

YOUR PATIENT'S TREATMENT PLAN IS SET...

DO YOU CLOSE YOUR EYES AND MAKE A WISH?

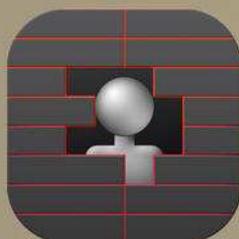


**THAT YOUR PATIENT WON'T CHANGE AT ALL?
AND THE DOSE WILL BE ON TARGET FOR
EVERY FRACTION?**

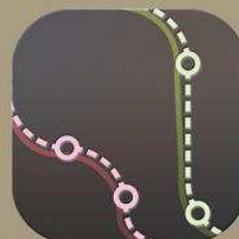
NOW YOU DON'T HAVE TO.



PRETREATMENT



IN VIVO



ADAPTIVE

ADAPTIVO™

THE FUTURE OF
PATIENT DOSIMETRY

Learn more about Adaptivo at:

www.standardimaging.com/adaptivo



Deterministic direct aperture optimization using multiphase piecewise constant segmentation

Dan Nguyen, Daniel O'Connor, Dan Ruan, and Ke Sheng^{a)}

Department of Radiation Oncology, University of Los Angeles California, Los Angeles, CA, USA

(Received 27 November 2016; revised 7 June 2017; accepted for publication 11 August 2017; published xx xxxx xxxx)

Purpose: Direct aperture optimization (DAO) attempts to incorporate machine constraints in the inverse optimization to eliminate the post-processing steps in fluence map optimization (FMO) that degrade plan quality. Current commercial DAO methods utilize a stochastic or greedy approach to search a small aperture solution space. In this study, we propose a novel deterministic direct aperture optimization that integrates the segmentation of fluence map in the optimization problem using the multiphase piecewise constant Mumford-Shah formulation.

Methods: The Mumford-Shah based direct aperture optimization problem was formulated to include an L2-norm dose fidelity term to penalize differences between the projected dose and the prescribed dose, an anisotropic total variation term to promote piecewise continuity in the fluence maps, and the multiphase piecewise constant Mumford-Shah function to partition the fluence into pairwise discrete segments. A proximal-class, first-order primal-dual solver was implemented to solve the large scale optimization problem, and an alternating module strategy was implemented to update fluence and delivery segments. Three patients of varying complexity—one glioblastoma multiforme (GBM) patient, one lung (LNG) patient, and one bilateral head and neck (H&N) patient with 3 PTVs—were selected to test the new DAO method. For each patient, 20 non-coplanar beams were first selected using column generation, followed by the Mumford-Shah based DAO (DAO_{MS}). For comparison, a popular and successful approach to DAO known as simulated annealing—a stochastic approach—was replicated. The simulated annealing DAO (DAO_{SA}) plans were then created using the same beam angles and maximum number of segments per beam. PTV coverage, PTV homogeneity ($\frac{D_{95}}{D_5}$), and OAR sparing were assessed for each plan. In addition, high dose spillage, defined as the 50% isodose volume divided by the tumor volume, as well as conformity, defined as the van't Riet conformation number, were evaluated.

Results: DAO_{MS} achieved essentially the same OAR doses compared with the DAO_{SA} plans for the GBM case. The average difference of OAR D_{max} and D_{mean} between the two plans were within 0.05% of the plan prescription dose. The lung case showed slightly improved critical structure sparing using the DAO_{MS} approach, where the average OAR D_{max} and D_{mean} were reduced by 3.67% and 1.08%, respectively, of the prescription dose. The DAO_{MS} plan substantially improved OAR dose sparing for the H&N patient, where the average OAR D_{max} and D_{mean} were reduced by over 10% of the prescription dose. The DAO_{MS} and DAO_{SA} plans were comparable for the GBM and LNG PTV coverage, while the DAO_{MS} plan substantially improved the H&N PTV coverage, increasing D_{99} by 6.98% of the prescription dose. For the GBM and LNG patients, the DAO_{MS} and DAO_{SA} plans had comparable high dose spillage but slightly worse conformity with the DAO_{MS} approach. For the H&N plan, DAO_{MS} was considerably superior in high dose spillage and conformity to the DAO_{SA}. The deterministic approach is able to solve the DAO problem substantially faster than the simulated annealing approach, with a 9.5- to 40-fold decrease in total solve time, depending on the patient case.

Conclusions: A novel deterministic direct aperture optimization formulation was developed and evaluated. It combines fluence map optimization and the multiphase piecewise constant Mumford-Shah segmentation into a unified framework, and the resulting optimization problem can be solved efficiently. Compared to the widely and commercially used simulated annealing DAO approach, it showed comparable dosimetry behavior for simple plans, and substantially improved OAR sparing, PTV coverage, PTV homogeneity, high dose spillage, and conformity for the more complex head and neck plan.

© 2017 American Association of Physicists in Medicine [https://doi.org/10.1002/mp.12529]

Key words: direct aperture optimization, piecewise constant segmentation

1. INTRODUCTION

Intensity modulated radiation therapy (IMRT) was theorized in the 1980s¹ and subsequently developed in several landmark papers.²⁻⁴ Since then, IMRT has been widely accepted

into radiation therapy clinics as the staple approach to radiotherapy. While IMRT has been proved exceptional in controlling of dose distributions, one major weakness lies in the fluence map optimization (FMO), which does not consider machine constraints in the inverse optimization and produces

optimal fluence maps that need to be converted to deliverable multileaf collimators (MLC) sequences. While several clever sequencing methods have been developed, such as sliding window or the reducing level method,⁵ this additional post processing step, separate from initial FMO inverse optimization, leads to varying degrees of degradation in the dose quality that often require plan reoptimization, which is ineffective to address the problem.

The process of converting a fluence map into deliverable apertures and be divided into two steps. First, a stratification step is performed, which bins the fluence's intensity values into patches of discrete levels. The patches are referred to as apertures in the context of IMRT. Typically, to reduce the discretization levels, the fluence maps need to be smoothed, more so for complicated and heavily modulated plans. In the second step, MLC motion trajectories are determined to efficiently deliver the apertures. Because of the modification to the optimized fluence map, the first step of stratification contributes most to the dosimetry degradation. Figure 1(a) illustrates the workflow of conventional fluence map optimization, followed by the stratification and the MLC sequencing step.

To alleviate the problem and produce simpler fluence maps that suffer a smaller loss in stratification, a total variation regularization term on the fluence maps was incorporated in the original optimization problem^{6–9} to encourage piecewise smoothness. Computationally, we recently showed that a first order primal dual method is efficient to solve the large scale non-differentiable optimization problem.⁹ However, as shown in Fig. 1(b), additional stratification is still needed to convert the piecewise smooth fluence map to apertures of uniform intensities, degrading the plan from its optimized version.

Direct aperture optimization (DAO) was invented to circumvent the stratification problem by optimizing the apertures instead of the beamlets [Fig. 1(c)]. One significant challenge with DAO is that the number of possible apertures

as the combination of beamlets is mathematically intractable. To manage the challenge, stochastic and greedy approaches have been implemented. The stochastic DAO method utilizes a simulated annealing process¹⁰ to iteratively update the aperture shapes and intensity values.^{11–18} Another stochastic approach of DAO employs a genetic algorithm to find a set deliverable segments with best fitness.¹⁹ Genetic algorithms, inspired by the process of natural selection, operates by having a population of solutions, and then performing genetic operations—such as crossover and mutations—on the fittest solutions in an attempt to increase the fitness value. The stochastic nature of the genetic algorithms stems from the random decisions made during the genetic operations. In the greedy approach, DAO plans are created by optimizing a predetermined library of apertures, and a column generation method to expand the aperture library until an acceptable treatment plan is acquired.^{20–22} These DAO methods result in plans that are directly deliverable by the machine without the need for a separate stratification step but there are significant limitations. Due to the computational cost, these methods can only search a small fraction of the possible apertures, which may be insufficient for complex IMRT plans.

In this study, we aim to overcome the limitations in existing fluence map and direct aperture optimization methods.^{6–9} Instead of incorporating the TV regularization that encourages piecewise smoothness, we incorporate a new regularizer where the fluence map stratification problem is formulated as a piecewise constant segmentation problem. The new regularizer would encourage piecewise constant patches that are equivalent to apertures to eliminate the need for post-optimization stratification [Figure 1(d)].

Piecewise constant segmentation is a well-researched image processing topic, aimed at approximating an original grayscale image by a few uniform-intensity patches. Following the original Mumford-Shah formulation,^{23–26} which was used to find

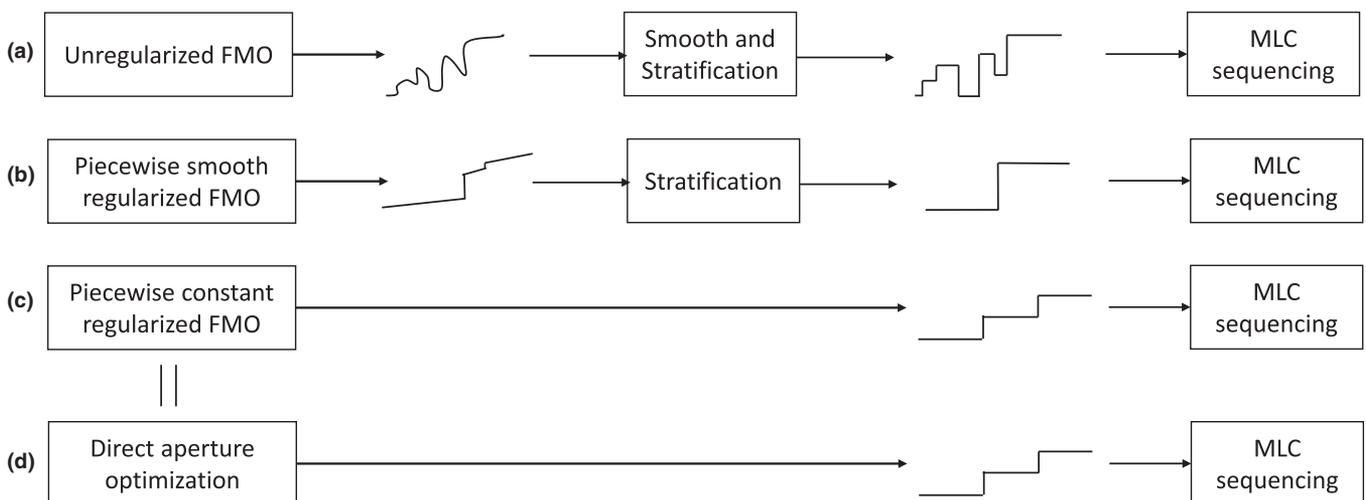


FIG. 1. Schematic of the fluence map optimization and MLC segmentation process. (a) Conventional fluence map optimization followed by smoothing and stratification. Because the smoothing is done post optimization, the plan quality degradation is often substantial. (b) Total variation regularization integrate in the fluence map optimization, resulting smaller losses in stratification and fewer intensity levels. (c) The current direct aperture optimization work flow. (d) The proposed piecewise constant Mumford-Shah segmentation regularized fluence map optimization resulting in a few discretized patches that bypass the stratification step.

