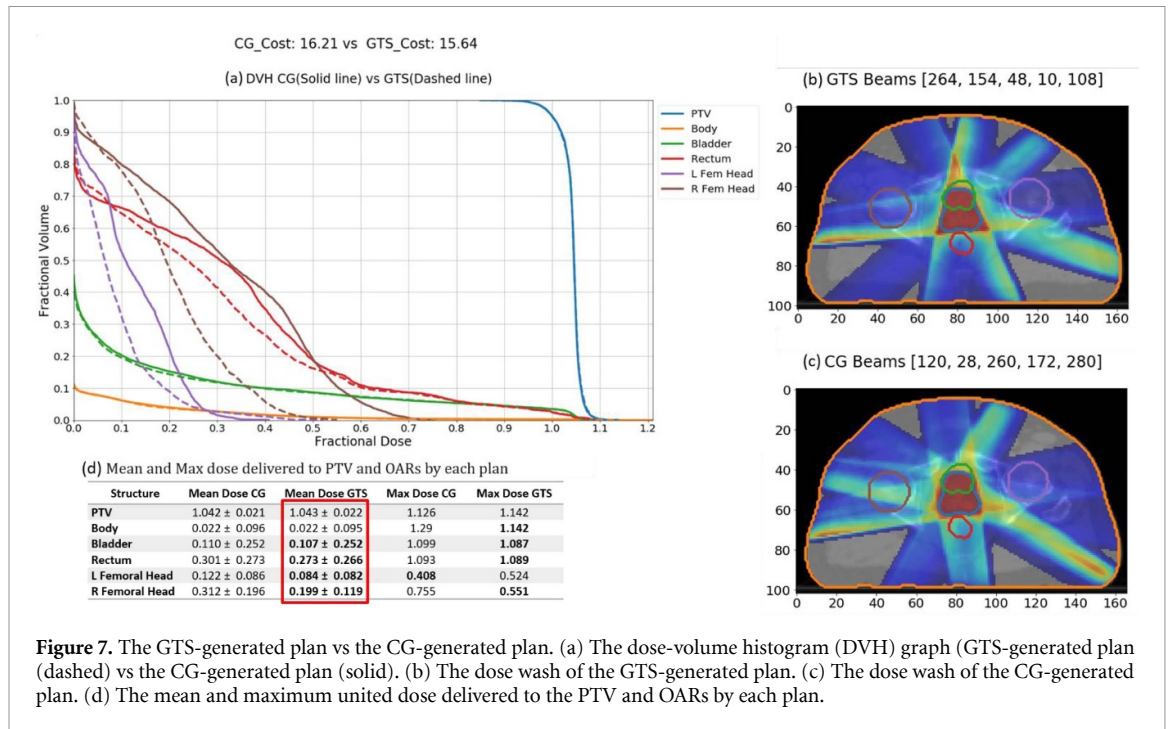


Figure 7. The GTS-generated plan vs the CG-generated plan. (a) The dose-volume histogram (DVH) graph (GTS-generated plan (dashed) vs the CG-generated plan (solid)). (b) The dose wash of the GTS-generated plan. (c) The dose wash of the CG-generated plan. (d) The mean and maximum united dose delivered to the PTV and OARs by each plan.

no knowledge of the problem at the beginning of the search, found solutions better than CG in 50% of the test cases. This shows the high potential to improve upon the solution found by CG.

We saw that GTS and GuidS both performed better than the other methods, which we expected because both of these methods use prior knowledge (a trained DNN) to explore the solution space. GTS outperforms even GuidS on average, because GTS combines GuidS with RTS, which means that adding a search method to GuidS can improve the quality of the solution. But considering the insignificant difference between the performances of RTS and RandS, adding any search method to GuidS may result in better solutions, and the improvement may not be directly related to RTS. This issue will be studied in future research. The poor performance of RTS may also suggest that using a uniform tree search is too slow to converge to the optimal selection of beams.

GTS found solutions with better objective function values than the other methods, but the dose spillage metrics, specifically the average dose received by the bladder in GTS plans, can be improved further. Considering the success of GTS in reducing the objective function and its potential for further improvement, we will continue exploring new methods and techniques to upgrade the quality of treatment planning with the help of artificial intelligence.

