Multicriteria optimization for volumetric-modulated arc therapy by decomposition into a fluence-based relaxation and a segment weight-based restriction

Rasmus Bokrantz[a)]
Optimization and Systems Theory, Department of Mathematics, KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden and RaySearch Laboratories, Sveavägen 25, SE-111 34 Stockholm, Sweden

(Received 5 March 2012; revised 3 July 2012; accepted for publication 8 September 2012; published 16 October 2012)

Purpose: To develop a method for inverse volumetric-modulated arc therapy (VMAT) planning that combines multicriteria optimization (MCO) with direct machine parameter optimization. The ultimate goal is to provide an efficient and intuitive method for generating high quality VMAT plans.
Methods: Multicriteria radiation therapy treatment planning amounts to approximating the relevant treatment options by a discrete set of plans, and selecting the combination thereof that strikes the best possible balance between conflicting objectives. This approach is applied to two decompositions of the inverse VMAT planning problem: a fluence-based relaxation considered at a coarsened gantry angle spacing and under a regularizing penalty on fluence modulation, and a segment weight-based restriction in a neighborhood of the solution to the relaxed problem. The two considered variable domains are interconnected by direct machine parameter optimization toward reproducing the dose-volume histogram of the fluence-based solution.
Results: The dose distribution quality of plans generated by the proposed MCO method was assessed by direct comparison with benchmark plans generated by a conventional VMAT planning method. The results for four patient cases (prostate, pancreas, lung, and head and neck) are highly comparable between the MCO plans and the benchmark plans: Discrepancies between studied dose-volume statistics for organs at risk were—with the exception of the kidneys of the pancreas case—within 1 Gy or 1 percentage point. Target coverage of the MCO plans was comparable with that of the benchmark plans, but with a small tendency toward a shift from conformity to homogeneity.
Conclusions: MCO allows tradeoffs between conflicting objectives encountered in VMAT planning to be explored in an interactive manner through search over a continuous representation of the relevant treatment options. Treatment plans selected from such a representation are of comparable dose distribution quality to conventionally optimized VMAT plans. [http://dx.doi.org/10.1118/1.4754652]

Key words: direct aperture optimization, direct machine parameter optimization, multicriteria optimization, Pareto optimality, total variation regularization, treatment planning, volumetric-modulated arc therapy

I. INTRODUCTION

Recent years have seen a surge in interest for rotational radiation therapy delivered using a conventional multileaf collimator (MLC) equipped linear accelerator. This interest has been sparked by the advance of rotational delivery with varying angular monitor unit (MU) level, which allows for delivery in a single—or very few—rotation(s) of the treatment gantry, see, e.g., the review.[1] This capability, referred to as volumetric-modulated arc therapy (VMAT), was first demonstrated on Varian[2] and Elekta[3] accelerators using control systems that allow variable dose rate, gantry speed, and MLC leaf positions during irradiation. VMAT delivery has more recently also been demonstrated on Siemens accelerators using statically collimated bursts over short arc intervals.[4] Numerous comparative studies report that VMAT is largely equivalent to fixed field intensity-modulated radiation therapy (IMRT) in terms of dose distribution quality, while it provides a substantial reduction in the number of MUs and total delivery time, see, e.g., Yu and Tang[1] and Teoh *et al.*[5] and references therein. Dynamic delivery, however, imposes additional machine constraints that make inverse planning for VMAT a more challenging problem than inverse planning for fixed field IMRT. This property is reflected in a marked comparative increase in manual planning time.[6,7]

Time-efficiency and transparency of IMRT planning have been addressed by a number of research groups from the viewpoint of a multicriteria decision problem, see, e.g., Refs. 8–13. The unifying theme of the cited papers is to first precompute a collection of treatment plans that give different emphasis to the considered planning objectives, and subsequently select the most preferred plan from this representation. A specific selection technique is in this paper considered that utilizes real-time navigation between discrete solutions by continuous interpolation, see Monz *et al.*[14] Recent comparative studies between multicriteria optimization (MCO) and standard inverse planning have for fixed field IMRT shown that MCO significantly reduces average planning time while

simultaneously providing plans that are ranked as superior by physicians.[15,16] It is plausible to hypothesize that VMAT planning due to its involved process would equally—if not more so—benefit from MCO. This paper therefore extends and evaluates the use of multicriteria planning to the VMAT modality.

Multicriteria planning for rotational therapy was first studied by Pardo-Montero and Fenwick.[17,18] These authors propose using linear combinations of a basis of anatomy-based arcs with MLC shapes defined by Boolean operations on the beam's eye view contours of the delineated regions of interests (ROIs). Craft *et al.*[19] have more recently proposed a method where navigation over a basis of fluence-based arcs is followed by unidirectional leaf-sequencing into dynamic MLC trajectories, and the tradeoff between delivery time and treatment quality examined using a mechanism for downsampling fluence maps.

The method for multicriteria VMAT optimization developed in this paper makes sequential use of a basis of fluence-based arcs and a basis of segment weight-based arcs with shared MLC shapes. The rationale of this approach is to first perform a global tradeoff between objectives in the fluence domain, and second fine-tune this tradeoff with respect to deliverable machine settings and final dose computed by a high-accuracy algorithm. The presented method is an extension of a previous technique for multicriteria planning with respect to fixed field step-and-shoot IMRT delivery, see Craft *et al.*[20]

A notable difference to previous work on multicriteria planning for rotational therapy is that MLC shapes are generated using direct machine parameter optimization (DMPO) (or direct aperture optimization as it is often synonymously called), thereby eliminating the need for postprocessing steps that may degrade treatment quality. There is for fixed field IMRT extensive empirical evidence that DMPO leads to a reduction in plan complexity without sacrificing dose distribution quality, see, e.g., Broderick *et al.*[21] for a review and references to the original literature. Rao *et al.*[22] report a similar result for VMAT in a comparative study of an anatomy-based, a fluence-based, and a machine parameter-based approach to treatment planning. In this study, the machine parameter-based approach produced plans that were slightly better than those generated by the fluence-based approach in terms of dose uniformity and conformity, most notably so for complex target geometries. The anatomy-based approach led to plans that required the least number of MUs, but at a cost of the least conformal dose distributions in general, and clinically unacceptable dose distributions for challenging target geometries.

The dosimetric quality of plans generated by the suggested MCO method is evaluated by comparison with benchmark plans generated by the conventional method of minimizing a weighted sum of objective functions. The purpose of the evaluation is to examine if MCO plans are of comparative dose distribution quality with plans generated by a standard method, disregarding the difficulty of finding an appropriate tradeoff between conflicting planning objectives with the conventional method. The MCO plans were for the purpose of a fair comparison generated toward showing a similar tradeoff between objectives as the benchmark plans. The evaluation can thereby be thought of as simulating a treatment planner that strives toward finding the best possible tradeoff between objectives, with the assumption being that the tradeoff of each benchmark plan is the correct one. The MCO plans should due to the design of this evaluation strategy not be expected to improve on the benchmark plans. The reported results should also be put in perspective of the referenced studies that point to the cognitive and time-efficiency benefits of multicriteria treatment planning, i.e., Refs. 15 and 16.

II. MATERIALS AND METHODS

II.A. Notation

For a given ROI, the minimum dose such that the associated isodose volume contains $x\%$ of the volume is denoted by D_x, the volume contained by the x Gy isodose volume by V_x, and the mean dose level by $\bar{\mathrm{D}}$. Volumes are to be understood in relative sense unless otherwise indicated. Planning target volumes (PTVs) are designated by their prescription level in subscript. The notation $\mathrm{PTV}_{\mathrm{low}}^{-x\,\mathrm{mm}} = \mathrm{PTV}_{\mathrm{low}} - (\mathrm{PTV}_{\mathrm{high}} + x\,\mathrm{mm})$ is used to denote the setwise difference between a low dose PTV and a high dose PTV, with magnitude of the expansion of the high dose PTV suppressed when zero.

II.B. Problem formulation

The method developed in this paper takes as input specification of a set of rotational arcs to be optimized with respect to n objectives f_i, with $n \geq 2$, and m constraints c_j. All objectives and constraints are assumed to map the dose distribution vector d to a nonnegative scalar. The constraints are in addition assumed to be of a form such that the feasible region is nonempty. An arc is represented as a sequence of control points at uniform gantry spacing and fixed collimator and couch angles. We consider optimization with respect to MLC leaf positions x, control point MUs μ, and arc delivery times t of the form

$$\begin{aligned} \underset{x,\mu,t}{\text{minimize}} \quad & [f_1(d(\psi(x,\mu))) \cdots f_n(d(\psi(x,\mu)))]^T \\ \text{subject to} \quad & c_j(d(\psi(x,\mu))) \leq 0, \qquad j = 1,\ldots,m, \\ & (x,\mu,t) \in \mathcal{F}, \end{aligned} \tag{1}$$

where ψ is the energy fluence vector and $\mathcal{F}$ the set of ordered triplets (x, μ, t) that correspond to feasible machine settings. The exact definition of $\mathcal{F}$ is determined by the commissioned machine specifications, see Ref. 23 for an explicit formulation. The representation of arc delivery times as scalar variables, as opposed to vectors of control point delivery times, implies that the gantry speed is assumed to be constant. This formulation is based on the recommendation by Bzdusek *et al.*[23] who found that variable gantry speed during

optimization does not contribute to any substantial plan improvements and often leads to large gantry accelerations that may introduce delivery errors.

