

A linear programming approach to inverse planning in Gamma Knife radiosurgery

J. Sjölund,^{a)} S. Riad, M. Hennix, and H. Nordström
Elekta Instrument AB, Kungstensgatan 18, Box 7593, SE-103 93 Stockholm, Sweden

(Received 10 June 2018; revised 19 December 2018; accepted for publication 15 January 2019; published 8 March 2019)

Purpose: Leksell Gamma Knife[®] is a stereotactic radiosurgery system that allows fine-grained control of the delivered dose distribution. We describe a new inverse planning approach that both resolves shortcomings of earlier approaches and unlocks new capabilities.

Methods: We fix the isocenter positions and perform sector-duration optimization using linear programming, and study the effect of beam-on time penalization on the trade-off between beam-on time and plan quality. We also describe two techniques that reduce the problem size and thus further reduce the solution time: dualization and representative subsampling.

Results: The beam-on time penalization reduces the beam-on time by a factor 2–3 compared with the naïve alternative. Dualization and representative subsampling each leads to optimization time-savings by a factor 5–20. Overall, we find in a comparison with 75 clinical plans that we can always find plans with similar coverage and better selectivity and beam-on time. In 44 of these, we can even find a plan that also has better gradient index. On a standard GammaPlan workstation, the optimization times ranged from 2.3 to 26 s with a median time of 5.7 s.

Conclusion: We present a combination of techniques that enables sector-duration optimization in a clinically feasible time frame. © 2019 The Authors. *Medical Physics* published by Wiley Periodicals, Inc. on behalf of American Association of Physicists in Medicine. [<https://doi.org/10.1002/mp.13440>]

Key words: inverse planning, Leksell Gamma Knife[®], optimization, stereotactic radiosurgery, treatment planning

1. INTRODUCTION

Stereotactic radiosurgery (SRS) is defined as the use of externally generated ionizing radiation to inactivate or eradicate defined targets, typically in the head or spine, without the need for a surgical incision.¹ Present-day neurosurgeons routinely use stereotactic radiosurgery for the management of a wide variety of brain disorders, including certain malignant and benign tumors,^{2–5} as well as cardiovascular and functional disorders within the brain.^{6–9}

Leksell Gamma Knife[®] (LGK) is a dedicated system for intracranial stereotactic radiosurgery. Its most recent incarnations, Perfexion[™] and Icon[™], use 192 ⁶⁰Co sources, each emitting gamma radiation. The radiation is collimated to create a focus where the radiation from every source converges. At the focus, both the radiation intensity and its gradient become very large. This makes it possible to deliver high radiation doses with minimal damage to surrounding healthy tissue.

The Perfexion[™] and Icon[™] systems enable two ways of tailoring the radiation dose according to the shape and size of the target. First, the patient can be precisely moved (robotically) in relation to the focus, effectively placing the focus in different *isocenters*. Second, the radiation sources are arranged in eight, independently controlled, sectors. Each sector can be in one of four different collimator states: the 4, 8, or 16 mm or in the beam-off state. For each isocenter position and collimator configuration (i.e., collimator size for each sector), the irradiation time can be specified. This composition is often referred to as a *shot*.

The large number of degrees of freedom allows sculpting of the dose distribution in unparalleled ways. At the same time, however, it is infeasible to explore them all by means of manual planning. Thus, an inverse planning method is required to make the full potential of LGK clinically accessible. Inverse planning methods only require the user to specify what objectives to strive for, and then uses mathematical optimization to search for the best possible treatment plan according to these objectives.

Inverse planning was introduced in Leksell GammaPlan[®] 10, and has since then become widely adopted. This inverse planner uses well-established metrics such as coverage, selectivity, and gradient index at a predetermined isodose level, together with a beam-on time (BOT) penalization. The optimization variables are the position of isocenters, collimator configurations, and irradiation times (beam weights more precisely). Unfortunately, the optimization problem is inherently difficult (nonconvex). The difficulty arises from two components: that the objectives use relative isodoses (instead of absolute doses) and that the positions of the shots can change. Besides, the direct use of relative isodoses makes it difficult to simultaneously manage multiple targets and to enforce criteria such as the maximum dose to an organ at risk (OAR).

Researchers have proposed optimization approaches that use absolute doses and shots that are allowed to move.^{10–15} This typically results in a so-called mixed-integer problem, which remains nonconvex. In practice, this implies that there is a compromise between computation time and the risk of ending up in a suboptimal solution. A consequence

is that it is difficult to explore what trade-offs are achievable — especially in complicated cases with multiple conflicting objectives.