We should note that CG is a greedy and deterministic algorithm, so using CG for the same problem always results in the same solution. This is of a completely different nature from our search methods, so it may not be fair to compare its performance with the four search procedures, which, given infinite time and resources, can act as brute force approaches and guarantee finding the optimal solution. However, our main goal is to find the best possible solution for the BOO problem and, in this work, we try to see which search algorithms can find the best solutions the fastest. We expect to see search algorithms outperform greedy algorithms. The resulting objective function values showed us that GTS, GuidS, RTS, RandS, and CG all perform similarly. Plans generated by DNN solutions may not be superior to CG, but the DNN can mimic the CG algorithm very efficiently [48] and explore the search space successfully, as shown by GuidS and GTS, especially when compared to CG, which can easily exhaust computational resources and is very slow to find a single solution for a problem. CG's good performance compared to RandS and RTS shows how powerful the CG method can be to find a solution, and the success of using DNN to explore the decision space represents the proper knowledge that can be achieved by learning from CG.

This work follows the feasibility study of using DNN [48], and is limited to the selection of five beam orientations; however, it can be applied to select a larger number of beam orientations as well, provided that the DNN (the guiding methods) is trained on cases with more than five beam orientations. Note that the current DNN can be used to predict 6th, 7th, etc beams, but the prediction error is expected to be high, since it is not trained in such examples. CG needs to be used to create training examples with six or more numbers of beam orientations, and that requires more computational memory and time. Furthermore, the size of the tree will grow exponentially by selecting more beam orientations, which increases the size of each branch,

and requires more space and time to be processed, and a time limit of 1000 s may not be applicable. Moreover, we study a coplanar BOO problem, which has a limited number of potential beam orientations to begin with (180 beams with 2° separation), but for non-coplanar BOO, the number of potential beams to select from may be a lot larger, for example 1162 in 4π non-coplanar beam selection [56]. Processing larger numbers of potential beam orientations requires even more computational time and space; although it may affect the training time and the architecture of the DNN, the structure of the tree-search proposed here can still be used with no change, but the depth of the tree will be highly affected by the size of the selected beam orientations. Hence, the GTS is suitable for the selection of limited numbers of beams; however, to select a larger set of beam orientations, more research is required.

Finally, GTS is a problem-specific search method that must be applied to each test case separately. To use the knowledge that we can get from the GTS performance, more advanced reinforcement algorithms can be trained to create a single general knowledge-based method that is not only very powerful in finding the best possible solutions, but also very fast in doing so. The advanced reinforcement learning method can then be easily applied to more sophisticated and challenging problems, such as proton and 4π radiation therapy. For future studies, we are working to develop a smart, fast and powerful tool to apply to these problems.

5. Conclusion

In this study, we proposed a Monte Carlo GTS to show that finding a solution with a better objective function compared to CG, in a reasonable amount of time, is feasible. GTS uses a pre-trained DNN to guide the MCTS. The DNN generates the beams' fitness values for nodes in the decision tree, where each node represents a set of selected beams. GTS continues to explore the decision tree for 1000 s. Along with GTS, we tested three other approaches: GuidS, which also uses a pre-trained DNN to select beams iteratively but does not update the probability distribution of beams during the search process; RTS, which is a simple tree search algorithm that starts by randomly sampling from a uniform distribution of beam orientations for each node and continues to update the beams' probability distribution based on the tree search approach presented for GTS; and RandS, which randomly selects beams—this is the most trivial and simple approach.

