AUTOMATED IMPROVEMENT OF RADIATION THERAPY TREATMENT PLANS BY OPTIMIZATION UNDER REFERENCE DOSE CONSTRAINTS

Albin FREDRIKSSON*

Abstract

A method is presented that automatically improves upon previous treatment plans by optimization under reference dose constraints. In such an optimization, a previous plan is taken as reference and a new optimization is performed towards some goal, such as minimization of the doses to healthy structures, under the constraint that no structure can become worse off than in the reference plan. Two types of constraints that enforce this are discussed: either each voxel or each dose-volume histogram of the improved plan must be at least as good as in the reference plan. These constraints ensure that the quality of the dose distribution cannot deteriorate, something that constraints on conventional physical penalty functions do not. To avoid discontinuous gradients, which may restrain gradient-based optimization algorithms, the positive part operators that constitute the optimization functions are regularized. The method was applied to a previously optimized plan for a C-shaped phantom and the effects of the choice of regularization parameter were studied. The method resulted in reduced integral dose and reduced doses to the organ at risk while maintaining target homogeneity. It could be used to improve upon treatment plans directly or as a means of quality control of plans.

1. Introduction

Studies show that the quality of intensity-modulated radiation therapy (IMRT) treatment plans is closely linked to the experience of the treatment planner [2,5]. This suggests that many treatment plans have room for improvement. Other causes of suboptimal plans include the use of inadequate planning methods, termination of the optimization process prior to optimality, or the measurement of plan quality with improper or blunt fitness functions. Examples of the latter are the semideviation penalties, which penalize voxels with doses above or below some threshold, that are often utilized in treatment plan optimization. When just voxels that receive doses higher than a prescribed dose level are penalized, the only incentive for the optimization algorithm is to reduce the dose to that very level, even though lower doses may well be attainable without any sacrifice of other treatment goals.

_____

*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. E-mail: `albfre@kth.se` and `albin.fredriksson@raysearchlabs.com`.

Consequently, even plans that are optimal with respect to such penalty functions can be improved upon. Attempts to solve this by maximum dose prescriptions of zero dose to organs at risk (OARs) tend to exacerbate the conflict between target coverage and low OAR doses, thereby making the treatment planning process more challenging and reducing the prospect of achieving satisfactory target coverage.

In the present paper, tailored optimization problems that exploit the possible leeway left by other treatment planning methods are formulated. To this end, the dose distribution of a previous plan is taken as a reference and an optimization is performed towards some goal while either the dose to each voxel or the dose-volume histogram (DVH) of each region of interest (ROI) is constrained to be at least as good as in the previous plan. The output from the optimization is a deliverable treatment plan. With this method, the aim is to automatically create better treatment plans and to find the criteria of a given plan that may be more strictly enforced. Thus, the method could be used in a clinical setting to improve upon plans and to assess plan quality as well as in the education of treatment planners to determine where their plans could be improved.

The potential suboptimality of IMRT treatment plans has inspired methods for plan quality assessment and improvement previously. Wu et al. [20] introduced a method for quality control of plans. In that method, the patient geometry of a newly optimized plan is matched against those in a database of previous plans. Structures for which the doses in the new plan are worse than the doses in previous plans with similar geometries are flagged as potential subjects for improvements. The plans are then reoptimized with stricter dose requirements for the flagged structures. The method improved the OAR sparing without sacrificing target coverage for a number of head-and-neck cases. Subsequently, Wu et al. [21] used the geometrical matching and the database of previous plans to generate suitable DVH objectives for new plans, which resulted in much improved planning efficiency. Moore et al. [17] used historical patient data to define a model to predict the achievable mean OAR dose as a function of the prescribed target dose and the overlap between the OAR and the planning target volume (PTV). Using the predictions, clinicians were able to improve on the normal tissue sparing, and the variability of plans between clinicians was reduced. For a review of quantitative metrics used for dose quality assessment in treatment planning, see Moore et al. [18]

Other methods are aimed at refining the optimization problem formulations in order to improve the quality of treatment plans. Cotrutz and Xing [6,7] used voxel-specific penalty weighting factors to convey the importance of different volumes in the patient geometry to the optimization algorithm. They performed optimizations iteratively to carefully balance the trade-offs between planning criteria. Wu et al. [22] also performed optimizations iteratively, in which they updated the weighting factors or the prescribed doses to each voxel. They found that the two types of updates are equivalent under certain conditions. In a similar vein, Lougovski et al. [14] constructed a scheme in which the prescribed dose levels were iteratively updated for voxels not satisfying the ideal dose prescriptions of uniform target dose and no dose to healthy structures. Holdsworth et al. [10] used a multiobjective evolutionary algorithm to explore the search space resulting when importance weights as well as

prescribed dose levels were allowed to vary for each voxel independently. They found that an enlarged search space can lead to higher quality plans, implying that too simplistic optimization functions may restrict the plan quality unnecessarily.

Similar to the methods discussed above, the method proposed in the present paper can determine whether plans can be improved upon (and also improves on the plans) and increases the flexibility of the optimization by applying an individual prescribed dose level to each voxel. However, it is designed to be applied only once, without any user input, and does not require multiple optimizations, nor does it require a database of previously optimized plans. The main contribution of this paper is the idea of improving plans by using constraints that enforce the dose to each voxel or the DVH of each ROI to be at least as good as in the reference plan while some other goal is optimized. An additional contribution is the use of a regularization that enables warm starting of the optimization from a solution that satisfies such constraints.