the piecewise smooth approximation of grayscale images, recent developments have led to multiphase piecewise constant versions of the formulation.^{24,25,27–30} In particular, the development of the Chan-Vese functional and algorithm,²⁷ which used level-set functions to define segments in the image, achieved the piecewise-constant property by only defining one value per segment. The piecewise constant version of the

Mumford-Shah based DAO and compare it to the simulated annealing DAO are explained.

2.A. Deterministic direct aperture optimization formulation

The master optimization formulation is written as

$$\begin{aligned}
 \underset{x, u, c}{\operatorname{argmin}} \quad & \underbrace{\frac{1}{2} \left\| W \left(\sum_{b=1}^{n_b} (A_b x_b) - d \right) \right\|_2^2}_{\text{Term 1: Dose Fidelity}} + \underbrace{\sum_{b=1}^{n_b} \left(\lambda_1 \|D_b^{\parallel} x_b\|_1 + \lambda_2 \|D_b^{\perp} x_b\|_1 \right)}_{\text{Term Set 2: Anisotropic TV on Fluence}} \\
 & + \gamma \sum_{b=1}^{n_b} \left(\underbrace{\frac{1}{2} \sum_{p=1}^{n_p} \sum_{s=0}^{n_s} u_{bps} |x_{bp} - c_{bs}|^2}_{\text{Term Set 3}} + \underbrace{\sum_{s=0}^{n_s} \left(\omega_1 \|D_b^{\parallel} u_{bs}\|_1 + \omega_2 \|D_b^{\perp} u_{bs}\|_1 \right)}_{\text{Term Set 4}} \right) \\
 & \underbrace{\hspace{15em}}_{\text{Multiphase Piecewise Constant Mumford-Shah Formulation}} \tag{1}
 \end{aligned}$$

subject to

$$\begin{aligned}
 x &\geq 0 \\
 u &\geq 0 \\
 \vec{1}^T u_{bp} &= 1,
 \end{aligned}$$

Mumford-Shah formulation, inspired by Chan-Vese formulation, utilizes a convexly relaxed labeling array instead of level set functions. By finding the minimizer of the Mumford-Shah formulation, an image can be segmented into multiple continuous regions, and a single value is assigned to each region during the optimization for the piecewise constant property. This is pertinent to the IMRT optimization problem, because the piecewise constant segmented regions are arbitrarily shaped, it provides a mathematically complete description of all possible MLC apertures for a specific fluence map.

In this study, we integrate the multiphase piecewise constant Mumford-Shah function with the fluence map optimization problem into a multiconvex formulation^{31,32}—a non-convex problem that yields a convex subproblem when all but one block of variables are held constant. This is commonly evaluated by an alternating module scheme. We then test the feasibility of this optimization on several patient cases and compare its performance to the DAO method using simulated annealing.

2. METHODS

The methods section is organized as follows. First, the novel deterministic DAO formulation is proposed. Second, a notations table of the important variables and data are presented. Third, the block alternating module approach to solving the DAO formulation is explained. Fourth, the simulated annealing DAO approach for comparison is described. Last, the evaluation details to test the performance to the

where the notations for the data and variables used in the DAO formulation are described in Table I.

Term 1 is the dose fidelity term which attempts to optimize the fluence, x , such that the projected dose, Ax , is penalized for deviations from the prescription dose, d . The structures of interest are weighted by the diagonal matrix, W . Term set 2 is the anisotropic total variation regularization on the fluence maps to promote piecewise continuous fluence maps.⁹ The matrices, D^{\parallel} and D^{\perp} , take derivatives parallel and orthogonal to the MLC leaf direction.

Term sets 3 and 4 belong to the multiphase piecewise constant Mumford-Shah formulation,^{23–25,27–30} which partitions each fluence map into n_s discrete segments. The function, for this study, utilizes the anisotropic version of total variation to account for the MLC leaf direction. The convexly relaxed segmentation labeling array, u , has been successfully applied to the Mumford-Shah formulation for piecewise constant segmentation.^{25,30} The rows of the labeling array, u_{bp} , are subject to the probability simplex, which is described by two constraints on u_{bp} —the elements of u_{bp} are non-negative and their sum is equal to 1. This probability simplex constraint is applied to push u_{bps} to equal 1 for the segment where $|x_{bp} - c_{bs}|^2$ is the smallest, but the total variation regularization term encourages each segment in u to be piecewise continuous. The index for segment, s , starts at 0, to account for a “segment” with zero fluence or completely closed MLC. By integrating MLC segmentation with FMO in a single objective function, the beam apertures become optimizable variables directly contributing to the dosimetry.

TABLE I. Important notations and data structures used in the study.

| Notation | Type | Description |
|-------------------------------|--------|--|
| Indices | | |
| B | Index | Index for beam. $b \in \{1, 2, \dots, n_b\}$ |
| P | Index | Index for beamlet. $p \in \{1, 2, \dots, n_p\}$ |
| S | Index | Index for segment. $s \in \{0, 1, 2, \dots, n_s\}$ |
| Optimization variables | | |
| X | Vector | All beamlet intensities from all selected beams. |
| x_b | Vector | All beamlet intensities for the b th beam. $x = \begin{bmatrix} x_{b=1} \\ \vdots \\ x_{b=n_b} \end{bmatrix}$ |
| x_{bp} | Scalar | Intensity value for the p th beamlet of the b th beam. $x_b = \begin{bmatrix} x_{b,p=1} \\ \vdots \\ x_{b,p=n_p} \end{bmatrix}$ |
| C | Matrix | Intensity values for all segments and beams. |
| c_b | Vector | Intensity values for all the segments of the b th beam. $c = \begin{bmatrix} c_{b=1}^T \\ \vdots \\ c_{b=n_b}^T \end{bmatrix}$ |
| c_{bs} | Scalar | Intensity value for the s th segment of the b th beam. $c = \begin{bmatrix} c_{b=1,s=1} & \cdots & c_{b=1,s=n_s} \\ \vdots & \ddots & \vdots \\ c_{b=n_b,s=1} & \cdots & c_{b=n_b,s=n_s} \end{bmatrix} c_b = \begin{bmatrix} c_{b,s=1} \\ \vdots \\ c_{b,s=n_s} \end{bmatrix}$ |
| U | Matrix | Segmentation labeling array for all selected beams. |
| u_b | Matrix | Segmentation labeling array for the b th beam. $u = \begin{bmatrix} u_{b=1} \\ \vdots \\ u_{b=n_b} \end{bmatrix}$ |
| u_{bs} | Vector | Probability that the beamlets from the b th beam belong to the s th segment. $u_b = [u_{b,s=1} \ \cdots \ u_{b,s=n_s}]$ |
| u_{bp} | Vector | Segment probabilities for the p th beamlet of the b th beam. Lives in probability simplex. $u_b = \begin{bmatrix} u_{b,p=1}^T \\ \vdots \\ u_{b,p=n_p}^T \end{bmatrix}$ |
| u_{bps} | Scalar | Probability that the p th beamlet of the b th beam belongs to the s th segment. $u_b = \begin{bmatrix} u_{b,p=1,s=1} & \cdots & u_{b,p=1,s=n_s} \\ \vdots & \ddots & \vdots \\ u_{b,p=n_p,s=1} & \cdots & u_{b,p=n_p,s=n_s} \end{bmatrix} u_{bs} = \begin{bmatrix} u_{b,p=1,s} \\ \vdots \\ u_{b,p=n_p,s} \end{bmatrix} u_{bp} = \begin{bmatrix} u_{b,p,s=1} \\ \vdots \\ u_{b,p,s=n_s} \end{bmatrix}$ |
| Other data | | |
| A | Matrix | Fluence-to-dose transformation matrix from all selected beams. |
| A_b | Matrix | Fluence-to-dose transformation matrix for the b th beam. $A = [A_{b=1} \ \cdots \ A_{b=n_b}]$ |
| W | Matrix | Weighting matrix for structures of interest. |
| D | Vector | Prescription dose the PTV and zero elsewhere. |
| D^{\parallel} | Matrix | Derivative matrix for all fluence maps from all beams. Takes derivative parallel to the MLC leaf motion. |
| D_b^{\parallel} | Matrix | Derivative matrix for fluence map of b th beam. Takes derivative parallel to the MLC leaf motion. $D^{\parallel} = \begin{bmatrix} D_{b=1}^{\parallel} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & D_{b=n_b}^{\parallel} \end{bmatrix}$ |
| D^{\perp} | Matrix | Derivative matrix for all fluence maps from all beams. Takes derivative perpendicular to the MLC leaf motion. |
| D_b^{\perp} | Matrix | Derivative matrix for fluence map of b th beam. Takes derivative perpendicular to the MLC leaf motion. $D^{\perp} = \begin{bmatrix} D_{b=1}^{\perp} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & D_{b=n_b}^{\perp} \end{bmatrix}$ |
| Q | Matrix | An intermediate variable that has the same structure and subscript notation as u^{\sim} . Is used in the appendix, section A3. $q_{bps} = \frac{1}{2} x_{bp} - c_{bs} ^2$ |