A remedy to most of these problems is to formulate a *convex* optimization problem. Convexity is a highly desirable property that allows optimization problems to be solved reliably and efficiently.¹⁶ One way to achieve convexity — which we also use — is to fix the isocenter positions and perform sector-duration optimization.¹⁷ This approach is inspired by optimization formulations, in particular fluence map optimization, that are common in intensity-modulated radiation therapy (IMRT).^{18,19} In sector-duration optimization, the collimator configurations are not packaged into shots during the optimization, instead the irradiation times of every collimator in every sector are treated independently during the optimization. The irradiation times are converted into deliverable shots after a solution has been found. Despite its promises, methods using sector-duration optimization has, so far, lacked an efficient way of controlling BOTs^{20,21} and have been too computationally costly for widespread clinical use.^{15,17,22} This is about to change.

We present the first sector-duration optimization that uses linear programming, which is possible thanks to a well-founded BOT penalization. Linear programming has a long history in radiotherapy,^{23–27} dating back at least to 1968.²⁸ Many of the early approaches constrained the doses to predetermined intervals and used the remaining freedom to maximize an objective function such as the mean²⁶ or minimum²⁹ dose in the target. But, in practice, this upfront specification of dose constraints often lead to infeasible problems. This motivated the use of elastic constraints,³⁰ where one instead minimizes the mean or maximum of the (weighted) constraint violations. Our formulation belongs to this category, but with an additional term for BOT penalization.

Linear programming has also been used in radiosurgery before, specifically in two-phase approaches to inverse planning both for Gamma Knife radiosurgery³¹ and for robotic radiosurgery.^{32–34} There, the decision variables are the weights of (deliverable) beams/shots that are generated in the first phase and optimized in the second. In Gamma Knife radiosurgery, the sector-duration optimization has 24 times the number of decision variables of the corresponding two-phase procedure. The equivalent of a sector-duration optimization for a system with a multileaf collimator would include a decision variable for the left and right positions of every leaf, resulting in a roughly hundredfold increase of decision variables. In this work, we present novel contributions that reduce the problem size of the sector-duration optimization while preserving the flexibility it offers. In summary, our method

- Manages multiple targets;
- Explicitly handles OARs;
- Reproducibly finds the optimum given a fixed set of isocenter locations;
- Runs in well under a minute;
- Efficiently optimizes for short BOTs;
- Allows hard constraints;
- Allows the exploration of achievable trade-offs.

2. MATERIALS AND METHODS

To achieve convexity, we divide the planning into three distinct phases: isocenter placement, optimization, and sequencing. The isocenters chosen in the first phase remain fixed throughout the rest of the planning. In the optimization phase, we formulate an optimization problem where competing objectives are combined as a weighted sum. By changing weights, it is straightforward to explore achievable trade-offs. Possible objectives include dose to target, sparing of OARs, and a new — highly efficient — BOT penalization. During the optimization, times for each sector and collimator are allowed to vary independently. In the sequencing phase, these times are converted into deliverable shots.

2.A. Isocenter placement

The first phase of the proposed inverse planner is to choose isocenter positions. These remain fixed throughout the subsequent phases, defining the search space of the optimization. The main objective of this phase is thus to provide enough freedom to find a high quality plan; the only harm in including extra isocenter positions is that it will take longer time to solve the optimization problem.

Although algorithms for automatic isocenter placement^{12,17,35–38} certainly have a role to play, we do not consider them to be the main focus of this work. Consequently, to remove this source of variability, we will reuse the isocenter positions from the corresponding (manual) reference plan in all examples that follow.

2.B. Optimization

As we will describe below, a major difference between the inverse planner we propose and previous approaches based on sector-duration optimization^{15,17,20–22} is that it uses *linear programming*. Such optimization problems are well behaved and well studied.^{16,39,40}

We recognize that there are multiple, possibly conflicting, objectives that are desirable. However, our working assumption is that the exact priority among these should be provided by the user. Consequently, we foresee that a number of options will be evaluated in the course of optimizing the treatment plan. This indeed reflects how it works presently, but the intention is that the proposed inverse planner will elucidate the achievable trade-offs. Formally, suppose we have m cost functions, $f = (f_1(x), \dots, f_m(x))^t$, representing, for example, target dose, dose to an OAR, and a penalization of BOT. This can be conveniently rearranged into a single cost function by forming a weighted sum, that is, by scalar multiplication with the weights $w = (w_1, \dots, w_m)^t$, where each weight quantifies the importance of the corresponding part of the cost function. We will now describe how we have defined the different parts of the optimization problem so that, taken together, it can be expressed as a linear programming problem.