2. Methods

Given the dose distribution d^{ref} of a treatment plan, an optimization is performed aiming at minimizing the dose under the constraint that no ROI be worse off than under d^{ref}, using either the dose distribution or the DVH curves as measure.

2.1. Optimization formulation

Assume that a treatment plan with a corresponding dose distribution d^{ref} is given. The ROIs of the treatment plan are indexed by the set $\mathcal{R}$, which is partitioned into the sets $\mathcal{O}$ and $\mathcal{T}$ for respectively organs at risk and target ROIs, with $\mathcal{O}$ including the external ROI. The optimization variables (e.g., the machine parameters) are denoted by the vector x, which is a member of the set $\mathcal{X}$ of feasible variables. One way to improve upon the given treatment plan is to optimize some criterion $f(x)$ while all dose criteria, including criteria that were not part of the previous plan, are enforced to be no worse than in the given plan. The latter can be administered by constraints that no voxel can receive a dose that is worse than in d^{ref}. Under the assumption that a uniform dose denoted by $\hat{d}_r$ is desired for each target ROI $r \in \mathcal{T}$, the optimization problem takes the form

$$
\begin{aligned}
&\underset{x \in \mathcal{X}}{\text{minimize}} \quad f(x) && \text{(2.1a)}\\
&\text{subject to} \quad \sum_{i \in \mathcal{V}_r} \Delta_i^r \left(d_i(x) - d_i^{\mathrm{ref}} \right)_+ \leq 0, && r \in \mathcal{O}, && \text{(2.1b)}\\
&\qquad\qquad \sum_{i \in \mathcal{V}_r} \Delta_i^r \left(\min\{d_i^{\mathrm{ref}}, \hat{d}_r\} - d_i(x) \right)_+ \leq 0, && r \in \mathcal{T}, && \text{(2.1c)}\\
&\qquad\qquad \sum_{i \in \mathcal{V}_r} \Delta_i^r \left(d_i(x) - \max\{d_i^{\mathrm{ref}}, \hat{d}_r\} \right)_+ \leq 0, && r \in \mathcal{T}, && \text{(2.1d)}
\end{aligned}
$$

where $\mathcal{V}_r$ enumerates the voxels of ROI r, $d_i(x)$ is the dose to voxel i as a function of the variables x, and Δ_i^r is the relative volume of voxel i that lies within ROI r, such that $\sum_{i \in \mathcal{V}_r} \Delta_i^r = 1$. The shorthand y_+ denotes the positive part $\max\{y, 0\}$ of

y. To avoid increasing plan complexity, constraints on the number of monitor units and, for volumetric-modulated arc therapy, the delivery time could additionally be enforced.

All the constraint functions of problem (2.1) are convex functions of dose. As in the case of conventional treatment plan optimization problems, the problem may still be nonconvex if the set $\mathcal{X}$ or the mapping $x \mapsto d$ is, such as when direct machine parameter optimization is performed. The constraints can be interpreted according to the following:

- (2.1b): No voxel of any OAR can receive higher dose than in the reference dose distribution.

- (2.1c): No voxel of any target can receive lower dose than the lowest of the reference dose in the voxel and the prescribed dose to the target.

- (2.1d): No voxel of any target can receive higher dose than the highest of the reference dose in the voxel and the prescribed dose to the target.

Constraints (2.1c) and (2.1d) thus enforce the dose of each voxel $i \in \mathcal{V}_r$ of each target ROI $r \in \mathcal{T}$ to lie within the interval $[\min\{d_i^{\mathrm{ref}}, \hat{d}_r\}, \max\{d_i^{\mathrm{ref}}, \hat{d}_r\}]$. The admissible DVH curves for an OAR and a target under a given reference dose distribution are illustrated by respectively the darkgray and lightgray areas in Figure 1.


Figure 1. Examples of the areas enclosing the admissible DVHs for an OAR (darkgray) and a target (lightgray) with prescribed dose level $\hat{d}$.

The objective function f in (2.1a) could be chosen to improve upon any feature, such as integral dose, target coverage, sparing of a subset of the OARs, number of monitor units, delivery time, or a combination thereof. In this paper, f will be aimed at reducing the OAR doses. For OARs, common dose prescriptions correspond to semideviation penalties, which are likely to evaluate to zero for some voxels, whereas for targets, all deviations from a narrow band of acceptable dose levels are usually penalized [9]. It can therefore be hypothesized that there is often more to be gained for OARs than for targets.

To ensure that the reference dose optimization does not leave room for improvement, the objective f should always provide incentive to approach better solutions.

A sufficient condition for this is strict convexity of f, which implies a unique optimal solution. When the goal is to minimize the OAR doses, the most important property of f is that it always provides incentive to reduce the doses, i.e., always has strictly positive gradient with respect to dose. This can be accomplished by a function f that is the sum of the equivalent uniform doses (EUDs) of the OARs:

$$f(x) = \sum_{r \in \mathcal{O}} \left(\sum_{i \in \mathcal{V}_r} \Delta_i^r d_i(x)^{a_r} \right)^{1/a_r} \tag{2.2}$$

for parameters $a_r \geq 1$ for $r \in \mathcal{O}$, which is a convex function of dose but, like the constraints, may be a nonconvex function of the optimization variables if the mapping $x \mapsto d$ is.