Term set 3 is not convex, rendering the whole objective formulation non-convex. Fortunately, the formulation is multiconvex^{31,32}—by updating one variable while holding the other two constant, we have a convex module. Similar to a previous study,³³ the algorithm is broken down into three modules, which are evaluated in an alternating block fashion.

$$\begin{aligned} \underset{u}{\operatorname{argmin}} \quad & \sum_{b=1}^{n_b} \left(\frac{1}{2} \sum_{p=1}^{n_p} \sum_{s=0}^{n_s} u_{bps} |x_{bp} - c_{bs}|^2 + \sum_{s=1}^{n_s} (\omega_1 \|D_b^\parallel u_{bs}\|_1 + \omega_2 \|D_b^\perp u_{bs}\|_1) \right) \\ \text{subject to} \quad & u \geq 0 \\ & \mathbf{1}^T u_{bp} = 1, \end{aligned} \quad (5)$$

2.A.1 Module 1

Module 1 updates x while holding u and c constant. The formulation for module 1 is written as

$$\begin{aligned} \underset{x}{\operatorname{argmin}} \quad & \frac{1}{2} \left\| W \left(\sum_{b=1}^{n_b} (A_b x_b) - d \right) \right\|_2^2 + \sum_{b=1}^{n_b} \left(\lambda_1 \|D_b^\parallel x_b\|_1 + \lambda_2 \|D_b^\perp x_b\|_1 \right) + \frac{\gamma}{2} \sum_{b=1}^{n_b} \sum_{p=1}^{n_p} \sum_{s=0}^{n_s} u_{bps} |x_{bp} - c_{bs}|^2 \\ \text{subject to} \quad & x \geq 0. \end{aligned} \quad (2)$$

The optimization formulation can be solved with a proximal-class, first-order primal-dual algorithm known as the Chambolle-Pock algorithm.^{25,34} This study uses the overrelaxed version of the algorithm³⁵ for faster convergence rate, as well as a preconditioning process³⁴ to select the step size. A detailed explanation of the Chambolle-Pock algorithm, as well as its application to module 1, can be found in the Appendix.

2.A.2 Module 2

Module 2 updates c , while holding x and u constant. Following the initial update of u , c is first initialized to be

$$c_{bs} = \frac{s}{n_s} \max(x_b) \quad \text{for all } b, s, \quad (3)$$

and the subsequent iterations will update c using a closed form solution²⁴

$$c_{bs} = \frac{u_{bs}^T x_b}{u_{bs}^T \mathbf{1}} \quad \text{for all } b, s. \quad (4)$$

In the case where $u_{b,s=k}^T \mathbf{1} = 0$, its corresponding $c_{b,s=k}$ takes on an average value from its neighbor values. In other words $c_{b,s=k} = \frac{1}{2} (c_{b,s=k-1} + c_{b,s=k+1})$. If this turned off segment happens to be the last segment, meaning that $u_{b,s=n_s}^T \mathbf{1} = 0$, then $c_{b,s=n_s} = 2 * c_{b,s=n_s-1}$. By default, $c_{b,s} = 0$ will always be set at zero to define an “off” segment. This strategy is implemented to assign an otherwise undefined $c_{b,s=k}$ a unique non-zero value, which may be effective in finding useful segments.

2.A.3 Module 3

Module 3 updates u , while holding x and c , constant. The terms in the objective setting that have u are simply the terms that belong to the Mumford-Shah function:

where γ weighting is ignored during module 3 since it does not influence the outcome of the optimization in Eq. (5). The formulation shown in Eq. (5) can also be efficiently evaluated using the preconditioned overrelaxed

Chambolle-Pock algorithm, which is described in the Appendix.

2.A.4 Alternating module schedule and other heuristics to handle non-convexity

To allow for a change in the fluence to immediately affect the average value of its aperture, module 1 and module 2 are alternated back-to-back, updating c as soon as x is updated. After the active modules 1 and 2 update x and c , module 3 will then update u . This routine would repeat until the aperture shapes converged to a constant shape. After analysis with a test case, the number of subiterations for each module was decided to be 100 for modules 1 and 2, and 1500 for module 3.

Once the aperture shapes have converged, a check is taken to see if u has converged to 1s and 0s. If not, then the u_{bps} with the highest value in its vector, $u_{b,p}$, is assigned to 1 and the rest are assigned to 0. A polishing step is added where the intensity values are directly optimized while the apertures are locked. The polishing optimization takes the form of

$$\begin{aligned} \underset{c}{\operatorname{argmin}} \quad & \frac{1}{2} \left\| W \left(\sum_{b=1}^{n_b} (A_b u_b c_b) - d \right) \right\|_2^2 \\ \text{subject to} \quad & c \geq 0, \end{aligned} \quad (6)$$

where the aperture intensities c are the only optimization variables. The polishing step allows for fine tuning of the aperture intensities that has been shown to improve the dosimetry. Ultimately, the aperture shapes and intensities are obtained

from u and c to reconstruct the piecewise constant fluence maps.

To ensure an acceptable local minimum that emphasizes the dose fidelity, a graduated weighting technique was implemented, where γ starts at 0 and is monotonically increased over the course of the optimization, until it reaches a max value γ_{max} . The idea behind this technique is to start with the convex part of the objective function that includes the dose fidelity term. During the optimization, the non-convexity is slowly introduced into the problem, ensuring that the optimized fluence segments still result in high dose fidelity. The γ was chosen to monotonically increase via the equation $\gamma = \gamma_{max} \cdot \min\left(1, \left(\frac{k}{\hat{k}}\right)^8\right)$, where k is the current subiteration number for the updates involving modules 1 and 2, and \hat{k} is a set subiteration number when $\gamma = \gamma_{max}$ for the first time. For all patients in this study, \hat{k} was set to 500, which would allow for the alternating module schedule to cycle through 5 times before $\gamma = \gamma_{max}$. Note that γ is not updated while module 3 is active, since γ does not directly influence the outcome of module 3. A pseudocode is provided in (7) order to provide clarity on the update scheme:

```

 $\hat{k} = 500$ 
 $k = 1$ 
while not converged do
  for 100 iterations
    update  $x$ (Module1)
    update  $c$ (Module2)
     $\gamma = \gamma_{max} \cdot \min\left(1, \left(\frac{k}{\hat{k}}\right)^8\right)$ 
     $k = k + 1$ 
  end for
  for 1500 iterations
    update  $u$ (Module3)
  end for
end while
polishing step

```

2.B. Simulated annealing DAO

Simulated annealing is a stochastic technique developed for non-convex optimization,¹⁰ and has been successfully used for DAO.^{11–17} We developed simulated annealing based DAO for comparison. Starting from an initial set of apertures such as the conformal aperture from beams' eye view, the method randomly selects either to change an aperture intensity or a leaf position for modification. A random number from a Gaussian distribution is sampled to determine the size and direction of the change. The standard deviation, SD , of the Gaussian is defined as

$$SD_{leaf} = 1 + (\vartheta - 1) \frac{1}{(n_{succ} + 1)^{1/T_0^{step}}}, \quad (8)$$

$$SD_{intensity} = (\vartheta - 1) \frac{1}{(n_{succ} + 1)^{1/T_0^{step}}},$$

where ϑ is the initial value, n_{succ} is the number of successes, and T_0^{step} defines the cooling rate.¹⁰ A success is defined as a change that results in a decrease of the cost function. Since the leaf position is quantized by beamlets, SD_{leaf} is defined such that the smallest possible value is 1 beamlet. On the other hand, $SD_{intensity}$ is allowed to diminish to 0.

Once either the aperture intensity or the leaf position is randomly modified, a cost function is evaluated. For this study, the cost function used is the dose fidelity term of Eq. (1):

$$\text{minimize } \frac{1}{2} \left\| W \left(\sum_{b=1}^{n_b} (A_b x_b) - d \right) \right\|_2^2, \quad (9)$$

where x_b a function of the leaf positions and aperture intensities—the variables that simulated annealing will be affecting directly. The theoretical leaf positions are constrained such that the two leaves on the same row will not overlap, and the aperture intensities are constrained to be non-negative. While the cost function in Eq. (9) is convex, the conversion of the leaf positions and aperture intensity information to x_b is non-convex. Any change that results in decreasing the cost value is accepted. If the change increases the cost value, the change can still be accepted with a probability defined as

$$P_{accept} = \phi \frac{1}{(n_{succ} + 1)^{1/T_0^{prob}}}, \quad (10)$$

where ϕ is the initial probability of acceptance, and T_0^{prob} defines the cooling rate. The motivation for allowing changes that increase the cost function is to allow for the jumping out of local minima.

2.C. Evaluation

Three planning cases—one glioblastoma multiforme (GBM) patient, one lung (LNG) patient, and one head and neck (H&N) patient with 3 PTVs at different prescription levels—were evaluated in this study to test the feasibility of the optimization formulation. Using a collapsed cone

TABLE II. Number of allowed segments per beam, prescription dose, and PTV volume for each patient.

| Patient | Number of allowed segments per beam (n_s) | Prescription dose (Gy) | PTV volume (cc) |
|-------------------------|---|------------------------|-----------------|
| Glioblastoma multiforme | 10 | 30 | 57.77 |
| Lung | 10 | 50 | 47.84 |
| Head & neck | 20 | 54 | 197.54 |
| | | 59.4 | 432.56 |
| | | 69.96 | 254.98 |

convolution/superposition code with a 6 MV polyenergetic kernel—with a dose array resolution of $(0.25 \text{ cm})^3$ and a beamlet size of $(0.5 \text{ cm})^2$ —the beamlet dose was calculated for 1162 beams, evenly spaced across the 4π steradian. After beam angles that would cause gantry to couch/patient collisions were removed, a column generation and pricing approach^{22,36–41} was utilized to automatically select 20 beam angles for each patient. The beamlet dose of the selected beams are stored in the dose array, A , for optimization. A shell structure around the PTV and skin structure were added

to every patient to reduce dose spillage outside the PTV. The thickness of the shell structure was calculated based on the equation: $thickness_{shell} = \sqrt[3]{\frac{3}{4\pi}TV}$, where TV is the target volume.