We briefly review some properties of the chain of transformations from optimization variables to dose. The energy fluence vector ψ is related to the vector of MLC leaf positions x by an integral over the fluence distributions of the bremsstrahlung target and the flattening filter. This integral is, if assuming Gaussian fluence distributions, described by a combination of error functions that are nonconvex due to their sigmoidal shape.[24] The components of the vector of control point MUs μ enter the fluence calculation as multiplicative factors. The function $\psi(x,\ \mu)$ is therefore nonlinear and nonconvex in x, and linear in μ. The dose distribution vector d is calculated from ψ by a superposition of weighted kernels, which is a linear operation. The function $d(\psi)$ is therefore linear.

We now turn to discussing the interpretation of a solution to problem (1). Solving a vector-valued optimization problem is most commonly understood in the sense of selecting the most preferred solution from the set of Pareto optimal solutions: feasible solutions such that none of the objectives can be improved without impairing at least one other objective, see, e.g., the monograph.[25] The current paradigm in multicriteria IMRT planning is to approximate the set of Pareto optimal solutions with a discrete set of points and their convex combinations. The theoretical justification of this method hinges on a convex problem formulation, i.e., that all objectives are convex functions of the optimization variables and the feasible region a convex set. If this criterion is satisfied, the set of Pareto optimal solutions forms a connected surface in the boundary of a convex set[26] (called the Pareto surface). Convex combinations of Pareto optimal points, moreover, remain feasible and with objective function values bounded by the corresponding convex combination of objective function value vectors.

II.C. Overview of the solution approach

The framework for convex MCO planning outlined in Sec. II.B cannot be directly applied to problem (1) due to its nonconvexity. This property is accounted for by first considering a convex relaxation of this problem, and then considering a convex restriction in a neighborhood of the solution to the relaxed problem. The cornerstones of this procedure is an algorithm for approximating convex Pareto surfaces, an algorithm for navigating a discrete Pareto surface representation, and a routine for single-objective VMAT optimization. These algorithms are applied to solve a sequence of optimization problems of the following form:

(a) *Fluence map optimization*
The relaxation of the initial problem given by considering energy fluence as a directly controllable variable.
(b) *Reference dose-volume histogram optimization*
The problem of reproducing the dose-volume histogram (DVH) of a fluence-based solution with respect to deliverable machine settings.
(c) *Segment weight optimization*
The restriction of the initial problem given by considering control point MUs and arc delivery times as variables while keeping the MLC leaf positions fixed.

The above procedure is summarized in Fig. 1. Because the optimization problems in (a) and (c) are vector-valued, solving these problems entails generating a discrete set of solutions using the Pareto surface approximation algorithm, and letting a decision maker select the most preferred convex combination from this representation using the navigation algorithm.

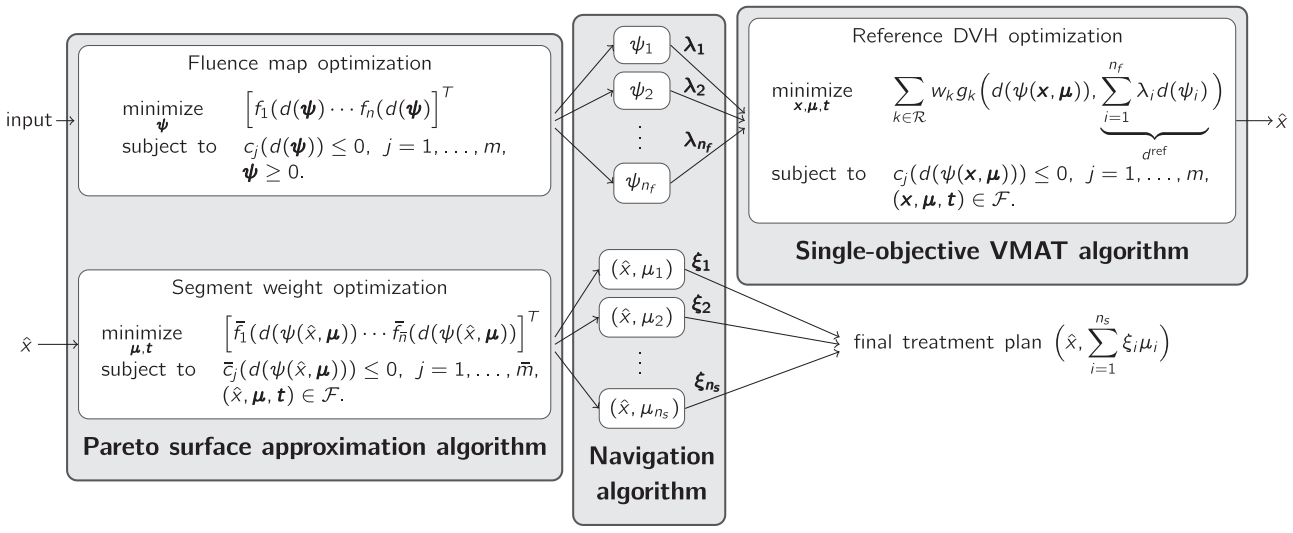

FIG. 1. Overview of the MCO method: A collection of n_f fluence-based plans are generated and combined into a single navigated plan by a convex combination with coefficient vector λ. The navigated plan is converted into a deliverable VMAT plan using reference DVH optimization, with the reference dose d^{ref} set equal to the dose distribution of the navigated solution. A collection of n_s segment weight-based plans in a neighborhood of the initial converted solution are generated and combined into a final treatment plan by a convex combination with coefficient vector ξ. Parameters that constitute variables at each considered step are indicated in bold. The mathematical notation used in the illustration is introduced in Secs. II.D–II.F.

The final solution to a vector-valued problem is thus always a single solution.

II.D. Fluence map optimization

The fluence map relaxation of (1) takes the form

$$\begin{aligned} \underset{\psi}{\text{minimize}} \quad & [f_1(d(\psi)) \cdots f_n(d(\psi))]^T \\ \text{subject to} \quad & c_j(d(\psi)) \leq 0, \qquad j = 1, \ldots, m, \\ & \psi \geq 0. \end{aligned} \tag{2}$$

This problem is a convex optimization problem whenever all objectives f_i and constraints c_j are convex in fluence since sublevel sets of convex functions are convex. A sufficient condition for a dose-based function to be convex in fluence is that the function is convex in dose. This fact follows from linearity of $d(\psi)$ and that composition with an affine function preserves convexity.

A large number of degrees of freedom makes the fluence map optimization problem a both computationally expensive problem and an inaccurate model of VMAT delivery if considered at the final gantry spacing and without restriction on fluence modulation. Problem (2) is therefore considered at a coarsened gantry spacing and under a regularizing penalty on a total variation measure of the fluence maps. This approach is supported by a number of observations by previous authors: Yu[27] has conjectured that plan quality is essentially a function of the number of strata, a quantity defined as the product between the number of beam angles and the number of intensity levels per beam. The number of intensity levels within a fluence modulated beam has been demonstrated by Zhu *et al.*[28] to be highly correlated with the total variation of the fluence vector. Finally, Tang *et al.*[29] report that dosimetric quality is insensitive to small angular displacements of the incident fluence distribution.

The total variation of the fluence vector is in this study defined as the L1-norm of its gradient with respect to the coordinate axes of the fluence planes. To aid the definition of the total variation functional $\mathrm{TV}(\cdot)$, denote by $\mathcal{B}$ an index set over the fluence planes, by $\Psi^{(b)}$ the transmission matrix associated with the bth fluence plane, and by $\mathcal{IJ}_b$ the set of bixel index pairs (i,j) associated with elements of $\Psi^{(b)}$ that are considered as optimization variables in Eq. (2). The total variation of the fluence vector is with this notation given by a sum of forward differences of the form

$$\mathrm{TV}(\psi) = \sum_{b \in \mathcal{B}} \left(\sum_{\substack{(i,j) \in \mathcal{IJ}_b: \\ (i+1,j) \in \mathcal{IJ}_b}} \left| \Psi^{(b)}_{i,j} - \Psi^{(b)}_{i+1,j} \right| + \sum_{\substack{(i,j) \in \mathcal{IJ}_b: \\ (i,j+1) \in \mathcal{IJ}_b}} \left| \Psi^{(b)}_{i,j} - \Psi^{(b)}_{i,j+1} \right| \right). \tag{3}$$

Total variation minimization is in addition to its use for smoothing of fluence maps in IMRT planning[28,30,31] widely used for image denoising where preserving edge information is important. It is for this application well-known to filter out high-frequency oscillations while preserving sharp discontinuities.[32] The scale of the image features that are smoothed out are inversely proportional to the penalty weight of the stabilizing functional.[33] The penalty weight of the total variation term was in this study determined by the empirical L-curve method, as described in Appendix A. The absolute values in Eq. (3) were made continuously differentiable using a conservative approximation of the form

$$|x| = \max\{x, -x\} \approx \frac{1}{\alpha} \ln(e^{\alpha x} + e^{-\alpha x}),$$

for some positive scalar α. This approximation is computationally inexpensive and can be made arbitrary close to the exact function by letting the parameter α tend to infinity. A more rigorous approach for handling the nondifferentiability of Eq. (3) would be to substitute linearly constrained auxiliary variables for the absolute values in this expression.

II.E. Reference dose-volume histogram optimization

The purpose of reference DVH optimization is to convert the navigated fluence-based solution to Eq. (2) into a deliverable VMAT plan. This conversion is carried out by DMPO toward minimizing the error in DVH resulting from the conversion. Discrepancies in the DVH domain are quantified using reference DVH functions g_k that map the current dose distribution d and a reference dose distribution d^{ref} into a one-sided penalty on the integral between the associated pair of DVH curves, see Fig. 2 for an illustration and Appendix C for a precise definition.