In theory, there is no need to aggregate the constraints of (2.1) per ROI: satisfying the aggregate constraints is equivalent to satisfying constraints $d_i(x) \leq d_i^{\text{ref}}$ for each voxel $i \in \mathcal{V}_r$ in each OAR $r \in \mathcal{O}$, and analogously for target voxels. Since, in general, nonlinear optimization solvers are necessary for this type of problem (because of the nonlinear relationship between machine parameters and dose unless pencil beam scanning is performed, in which case linear programming could be used), the large number of voxels makes such individual handling too time-consuming to be practical.

For many nonlinear programming solvers, it is advantageous to square the optimization functions in (2.1), or their summands, in order to introduce curvature into the problem. For the objective function (2.2), squaring each individual EUD changes the optimal solution of the optimization problem into one that emphasizes the importance of reducing higher EUDs. If instead the entire objective is squared and at the same time either the summands or the entire sums of the constraints are squared, the optimal solution remains the same as the one to (2.1) with objective (2.2).

2.2. Reference DVH constraints

Problem (2.1) enforces each voxel to be no worse off than in the reference dose distribution. This might be an overly restrictive request since in reality, the DVHs may be of more interest than the individual voxels. One can thus modify the constraints as not to penalize the deviation of each voxel from the reference dose, but the deviation of each DVH from the corresponding DVH of the reference dose distribution. The optimization functions then becomes nonconvex (since not even the common minimum and maximum DVH functions are convex). The DVH constraints depend on the functions $D_r(v; x)$ and $D_r^{\text{ref}}(v)$, which parameterize respectively the planning DVH, given the optimization variables $x \in \mathcal{X}$, and the reference DVH as functions of the volume $v \in (0, 1]$. The function $D_r(v; x)$ is defined as the highest dose level $\hat{d}$ in the planning dose distribution $d(x)$ such that a fraction v of the volume of ROI $r \in \mathcal{R}$ receives a dose greater than or equal to $\hat{d}$. The function $D_r^{\text{ref}}(v)$ is similar but for the reference dose distribution d^{ref}. Instead of the maximum reference dose constraint (2.1b), a maximum reference DVH constraint is now used, which does not

allow any part of the DVH to exceed the reference DVH, and analogously for the target constraints (2.1c) and (2.1d). The reference DVH constraints are formulated as

$$\int_0^1 \left(D_r(v; x) - D_r^{\text{ref}}(v) \right)_+ dv \leq 0, \qquad\qquad r \in \mathcal{O}, \qquad (2.3a)$$

$$\int_{\max\{v:D_r^{\text{ref}}(v)=\hat{d}_r\}}^1 \left(D_r^{\text{ref}}(v) - D_r(v; x) \right)_+ dv \leq 0, \quad r \in \mathcal{T}, \qquad (2.3b)$$

$$\int_0^{\min\{v:D_r^{\text{ref}}(v)=\hat{d}_r\}} \left(D_r(v; x) - D_r^{\text{ref}}(v) \right)_+ dv \leq 0, \quad r \in \mathcal{T}. \qquad (2.3c)$$

Note that the limits of integration for the target constraints (2.3b) and (2.3c) are modified to exclude part of the interval $(0, 1]$ and that the min and max operators of the target constraints (2.1c) and (2.1d) are not needed. Provided that the above constraints are satisfied, monotonicity of the DVH functions ensures that similar constraints, but with respectively $\min\{D_r^{\text{ref}}(v), \hat{d}_r\}$ and $\max\{D_r^{\text{ref}}(v), \hat{d}_r\}$ substituted for $D_r^{\text{ref}}(v)$, are satisfied for the full interval.

An alternative to the reference DVH-based optimization functions, which still is less restrictive than the voxel-specific requirements but preserves convexity, is to enforce the reference dose distribution in voxel clusters. The mean doses of the clusters are then constrained to be no worse than the corresponding mean doses of the reference dose distribution.

2.3.   Regularization of positive part functions

A difficulty with optimization functions like those in the constraints in (2.1) and (2.3) is that the gradients vanish for feasible solutions: For all points satisfying the constraints, all the positive part functions evaluate to zero, which implies that the gradients are zero in the interior of the feasible region. Moreover, since positive part functions have discontinuous derivatives at zero, the gradients of the constraints have discontinuities. Given a feasible point, a gradient-based optimization algorithm thus has a hard time predicting how large a step can be made. For the considered optimization problem, this is mainly troublesome for the target, since the target constraints provide the only incentive to maintain target coverage, while the objective and all other constraints provide incentive to reduce the dose. The vanishing gradients can be circumvented through regularization of the positive part functions into smooth functions with everywhere nonzero gradients. A common smooth and convex regularization of max functions is the log-sum-exp function [11], which is given by

$$\text{lse}_\epsilon(x_1, \ldots, x_n) = \epsilon \log \left(\sum_{i=1}^n \exp(x_i/\epsilon) \right),$$