Patient plans were created using Mumford-Shah based DAO (DAO_{MS}) and simulated annealing DAO (DAO_{SA}). The same beam angles and number of allowed segments per beam are used for both DAO_{MS} and DAO_{SA} for the same patient. The allowed number of segments, prescription doses and PTV volumes for three patients are shown in Table II. To ensure an

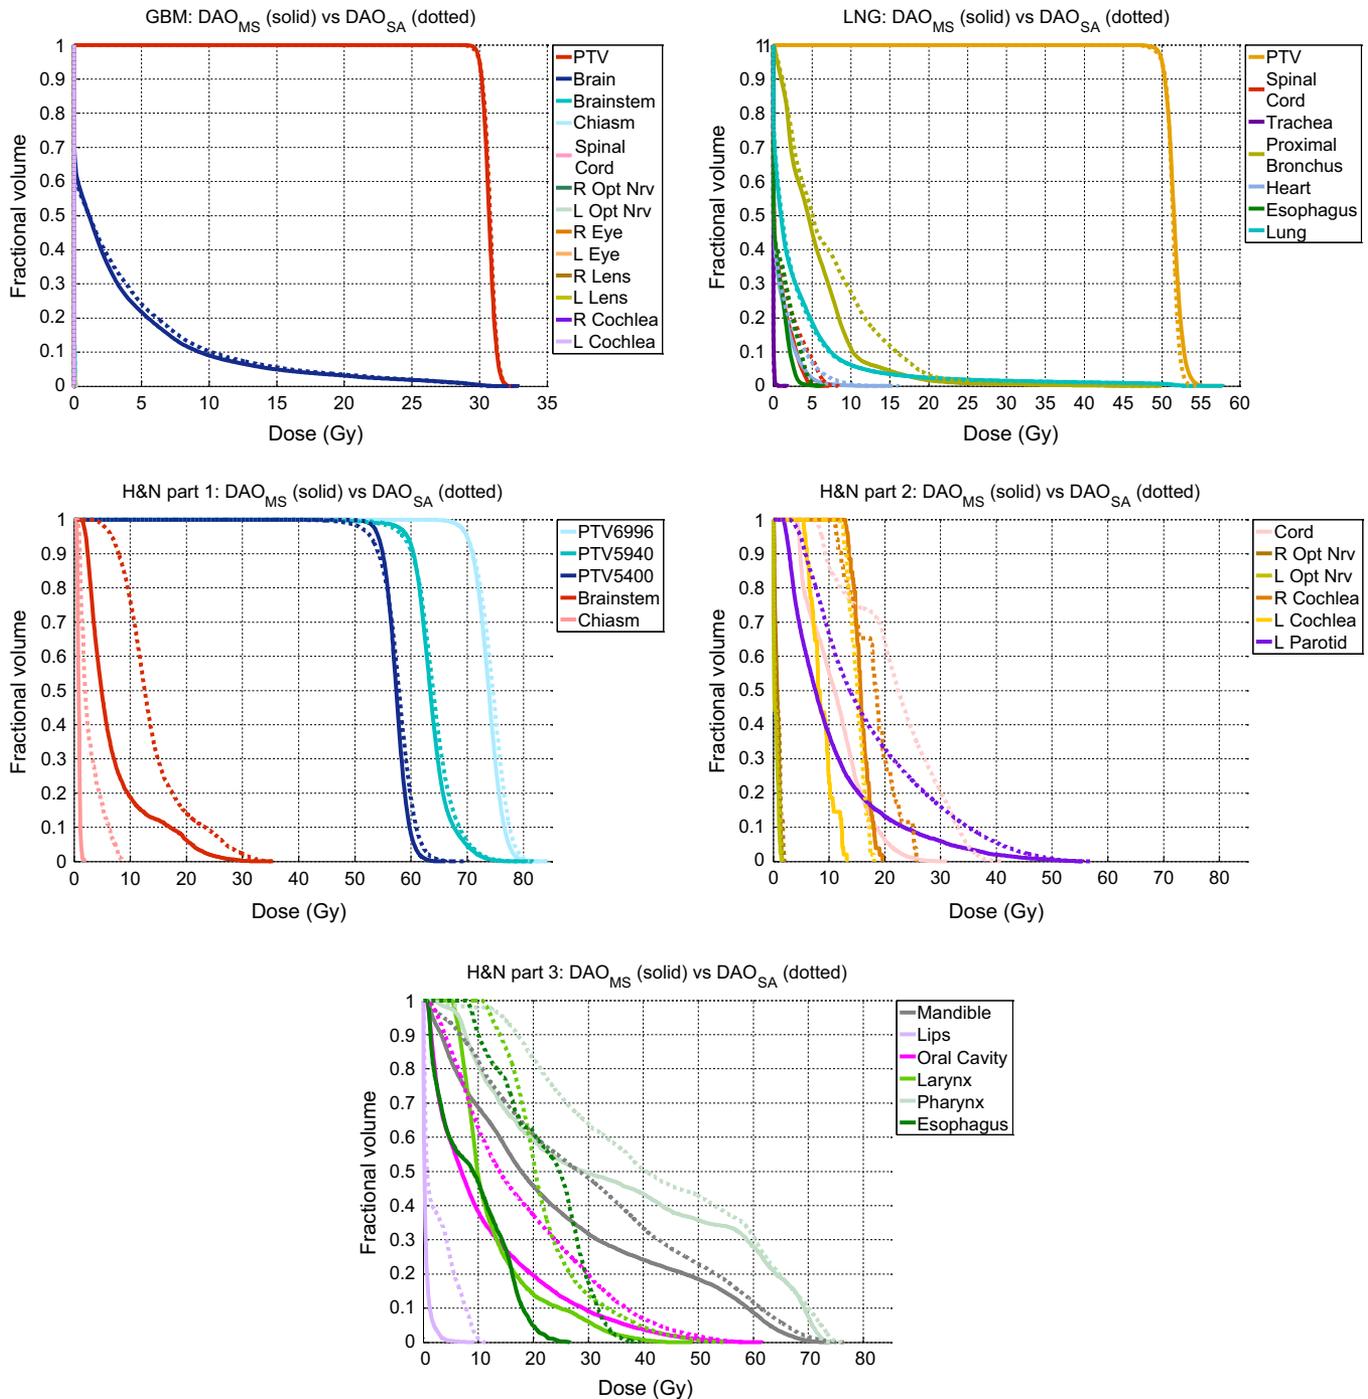


FIG. 2. DVH comparisons of the GBM, LNG, and H&N cases.

unbiased plan comparison, the weighting matrices W of DAO_{MS} and DAO_{SA} plans are iteratively updated to outperform each other in terms of OAR sparing and PTV coverage and homogeneity until further improvement in one aspect is impossible without sacrificing the other. At the end, with the same fidelity term, the structure weightings used by both the DAO_{MS} plans and the DAO_{SA} plans were found to be similar.

For evaluation, all treatment plans were normalized such that 100% of the prescription dose is delivered to 95% of the PTV. For the H&N case, the 69.96 Gy PTV was used for normalization. PTV homogeneity ($\frac{D_{95}}{D_5}$), D98, D99, and D_{max} were evaluated to assess the PTV coverage, hot spots, and homogeneity. OAR D_{max} and D_{mean} were assessed to determine dose sparing to the critical structures. D_{max} is defined as the dose at 2% of the structure volume, D2, which is recommended by the ICRU-83 report.⁴² To quantitate the amount of high dose spillage to the normal tissue, R50, defined as the 50% isodose volume divided by the target volume, was determined. Last, the van't Reit conformation number (VRCN)⁴³ was also assessed. VRCN is defined as $\frac{TV_{RI}}{TV} \times \frac{TV_{RI}}{V_{RI}}$, where TV_{RI} is the target volume covered by the reference isodose, TV is the target volume, and V_{RI} is the volume of the reference isodose. van't Reit conformation number is global conformity index that takes into account both the irradiation of the target volume and healthy tissues, and ranges from 0 to 1, with 1 being the perfect case. For the H&N case, the R50 and VRCN values were calculated using the sum of all 3 PTVs as the target volume and 54 Gy as the reference dose.

The computer used to solve the optimizations has 32 GB RAM, an NVIDIA GeForce GTX 690, and an Intel Core i7-3960X CPU, with six physical cores overclocked to 4.00 GHz.

3. RESULTS

Figure 2 shows the DVHs of all patients. The GBM DAO_{MS} and DAO_{SA} plans are almost identical in OAR doses with difference less than 0.05% of the prescription dose. For the LNG case, the DVH shows that the DAO_{MS} method was able to better spare the proximal bronchus dose, but otherwise, the two methods performed similarly. On average, for the LNG case, the DAO_{MS} method was able to spare proximal bronchus D_{max} by 3.67% and D_{mean} by 1.08% of the prescription dose. For the more complicated H&N plan, the DAO_{MS} plan evidently improved OAR sparing compared to the DAO_{SA} plan, with as much as a 13 Gy reduction to the spinal cord. The improvement is consistent for all OARs in the H&N DAO_{MS} plan. On average, D_{max} and D_{mean} were reduced by 10.91% and 10.58% of the prescription dose. The OAR sparing is summarized in Table III.