2.B.1. Dose-based objectives

Given a fixed set of isocenter positions, our optimization variables are the irradiation times t_{isc} corresponding to every isocenter i , collimator state c , and sector s . The dose D at a position r is linear in terms of these irradiation times,

$$D(r, t_{isc}) = \sum_{i=1}^{N_{iso}} \sum_{s=1}^8 \sum_{c=1}^3 \Phi_{isc}(r) t_{isc}. \quad (1)$$

Often, we are interested in the doses D_n given to a discrete set of voxels at positions r_n , where $n = 1, \dots, N$. Then, after appropriate rearrangements, we may evaluate the dose as a matrix-vector multiplication: $D_n = (\Phi t)_n$. Thus, the $(N \times 24N_{iso})$ -matrix Φ , which we refer to as the dose rate kernel, maps irradiation times to doses.

In regions of interest, each voxel can be assigned dose-based objectives that reflect the amount of underdosage or overdosage it receives. Typically, the regions of interest include all targets, a volume of healthy tissue surrounding each target^{10,17} and clinically relevant OARs. We express the dose-based objectives using one or several hinge functions,

$$(D_n - \hat{D}_n)_+ = \max(D_n - \hat{D}_n, 0), \quad (2)$$

where D_n is the dose in the voxel in question and \hat{D}_n is a reference dose. This function, or its square, is commonly used as a dose-based objective.^{17,41} In principle, every voxel could have a unique dose-based objective, which makes it possible to, for instance, perform full-fledged dose painting by numbers.^{42–44} By introducing auxiliary variables, piecewise linear convex functions such as the hinge function above, can be recast as linear programming problems, cf. Appendix A.

2.B.2. Beam-on time penalization

To control the treatment time, we use a highly efficient penalization function which we refer to as the idealized beam-on time (iBOT).^{20,21} It is defined as

$$\Theta(t) = \sum_{i=1}^{N_{iso}} \max_s \sum_{c=1}^3 t_{isc}, \quad (3)$$

where N_{iso} is the number of isocenters; s and c , respectively, correspond to the eight sectors and three collimators. We emphasize that iBOT is *not* an L_p -norm nor does it promote sparsity — in fact, it is quite the opposite: it encourages the total BOT of each sector to be equally long, which is advantageous since they can irradiate simultaneously. By introducing auxiliary variables, the iBOT function can be expressed using linear programming,⁴⁵ cf. Appendix A.

2.B.3. Constraints

In a linear programming problem, it is possible to include *hard constraints*, that is, conditions that the solution must

fulfill. Physics only dictates one such hard constraint, namely, that all times must be non-negative, $t \geq 0$, and therefore, this is the only hard constraint that is strictly necessary — exempting the “artificial” constraints coming from the auxiliary variables. One can, optionally, include hard constraints on doses in various regions.

2.B.4. Normal tissue sparing

We promote normal tissue sparing by encompassing the target within two nonoverlapping thick “shells” shaped according to the target surface. By penalizing high doses inside the inner and outer shells, we can control selectivity and gradient index (cf. Section 2.D), respectively. These two shells may seem similar to the widely used concept of margins,⁴⁶ but unlike margins — which have well-defined clinical meanings — our shells are merely used to steer the optimization. We construct the inner and outer shells as geometric expansions, defined via the Euclidean distance transform, of the target. The inner shell is expanded from the target surface until its total volume is half that of the target volume. The outer shell is expanded from the outer surface of the inner shell until its volume is twice that of the target.

2.B.5. Illustrative example — primal formulation

We will base our exposition on a minimal, yet illustrative, example of the optimization problem when there is a single target and an OAR where we want to limit the maximum dose to at most D_O . In this case, the cost function has four components that control the target dose, selectivity, gradient index and BOT, respectively. We denote the prescription dose to the target D_T and the dose thresholds for selectivity and gradient index by D_S and D_G , respectively. We thus arrive at the optimization problem

$$\begin{aligned} \underset{t}{\text{minimize}} \quad & \frac{w_T}{D_T N_T} \sum_{i=1}^{N_T} (D_T - (\Phi t)_i)_+ \\ & + \frac{w_S}{D_S N_S} \sum_{i=1}^{N_S} ((\Phi_S t)_i - D_S)_+ \\ & + \frac{w_G}{D_G N_G} \sum_{i=1}^{N_G} ((\Phi_G t)_i - D_G)_+ \\ & + \frac{w_{BOT}}{D_T / \varphi_{cal}} \Theta(t) \end{aligned} \quad (4)$$

subject to $\Phi_O t \leq D_O$

$$t \geq 0,$$

where φ_{cal} is the calibration dose rate, N_α is the number of voxels in the structures $\alpha \in \{T, S, G, O\}$ and the weights w_α govern the relative importance of each term. Typically, we let $D_S = D_T$ to promote selectivity and $D_G = D_T/2$ to reduce the gradient index. In Appendix A, we give the explicit representation of the problem in Eq. (4) as a linear programming problem on standard form.