where $\epsilon > 0$ is a parameter that can be used to determine the exactness of the regularization. As $\epsilon \to 0$, the regularized function converges uniformly to the corresponding max function. To avoid overflow, the function is implemented as

$$\text{lse}_\epsilon(x_1, \ldots, x_n) = x_{\max} + \epsilon \log \left(\sum_{i=1}^n \exp((x_i - x_{\max})/\epsilon) \right),$$

where $x_{\max} = \max\{x_i : i = 1, \ldots, n\}$. In this paper, it is always the case that $n = 2$ with one of the arguments 0; the regularization of the positive part function is given by

$$x_+ \approx \mathrm{lse}_\epsilon(x, 0). \tag{2.4}$$

This approximation satisfies the inequalities $x_+ < \mathrm{lse}_\epsilon(x, 0) \le x_+ + \epsilon \log 2$ for all $x \in \mathbb{R}$. The positive part function and the regularized positive part function are illustrated in Figure 2. The maximum error of $\epsilon \log 2$ is attained at $x = 0$. Therefore,


Figure 2. The log-sum-exp function (solid) and the positive part function (dashed).

when all constraints of the optimization problem (2.1) or a similar problem with constraints of the type (2.3) have been regularized, the corresponding upper bounds must be adjusted to match the errors of the regularizations. For an OAR $r \in \mathcal{O}$, the regularized constraint corresponding to (2.1b) takes the form

$$\sum_{i \in \mathcal{V}_r} \Delta_i^r \mathrm{lse}_\epsilon\left(d_i(x) - d_i^{\mathrm{ref}}, 0\right) \le \epsilon \log 2. \tag{2.5}$$

For target ROIs, the constraints can be defined analogously, but the right-hand sides must be modified to correspond to the regularized functions evaluated in the reference dose distribution. For a target ROI $r \in \mathcal{T}$, the regularized constraints corresponding to (2.1c) and (2.1d) require upper bounds that are respectively given by

$$\sum_{i \in \mathcal{V}_r} \Delta_i^r \mathrm{lse}_\epsilon\left(-\left(d_i^{\mathrm{ref}} - \hat{d}_r\right)_+, 0\right) \qquad \text{and} \qquad \sum_{i \in \mathcal{V}_r} \Delta_i^r \mathrm{lse}_\epsilon\left(-\left(\hat{d}_r - d_i^{\mathrm{ref}}\right)_+, 0\right).$$

When the regularization (2.4) is applied to the positive part functions of a reference DVH constraint of the type (2.3), the resulting constraint corresponding to (2.3a) takes the form

$$\int_0^1 \mathrm{lse}_\epsilon\left(D_r(v; x) - D_r^{\mathrm{ref}}(v), 0\right) dv \le \epsilon \log 2. \tag{2.6}$$

For the target ROIs with regularized reference DVH constraints corresponding to (2.3b) and (2.3c), the upper bounds are respectively given by

$$(\epsilon \log 2)\left(1 - \max\{v : D_r^{\mathrm{ref}}(v) = \hat{d}_r\}\right) \qquad \text{and} \qquad (\epsilon \log 2)\min\{v : D_r^{\mathrm{ref}}(v) = \hat{d}_r\}.$$

It should be noted that a dose distribution satisfying the regularized constraints does not necessarily satisfy the constraints in (2.1) or (2.3). Since the regularized positive part function is everywhere strictly increasing in the argument, it is possible to remain feasible while delivering higher dose than the reference dose to a voxel provided the dose to another voxel is lower than its reference dose. However, the exchange rate for doing so is bad: the deficiency must be considerably larger than the excess. One way to reduce the risk of exceeding the reference dose is to separate the voxels of each ROI into a number of subsets, for each of which a constraint is added to the optimization problem. If the voxels of the OAR $r \in \mathcal{O}$ are separated into subsets $\mathcal{V}_r(s)$ for s indexed by $\mathcal{S}_r$, the new constraints corresponding to constraint (2.5) for OAR r take the form

$$\sum_{i \in \mathcal{V}_r(s)} \Delta_i^r \mathrm{lse}_\epsilon \left(d_i(x) - d_i^{\mathrm{ref}}, 0 \right) \leq \epsilon \sum_{i \in \mathcal{V}_r(s)} \Delta_i^r \log 2, \qquad s \in \mathcal{S}_r, \tag{2.7}$$

and analogously for the target ROIs and DVH constraints. The separation into subsets can be performed in a number of ways. For instance, proximate voxels or voxels with reference doses within some given interval can be chosen to belong to the same set. As the separation into subsets is refined, the feasible region of the regularized problem converges to that of the original problem. When the separation is so fine that there is one singleton subset for each voxel, the feasible regions coincide.

As in the non-regularized case, it is often advantageous to square the regularized optimization functions or their summands in order to introduce curvature to the problem. In the regularized case, the choice results in optimization problems with different optima because of the nonzero right-hand sides.

For many other types of optimization problems, it is not suitable to use the log-sum-exp regularization, since it does not guarantee that accumulated max constraints are satisfied. If it is crucial that no constraint is violated, the regularization is thus not advisable. For the IMRT problem considered here, small constraint violations may however be tolerated in some regions, especially when they are compensated by slack in other regions.