Table IV shows PTV dose statistics, as well the dose spillage (R50) and the conformity number (VRCN). On average, the PTV was better covered using DAO_{MS} . D98 and D99 was increased by 1.66% and 2.21% of the prescription dose. The H&N plan had the largest improvement in dose coverage, with an increase as large as 3.77 Gy in D99 for the 54 Gy PTV. The average homogeneity of the two planning methods

TABLE III. Largest, smallest, and average values found for ($D_{DAO_{MS}} - D_{DAO_{SA}}$ (Gy)) dose differences for D_{max} and D_{mean} . Negative values represent dose sparing for DAO_{MS} , while positive values represent dose sparing for DAO_{SA} . OARs that received 0 Gy in both cases are excluded in the evaluation.

| Dose difference ($D_{DAO_{MS}} - D_{DAO_{SA}}$ (Gy)) | D_{max} | | | D_{mean} | | |
|--|---------------------|------------------------|---------------|---------------------|----------------------|---------------|
| | Largest value | Smallest value | Average value | Largest value | Smallest value | Average value |
| GBM | +0.021 Chiasm | -0.031 L Eye | -0.004 | +0.003 Chiasm | -0.003 L Eye | -0.0004 |
| LNG | +0.036 Trachea | -3.000 ProxBronch | -1.835 | +0.004 Trachea | -1.593 ProxBronch | -0.541 |
| H&N | -0.162 L Opt Nrv | -13.067 Spinal Cord | -5.891 | -0.046 R Opt Nrv | -13.073 Esophagus | -5.714 |

TABLE IV. Comparison of the PTV homogeneity, D98, D99, and D_{max} , as well as R50 and VRCN.

| Patient Case | PTV Statistics | | | | | | | | | |
|--------------|----------------|------------|--------|--------|-----------|----------------|----------------|------------|------------|-------|
| | Homogeneity | | D98 | D99 | D_{max} | R50 | | VRCN | | |
| | DAO_{MS} | DAO_{SA} | | | | $D_{DAO_{MS}}$ | $D_{DAO_{SA}}$ | DAO_{MS} | DAO_{SA} | |
| GBM | 0.954 | 0.952 | +0.089 | +0.174 | -0.047 | 2.553 | 2.712 | 0.816 | 0.847 | |
| LNG | 0.938 | 0.951 | +0.294 | +0.494 | +0.935 | 2.840 | 2.680 | 0.831 | 0.892 | |
| H&N | 54 | 0.895 | 0.853 | +2.626 | +3.771 | -1.595 | 2.248 | 3.074 | 0.748 | 0.607 |
| | 59.4 | 0.850 | 0.833 | +1.237 | +1.448 | -0.483 | | | | |
| | 69.96 | 0.904 | 0.892 | +0.153 | -0.101 | -1.054 | | | | |
| Average | | 0.908 | 0.896 | +0.880 | +1.157 | -0.449 | 2.547 | 2.822 | 0.798 | 0.782 |

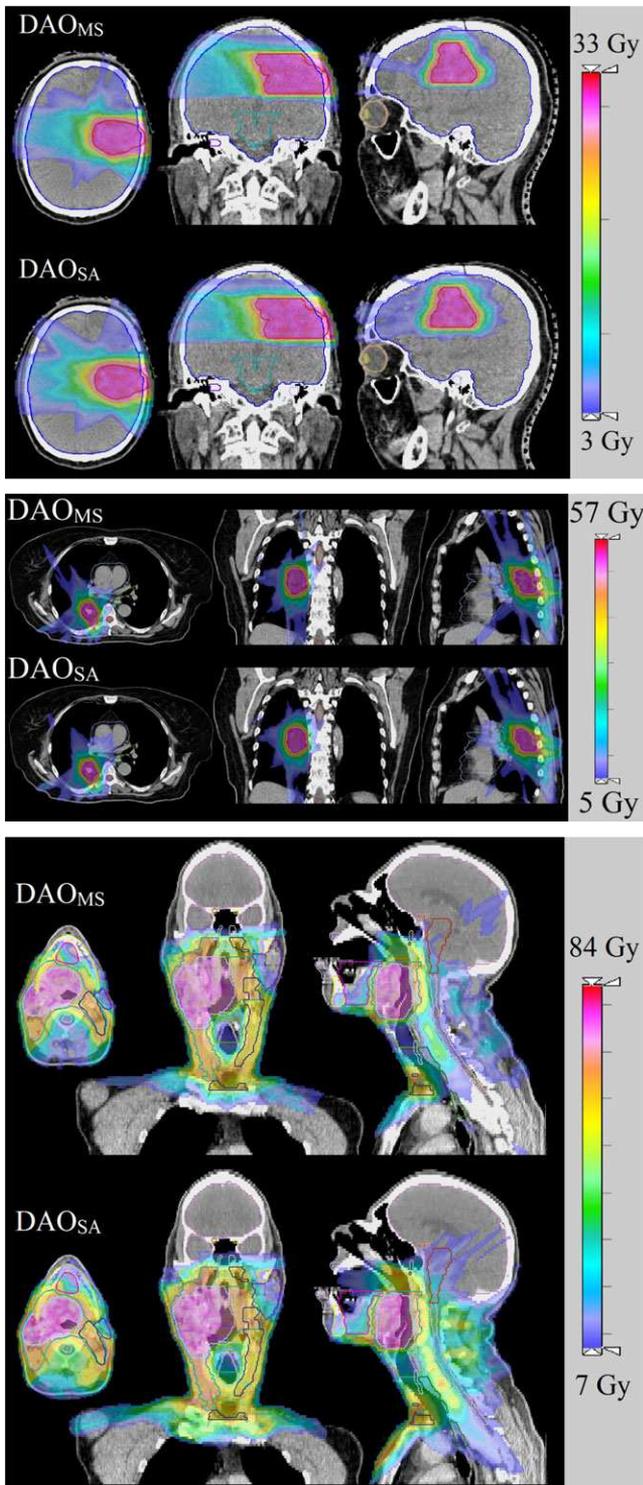


FIG. 3. Isodose of the GBM, LNG, and H&N patients. The low dose cutoff for viewing was set to be 10% of the prescription dose.

is similar with DAO_{MS} being slightly better by 0.012. The average R50 and VRCN was more clearly improved in the H&N case using DAO_{MS} . The R50 for the GBM and LNG cases were similar for the two plan types, differing by 0.16. For the H&N plan, R50 was lowered by 0.826 and VRCN

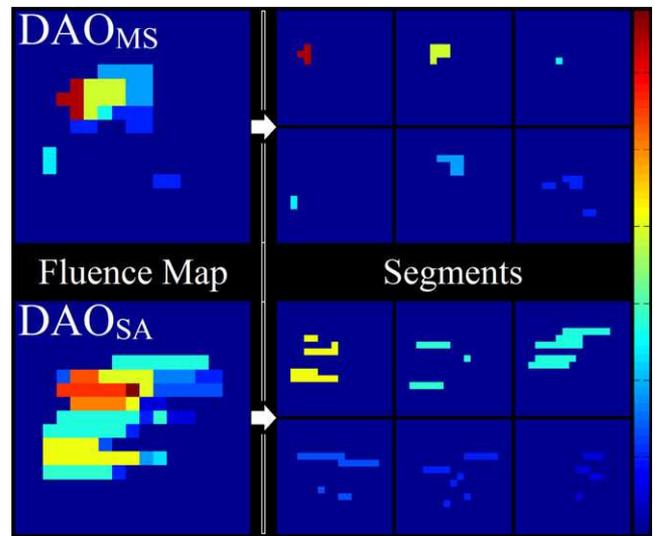


FIG. 4. Schematic of an optimized fluence from each plan type, and the breakdown of the fluences to their segments.

TABLE V. Aperture statistics for the DAO_{MS} and the DAO_{SA} methods for the patient cases.

| | Allowed number of segments per beam (n_s) | Average number of solved segments per beam | | Mean number of beamlets in a segment | |
|-----|---|--|------------|--------------------------------------|------------|
| | | DAO_{MS} | DAO_{SA} | DAO_{MS} | DAO_{SA} |
| GBM | 10 | 6.10 | 7.80 | 6.72 | 17.51 |
| LNG | 10 | 9.40 | 9.15 | 4.93 | 14.93 |
| H&N | 20 | 16.00 | 10.20 | 8.84 | 50.50 |

was increased by 1.355 using DAO_{MS} , indicating substantially improved dose compactness and conformality.

Figure 3 shows the isodose comparison. The low dose cutoff for viewing was set at 10% of the prescription dose. Qualitatively, the dose distributions produced by the DAO_{MS} and the DAO_{SA} methods for GBM and LNG plans are similar. On the other hand, the dose distribution differences of the H&N plans are more compelling. The dose is visually more compact in the DAO_{MS} plan, with lower dose to the brainstem, spinal cord, and posterior neck region in the sagittal slice, as well as the larynx and left parotid in the coronal slice.

Figure 4 shows an example of MLC segments and resultant fluence map for the GBM patient. The pairwise disjoint nature of the piecewise constant Mumford-Shah formulation is apparent in the segmentation results, as the DAO_{MS} segments do not overlap. The DAO_{SA} MLC segments, on the other hand, can overlap, resulting in more complex fluence. On average, the size of DAO_{MS} segments is smaller than that of DAO_{SA} . Interestingly, Table V shows that none of the planning methods utilized the maximal allowed segments for delivery.