A natural formulation of the reference DVH optimization problem would be to minimize a vector of DVH penalties with respect to the navigated solution to Eq. (2), subject to the constraints of problem (1). The resulting problem would, however, be a multiobjective program and therefore as difficult as problem (1) in general. Instead, positive weights w_k defined over an index set $\mathcal{R}$ over the considered ROIs are

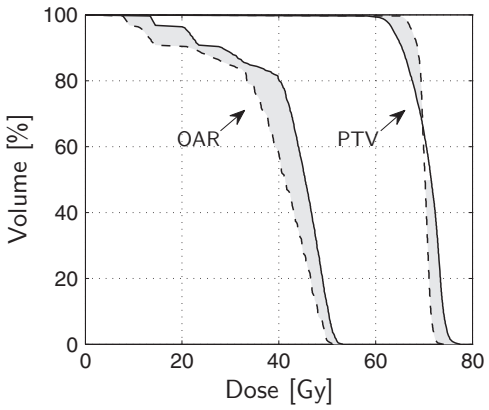

FIG. 2. Penalty imposed by a max reference DVH function assigned to an OAR and a one-sided uniform reference DVH function assigned to a PTV. The current DVH is indicated by solid lines and the reference DVH by dashed lines. The shaded regions indicate cumulative volumes where a quadratic penalty is imposed on differences along the dose axis.

introduced to yield a scalar-valued problem of the form

$$\begin{aligned} \underset{x,\mu,t}{\text{minimize}} \quad & \sum_{k\in\mathcal{R}} w_k g_k(d(\psi(x,\mu)), d^{\text{ref}}) \\ \text{subject to} \quad & c_j(d(\psi(x,\mu))) \leq 0, \qquad j = 1,\ldots,m, \\ & (x,\mu,t) \in \mathcal{F}, \end{aligned} \tag{4}$$

where d^{ref} is the dose distribution vector associated with the navigated solution to Eq. (2). Scalarization with *a priori* weights in Eq. (4) can be motivated by that this problem is expected to have feasible solutions where the functions g_k evaluate to values close their lower bound of zero. This property makes Eq. (4) less sensitive to the choice of weighting coefficients than the problem given by substituting a non-negatively weighted sum of the objectives for the vector-valued objective function in Eq. (1). Validity of this statement relies on that the DVH associated with the navigated fluence-based solution is close to a DVH that can be realized by a feasible VMAT plan.

II.F. Segment weight optimization

The segment weight optimization problem can be derived from Eq. (1) by fixating the MLC leaf positions at the output values $\hat{x}$ from the reference DVH optimization step. It is typically not necessary to consider all initial objectives during segment weight optimization, but rather only those where an improvement is prioritized. A set of $\bar{n}$ objectives $\bar{f}_i$, where $\bar{n} \geq 2$, is therefore substituted for the objectives in Eq. (1). A set of $\bar{m}$ constraints $\bar{c}_j$ chosen such that the solution to the reference DVH optimization problem is feasible is similarly substituted for the constraints in Eq. (1). The resulting optimization problem is of the form

$$\begin{aligned} \underset{\mu,t}{\text{minimize}} \quad & [\bar{f}_1(d(\psi(\hat{x},\mu))) \cdots \bar{f}_{\bar{n}}(d(\psi(\hat{x},\mu)))]^T \\ \text{subject to} \quad & \bar{c}_j(d(\psi(\hat{x},\mu))) \leq 0, \qquad j = 1,\ldots,\bar{m}, \\ & (\hat{x},\mu,t) \in \mathcal{F}. \end{aligned} \tag{5}$$

The set $\mathcal{F}$ is defined by a set of linear inequalities and therefore convex. This property together with the linear relation between MU and dose implies that Eq. (5) is a convex problem whenever all functions $\bar{f}_i$ and all $\bar{c}_j$ are convex in dose. The linearity between MU and dose also implies that dose computed by a high-accuracy algorithm prior to solving Eq. (5) remains valid during segment weight optimization.

II.G. Treatment planning system

The suggested MCO method was implemented in a research version of the RayStation treatment planning system version 2.6 (RaySearch Laboratories, Stockholm, Sweden). The single-objective VMAT algorithm used in this system is described in Bzdusek *et al.*[23] Briefly, fluence map optimization is first performed at a gantry spacing of 24°. MLC shapes are subsequently generated to approximate the optimized fluence profiles, distributed over adjacent gantry angles, and further optimized using DMPO. Segment weight optimization can optionally be performed as a final step. Dose calculations during fluence map optimization and DMPO is performed using a pencil beam convolution technique based on singular value decomposition, similar to Bortfeld *et al.*[34] Accurate dose is computed using a collapsed cone (CC) convolution algorithm, see, e.g., Ahnesjö.[35] If an intermediate CC dose calculation is performed, it is used as a background dose for subsequent dose calculations. All nonlinear programming tasks are performed using a quasi-Newton sequential quadratic programming (SQP) method, see, e.g., Gill *et al.*[36] for a review of a similar algorithm. The considered version of RayStation does not support nonlinear constraints for rotational arcs. All dose-based constraints in Eqs. (4) and (5) were therefore relaxed into penalty terms in the objective function. Multiobjective programs are solved by approximating the Pareto surface with a predetermined number of plans. Each plan is generated by minimizing a non-negatively weighted sum of objectives subject to the initial constraints. The weighting coefficients are chosen toward optimally limiting the approximation error of the current representation of the Pareto surface, as detailed in Bokrantz and Forsgren.[37] A graphical user interface allows for selection of the most preferred combination of basis plans, see Craft *et al.*[13] for a description of a prototype version.

II.H. Patient cases

The developed method is evaluated by retrospective planning for four patient cases:

- A prostate case with a prescribed dose of 72.2 Gy over 39 fractions to a PTV that encompasses the prostate and seminal vesicles, with a simultaneous boost of 78 Gy in total to the prostate. Considered critical structures are the bladder, femurs, and rectum.
- A pancreas case with prescribed dose to the PTV of 50.4 Gy over 28 fractions. Considered critical structures are the kidneys, liver, and stomach.
- A lung case with a PTV that encompasses a primary tumor situated in middle lobe of the right lung, and surrounding lymph nodes. The PTV is composed of two disjoint subvolumes that are about 5 mm apart. Its prescribed dose was 66 Gy over 24 fractions. The right lung of the patient was partially collapsed and filled with fluid. Considered critical structures are the esophagus, heart, lungs, and spinal cord.
- A head and neck case with a prescribed dose of 72 Gy over 36 fractions to the PTV of the primary tumor and a prescribed dose of 59.4 Gy over 36 fractions to the elective nodal PTV. Considered critical structures are the brainstem, esophagus constrictors, larynx, mandible, oral cavity, parotids, and spinal cord.

The objectives and constraints used during treatment plan optimization were of the form of least-squares penalties on deviation in voxel dose or equivalent uniform dose (EUD) from a reference level. A complete summary of objectives and constraints per patient case is provided in Appendix B. Mathematical definition of the used optimization functions is given in Appendix C.

The DVH-based functions that were used for a subset of the target constraints are nonconvex in dose[38] and therefore violate the convexity requirement stated in Sec. II.B. Empirical evidence, however, indicates that this class of functions do not impose any severe local minima effects, and that they therefore are compatible with gradient-based optimization algorithms.[39,40] DVH-based functions were included in the present study since functions of this form are in standard use in clinical practice.

II.I. Machine model and algorithm settings

All patient cases were planned for delivery with a Varian 2100 linear accelerator (Varian Medical Systems, Palo Alto, California) equipped with a 120-leaf interdigitating MLC system and at a nominal energy of 6 MV. The treatment machine was modeled with continuously variable dose rate in the range 50–600 MU/min, a maximum gantry speed of 5.54 °/s, and a maximum leaf speed of 22.5 mm/s. An upper bound on the maximum delivery time was introduced at 120 s for each arc, the control point spacing was set to 3°, the collimator angle set to 45°, and the couch angle set to 0°, with the latter two being defined according to IEC convention. Cases with close to convex target geometries (prostate, pancreas) were planned for delivery with a 359° single arc while cases with highly nonconvex target geometries (lung, head and neck) were planned for delivery with a 359° dual arc. Dual arcs were preferred over single arc delivery at a refined control point spacing and with an increased bound on the arc delivery time because the considered implementation of DMPO requires a dose calculation at every considered control point angle. Dose calculations for a single arc plan with halved control point spacing are therefore approximately twice as costly as those for a dual arc plan because the fluence distributions of each arc in the dual arc plan can be superimposed prior to making the transformation from fluence to dose. A discretization of the fluence planes into $5 \times 5\,\mathrm{mm}^2$ bixels and a discretization of the patient volume into $3 \times 3 \times 3\,\mathrm{mm}^3$ voxels were used throughout during optimization. The maximum number of SQP iterations during fluence map optimization, DMPO, and segment weight optimization was kept at 50, 45, and 30, respectively. An intermediate CC dose calculation at iteration 30 and a final CC dose calculation at completion was throughout performed during DMPO.

III. RESULTS

III.A. Treatment plan generation

A single benchmark plan per patient case was generated by minimizing a positively weighted sum of the objectives in Eq. (1) subject to the initial constraints of this problem. Objective weights were manually adjusted and the plan reoptimized until a treatment plan that was deemed satisfactory had been obtained. MCO plans were subsequently generated with navigation manually performed toward finding a plan with a similar tradeoff between objectives as that of the corresponding benchmark plan, this in order to facilitate a fair sideways comparison. The number of treatment plans in the Pareto surface representations was 50 during fluence map optimization and 25 during segment weight optimization. Reference DVH optimization was, unless otherwise stated, performed with respect to all ROIs assigned with an objective and a 10 mm shell isotropically expanded from the union of all targets. Target structures were assigned with a one-sided uniform reference DVH function, critical structures with a max reference DVH function, and unclassified tissue with a reference dose fall-off function. The weights w_k in problem (4) were set to 30 for functions associated with targets or shell structures and set to unity otherwise. The objective function of all segment weight optimization problems was augmented with a term identical to the composite objective function of the preceding reference DVH optimization problem at a relative weight of 10^{-4} in order to avoid radically altering the dose distribution.