When ϵ is incorporated into the optimization and suitable changes to the penalty function are made so that the minimum occurs for $\epsilon = 0$, the minimum of the original problem (2.1) can be found, see Huyer and Neumaier [11]. In this paper, the value of ϵ is fixed during the optimizations for the ease of using the method with existing software.

3.   Results

The proposed method is applied to a reference plan for a phantom case and the resulting plan is compared to the reference. Further, the effects of the regularization are studied.

3.1.   Computational study

The proposed method was implemented in the RayStation 2.4 treatment planning system (RaySearch Laboratories, Stockholm, Sweden). The reference dose-based

and reference DVH-based optimization functions of this system, which have been previously used to convert plans between different treatment modalities, were modified with the log-sum-exp regularization discussed in the present paper. Plans were optimized using direct step-and-shoot optimization. The optimization algorithm of RayStation is a quasi-Newton sequential quadratic programming algorithm that uses Broyden-Fletcher-Goldfarb-Shanno updates of an approximation of the Hessian of the Lagrangian. During the optimization, a fast dose computation algorithm based on singular value decomposition of pencil beam kernels [3] is used for dose computation. At intermediate iterations, accurate dose is computed using a collapsed cone dose computation algorithm [1], which is also used to compute final dose for the plan comparisons. The optimization is performed on the dose computed using the fast algorithm incremental from that of the accurate algorithm.

The proposed plan improvement method was evaluated on a nine-field IMRT plan for the C-shaped phantom described by the AAPM Task Group 119 [9]. A slice of the geometry is shown in Figure 3. The considered structures of the geometry were the PTV, the OAR, and the external ROI. Since one of the problems the proposed method is intended to solve is that associated with semideviation penalties (i.e., that voxels with doses below the prescribed dose level provide no incentive for the optimization algorithm to reduce doses further), the optimization problem for the reference plan was formulated with such penalties. The goals of the optimization were derived from Ezzel et al. [9] and a standard optimization of a weighted sum of functions representing the treatment goals subject to constraints on other goals was pursued; the optimization problem is given in Table 1 and the optimization functions are formulated mathematically in Appendix A.

The optimization was allowed to run for 80 iterations, the first 7 of which optimized the fluence maps, which were then converted into segments. After that, direct step-and-shoot optimization was performed. Accurate dose was computed after iterations 7, 40, and 80.


Figure 3. A transversal slice of the C-shaped phantom geometry.

For the proposed improvement optimizations, the optimization problem was similar to (2.1) using the EUD objective (2.2) with $a_r = 1$ for all $r \in \mathcal{O}$: to minimize the sum of the mean doses to the OAR and the external ROI with the PTV subtracted under constraints preventing the dose distribution from becoming worse than in the reference plan. Alternatively, DVH constraints according to (2.3) were used. All

Table 1. Optimization functions of the reference plan. The weight ∞ indicates a constraint.

Structure	Function	Dose level [Gy]	Weight
External$-$PTV	Max dose	15	1
OAR	Max dose	25	10
OAR	Max 5 % dose	25	∞
PTV	Max 10 % dose	55	∞
PTV	Max dose	55	100
PTV	Min 95 % dose	50	∞
PTV	Min dose	50	100

constraints were regularized using the log-sum-exp regularization, and thus formulated respectively like (2.5) and (2.6), and analogously for the target constraints. The target reference dose level $\hat{d}$ was set to D_{50} of the reference dose distribution, where D_x of a ROI denotes the minimum dose level with an isodose volume containing x % of the ROI. The summands and integrands of the reference dose and DVH functions were squared, as were the individual EUDs of the OARs in the objective. Since the improvement optimizations were warm started from the reference plan, they used the same beam directions and number of segments as the reference plan.

The target constraints were split up according to (2.7). When reference dose constraints were used, the regions of the split constraints were selected as equally large subintervals of $[D_{99.25}, D_{0.75}]$ of the target in the reference dose distribution. When reference DVH constraints were used, the regions of min reference DVH constraints were selected as subintervals of $[D_{99.25}, \hat{d}]$ and those of max reference DVH constraints were selected as subintervals of $[\hat{d}, D_{0.75}]$. The subinterval containing $D_{99.25}$ was extended to 0 and that containing $D_{0.75}$ was extended to ∞. Each new constraint took into account the voxels with reference doses in one of the resulting intervals.

3.2.   Plan comparison

The proposed method was applied to the reference plan of the phantom case to create improved plans with clinical doses. The improvement optimizations were allowed to run for 40 iterations with accurate dose computations after iterations 20 and 40. The regularization parameter was set to 1/8 (with doses in Gy) and each target constraint was split into 3 constraints.