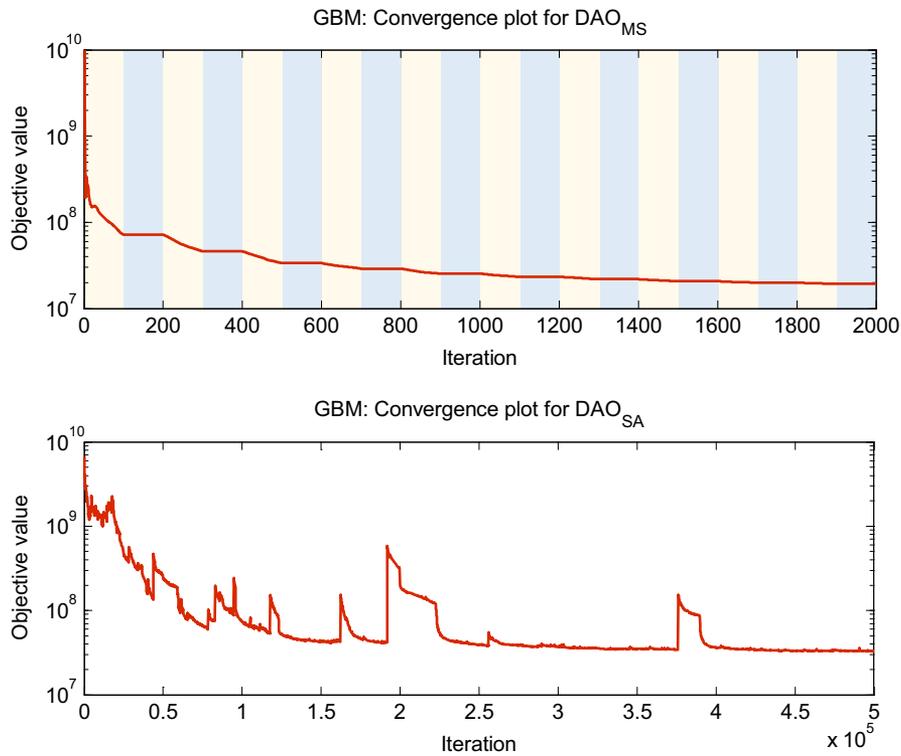


FIG. 5. Example convergence plots for the GBM case. For DAO_{MS} , the light-shaded regions represent where module 1 and module 2 are updating x and c at each subiteration. The dark-shaded region represents where module 3 is updating u , and each data point in the blue region represents 15 subiterations. Objective value calculation for DAO_{MS} uses the γ_{max} value for all iterations.

DAO_{MS} is substantially faster than DAO_{SA} , with a 9.5 to 40 fold decrease in total solve time. The number of times retuning the weights was greater for the DAO_{MS} approach because it was run 2–3 times in the beginning before the first iteration of the DAO_{SA} approach, which then adopted the DAO_{MS} weights as its starting point. Figure 5 shows example convergence plots for the GBM case. The DAO_{MS} approach is able to consistently lower the cost function, since each module is a convex subproblem. The stochastic property of simulated annealing causes the cost function to randomly increase at times, while attempting to escape from local minima. Figure 5 and Table VI highlight the strengths of the multiconvex approach to quickly and deterministically reach a minimum that contains a high quality dosimetric plan.

TABLE VI. Solve time for one optimization run for each case and plan type. DAO_{MS} solve times include the polishing step.

| | Solve time (s) | | Number of times retuning structure weights | |
|-----|----------------|------------|--|------------|
| | DAO_{MS} | DAO_{SA} | DAO_{MS} | DAO_{SA} |
| GBM | 599.8 | 5719.6 | 6 | 4 |
| LNG | 647.1 | 25747.7 | 7 | 4 |
| H&N | 2170.2 | 35682.8 | 11 | 10 |

4. DISCUSSION

In this study, we developed a novel direct aperture optimization technique that integrates fluence map optimization—employing L2-norm fidelity and L1-norm regularization terms—with the multiphase piecewise constant Mumford-Shah method. We utilized a first-order primal-dual proximal-class algorithm, known as the Chambolle-Pock algorithm, to solve the multiconvex direct aperture optimization problem. This novel method simultaneously solves for the fluence segments while minimizing the dose fidelity term offers a number of advantages.

Ultimately, DAO_{MS} result in perfectly piecewise constant fluence maps that are equivalent to apertures without additional stratification. The piecewise constancy cannot be achieved with total variation regularization alone^{6–9}, because total variation energies do not differentiate piecewise smooth and piecewise constant functions.

The optimization formulation offers an elegant theoretical description of the physical problem. The fidelity term penalizes deviations from the prescription dose, the total variation term promotes piecewise smoothness on the fluence map, and the piecewise constant Mumford-Shah formulation defines segments that constitute the fluence map. As a non-convex optimization, this method cannot guarantee to find the global minima of the problem. However, with the graduated weighting technique on the non-convex term, the optimization is able to find a local minimum that is biased to

satisfying to dose fidelity term. By slowly introducing the non-convexity, optimizer can begin to form segments for the fluence map that indirectly penalize deviations between the projected dose and the prescription dose.

In addition to its theoretical appeal, the deterministic multiconvex approach offers practical advantages in computational speed and dosimetric performance, compared to the existing stochastic simulated annealing method for DAO. The DAO_{SA} method struggled to include more apertures in the optimization but the explicit addition is heavily penalized by computational time. Given the vast number of possible apertures as a combination of beamlets, the previous methods are inherently inefficient. Our method, on the other hand, segments the MLC apertures on the fly, thus avoiding being limited to a small subset of possible apertures. As hypothesized, the advantage is shown more evident for the multilevel head and neck plan requiring complex modulation.

The Chambolle-Pock algorithm was selected to solve the optimization problem for several reasons. As a proximal-class algorithm, it can solve many types of non-differentiable optimization formulations exactly, such as formulations involving the L1-norm. The algorithm is highly efficient on memory usage and computation cost because it does not require a to solve system of linear equations involving the fluence to dose transformation matrix at every iteration, contrasting to other methods such as alternating direction method of multipliers (ADMM).⁴⁴ The computation cost of Chambolle-Pock relies on the simple multiplication of a large matrix and its transpose at each iteration.

While the piecewise constant Mumford-Shah term promotes apertures to be MLC deliverable, the DAO_{MS} method does not explicitly guarantee that the resulting aperture is MLC deliverable in 1 segment. For example, in Fig. 3, the last segment of the DAO_{MS} plan is not deliverable for horizontal MLCs. The simplest solution is to break down the segment into two delivery segments, which should minimally add to the treatment time since only a small fraction of the solved segments need to be further broken down.

A drawback of the current approach is that the Mumford-Shah formulation is designed to describe non-overlapping segments. The ability to have overlapping segments would allow for larger and fewer segments. Although the drawback did not prevent DAO_{MS} from outperforming the stochastic approach, particularly in complex cases, there is clearly space for improvement. As a topic of further investigation, the smaller segment size problem may be alleviated by devising an improved segmentation algorithm to divide the DAO_{MS} fluence into fewer and larger overlapping segments, without modifying the current fluence.

Another interesting observation one can make from the result is that both methods did not use the maximal allowed number of segments. For the DAO_{MS} plans, this is due to the total variation term that reduces the fluence map complexity and subsequently fewer segments. The DAO_{SA} method, on the other hand, is not regularized to penalize more segments. Given that DAO_{SA} plans are suboptimal compared to the DAO_{MS} plans, this reflects the weakness of

using simulated annealing method to search a larger aperture solution space.

The run time of the DAO_{SA} method used in this study is noticeably longer than the simulated annealing implemented commercial software. The longer computational time can be attributed to the following reasons. First, the plans used 20 beams, which increased the fluence-to-dose transformation matrix size. Second, the fluence-to-dose transformation matrices in this study were not downsampled. Last, the plans allowed many more segments than what would be allowed in a commercial planning software.

5. CONCLUSION

A novel deterministic direct aperture optimization formulation combining fluence map optimization and multiphase piecewise constant Mumford-Shah segmentation into a unified framework was proposed and evaluated. The new approach enables generating MLC segments on the fly without being limited to a small subset of possible apertures as previous methods did. The non-convex optimization formulation was split into multiple convex modules, and solved alternately using a first-order primal-dual proximal-class algorithm. The new deterministic method solved the DAO problem is considerably faster than the simulated annealing method and is dosimetrically superior, particularly for the complex head and neck case.

ACKNOWLEDGEMENTS

This research was supported by NIH grants R43CA183390, R44CA183390, R01CA188300 and DOE grant 0000227808.

CONFLICT OF INTEREST

There is no conflict of interest involved in the study.

APPENDIX

A1 The overrelaxed Chambolle-Pock algorithm

The Chambolle-Pock algorithm, a first-order primal dual algorithm,^{25,26} has previously been successfully implemented for fluence map optimization problems.^{9,33,36,45} The algorithm solves a primal problem of the canonical form

$$\text{minimize } F(Kx) + G(x), \quad (\text{A1})$$

where F and G are lower semicontinuous functions and K is a matrix. To solve an optimization problem with Chambolle-Pock, the optimization must be rewritten to fit this canonical form. It will be seen later that all optimizations in this study that use the Chambolle-Pock algorithm can easily be expressed in this canonical form. The overrelaxed version of the Chambolle-Pock algorithm³⁵ solves the primal problem via the iteration

$$\begin{aligned}
\bar{x}^{n+1} &= \text{prox}_{\tau G}(x^n - \tau K^T z^n) \\
\bar{z}^{n+1} &= \text{prox}_{\sigma F^*}(z^n + \sigma K(2\bar{x}^{n+1} - x^n)) \\
x^{n+1} &= \rho \bar{x}^{n+1} + (1 - \rho)x^n \\
z^{n+1} &= \rho \bar{z}^{n+1} + (1 - \rho)z^n
\end{aligned} \tag{A2}$$

where z variable of the dual problem: *minimize* $G^*(-K^T z) + F^*(z)$. The Chambolle-Pock algorithm is solving both the primal and dual problem simultaneously. The function, F^* , is the convex conjugate of F , defined as $F^*(z) = \sup_y (z^T y - F(y))$. The overrelaxation parameter, $\rho \in [0, 2]$, can be adjusted to control the convergence. The algorithm degenerates to the original algorithm when $\rho = 1$. As a proximal-class algorithm,⁴⁶ the Chambolle-Pock algorithm relies on the ‘‘proximal mapping’’, or ‘‘prox operator’’, defined as