III.B. Evaluation of plan quality

Plan quality was assessed with respect to target coverage, dose conformity, and sparing of organs at risk (OARs). Target coverage was quantified by a homogeneity index (HI) (Ref. 41) according to

$$\mathrm{HI} = (\mathrm{D}_2 - \mathrm{D}_{98}) / \mathrm{D}_{50}.$$

Dose conformity was quantified by a conformity index (CI) (Ref. 42) being the ratio between treated volume at 95% of the prescription level and target volume according to

$$\mathrm{CI} = \mathrm{V}_{95\%}^{\mathrm{External}} / \mathrm{V}^{\mathrm{PTV}},$$

with volume interpreted in absolute sense. Note that the ideal values for the HI and CI metrics are 0% and 100%, respectively. Sparing of OARs was quantified in terms of dose-volume statistics. Plan quality was additionally assessed with respect to DVH and planned dose distribution. The isodose lines of all depicted dose distributions follow a color table that is defined relative to the minimal prescription level taken over all target structures of the considered patient case.

III.B.1. Prostate

The results for the prostate case are summarized in Fig. 3 and Table I. The main dosimetric challenge was to obtain a uniform dose to both target volumes at their respective prescription level. This goal was in sharp conflict with sparing of the rectum. The treatment plan selected during fluence-based navigation is slightly better than the benchmark plan with respect to all objectives, most notably so by an improved average rectal dose of about 3 Gy. The plan obtained after reference DVH optimization shows mean dose levels for the bladder and rectum that are within 1 Gy of those of the benchmark plan. A comparison of the DVH curves of the target structures reveals a shift toward the low dose region of about 0.8 Gy for the MCO plan. Segment weight optimization was performed with respect to the subset of initial objectives that were associated with the PTVs and rectum. The treatment plan selected during segment weight-based navigation is essentially equivalent to the benchmark plan in terms

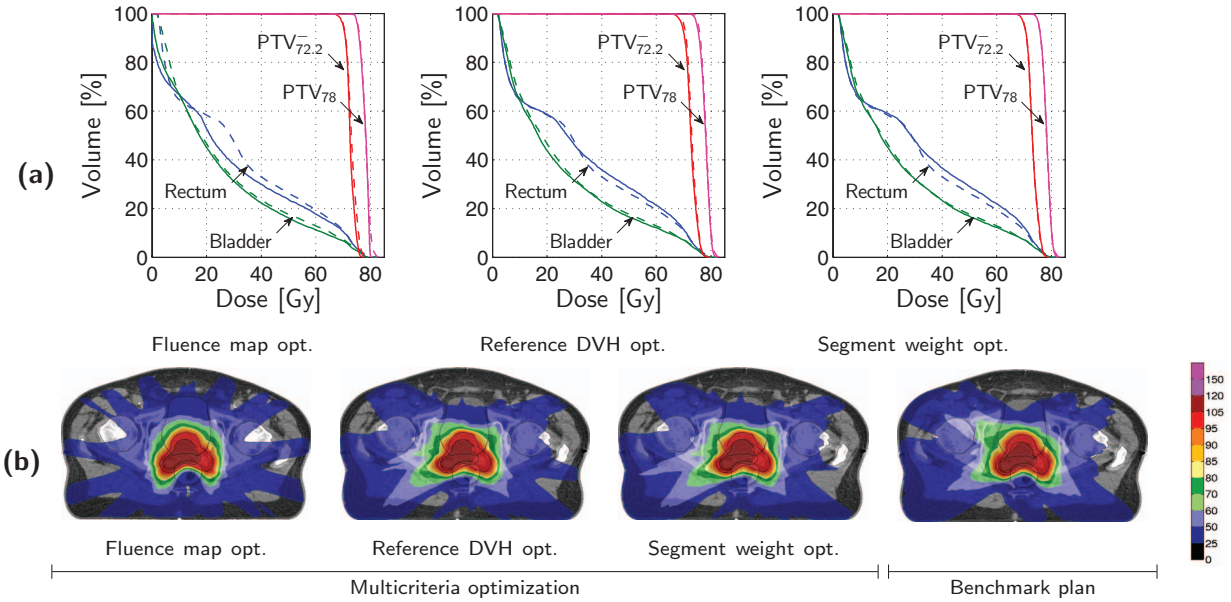

FIG. 3. (a) DVH results for the prostate case. Incremental stages of the MCO method are indicated by solid lines while the benchmark plan is indicated by dashed lines. (b) Transversal slices of dose distributions for the prostate case at incremental stages of the MCO method and for the benchmark method. ROI contours are indicated in black.

of target homogeneity and dose conformity, while showing a lower bladder mean dose of about 0.5 Gy, and a higher rectum mean dose of about 1 Gy.

III.B.2. Pancreas

The results for the pancreas case are summarized in Fig. 4 and Table II. Prioritized goals during fluence-based navigation was to limit kidney dose to $V_{10} \leq 60\%$ and $V_{20} \leq 30\%$. The discrepancy to the benchmark plan of 3–4 percentage points (pp) in the right kidney DVH for cumulative volumes below 15% could not be improved without compromising either sparing of the left kidney or target coverage. This discrepancy can however, at least in part, be attributed to that the min dose constraint at 45.4 Gy for the PTV was handled rigorously for the MCO plan, while relaxed into a penalty term for the benchmark plan. As a consequence of this constraint handling, 0.37% of the PTV volume violated this constraint, as compared to 0.01% for the fluence-based MCO plan. The treatment plan obtained after reference DVH optimization is qualitatively similar the benchmark plan. The most notable differences is a shift of the target DVH of about 0.5 Gy and the discrepancy for the right kidney that was already present in the fluence domain. Segment weight optimization was performed toward the subset of the initial objectives that were associated with the PTV and right kidney. The treatment plan selected during segment weight-based navigation shows comparable target coverage with the benchmark plan whereas sparing of the right kidney could not be simultaneously improved. The dose statistics in Table II indicates that the MCO plan is essentially equivalent to the benchmark plan in terms of target homogeneity and sparing of the liver and stomach, while slightly less conformal, and has higher V_{20} levels for both kidneys of about 3 pp–4 pp.

III.B.3. Lung

The results for the lung case are summarized in Fig. 5 and Table III. The primary consideration during treatment plan optimization was to obtain a uniform dose to the PTV at the prescription level while maintaining a sharp dose fall-off outside this structure. Secondary goals were to reduce lung dose and to avoid dose in the PTV overlap volume of the esophagus and heart beyond the prescription level. Sparing of the

TABLE I. Dose statistics for the prostate case.

	PTV_{78}		$PTV^{-}_{72.2}$	$PTV_{72.2}$	Bladder		Rectum	
Plan	HI (%)	CI (%)	HI (%)	CI (%)	D_{10} (Gy)	$\bar{D}$ (Gy)	D_{10} (Gy)	$\bar{D}$ (Gy)
Fl. map	6.6	113.9	9.4	122.1	62.8	24.2	70.4	28.0
Ref. DVH	8.6	119.6	11.7	121.3	64.0	24.8	71.0	31.2
Seg. wt.	8.0	125.1	11.0	124.1	64.3	25.2	71.5	31.8
Benchmark	8.5	124.8	10.7	125.8	65.5	25.7	70.9	30.7

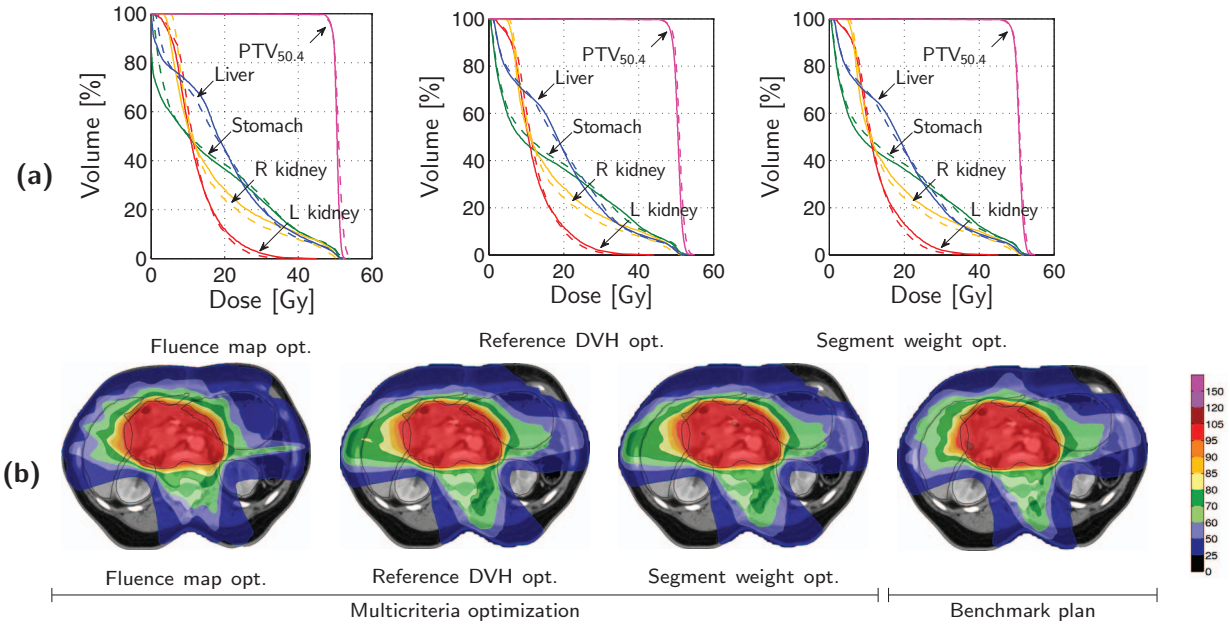

FIG. 4. (a) DVH results for the pancreas case. Incremental stages of the MCO method are indicated by solid lines while the benchmark plan is indicated by dashed lines. (b) Transversal slices of dose distributions for the pancreas case at incremental stages of the MCO method and for the benchmark method. ROI contours are indicated in black.