Data availability statement

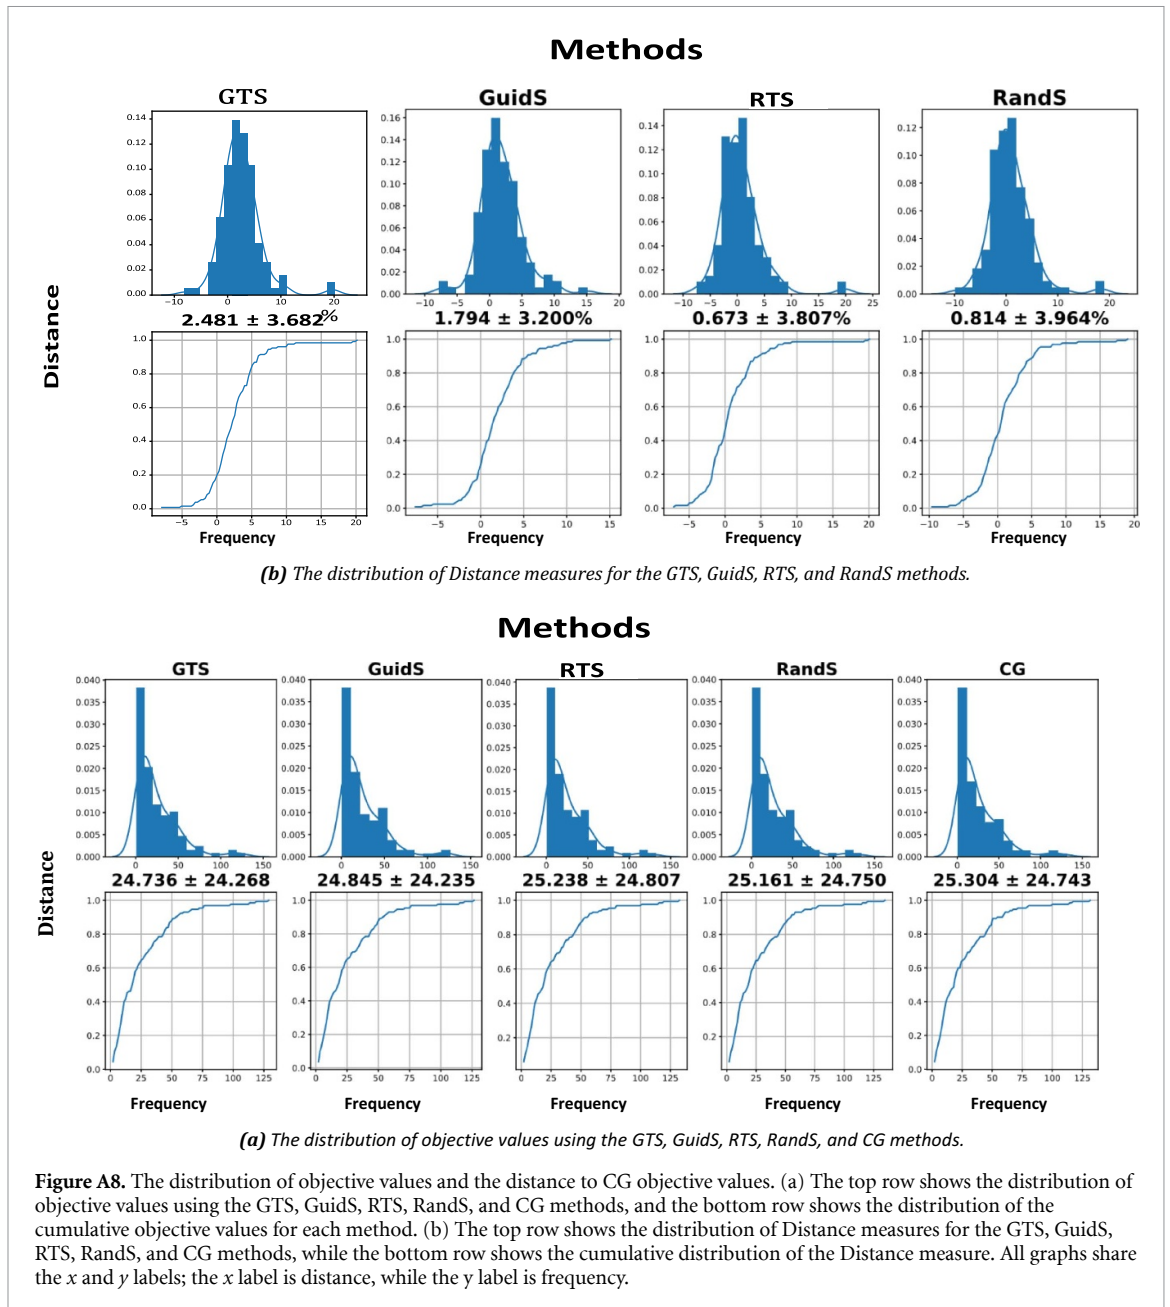
The data that support the findings of this study are available upon reasonable request from the authors.

Acknowledgments

The authors would like to thank Jonathan Feinberg for editing the manuscript. This work was sponsored by NIH Grant No. R01CA237269 Foundation for the National Institutes of Health and the Cancer Prevention and Research Institute of Texas (CPRIT) (IIRA RP150485, MIRA RP160661).