$$\text{prox}_{th}(x) = \underset{v}{\text{argmin}} \left(h(v) + \frac{1}{2t} \|v - x\|_2^2 \right), \tag{A3}$$

where h is a lower semicontinuous function and t is a parameter that serves as a step size. Intuitively, $\text{prox}_{th}(x)$ will try to find an argument that minimizes h , without straying too far from x . The efficacy of proximal algorithms relies on that the function h has a simple evaluation with the prox operator. A useful property of the prox operator is the separable sum rule,

$$\begin{aligned}
&\underset{x}{\text{argmin}} \quad \frac{1}{2} W \left\| \left(\sum_{b=1}^{n_b} (A_b x_b) - d \right) \right\|_2^2 + \sum_{b=1}^{n_b} \left(\lambda_1 \|D_b^+ x_b\|_1 + \lambda_2 \|D_b^- x_b\|_1 \right) + \frac{\gamma}{2} \sum_{b=1}^{n_b} \sum_{s=0}^{n_s} \left\| \sqrt{\text{diag}(u_{bs})} (x_b - c_{bs} \mathbf{1}) \right\|_2^2 \\
&\text{subject to} \quad x \geq 0,
\end{aligned} \tag{A7}$$

$$\text{prox}_{\sigma F^*} \left(\begin{bmatrix} z_1 \\ \vdots \\ z_n \end{bmatrix} \right) = \begin{bmatrix} \text{prox}_{\sigma f_1^*}(z_1) \\ \vdots \\ \text{prox}_{\sigma f_n^*}(z_n) \end{bmatrix}, \tag{A4}$$

where $F \left(\begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix} \right) = f_1(y_1) + \dots + f_n(y_n)$, and this allows us to break down $F(Kx)$ into a sum of convex functions. Another valuable property is the Moreau decomposition,

$$\text{prox}_{\sigma F^*}(Z) = Z - \sigma \cdot \text{prox}_{\sigma^{-1} F}(\sigma^{-1} z), \tag{A5}$$

which allows us to calculate the prox operator of the convex conjugate of a function without having to evaluate the convex conjugate itself.

In the algorithm, we have the parameters τ and σ , which serve as our step sizes during the optimization. These can be selected based on a diagonal preconditioning approach³⁴:

$$\begin{aligned}
\tau_j &= \frac{1}{\sum_{i=1}^{i=n_{\text{rows}}} K(i,j)} & \sigma_i &= \frac{1}{\sum_{j=1}^{j=n_{\text{cols}}} K(i,j)} \\
\tau &= \begin{bmatrix} \tau_{j=1} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \tau_{j=n_{\text{cols}}} \end{bmatrix} & \sigma &= \begin{bmatrix} \sigma_{i=1} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \sigma_{i=n_{\text{rows}}} \end{bmatrix},
\end{aligned} \tag{A6}$$

where τ and σ are diagonal matrices with the same number of diagonal elements as the length of x and z , respectively. The diagonal preconditioning guarantees convergence under the assumption that $F(Kx)$ and $G(x)$ have sum separable operations. In the case where these functions are not sum separable, but are block separable, it is possible to assign the block of variables the smallest step size—found during the preconditioning process—within its block. Otherwise, τ and σ can be scalars and follow the relation, $\tau \sigma \|K\|^2 \leq 1$, to guarantee convergence. The operator norm, K , can be estimated via power iteration.⁴⁷

A2 Solving Module 1 with the Chambolle-Pock Algorithm

Module 1 updates x , while holding u and c constant. To solve module 1 with Chambolle-Pock, which solves the optimization of the form *minimize* $F(Kx) + G(x)$, we first rewrite problem 2 as:

where the third term is an equivalent expression to $\frac{\gamma}{2} \sum_{b=1}^{n_b} \sum_{p=1}^{n_p} \sum_{s=0}^{n_s} u_{bps} |x_{bp} - c_{bs}|^2$. It is written this way as a more convenient expression for updating x . We then define the matrix, K , as well as the functions, F and G , as

$$\begin{aligned}
K &= \begin{bmatrix} WA \\ D_1 \\ D_2 \\ \sqrt{\text{diag}(u_{b=1,s=0})} \\ \vdots \\ \sqrt{\text{diag}(u_{b=n_b,s=n_s})} \end{bmatrix} \\
F(y) &= \underbrace{\frac{1}{2} \|y_1 - Wd\|_2^2}_{f_1(y_1)} + \underbrace{\lambda_1 \|y_2\|_1}_{f_2(y_2)} + \underbrace{\lambda_2 \|y_3\|_1}_{f_3(y_3)} \\
&\quad + \underbrace{\sum_{b=1}^{n_b} \sum_{s=0}^{n_s} \frac{1}{2} \|y_{bs} - \sqrt{\text{diag}(u_{bs})} c_{bs} \mathbf{1}\|_2^2}_{\sum_b \sum_{s=0}^{n_s} f_{bs}(y_{bs})} \\
G(x) &= \begin{cases} 0 & \text{if } f \geq 0 \\ \infty & \text{otherwise} \end{cases}.
\end{aligned} \tag{A8}$$

Once K , F , and G are defined, the Chambolle-Pock algorithm can be applied to solve the problem. The prox operator evaluations of the functions, F^* and G , yield simple low cost expressions:

$$\begin{aligned} \text{prox}_{\sigma f_1^*}(\widehat{z}_1) &= \frac{\widehat{z}_1 - \sigma Wd}{1 + \sigma} \\ \text{prox}_{\sigma f_2^*}(\widehat{z}_2) &= P_{\lambda_1 B}(\widehat{z}_2) \\ \text{prox}_{\sigma f_3^*}(\widehat{z}_3) &= P_{\lambda_2 B}(\widehat{z}_3) \\ \text{prox}_{\sigma f_{bs}^*}(\widehat{z}_{bs}) &= \gamma \left(\frac{\widehat{z}_{bs} - \sigma \sqrt{\text{diag}(u_{bs})} c_{bs} \mathbf{1}}{\gamma + \sigma} \right)_+, \\ \text{prox}_{\tau G}(\widehat{x}) &= (\widehat{x})_+ \end{aligned} \quad (\text{A9})$$

where $P_{\lambda B}(\cdot)$ projects its argument onto the norm ball $\lambda B = \{\psi \mid \|\psi\|_\infty \leq \lambda\}$, and $(\cdot)_+$ projects its argument onto the non-negative orthant. These expressions are used at each iteration of the algorithm to update x .

A3 Solving Module 3 with the Chambolle-Pock Algorithm

Module 3 updates u , while holding x and c constant. Since x and c are constants, Eq. (5) can be rewritten as

$$\begin{aligned} \underset{u}{\text{argmin}} \quad & \langle q, u \rangle + \sum_{b=1}^{n_b} \pm \sum_{s=1}^{n_s} (\omega_1 \|D_b^\parallel u_{bs}\|_1 + \omega_2 \|D_b^\perp u_{bs}\|_1) \\ \text{subject to} \quad & u > 0 \\ & \mathbf{1}^T u_{bp} = 1, \end{aligned} \quad (\text{A10})$$

where $q_{bps} = \frac{1}{2} |f_{bp} - c_{bs}|^2$ for all b, p , and s . The intermediate variable, q , has the same dimensions and subscript structure as u . The inner product, $\langle \cdot, \cdot \rangle$, takes an element-wise product between its two arguments and sums the result. To solve problem 20 with Chambolle-Pock, which solves an optimization of the form $F(Ku) + G(u)$, we first define our matrix, K , as well as our functions F and G as

$$\begin{aligned} K &= \begin{bmatrix} D^\parallel \\ D^\perp \end{bmatrix} \\ F(y) &= \underbrace{\omega_1 \|y_1\|_1}_{f_1(y_1)} + \underbrace{\omega_2 \|y_2\|_1}_{f_2(y_2)} \\ G(u) &= \langle q, u \rangle + \sum_{b=1}^{n_b} \sum_{p=1}^{n_p} I_{sp}(u_{bp}), \end{aligned} \quad (\text{A11})$$

where I_{sp} is the indicator function for the probability simplex. In other words,

$$\begin{aligned} I_{sp}(u_{bp}) &= I_+(u_{bp}) + I_1(u_{bp}) \\ I_+(u_{bp}) &= \begin{cases} 0 & \text{if } u_{bp} \geq 0 \\ \infty & \text{otherwise} \end{cases} \\ I_1(u_{bp}) &= \begin{cases} 0 & \text{if } u_{bp} \geq 0 \\ \infty & \text{otherwise} \end{cases} \end{aligned} \quad (\text{A12})$$

Evaluation of the prox operators of F^* and G yield the following expressions.

$$\begin{aligned} \text{prox}_{\sigma f_1^*}(\widehat{z}_1) &= P_{\omega_1 B}(\widehat{z}_1) \\ \text{prox}_{\sigma f_2^*}(\widehat{z}_2) &= P_{\omega_2 B}(\widehat{z}_2) \\ \text{prox}_{\tau G}(\widehat{u}) &= P_{S_p \text{row}}(\widehat{u} - \tau q) \end{aligned} \quad (\text{A13})$$

where $P_{\omega B}(\cdot)$ projects its argument onto the norm ball $\omega B = \{\psi \mid \|\psi\|_\infty \leq \omega\}$, and $P_{S_p \text{row}}(\cdot)$ projects each row of its argument onto the probability simplex. Recall that each row in u is a vector u_{bp} . These simple expressions are used at each iteration of the Chambolle-Pock algorithm to update u .