left lung was prioritized over sparing of the right lung due to the location of the primary tumor. The treatment plan selected during fluence-based navigation is better than the benchmark plan with respect to all objectives, in particular with respect to those associated with OAR sparing. Reference DVH optimization was performed with respect to all ROIs assigned with an objective, except for the ROI derived from the union of the lungs. A 70 mm shell expanded from the PTV was also included in the optimization. Expansion beyond 10 mm was for this structure exclusively performed in the gantry rotation plane. The treatment plan obtained after reference DVH optimization remains better than the benchmark plan in terms of all objectives associated with OARs, whereas it is slightly worse than the benchmark plan in terms of conformity and homogeneity. The improvement in mean dose level is less than 1 Gy for all OARs. Segment weight optimization was performed with respect to the subset of initial objectives that were associated with the PTV and lungs. The treatment plan selected during segment weight-based navigation shows target uniformity that is comparable with the benchmark plan and OAR sparing that is within 1 Gy in terms of mean dose. The benchmark plan is marginally more conformal than the MCO plan.

III.B.4. Head and neck

The results for the head and neck case are summarized in Fig. 6 and Table IV. Treatment plan optimization was performed toward obtaining a uniform dose to both target volumes at their respective prescription level while sparing the parotid glands and avoiding hot spots. Sparing of the left parotid was prioritized over sparing of the right parotid as the left gland was more remotely located from the primary tumor volume. In comparison to the benchmark plan, the treatment plan selected during fluence-based navigation shows a higher degree of homogeneity for both targets, and improved parotid sparing in the low dose region (<20 Gy). Reference DVH optimization was performed with respect to all ROIs assigned with an objective and two shells expanded 50 and 70 mm, respectively, from the union of the PTVs. Expansion beyond 10 mm was for these structures exclusively performed in the gantry rotation plane. The resulting treatment plan shows target coverage for PTV_{72} that is comparable with the benchmark plan, slightly worse target coverage for PTV_{72}, and improved parotid sparing in the low dose region. Segment weight optimization was performed with respect to the subset of initial objectives that were associated

TABLE II. Dose statistics for the pancreas case.

	$PTV_{50.4}$		L kidney		R kidney		Liver	Stomach
Plan	HI (%)	CI (%)	V_{20} (%)	$\bar{D}$ (Gy)	V_{20} (%)	$\bar{D}$ (Gy)	$\bar{D}$ (Gy)	$\bar{D}$ (Gy)
Fl. map	7.5	103.8	12.1	12.3	28.3	17.0	15.7	19.1
Ref. DVH	10.0	105.8	11.7	12.3	28.5	17.3	16.2	18.9
Seg. wt.	9.7	108.5	13.2	12.8	27.8	17.1	16.2	19.1
Benchmark	10.9	105.5	10.4	12.5	24.1	16.6	16.7	19.1

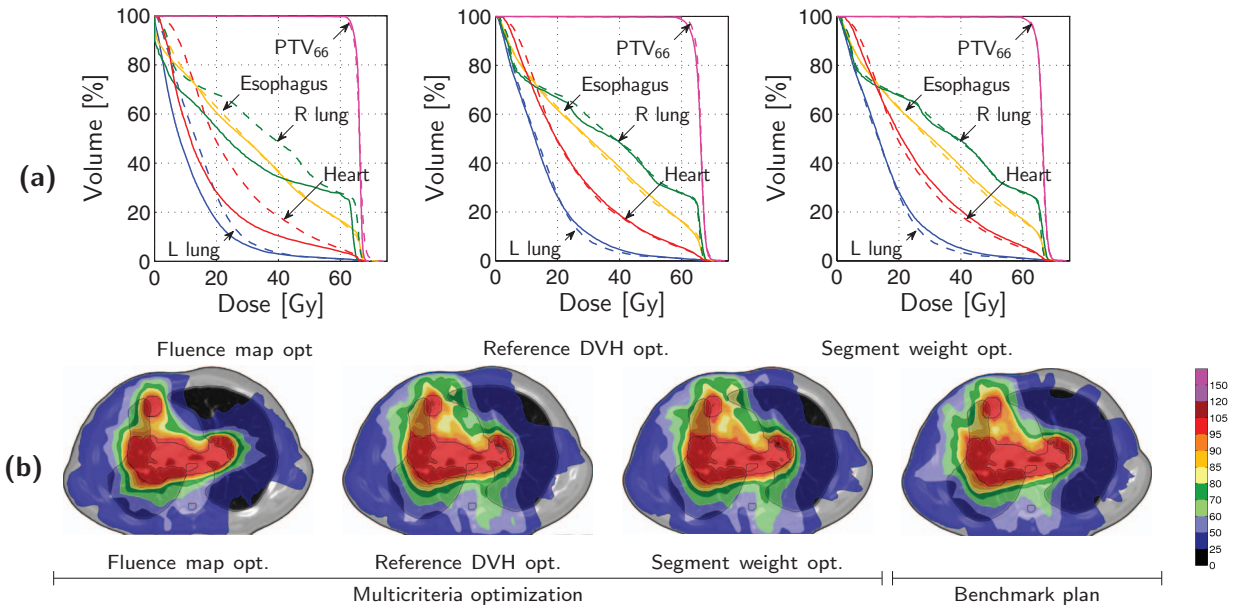

FIG. 5. (a) DVH results for the lung case. Incremental stages of the MCO method are indicated by solid lines while the benchmark plan is indicated by dashed lines. (b) Transversal slices of dose distributions for the lung case at incremental stages of the MCO method and for the benchmark method. ROI contours are indicated in black.

with the PTVs, a max dose objective for a 10 mm PTV shell with reference dose level at 56.5 Gy, and a max dose objective for the structure given by subtracting $PTV_{59.4}$ with a margin of 10 mm from the external ROI, with reference dose level at 35 Gy. The treatment plan selected during segment weight-based navigation has better homogeneity for both PTVs than the benchmark plan, but is slightly less conformal. The MCO plan shows a lower parotid mean dose of about 0.7 Gy for both glands.

III.C. Computational cost

Computational times for generation of the conventionally optimized benchmark plans were 15–20 min for the single arc plans and 40–50 min for the dual arc plans. These values are the total time required for a single optimization that includes two CC dose calculations and do not take into account the time required to manually identify suitable objective weights. Generating a single fluence-based MCO plan took about 2 min for the single arc plans and about 5 min for the dual arc plans, totaling to a time of about 1.5 h and 4 h, respectively, for complete generation of the fluence-based Pareto plan databases. Computational times for reference DVH optimization were comparable with those reported for the benchmark plans. Generating a single segment weight-based plan took about 1.2 min for the single arc plans and about twice that time for the dual arc plans, totaling to a time of about 30 min and 1 h, respectively, for complete generation of the segment weight-based Pareto plan databases.

IV. DISCUSSION

The main difficulty in applying MCO to VMAT planning is modeling the initially nonconvex VMAT optimization problem in a convex setting. We have studied two convex decompositions: a fluence-based relaxation and a segment weight-based restriction. The fluence-based formulation was defined at a coarsened gantry spacing and under a regularizing penalty on total fluence modulation, with the penalty weight of the total variation functional determined by empirical L-curve analysis. A requirement on identical MLC shapes was imposed on the segment weight-based formulation to allow for continuous interpolation between treatment plans. This formulation has the additional advantage that dose computed by a high-accuracy algorithm remains valid during navigation. The two

TABLE III. Dose statistics for the lung case.