Appendix

Although statistically, with 130 test cases, we can assume that our metrics approximately follow a normal distribution, the graphs in figure 8(a) suggest that this assumption may not be practical. Because of this, we introduced *Distance* measures to normalize our metrics. The probability distribution and cumulative probability mass function of *Distance* measures are presented in figure 8(b). From this graph, we can verify that this measure is approximately normally distributed with a similar standard deviation.



ORCID iDs

Azar Sadeghnejad-Barkousaraie  <https://orcid.org/0000-0002-4331-4988>

Dan Nguyen  <https://orcid.org/0000-0002-9590-0655>

References

- [1] Baskar R, Lee K A, Yeo R and Yeoh K-W 2012 Cancer and radiation therapy: current advances and future directions *Int. J. Med. Sci.* **9** 193–9
- [2] Taylor A and Powell M E B 2004 Intensity-modulated radiotherapy—what is it? *Cancer Imaging* **4** 68–73
- [3] Breedveld S, Storchi P R M, Voet P W J and Heijmen B J M 2012 iCycle: integrated, multicriterial beam angle, and profile optimization for generation of coplanar and noncoplanar IMRT plans *Med. Phys.* **39** 951–63
- [4] Ehr Gott M, Holder A and Reese J 2008 Beam selection in radiotherapy design *Linear Algebr. Appl.* **428** 1272–312
- [5] Bangert M and Oelfke U 2010 Spherical cluster analysis for beam angle optimization in intensity-modulated radiation therapy treatment planning *Phys. Med. Biol.* **55** 6023
- [6] Cabrera G G, Ehr Gott M, Mason A J and Raith A 2018 A metaheuristic approach to solve the multiobjective beam angle optimization problem in intensity-modulated radiation therapy *Int. Trans. Oper. Res.* **25** 243–68
- [7] Bortfeld T and Schlegel W 1993 Optimization of beam orientations in radiation therapy: some theoretical considerations *Phys. Med. Biol.* **38** 291
- [8] Rowbottom C G, Webb S and Oldham M 1999 Beam-orientation customization using an artificial neural network *Phys. Med. Biol.* **44** 2251
- [9] Pugachev A, Li J G, Boyer A L, Hancock S L, Le Q-T, Donaldson S S and Xing L 2001 Role of beam orientation optimization in intensity-modulated radiation therapy *Int. J. Radiat. Oncol. Biol. Phys.* **50** 551–60
- [10] Djajaputra D, Wu Q, Wu Y and Mohan R 2003 Algorithm and performance of a clinical IMRT beam-angle optimization system *Phys. Med. Biol.* **48** 3191
- [11] Yongjie L, Jonathan Y and Dezhong Y 2004 Automatic beam angle selection in IMRT planning using genetic algorithm *Phys. Med. Biol.* **49** 1915
- [12] Li Y, Yao D, Yao J and Chen W 2005 A particle swarm optimization algorithm for beam angle selection in intensity-modulated radiotherapy planning *Phys. Med. Biol.* **50** 3491
- [13] Romeijn H E, Ahuja R K, Dempsey J F and Kumar A 2005 A column generation approach to radiation therapy treatment planning using aperture modulation *SIAM J. Optim.* **15** 838–62
- [14] Schreibmann E and Xing L 2005 Dose–volume based ranking of incident beam direction and its utility in facilitating IMRT beam placement *Int. J. Radiat. Oncol. Biol. Phys.* **63** 584–93
- [15] Aleman D M, Kumar A, Ahuja R K, Romeijn H E and Dempsey J F 2008 Neighborhood search approaches to beam orientation optimization in intensity modulated radiation therapy treatment planning *J. Glob. Optim.* **42** 587–607
- [16] Lim G J, Choi J and Mohan R 2008 Iterative solution methods for beam angle and fluence map optimization in intensity modulated radiation therapy planning *OR Spectrum* **30** 289–309