DVHs for the reference plan and the improved plans are shown in Figure 4. The proposed method leads to decreased OAR and external ROI doses for both considered types of constraints. Since the reference DVH constraints are less restrictive than the reference dose constraints, the plan with reference DVH constraints achieves lower healthy structure doses. The target coverage and homogeneity are similar in the three plans. These effects are also reflected in the dose statistics, shown in Table 2.
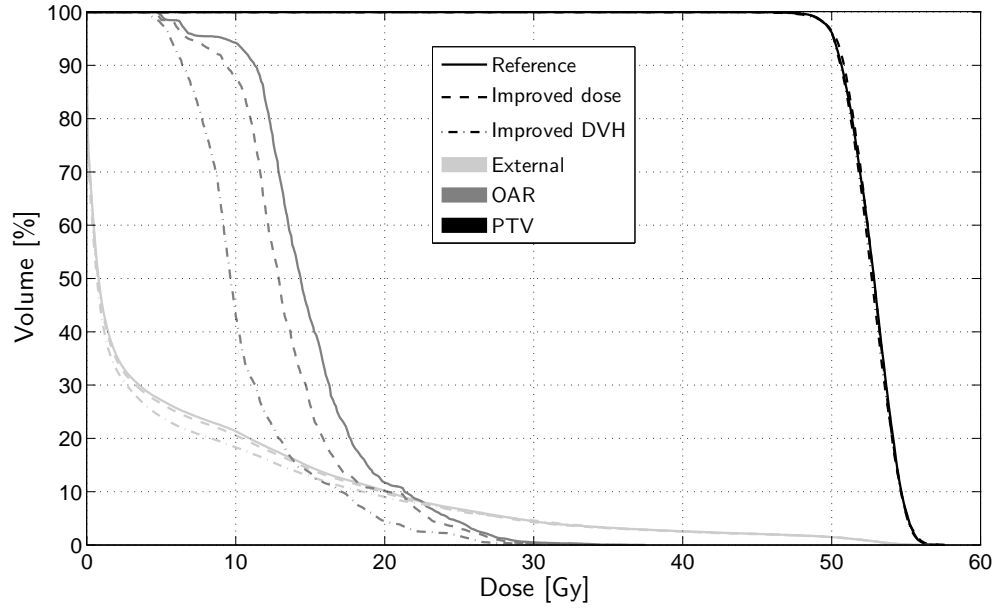
Figure 4. DVHs of the reference plan (solid), the improved dose plan (dashed) and the improved DVH plan (dot-dashed).

Table 2. Dose statistics for the reference plan, the improved dose plan, and the improved DVH plan. For each given ROI, $\bar{D}$ denotes the mean dose.

Structure	Statistic	Reference	Improved dose	Improved DVH
External	$\bar{D}$ [Gy]	5.9	5.7	5.4
OAR	$\bar{D}$ [Gy]	15.1	13.7	10.7
PTV	$\bar{D}$ [Gy]	52.7	52.8	52.7
PTV	D_1 [Gy]	55.8	55.9	55.8
PTV	D_2 [Gy]	55.6	55.6	55.5
PTV	D_{98} [Gy]	49.5	49.5	49.5
PTV	D_{99} [Gy]	49.0	48.9	48.9

3.3. Regularization error

The error introduced by the regularization was studied for improvement optimizations with regularization parameter $\epsilon = 2^{-k}$ for $k = 0, \ldots, 4$ (with doses in Gy) and with each target constraint split into one, three, or five constraints. With reference dose constraints, numerical difficulties prevented the optimization from proceeding when $k \geq 8$, and with reference DVH constraints, when $k \geq 6$. Since the goal was to determine the effect of the regularization alone, accurate dose was not computed during or after these optimizations, and they were allowed to run for 80 iterations.

The mean violation and D_{99} of the target for the resulting plans are shown in Figure 5. For the plan with reference dose constraints, the mean dose distribution violation was considered, which is the sum of the left-hand sides of constraints (2.1c) and (2.1d). For the plan with reference DVH constraints, the mean DVH violation was considered, which is the sum of the left-hand sides of constraints (2.3b) and (2.3c).
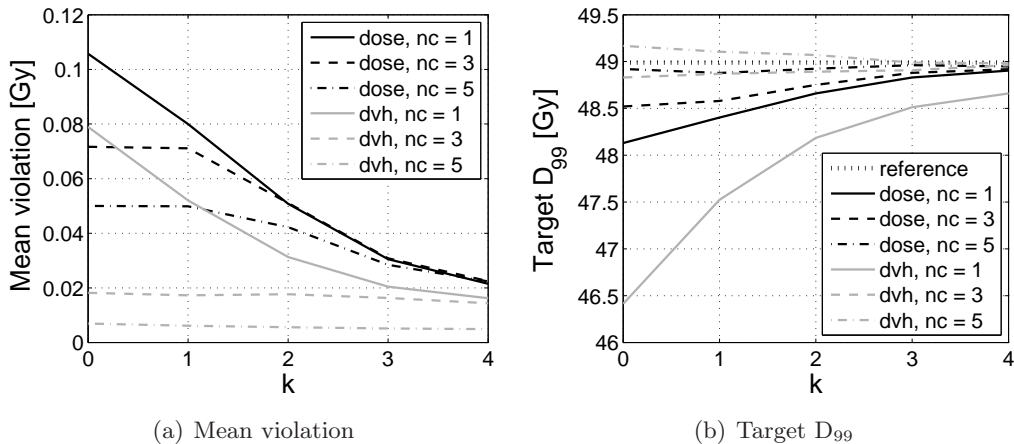

(a) Mean violation   (b) Target D_{99}

Figure 5. The effects on the target dose distribution and DVH of the choice of regularization parameter $\epsilon = 2^{-k}$ for $k = 0, \ldots, 4$ and the number of constraints $\mathsf{nc} = 1, 3, 5$ that each target constraint was split into. Plans using reference dose constraints are indicated by "dose" and those using reference DVH constraints by "dvh" in the legends.