	PTV_{66}		Esophagus	Heart	L lung		R lung	
Plan	HI (%)	CI (%)	$\bar{D}$ (Gy)	$\bar{D}$ (Gy)	V_{20} (%)	$\bar{D}$ (Gy)	V_{20} (%)	$\bar{D}$ (Gy)
Fl. map	7.0	107.2	30.1	17.5	16.6	11.3	60.5	30.4
Ref. DVH	11.9	102.8	35.7	24.6	26.2	15.0	61.8	31.5
Seg. wt.	10.9	108.9	36.4	25.6	28.3	15.7	62.3	31.8
Benchmark	10.4	105.9	36.6	24.9	27.0	15.0	62.8	31.3

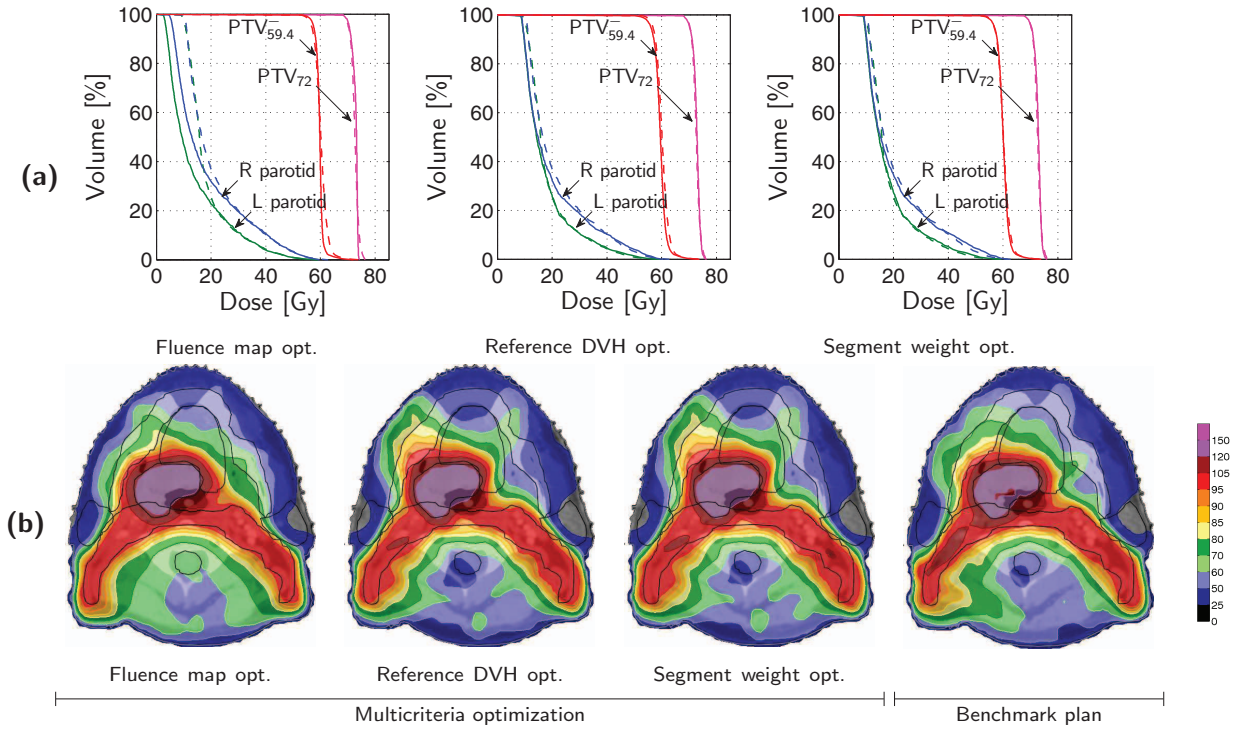

FIG. 6. (a) DVH results for the head and neck case. Incremental stages of the MCO method are indicated by solid lines while the benchmark plan is indicated by dashed lines. (b) Transversal slices of dose distributions for the head and neck case at incremental stages of the MCO method and for the benchmark method. ROI contours are indicated in black.

convex domains were interconnected by DMPO toward reproducing the DVH distribution of the fluence-based solution.

The MCO method was evaluated with respect to treatment planning for prostate, pancreas, lung, and head and neck by direct comparison with benchmark plans generated using a standard inverse planning method. The final MCO plans were consistently qualitatively very similar to their corresponding benchmark plan. Dose-volume statistics associated with OAR sparing was within 1 Gy or 1pp of those of the benchmark plan. The single exception to this observation was a discrepancy of 3 pp–4 pp for the V_{10} level of the kidneys for the pancreas case that was in favor of the benchmark plan. It was argued that this is an effect of approximate constraint handling during optimization of the benchmark plan. Dose uniformity was quantitatively comparable for all target structures across all patient cases. Dose conformity was comparable for the pancreas and prostate case, whereas the plans for the lung and head and neck case were slightly less conformal than their corresponding benchmark plan.

The goal of the evaluation was to assess the dose distribution quality of the MCO plans. The tradeoff between objectives of the benchmark plan was for the purpose of the evaluation assumed to be the best possible from a clinical perspective. This is a strong assumption that may not be fulfilled in practice. An inherent difficulty in conventional inverse planning is to assess whether a given plan is the best possible for the considered patient case. It is in general also not clear how to adjust the optimization problem formulation in order to introduce a desired modification to the optimized dose distribution. Accordingly, multiple studies show that the treatment quality of plans generated by conventional methods varies with the level of experience of the treatment planner.[43,44] Physicians have furthermore in blinded tests been shown to judge plans generated by MCO as superior to conventionally optimized plans,[15,16] a plausible explanation being that MCO enables treatment planners to find a more appropriate balance between conflicting treatment goals.

TABLE IV. Dose statistics for the head and neck case.

	PTV_{72}		$PTV_{59.4}^{-}$	$PTV_{59.4}$	L parotid		R parotid	
Plan	HI (%)	CI (%)	HI (%)	CI (%)	V_{30} (%)	$\bar{D}$ (Gy)	V_{30} (%)	$\bar{D}$ (Gy)
Fl. map	5.8	134.2	11.1	153.1	11.1	14.6	19.6	19.6
Ref. DVH	8.0	139.5	16.7	152.2	10.6	18.1	18.6	21.3
Seg. wt.	7.8	139.3	15.7	160.2	11.2	18.5	19.1	21.9
Benchmark	8.4	137.2	16.1	145.7	11.1	19.2	20.0	22.6

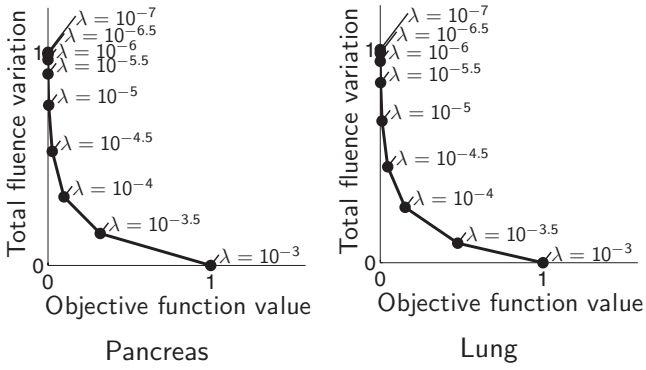

FIG. 7. L-curves of total variation of the fluence vector versus composite function value of a weighted sum of the objectives, with both quantities normalized to [0, 1]. The L-curves are parameterized by the penalty weight λ of the total variation functional.

Likewise, the reported computational times should be viewed in context of that conventional inverse planning methods typically require multiple reoptimizations to yield acceptable treatment plans. A recent example is an average number of 27.6 ± 10.4 reoptimizations reported in a study on head and neck planning.[45] Generation of plans that form a discrete Pareto surface representation—the major computational expense in MCO procedure—can also be performed in fully automated fashion and in parallel over multiple workstations.

V. CONCLUSIONS

The developed method allows the dosimetric tradeoffs between conflicting objectives encountered in VMAT planning to be explored in an intuitive manner without sacrificing treatment quality. The fluence-based relaxation is the primary domain of interest for trading sparing of sensitive structures against target coverage. The segment weight-based restriction is useful in retrieving loss of plan quality due to conversion into a deliverable VMAT plan, and for compensating for inaccuracies in an approximate optimization dose algorithm. The potential of the developed method is best realized with a clinical implementation that exploits that generation of Pareto surface representations can be performed without manual interaction and is well suited for distributed computing.

ACKNOWLEDGMENTS

The author is grateful to Anders Forsgren, Björn Hårdemark, and Erik Traneus for constructive comments on earlier drafts, and to Kjell Eriksson for discussions on VMAT optimization and help with the patient cases.

APPENDIX A: PRESTUDY ON CHOICE OF REGULARIZATION PARAMETER

A suitable penalty weight for the total variation term (3) was determined empirically by L-curve analysis, see, e.g. Hansen,[46] for a review. This technique selects the regularization parameter from the high-curvature region of a plot of smoothing norm (total variation) as a function of residual norm (objective function value). For the sake of brevity, results are only presented for the pancreas and lung case.

L-curves over the tradeoff between dosimetric quality and fluence modulation are shown in Fig. 7. The curves were plotted on a linear scale, to the contrary of on a logarithmic scale as is the convention for L-curves, in line with the recommendation by Chvetsov[47] regarding L-curve analysis of IMRT plans. The depicted results were generated by fluence map optimization for a single-objective problem constructed by replacing the vector-valued objective function in Eq. (2) by a positively weighted sum of its components. Fifty SQP iterations per treatment plan were used throughout. The results indicate that the relevant range of regularization parameter values is $\lambda \in [10^{-5}, 10^{-3}]$. The relationship between regularization parameter and fluence modulation for values in this range is illustrated in Figs. 8 and 9. A regularization parameter of $\lambda = 10^{-4}$ was used in all subsequent numerical experiments based on the depicted results.

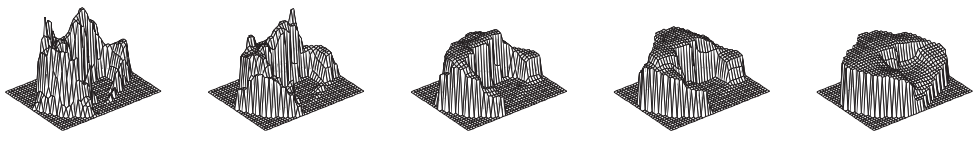

FIG. 8. Optimized fluence profile at a fixed gantry angle as function of penalty weight λ of the total variation functional for the pancreas case.