- [17] Breedveld S, Storchi P and Heijmen B 2009 The equivalence of multi-criteria methods for radiotherapy plan optimization *Phys. Med. Biol.* **54** 7199
- [18] Lim G J, Holder A and Reese J 2009 A clustering approach for optimizing beam angles in IMRT planning *IIE Annual Conf.*
- [19] Craft D and Monz M 2010 Simultaneous navigation of multiple Pareto surfaces, with an application to multicriteria IMRT planning with multiple beam angle configurations *Med. Phys.* **37** 736–41
- [20] Rocha H, Dias J M, Ferreira B C and Lopes M C 2013 Beam angle optimization for intensity-modulated radiation therapy using a guided pattern search method *Phys. Med. Biol.* **58** 2939–53
- [21] Yuan L, Wu Q J, Yin F, Li Y, Sheng Y, Kelsey C R and Ge Y 2015 Standardized beam bouquets for lung IMRT planning *Phys. Med. Biol.* **60** 1831–43
- [22] Amit G, Purdie T G, Levinshstein A, Hope A J, Lindsay P, Marshall A, Jaffray D A and Pekar V 2015 Automatic learning-based beam angle selection for thoracic IMRT *Med. Phys.* **42** 1992–2005
- [23] Liu H, Dong P and Xing L 2017 A new sparse optimization scheme for simultaneous beam angle and fluence map optimization in radiotherapy planning *Phys. Med. Biol.* **62** 6428
- [24] Cabrera-Guerrero G, Lagos C, Cabrera E, Johnson F, Rubio J M and Paredes F 2018 Comparing local search algorithms for the beam angles selection in radiotherapy *IEEE Access* **6** 23701–10
- [25] Rocha H et al 2018 Comparison of combinatorial and continuous frameworks for the beam angle optimization problem in IMRT *Int. Conf. on Computational Science and Its Applications* (Springer) (https://doi.org/10.1007/978-3-319-95165-2_42)
- [26] O'Connor D, Yu V, Nguyen D, Ruan D and Sheng K 2018 Fraction-variant beam orientation optimization for non-coplanar IMRT *Phys. Med. Biol.* **63** 045015
- [27] Cabrera-Guerrero G, Mason A J, Raith A and Ehr Gott M 2018 Pareto local search algorithms for the multi-objective beam angle optimisation problem *J. Heuristics* **24** 205–38
- [28] Ramar N, Meher S R, Ranganathan V, Perumal B, Kumar P, Anto G J and Etti S H 2020 Objective function based ranking method for selection of optimal beam angles in IMRT *Phys. Med.* **69** 44–51
- [29] Ventura T, Rocha H, da Costa Ferreira B, Dias J and Lopes M D C 2019 Comparison of two beam angular optimization algorithms guided by automated multicriterial IMRT *Phys. Med.* **64** 210–21
- [30] Potrebko P S, McCurdy B M C, Butler J B and El-Gubtan A S 2008 Improving intensity-modulated radiation therapy using the anatomic beam orientation optimization algorithm *Med. Phys.* **35** 2170–9
- [31] Jorge L, Li S, Agazaryan N, Promberger C and Solberg T D 2009 Non-coplanar automatic beam orientation selection in cranial IMRT: a practical methodology *Phys. Med. Biol.* **54** 1337
- [32] Yu V Y, Landers A, Woods K, Nguyen D, Cao M, Du D, Chin R K, Sheng K and Kaprelian T B 2018 A prospective 4 π radiation therapy clinical study in recurrent high-grade glioma patients *Int. J. Radiat. Oncol. Biol. Phys.* **101** 144–51
- [33] Yarmand H and Craft D 2013 Effective heuristics for beam angle optimization in radiation therapy *Simulation* **94** 1041–9
- [34] Lulin Y, Zhu W, Ge Y, Jiang Y, Sheng Y, Yin F-F and Wu Q J 2018 Lung IMRT planning with automatic determination of beam angle configurations *Phys. Med. Biol.* **63** 135024
- [35] Rocha H et al 2018 Beam angle optimization in IMRT: are we really optimizing what matters? *Int. Trans. Oper. Res.* **26** 908–28