The mean violation decreases as the regularization parameter decreases and as the number of constraints increases. For the plans with each target constraint split into 5 constraints, the mean violations of the optimizations with reference dose and reference DVH constraints are less than respectively 0.05 Gy and 0.01 Gy for all considered values of ϵ, and their D_{99} levels are at most respectively 0.12 Gy and 0.01 Gy below the level of the reference plan.

4. Discussion

The results show that the reference plan had not been optimized to its full potential and that the proposed method could reduce the doses to the OAR and the external

ROI while maintaining target coverage. While there are criteria that can be used to evaluate plan fitness other than the dose distribution and the DVHs (e.g., plan complexity measures such as number of monitor units), the large magnitudes of the dose decrements make it probable that there was room left for improvements on some of the criteria of the plan. The main cause of this is that the reference plan was created using semideviation penalty functions that the optimization algorithm could satisfy in many voxels for the OAR and the external ROI, leaving little incentive to reduce the doses further, even though it was possible without sacrificing target coverage. Introducing penalties on the mean doses to the healthy structures often leads to deteriorated target coverage. Most likely, there exist weights for mean dose penalties that lead to a solution similar to the one found by the proposed method, but these weights are a priori unknown. The proposed method avoids the problem of selecting these weights by optimizing the goals in a post-processing step.

The regularization parameter ϵ should be chosen as small as possible—otherwise, it is possible that the "improved" plan does not really improve on the reference plan. However, too small values can lead to numerical difficulties for gradient-based optimizers. When the constraints are separated into multiple constraints, the choice of the regularization parameter becomes less crucial. Since the objective favors dose decrements, the regularization is likely of less importance for the OARs than for the target. Reducing ϵ—or dropping the regularization altogether—for the OAR constraints could make additional constraints for the OARs unnecessary.

Optimizing with constraints dependent on a reference plan is related to the multicriteria optimization concept of lexicographic ordering (LO) [4, 12, 13, 15, 19]. In LO, the treatment goals are ordered into a number of priority levels. The goals of highest priority are first optimized with the other goals neglected. The attained function values of these first goals are then introduced as constraints for the corresponding functions (possibly with some slack), and a new optimization of the goals of second highest priority is performed under these constraints. This procedure is repeated for the goals of all priority levels.

The constraints of the proposed method differ from those used in LO. In the proposed method, the functions that were used to acquire the reference plan are neglected, and the subsequent optimization enforces the individual voxel doses or the DVHs to become no worse than in the reference plan. The new dose distribution is thereby constrained to be at least as good as the previous one. In LO, the only guarantee is that the values of the functions used to acquire the reference plan do not deteriorate. Since different dose distributions can yield the same function values, this does not guarantee maintained dose distribution quality. It is thus feasible in LO to move hot spots, or even create new ones when constraints on DVH points are used. However, if in the last step of LO the doses to all OARs are minimized under constraints on all previous goals, this last step can be seen as a relaxation of the proposed method. (To see that it is a relaxation, note that a function satisfying the constraints of the proposed method will also satisfy the DVH and EUD constraints typically used in LO, but not the other way around.) Conversely, the proposed method can be seen as providing a new and more stringent means of preserving the plan quality in LO.

In theory, the problem of plans that can be improved upon could also be solved by multicriteria optimization with Pareto surface navigation [16], provided mean dose objectives are introduced for the OARs. In practice, the Pareto surface is approximated by a discrete set of points and a navigated IMRT plan must be converted before it becomes deliverable [8]. The approximation and conversion errors may lead to plans that can still be improved upon.

Studying the proposed method applied to clinical cases is a delicate issue. If the method is able to improve upon the plans, it indicates that patients are suboptimally treated. This may be interpreted as resulting from a treatment planner performing poorly or as an indication of an unintuitive treatment planning system. Nevertheless, the method provides a means to solve these possible issues. In future studies, the method should be used to determine the quality of clinical plans.

5. Conclusion

A method for automatically improving upon any given treatment plan was presented. It takes a reference plan as input and optimizes some criterion, such as the doses to healthy structures, while enforcing that no ROI becomes worse off than in the reference plan—either in the sense of voxel doses or in the sense of DVH curves. Numerical problems associated with the enforcing constraints are dealt with by regularization. Due to the regularization, the constraints are not always as strictly enforced as desired. With an appropriately low regularization parameter or with each constraint split into multiple constraints, this problem is alleviated. The proposed method was applied to a phantom case with a previously optimized plan, for which it resulted in substantial decrements in the doses to healthy tissues while preserving the target coverage and homogeneity of the original plan. This indicates that the method could be used in a clinical setting to improve upon plans prior to delivery or as a quality control of plans, and for educational purposes, since it shows where the treatment planner could enforce stricter requirements in the optimization.