A4 The probability simplex projection

Projection onto the probability simplex is defined as

$$P_{S_p}(v) = \left(v - \eta \mathbf{1} \right)_+, \quad (\text{A14})$$

where v is a vector, $(\cdot)_+$ projects its argument onto the non-negative orthant, and η is a scalar that is the solution to the equation

$$\mathbf{1}^T \left(v - \eta \mathbf{1} \right)_+ = 1. \quad (\text{A15})$$

Since this projection occurs at every iteration, it is imperative to use a fast, low-cost algorithm to project a vector onto the probability simplex. Wang et al.,⁴⁸ provided an efficient algorithm for the projection:

$$\begin{aligned} \text{Input : } & v_{in} \in R^D \\ \text{Sort } & u_{in} \text{ into } \tilde{v} : \tilde{v}_1 \geq \tilde{v}_2 \geq \dots \geq \tilde{v}_D \\ \text{Find } & j^* = \max \left\{ 1 \leq j \leq D \mid \tilde{v}_j + \frac{1}{j} \left(1 - \sum_{i=1}^j \tilde{v}_i \right) > 0 \right\} \\ \text{Define } & \eta = -\frac{1}{j^*} \left(1 - \sum_{i=1}^{j^*} \tilde{v}_i \right) \\ \text{Output : } & v_{out} = \left(v_{in} - \eta \mathbf{1} \right)_+ \end{aligned} \quad (\text{A16})$$

Essentially, this algorithm efficiently solves for η to properly project the argument onto the probability simplex.

^{a)} Author to whom correspondence should be addressed. Electronic mail: ksheng@mednet.ucla.edu.

REFERENCES

1. Brahme A. Optimization of stationary and moving beam radiation therapy techniques. *Radiother Oncol.* 1988;12:129–140.
2. Bortfeld T, Bürkelbach J, Boesecke R, Schlegel W. Methods of image reconstruction from projections applied to conformation radiotherapy. *Phys Med Biol.* 1990;35:1423.

3. Webb S. Optimisation of conformal radiotherapy dose distribution by simulated annealing. *Phys Med Biol.* 1989;34:1349.
4. Convery D, Rosenbloom M. The generation of intensity-modulated fields for conformal radiotherapy by dynamic collimation. *Phys Med Biol.* 1992;37:1359.
5. Xia P, Verhey LJ. Multileaf collimator leaf sequencing algorithm for intensity modulated beams with multiple static segments. *Med Phys.* 1998;25:1424–1434.
6. Zhu L, Lee L, Ma Y, Ye Y, Mazzeo R, Xing L. Using total-variation regularization for intensity modulated radiation therapy inverse planning with field-specific numbers of segments. *Phys Med Biol.* 2008;53:6653.
7. Zhu L, Niu T, Choi K, Xing L. Total-variation regularization based inverse planning for intensity modulated arc therapy. *Technol Cancer Res Treat.* 2012;11:149–162.
8. Zhu L, Xing L. Search for IMRT inverse plans with piecewise constant fluence maps using compressed sensing techniques. *Med Phys.* 2009;36:1895–1905.
9. Nguyen D, O'Connor D, Yu VY, et al. Dose domain regularization of MLC leaf patterns for highly complex IMRT plans. *Med Phys.* 2015;42:1858–1870.
10. Kirkpatrick S. Optimization by simulated annealing: quantitative studies. *J Stat Phys.* 1984;34:975–986.
11. Earl M, Shepard D, Naqvi S, Li X, Yu C. Inverse planning for intensity-modulated arc therapy using direct aperture optimization. *Phys Med Biol.* 2003;48:1075.
12. Shepard DM, Earl MA, Li XA, Naqvi S, Yu C. Direct aperture optimization: a turnkey solution for step-and-shoot IMRT. *Med Phys.* 2002;29:1007–1018.
13. Earl M, Afghan M, Yu C, Jiang Z, Shepard D. Jaws-only IMRT using direct aperture optimization. *Med Phys.* 2007;34:307–314.
14. Bergman AM, Bush K, Milete M-P, Popescu IA, Otto K, Duzenli C. Direct aperture optimization for IMRT using Monte Carlo generated beamlets. *Med Phys.* 2006;33:3666–3679.
15. Sui H, PIC CA. Direct aperture optimization.
16. Zhang G, Jiang Z, Shepard D, Zhang B, Yu C. Direct aperture optimization of breast IMRT and the dosimetric impact of respiration motion. *Phys Med Biol.* 2006;51:N357.
17. Jiang Z, Earl MA, Zhang GW, Yu CX, Shepard DM. An examination of the number of required apertures for step-and-shoot IMRT. *Phys Med Biol.* 2005;50:5653–5663.
18. Mestrovic A, Milete MP, Nichol A, Clark BG, Otto K. Direct aperture optimization for online adaptive radiation therapy. *Med Phys.* 2007;34:1631–1646.
19. Cotrutz C, Xing L. Segment-based dose optimization using a genetic algorithm. *Phys Med Biol.* 2003;48:2987.
20. Men C, Romeijn HE, Taşkın ZC, Dempsey JF. An exact approach to direct aperture optimization in IMRT treatment planning. *Phys Med Biol.* 2007;52:7333.
21. Salari E, Unkelbach J. A column-generation-based method for multi-criteria direct aperture optimization. *Phys Med Biol.* 2013;58:621.
22. Romeijn H, Ahuja R, Dempsey J, Kumar A. A column generation approach to radiation therapy treatment planning using aperture modulation. *SIAM J Optim.* 2005;15:838–862.
23. Mumford D, Shah J. Optimal approximations by piecewise smooth functions and associated variational problems. *Commun Pure Appl Math.* 1989;42:577–685.
24. Vese LA, Chan TF. A multiphase level set framework for image segmentation using the Mumford and Shah model. *Int J Comput Vision.* 2002;50:271–293.
25. Chambolle A, Pock T. A first-order primal-dual algorithm for convex problems with applications to imaging. *J Math Imaging Vis.* 2011;40:120–145.
26. Pock T, Cremers D, Bischof H, Chambolle A. An algorithm for minimizing the Mumford-Shah functional. *2009 IEEE 12th International Conference on Computer Vision, Kyoto, 2009*;1133–1140.
27. Chan TF, Vese LA. Active contours without edges. *IEEE Trans Image Process.* 2001;10:266–277.
28. Esedog S, Tsai Y-HR. Threshold dynamics for the piecewise constant Mumford-Shah functional. *J Comput Phys.* 2006;211:367–384.
29. Chan TF, Vese LA. A level set algorithm for minimizing the Mumford-Shah functional in image processing. In: *Variational and Level Set Methods in Computer Vision, 2001. Proceedings. IEEE Workshop on,* 161–168. IEEE, 2001.
30. Chan TF, Esedoglu S, Nikolova M. Algorithms for finding global minimizers of image segmentation and denoising models. *SIAM Appl Math.* 2006;66:1632–1648.
31. Shen X, Diamond S, Udell M, Gu Y, Boyd S. Disciplined multi-convex programming; 2016.
32. Xu Y, Yin W. A block coordinate descent method for regularized multi-convex optimization with applications to nonnegative tensor factorization and completion. *SIAM J Imaging Sci.* 2013;6:1758–1789.
33. Nguyen D, Lyu Q, Ruan D, O'Connor D, Low DA, Sheng K. A comprehensive formulation for volumetric modulated arc therapy planning. *Med Phys.* 2016;43:4263–4272.
34. Pock T, Chambolle A. Diagonal preconditioning for first order primal-dual algorithms in convex optimization. *2011 International Conference on Computer Vision, Barcelona, 2011*;1762–1769.
35. Condat L. A primal-dual splitting method for convex optimization involving lipschitzian, proximable and linear composite terms. *J Optim Theory Appl.* 2013;158:460–479.
36. Nguyen D, Thomas D, Cao M, O'Connor D, Lamb J, Sheng K. Computerized triplet beam orientation optimization for MRI-guided Co-60 radiotherapy. *Med Phys.* 2016;43:5667–5675.
37. Dong P, Lee P, Ruan D, et al. 4 π Noncoplanar stereotactic body radiation therapy for centrally located or larger lung tumors. *Int J Radiat Oncol Biol Phys.* 2013;86:407–413.
38. Dong P, Lee P, Ruan D, et al. 4 π Non-Coplanar Liver SBRT: a novel delivery technique. *Int J Radiat Oncol Biol Phys.* 2013;85:1360–1366.
39. Nguyen D, Rwigema J-CM, Yu VY, et al. Feasibility of extreme dose escalation for glioblastoma multiforme using 4 π radiotherapy. *Radiat Oncol.* 2014;9:1–9.
40. Dong P, Nguyen D, Ruan D, et al. Feasibility of prostate robotic radiation therapy on conventional C-arm linacs. *Pract Radiat Oncol.* 2014;4:254–260.
41. Dong P, Yu V, Nguyen D, et al. Feasibility of using intermediate x-ray energies for highly conformal extracranial radiotherapy. *Med Phys.* 2014;41:041709.
42. Grégoire V, Mackie TR. State of the art on dose prescription, reporting and recording in intensity-modulated radiation therapy (ICRU report No. 83). *Cancer/Radiothér.* 2011;15:555–559.
43. Feuvret L, Noël G, Mazeron J-J, Bey P. Conformity index: a review. *Int J Radiat Oncol Biol Phys.* 2006;64:333–342.
44. Boyd S, Parikh N, Chu E, Peleato B, Eckstein J. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations Trends® Mach Learn.* 2011;3:1–122.
45. Nguyen D, Ruan D, O'Connor D, et al. A novel software and conceptual design of the hardware platform for intensity modulated radiation therapy. *Med Phys.* 2016;43:917–929.
46. Parikh N, Boyd S. Proximal algorithms. *Foundations Trends Optim.* 2013;1:123–231.
47. Golub GH, Van Loan CF. *Matrix Computations.* Baltimore: Johns Hopkins University Press; 1989.
48. Wang W, Carreira-Perpinán MA. Projection onto the probability simplex: an efficient algorithm with a simple proof, and an application; 2013.