- [36] Bedford J L, Ziegenhein P, Nill S and Oelfke U 2019 Beam selection for stereotactic ablative radiotherapy using Cyberknife with multileaf collimation *Med. Eng. Phys.* **64** 28–36
- [37] Haseai S, Arimura H, Asai K, Yoshitake T and Shioyama Y 2020 Similar-cases-based planning approaches with beam angle optimizations using water equivalent path length for lung stereotactic body radiation therapy *Radiol. Phys. Technol.* **13** 119–27
- [38] Oelfke U and Bortfeld T 2001 Inverse planning for photon and proton beams *Med. Dosim.* **26** 113–24
- [39] Gu W, O'Connor D, Nguyen D, Yu V Y, Ruan D, Dong L and Sheng K 2018 Integrated beam orientation and scanning-spot optimization in intensity-modulated proton therapy for brain and unilateral head and neck tumors *Med. Phys.* **45** 1338–50
- [40] Shirato H, Le Q-T, Kobashi K, Prayongrat A, Takao S, Shimizu S, Giaccia A, Xing L and Umegaki K 2018 Selection of external beam radiotherapy approaches for precise and accurate cancer treatment *J. Radiat. Res.* **59** i2–i10
- [41] Gu W et al 2019 Robust beam orientation optimization for intensity-modulated proton therapy *Med. Phys.* **46** 3356–70
- [42] Taasti V T, Hong L, Shim J S, Deasy J O and Zarepisheh M 2020 Automating proton treatment planning with beam angle selection using Bayesian optimization *Med. Phys.* **47** 3286–96
- [43] Gu W, O'Connor D, Ruan D, Zou W, Dong L and Sheng K 2020 Fraction-variant beam orientation optimization for intensity-modulated proton therapy *Med. Phys.* **47** 3826–34
- [44] Azizi-Sultan A S 2006 Optimization of beam orientation in intensity modulated radiation therapy planning *PhD Thesis* (Technische Universität Kaiserslautern)
- [45] Dong P, Lee P, Ruan D, Long T, Romeijn E, Yang Y, Low D, Kupelian P and Sheng K 2013 4π non-coplanar liver SBRT: a novel delivery technique *Int. J. Radiat. Oncol. Biol. Phys.* **85** 1360–6
- [46] Rwigema J C M et al 2015 4π noncoplanar stereotactic body radiation therapy for head-and-neck cancer: potential to improve tumor control and late toxicity *Int. J. Radiat. Oncol. Biol. Phys.* **91** 401–9
- [47] Nguyen D, Thomas D, Cao M, O'Connor D, Lamb J and Sheng K 2016 Computerized triplet beam orientation optimization for MRI-guided Co-60 radiotherapy *Med. Phys.* **43** 5667–75
- [48] Sadeghnejad Barkousaraie A et al 2019 A fast deep learning approach for beam orientation optimization for prostate cancer treated with intensity modulated radiation therapy *Med. Phys.* **47** 880–97
- [49] Chambolle A and Pock T 2011 A first-order primal-dual algorithm for convex problems with applications to imaging *J. Math. Imaging Vis.* **40** 120–45
- [50] Kuhn H W and Tucker A W 1951 Nonlinear programming *Proc. Second Berkeley Symposium on Mathematical Statistics and Probability* ed J Neyman (Berkeley: Univ. of California Press) pp 481–92
- [51] Karush W 2014 Minima of functions of several variables with inequalities as side conditions *Traces and Emergence of Nonlinear Programming* 1 Giorgi Giorgio and Kjeldsen Tinne Hoff (Basel: Birkhäuser) pp 217–45
- [52] Browne C B, Powley E, Whitehouse D, Lucas S M, Cowling P I, Rohlfshagen P, Tavener S, Perez D, Samothrakis S and Colton S 2012 A survey of Monte Carlo tree search methods *IEEE Trans. Comput. Intell. AI Games* **4** 1–43
- [53] Riet A V T, Mak A C A, Moerland M A, Elders L H and van der Zee W 1997 A conformation number to quantify the degree of conformality in brachytherapy and external beam irradiation: application to the prostate *Int. J. Radiat. Oncol. Biol. Phys.* **37** 731–6
- [54] Paddick I 2000 A simple scoring ratio to index the conformity of radiosurgical treatment plans *J Neurosurg* **93** 219–22
- [55] Hodapp N 2012 The ICRU Report No. 83: prescribing, recording and reporting photon-beam intensity-modulated radiation therapy (IMRT) *Strahlentherapie und Onkologie* (Springer) pp 97–99
- [56] Nguyen D, Dong P, Long T, Ruan D, Low D A, Romeijn E and Sheng K 2014 Integral dose investigation of non-coplanar treatment beam geometries in radiotherapy *Med. Phys.* **41** 011905