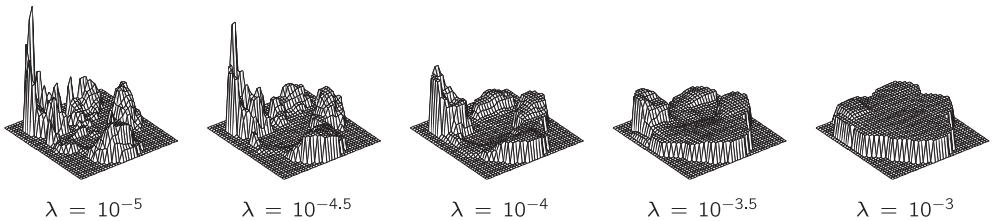

FIG. 9. Optimized fluence profile at a fixed gantry angle as function of penalty weight λ of the total variation functional for the lung case.

TABLE V. Problem formulation for the prostate case.

Objectives			Constraints		
ROI	Function	$\hat{d}$ (Gy)	ROI	Function	$\hat{d}$ (Gy)
PTV_{78}	Min dose	78.0	PTV_{78}	Min dose	70.2
	Uniform dose	78.0		Min 98% DVH	74.1
$PTV^{-}_{72.2}$	Min dose	72.2		Max dose	81.9
	Uniform dose	72.2	$PTV^{-}_{72.2}$	Min dose	65.0
Bladder	Max EUD $a = 2$	0.0		Min 98% DVH	68.6
Rectum	Max EUD $a = 2$	0.0	$PTV^{-4\,\mathrm{mm}}_{72.2}$	Max dose	75.8
External	Dose fall-off 10 mm	30.0–78.0	R femoral head	Max dose	38.0
			L femoral head	Max dose	38.0
			External	Max dose	81.9

APPENDIX B: PROBLEM FORMULATIONS

The objectives and constraints using during treatment plan optimization of the patient cases described in Sec. II.H are summarized in Tables V–VIII. Reference dose and reference EUD levels are both denoted by $\hat{d}$.

APPENDIX C: OPTIMIZATION FUNCTIONS

The optimization functions used in this study can be broadly categorized into dose-volume functions, EUD functions, and reference dose-based functions. All functions act on the restriction of the dose distribution to the voxels of a specified ROI.

1. Dose-volume functions

The dose-volume functions considered in this study are min dose, max dose, uniform dose, min DVH, max DVH, and dose fall-off functions. Min DVH and max DVH functions are specified by a reference dose level $\hat{d}$ and a reference volume level $\hat{v}$. A max DVH function imposes a penalty according to the cumulative volume integral

$$f(d) = \int_{\hat{v}}^{1} \max\{D(v, d) - \hat{d}, 0\}^2 \, dv, \tag{C1}$$

where $D(\cdot, d)$ is the function that takes the dose distribution vector d to parameterizations of the DVH curve for the considered ROI along the cumulative volume axis. A min DVH functions is obtained if reversing the sign of the two terms

TABLE VI. Problem formulation for the pancreas case.

Objectives			Constraints		
ROI	Function	$\hat{d}$ (Gy)	ROI	Function	$\hat{d}$ (Gy)
$PTV_{50.4}$	Min dose	50.4	$PTV_{50.4}$	Min dose	45.4
	Uniform dose	50.4		Min 95% DVH	47.9
L kidney	Max EUD $a = 2$	0.0		Max dose	53.9
R kidney	Max EUD $a = 2$	0.0	Spinal cord	Max dose	35.0
Liver	Max EUD $a = 2$	0.0	External	Max dose	53.9
Stomach	Max EUD $a = 2$	0.0			
External	Dose fall-off 10 mm	25.2–50.4			

TABLE VII. Problem formulation for the lung case.

Objectives			Constraints		
ROI	Function	$\hat{d}$ (Gy)	ROI	Function	$\hat{d}$ (Gy)
PTV_{66}	Min dose	66.0	PTV_{66}	Min dose	59.4
	Uniform dose	66.0		Min 95% DVH	62.7
Esophagus	Max EUD $a = 2$	0.0		Max dose	70.6
Heart	Max EUD $a = 2$	0.0	Esophagus	Max dose	66.5
L & R lungs	Max EUD $a = 2$	0.0	Heart	Max dose	66.5
L lung	Max EUD $a = 2$	0.0	Spinal cord	Max dose	38.0
PTV_{66} shell 10 mm	Max dose	56.0	External	Max dose	70.6
External	Dose fall-off 10 mm	30.0–66.0			

TABLE VIII. Problem formulation for the head and neck case.

Objectives			Constraints		
ROI	Function	$\hat{d}$ (Gy)	ROI	Function	$\hat{d}$ (Gy)
PTV_{72}	Min dose	72.0	PTV_{72}	Min dose	64.8
	Uniform dose	72.0		Min 98% DVH	68.4
$PTV^{-}_{59.4}$	Min dose	59.4		Min 90% DVH	71.3
$PTV_{59.4}$	Uniform dose	59.4		Max dose	77.0
L parotid	Max EUD $a = 2$	0.0	$PTV^{-}_{59.4}$	Min dose	53.5
R parotid	Max EUD $a = 2$	0.0		Min 90% DVH	58.1
$PTV_{59.4}$ shell 10 mm	Max dose	49.4	$PTV^{-6\,\mathrm{mm}}_{59.4}$	Max 10% DVH	61.8
External	Dose fall-off 10 mm	35.0–59.4		Max 1% DVH	65.3
			Brainstem	Max dose	45.0
			Esophagus constr.	Max EUD $a = 2$	45.0
			Larynx	Max EUD $a = 2$	45.0
			Mandible	Max EUD $a = 2$	45.0
				Max dose	65.0
			Oral cavity	Max EUD $a = 2$	45.0
			Spinal cord	Max dose	40.0
			External	Max dose	77.0

in the first argument of the max function in Eq. (C1) and integrating over $(0, \hat{v}]$. Min dose and max dose functions can be derived from their DVH counterparts by specializing $\hat{v}$ to unity and zero, respectively. A uniform dose function is given by the direct sum of a min dose function and a max dose function. A dose fall-off function is of the form of a max dose function with spatially variable reference dose level according to

$$\hat{d}(r) = \begin{cases} \infty & \text{if } r = 0, \\ d^{\text{low}} + (d^{\text{high}} - d^{\text{low}})(1 - \frac{r}{\delta}) & \text{if } 0 < r < \delta, \\ d^{\text{low}} & \text{otherwise,} \end{cases}$$

where r is the distance from the closest target structure, d^{low} a low dose level, d^{high} a high dose level, and δ a dose fall-off distance. Min DVH and max DVH functions are in general nonconvex functions of the dose distribution, whereas min dose, max dose, uniform dose, and dose fall-off functions are convex in dose.

2. Equivalent uniform dose functions

The only EUD function that is considered is max EUD. This function is specified by a reference EUD level $\hat{d}$ and an exponent a, where $a \geq 1$. The penalty imposed by a max EUD function is given by

$$f(d) = \max\left\{\left(\int_0^1 d^a(v)\,dv\right)^{1/a} - \hat{d}, 0\right\}^2,$$

where the integral is taken over relative volume. The exponent a is a tissue-specific parameter that allows continuous scaling between a penalty on mean dose ($a = 1$) and a penalty on maximum dose ($a \to \infty$). Max EUD functions are convex in dose.

3. Reference dose-based functions

The reference dose-based functions considered in this study are min reference DVH, max reference DVH, one-sided uniform reference DVH, and reference dose fall-off. All reference dose-based functions are specified by a reference dose distribution d^{ref}. Min reference DVH and max reference DVH functions can be derived from their DVH function counterparts by substituting $D(v, d^{\text{ref}})$ for $\hat{d}$ in Eq. (C1) and integrating over (0, 1]. A one-sided uniform reference DVH function is of the form of a max reference DVH function for cumulative volumes in the range (0, 1/2], and of the form of a min reference DVH function for cumulative volumes in the range (1/2, 1]. A reference dose fall-off function is defined analogously with a dose fall-off function, but with $\hat{d}(r)$ determined from the reference dose distribution. The components of d^{ref} are partitioned into subsets according to shortest distance for the associated voxel to a target structure. The partitioning is performed so that each subset corresponds to a distance interval with length equal to the shortest side length of a voxel. The function $\hat{d}(r)$ is made continuous by linear interpolation and extrapolation over the 95th percentile levels of the subsets. Min reference DVH, max reference DVH, and one-sided uniform reference DVH functions are nonconvex in dose, while reference dose fall-off functions are convex in dose.