Acknowledgements

A. Optimization functions

Given $D(v; x)$ parameterizing the DVH of some ROI as a function of the volume $v \in (0, 1]$ and the optimization variables $x \in \mathcal{X}$, a max DVH function with dose level $\hat{d}$ and volume parameter $\hat{v}$ is given by

$$\int_{\hat{v}}^{1} \left(D(v; x) - \hat{d} \right)_{+}^{2} dv.$$

Min DVH functions are similar, but with the signs of $D(v; x)$ and $\hat{d}$ reversed and the integration taken over $(0, \hat{v}]$. Max and min dose functions are derived from the

corresponding DVH functions with $\hat{v}$ set to respectively 0 and 1.

References

[1] A. Ahnesjö. Collapsed cone convolution of radiant energy for photon dose calculation in heterogeneous media. *Medical Physics*, 16(4):577–592, 1989.

[2] J. Bohsung, S. Gillis, R. Arrans, A. Bakai, C. De Wagter, T. 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. *Radiotherapy and Oncology*, 76(3):354–361, 2005.

[3] T. Bortfeld, W. Schlegel, and B. Rhein. Decomposition of pencil beam kernels for fast dose calculations in three-dimensional treatment planning. *Medical Physics*, 20(2):311–318, 1993.

[4] S. Breedveld, P. R. M. Storchi, M. Keijzer, A. W. Heemink, and B. J. M. Heijmen. A novel approach to multi-criteria inverse planning for IMRT. *Physics in Medicine and Biology*, 52(20):6339, 2007.

[5] H. T. Chung, B. Lee, E. Park, J. 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. *International Journal of Radiation Oncology*Biology*Physics*, 71(4):1167–1174, 2008.

[6] C. Cotrutz and L. Xing. Using voxel-dependent importance factors for interactive DVH-based dose optimization. *Physics in Medicine and Biology*, 47(10):1659–1669, 2002.

[7] C. Cotrutz and L. Xing. IMRT dose shaping with regionally variable penalty scheme. *Medical Physics*, 30(4):544–551, 2003.

[8] D. Craft, F. Carlsson, T. Bortfeld, and H. Rehbinder. Multi-objective IMRT planning which produces deliverable plans. *Medical Physics*, 35(6):2846–2846, 2008.

[9] G. A. Ezzell, J. W. Burmeister, N. Dogan, T. J. LoSasso, J. G. Mechalakos, D. Mihailidis, A. Molineu, J. R. Palta, C. R. Ramsey, B. J. Salter, J. Shi, P. Xia, N. J. Yue, and Y. Xiao. IMRT commissioning: multiple institution planning and dosimetry comparisons, a report from AAPM task group 119. *Medical Physics*, 36(11):5359–73, 2009.

[10] C. Holdsworth, M. Kim, J. Liao, and M. Phillips. The use of a multiobjective evolutionary algorithm to increase flexibility in the search for better imrt plans. *Medical Physics*, 39(4):2261–2274, 2012.

[11] W. Huyer and A. Neumaier. A new exact penalty function. *SIAM Journal on Optimization*, 13:1141–1158, October 2002.

[12] K.-W. Jee, D. L. McShan, and B. A. Fraass. Lexicographic ordering: intuitive multicriteria optimization for IMRT. *Physics in Medicine and Biology*, 52(7):1845, 2007.

[13] J. Löf. *Development of a general framework for optimization of radiation therapy.* PhD thesis, Stockholm University, 2000.

[14] P. Lougovski, J. LeNoach, L. Zhu, Y. Ma, Y. Censor, and L. Xing. Toward truly optimal IMRT dose distribution: inverse planning with voxel-specific penalty. *Technology in Cancer Research and Treatment*, 9(6):629–636, 2010.

[15] K. Miettinen. *Nonlinear Multiobjective Optimization*, volume 12 of *International Series in Operations Research and Management Science*. Kluwer Academic Publishers, Dordrecht, 1999.

[16] M. Monz, K. H. Kufer, T. R. Bortfeld, and C. Thieke. Pareto navigation—algorithmic foundation of interactive multi-criteria IMRT planning. *Physics in Medicine and Biology*, 53(4):985–998, 2008.

[17] K. Moore, R. Brame, D. Low, and S. Mutic. Experience-based quality control of clinical intensity-modulated radiotherapy planning. *International Journal of Radiation Oncology*Biology*Physics*, 81(2):545–551, 2011.

[18] K. Moore, R. Brame, D. Low, and S. Mutic. Quantitative metrics for assessing plan quality. *Seminars in Radiation Oncology*, 22(1):62–69, 2012.

[19] J. J. Wilkens, J. R. Alaly, K. Zakarian, W. L. Thorstad, and J. O. Deasy. IMRT treatment planning based on prioritizing prescription goals. *Physics in Medicine and Biology*, 52(6):1675, 2007.

[20] B. Wu, F. Ricchetti, G. Sanguineti, M. Kazhdan, P. Simari, M. Chuang, R. Taylor, R. Jacques, and T. McNutt. Patient geometry-driven information retrieval for IMRT treatment plan quality control. *Medical Physics*, 36(12):5497–5505, 2009.

[21] 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. *International Journal of Radiation Oncology*Biology*Physics*, 79(4):1241–1247, 2011.

[22] C. Wu, G. H. Olivera, R. Jeraj, H. Keller, and T. R. Mackie. Treatment plan modification using voxel-based weighting factors/dose prescription. *Physics in Medicine and Biology*, 48(15):2479–2491, 2003.