[a]Electronic addresses: bokrantz@kth.se and rasmus.bokrantz@raysearchlabs.com

[1]C. Yu and G. Tang, "Intensity-modulated arc therapy: Principles, technologies and clinical implementation," Phys. Med. Biol. **56**(5), R31–R54 (2011).
[2]C. Ling, P. Zhang, Y. Archambault, J. Bocanek, G. Tang, and T. Losasso, "Commissioning and quality assurance of RapidArc radiotherapy delivery system," Int. J. Radiat. Oncol., Biol., Phys. **72**(2), 575–581 (2008).
[3]J. Bedford and A. Warrington, "Commissioning of volumetric modulated arc therapy (VMAT)," Int. J. Radiat. Oncol., Biol., Phys. **73**(2), 537–545 (2009).
[4]B. Salter, V. Sarkar, B. Wang, H. Shukla, M. Szegedi, and P. Rassiah-Szegedi, "Rotational IMRT delivery using a digital linear accelerator in

very high dose rate 'burst mode'," Phys. Med. Biol. **56**(7), 1931–1946 (2011).
[5]M. Teoh, C. Clark, K. Wood, S. Whitaker, and A. Nisbet, "Volumetric modulated arc therapy: A review of current literature and clinical use in practice," Br. J. Radiol. **84**(1007), 967–996 (2011).
[6]M. Oliver, W. Ansbacher, and W. Beckham, "Comparing planning time, delivery time and plan quality for IMRT, RapidArc and Tomotherapy," J. Appl. Clin. Med. Phys. **10**(4), 117–131 (2009).
[7]M. Rao, W. Yang, F. Chen, K. Sheng, J. Ye, V. Mehta, D. Shepard, and D. Cao, "Comparison of Elekta VMAT with helical tomotherapy and fixed field IMRT: Plan quality, delivery efficiency and accuracy," Med. Phys. **37**(3), 1350–1359 (2010).
[8]C. Cotrutz, M. Lahanas, C. Kappas, and D. Baltas, "A multiobjective gradient-based dose optimization algorithm for external beam conformal radiotherapy," Phys. Med. Biol. **46**(8), 2161–2175 (2001).
[9]M. Lahanas, E. Schreibmann, and D. Baltas, "Multiobjective inverse planning for intensity modulated radiotherapy with constraint-free gradient-based optimization algorithms," Phys. Med. Biol. **48**(17), 2843–2871 (2003).
[10]K.-H. Küfer, A. Scherrer, M. Monz, F. Alonso, H. Trinkaus, T. Bortfeld, and C. Thieke, "Intensity-modulated radiotherapy—a large scale multi-criteria programming problem," OR Spectrum **25**, 223–249 (2003).
[11]D. Craft, T. Halabi, and T. Bortfeld, "Exploration of tradeoffs in intensity-modulated radiotherapy," Phys. Med. Biol. **50**(24), 5857–5868 (2005).
[12]A. Hoffmann, A. Siem, D. den Hertog, J. Kaanders, and H. Huizenga, "Derivative-free generation and interpolation of convex Pareto optimal IMRT plans," Phys. Med. Biol. **51**(24), 6349–6369 (2006).
[13]D. Craft, T. Halabi, H. Shih, and T. Bortfeld, "An approach for practical multiobjective IMRT treatment planning," Int. J. Radiat. Oncol., Biol., Phys. **69**(5), 1600–1607 (2007).
[14]M. Monz, K.-H. Kufer, T. R. Bortfeld, and C. Thieke, "Pareto navigation—algorithmic foundation of interactive multi-criteria IMRT planning," Phys. Med. Biol. **53**(4), 985–998 (2008).
[15]D. Craft, T. Hong, H. Shih, and T. Bortfeld, "Improved planning time and plan quality through multi-criteria optimization for intensity modulated radiation therapy," Int. J. Radiat. Oncol., Biol., Phys. **82**(1), e83–e90 (2012).
[16]J. Wala, D. Craft, J. Paly, and J. Efstathiou, "Multi-criteria optimization of IMRT plans in prostate cancer," *Poster at the American Society for Radiation Oncology 53rd Annual Meeting, Miami Beach, Florida, October 2011*.
[17]J. Pardo-Montero and J. Fenwick, "An approach to multiobjective optimization of rotational therapy," Med. Phys. **36**(7), 3292–3303 (2009).
[18]J. Pardo-Montero and J. Fenwick, "An approach to multiobjective optimization of rotational therapy. II. Pareto optimal surfaces and linear combinations of modulated blocked arcs for a prostate geometry," Med. Phys. **37**(6), 2606–2616 (2010).
[19]D. Craft, D. McQuaid, J. Wala, W. Chen, E. Salari, and T. Bortfeld, "Multicriteria VMAT optimization," Med. Phys. **39**(2), 686–696 (2012).
[20]D. Craft, F. Carlsson, T. Bortfeld, and H. Rehbinder, "Multi-objective IMRT planning which produces deliverable plans," Med. Phys. **35**(6), 2846–2846 (2008).
[21]M. Broderick, M. Leech, and M. Coffey, "Direct aperture optimization as a means of reducing the complexity of intensity modulated radiation therapy plans," Radiat. Oncol. (2009) **4**:8.
[22]M. Rao, D. Cao, F. Chen, J. Ye, V. Mehta, T. Wong, and D. Shepard, "Comparison of anatomy-based, fluence-based and aperture-based treatment planning approaches for VMAT," Phys. Med. Biol. **55**(21), 6475–6490 (2010).
[23]K. Bzdusek, H. Friberger, K. Eriksson, B. Hårdemark, D. Robinson, and M. Kaus, "Development and evaluation of an efficient approach to volumetric arc therapy planning," Med. Phys. **36**(6), 2328–2339 (2009).
[24]M. Fippel, F. Haryanto, O. Dohm, F. Nüsslin, and S. Kriesen, "A virtual photon energy fluence model for Monte Carlo dose calculation," Med. Phys. **30**(3), 301–311 (2003).
[25]K. Miettinen, *Nonlinear Multiobjective Optimization* (Kluwer Academic, Boston, MA, 1999).
[26]E. Romeijn, J. Dempsey, and J. Li, "A unifying framework for multi-criteria fluence map optimization models," Phys. Med. Biol. **49**(10), 1991–2013 (2004).
[27]C. Yu, "Intensity-modulated arc therapy with dynamic multileaf collimation: An alternative to tomotherapy," Phys. Med. Biol. **40**(9), 1435–1449 (1995).
[28]L. Zhu, L. Lee, Y. Ma, Y. Ye, R. Mazzeo, and L. Xing, "Using total-variation regularization for intensity modulated radiation therapy inverse planning with field-specific numbers of segments," Phys. Med. Biol. **53**(23), 6653–6672 (2008).
[29]G. Tang, M. Earl, S. Luan, S. Naqvi, and C. Yu, "Converting multiple-arc intensity modulated arc therapy into a single arc for efficient delivery," Int. J. Radiat. Oncol., Biol., Phys. **69**(3), S673 (2007).
[30]L. Zhu and L. Xing, "Search for IMRT inverse plans with piecewise constant fluence maps using compressed sensing techniques," Med. Phys. **36**(5), 1895–1905 (2009).
[31]T. Kim, L. Zhu, T.-S. Suh, S. Geneser, B. Meng, and L. Xing, "Inverse planning for IMRT with nonuniform beam profiles using total-variation regularization (TVR)," Med. Phys. **38**(1), 57–66 (2011).
[32]L. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," Physica D **60**, 259–268 (1992).
[33]D. Strong and T. Chan, "Edge-preserving and scale-dependent properties of total variation regularization," Inverse Probl. **19**(6), S165–S187 (2003).
[34]T. Bortfeld, W. Schlegel, and B. Rhein, "Decomposition of pencil beam kernels for fast dose calculations in three-dimensional treatment planning," Med. Phys. **20**(2), 311–318 (1993).
[35]A. Ahnesjö, "Collapsed cone convolution of radiant energy for photon dose calculation in heterogeneous media," Med. Phys. **16**(4), 577–592 (1989).
[36]P. Gill, W. Murray, and M. Saunders, "SNOPT: An SQP algorithm for large-scale constrained optimization," SIAM Rev. **47**(1), 99–131 (2005).
[37]R. Bokrantz and A. Forsgren, "An algorithm for approximating convex Pareto surfaces based on dual techniques," INFORMS J. Comput. ijoc.1120.0508, published online before print June 6, 2012.
[38]J. Deasy, "Multiple local minima in radiotherapy optimization problems with dose-volume constraints," Med. Phys. **24**(7), 1157–1161 (1997).
[39]J. Llacer, J. Deasy, T. Bortfeld, T. Solberg, and C. Promberger, "Absence of multiple local minima effects in intensity modulated optimization with dose volume constraints," Phys. Med. Biol. **48**(2), 183–210 (2003).
[40]Q. Wu and R. Mohan, "Multiple local minima in IMRT optimization based on dose-volume criteria," Med. Phys. **29**(7), 1514–1527 (2002).
[41]International Commission on Radiation Units and Measurements, "Prescribing, recording, and reporting photon-beam intensity-modulated radiation therapy (IMRT)," Report No. 83 (ICRU Publication, Oxford University Press, Oxford, UK, 2010), Vol. 11, Issue (1).
[42]International Commission on Radiation Units and Measurements, "Prescribing, recording and reporting photon beam therapy (supplement to ICRU Report 50)," ICRU Report 62 (ICRU Publication, Bethesda, MD, 1999).
[43]H. Chung, B. Lee, E. Park, J. Lu, and P. Xia, "Can all centers plan intensity-modulated radiotherapy (IMRT) effectively? An external audit of dosimetric comparisons between three-dimensional conformal radiotherapy and IMRT for adjuvant chemoradiation for gastric cancer," Int. J. Radiat. Oncol., Biol., Phys. **71**(4), 1167–1174 (2008).
[44]J. Bohsung, S. Gillis, R. Arrans, A. Bakai, C. de Wagter, Tommy Knöös, B. Mijnheer, M. Paiusco, B. Perrin, H. Welleweerd, and P. Williams, "IMRT treatment planning—a comparative inter-system and inter-centre planning exercise of the ESTRO QUASIMODO group," Radiother. Oncol. **76**(3), 354–361 (2005).
[45]B. Wu, F. Ricchetti, G. Sanguineti, M. Kazhdan, P. Simari, R. Jacques, R. Taylor, and T. McNutt, "Data-driven approach to generating achievable dose-volume histogram objectives in intensity-modulated radiotherapy planning," Int. J. Radiat. Oncol. **79**(4), 1241–1247 (2011).
[46]P. Hansen, "Analysis of discrete ill-posed problems by means of the L-curve," SIAM Rev. **34**(4), 561–580 (1992).
[47]A. Chvetsov, "L-curve analysis of radiotherapy optimization problems," Med. Phys. **32**(8), 2598–2605 (2005).