The example given here can easily be extended in several ways. For instance, additional structures could be incorporated as additional terms in the cost function, and homogeneous (target) dose distributions could be promoted by including a hinge function that penalizes overdosage.

2.B.6. Representative subsampling

The introduction of one auxiliary variable for every voxel in a relevant structure vastly increases the size of the optimization problem. In realistic cases, there could be, say, 10^5 variables. Naïvely solving the optimization problem thus becomes time-consuming at best, but it could even be impossible due to memory limitations. Thus, it is essential to reduce the size of the problem.

As a first step, we propose to use an approximation we refer to as *representative subsampling*. Representative subsampling is based on the observation that the dose tends to vary quite smoothly from one position to another. True enough, there are regions like the penumbra region of the Gamma Knife (and flattening filter-free accelerators) that may exhibit stronger variability than others; however, for practical purposes, it is redundant to take the dose in every voxel into account during the optimization. We exploit this realization by sampling only a representative fraction of the voxels in each structure that we use in the optimization.

Formally, we may understand this from the observation that most of the criteria involving a structure Ω can be phrased as an integral

$$J = \int_{\Omega} f(r, D(r, t)) dr, \quad (5)$$

where f is a function that assigns a cost to the dose, $D(r, t) = \Phi(r) t$, at position r . For instance, a dose-based hinge functions, as in Eq. (2), corresponds to $f(r, D(r, t)) = (D(r, t) - \hat{D}(r))_+$ and a hard constraint on the maximum dose corresponds to

$$f(r, D(r, t)) = \begin{cases} 0 & \text{if } D(r, t) \leq \hat{D}(r), \\ \infty & \text{otherwise.} \end{cases} \quad (6)$$

If we discretize the integral (5) uniformly (i.e., at the sampling points), we recover the conventional expression

$$J \approx \frac{V(\Omega)}{N} \sum_{n=1}^N f(r_n, \Phi(r_n) t), \quad (7)$$

where $V(\cdot)$ is the volume operator. However, it is well known that this scales poorly with dimension — the number of samples required to reach a given precision increases exponentially with the dimension. To mitigate this, researchers have proposed to subsample voxels on a regular grid¹³ or to aggregate voxels into clusters.^{47,48} Another option is to use a stochastic gradient descent algorithm that resamples voxels at each iteration,⁴⁹ but stochastic gradient descent is often ineffective on constrained problems, since every iteration requires a projection onto the feasible set.⁵⁰

However, randomness is powerful. Instead of the deterministic sampling schemes mentioned above, we build upon the well-established efficiency of Monte Carlo integration to approximate Eq. (5) by sampling positions uniformly at random in the structure. In addition, because we want to subsample also when using minimum and maximum dose constraints in the target or maximum dose constraints in organs at risk, we include a separate term corresponding to positions sampled at random on the tessellated surface of the volume.

2.B.7. Dualization

The introduction of auxiliary variables in the primal formulation increases the number of variables — often severalfold. However, as can be appreciated from the explicit formulation in Appendix A, the resulting matrix is highly structured; it could, for example, be decomposed as the sum of a low-rank matrix and a sparse matrix. Exploiting this structure makes it possible to reduce computation times drastically. We have found that dualization reduces the computation time by a factor 5–20 depending on the features of the problem. Since strong duality holds for linear programming problems, the primal and dual problems are equivalent.

In Appendix B, we revisit our earlier example and give the explicit formulas for the corresponding dual problem. Importantly, however, the constraints introduced when rewriting the hinge function using auxiliary variables become trivial thanks to the dualization. The resulting size reduction leads to dramatic performance improvements when solving the dual compared to the original, primal, problem. Moreover, the computational gain due to dualization is entirely complementary to that of representative subsampling.

2.C. Sequencing algorithm for shot composition

To deliver the treatment the Gamma Knife system requires composite shots, that is, collimator and sector configurations for each shot. The conversion from irradiation times for each isocenter position t_{isc} to such shots is referred to as sequencing. The sequencing can be done in several different ways, and although we will not go into details here,⁵¹ it is worth noting that the result of the optimization will often lead to multiple shots in the same isocenter position.

2.D. Evaluation

In Fig. 1, we define the target volume (TV) and the planning isodose volume (PIV), which are used, together with the volume operator $V(\cdot)$, to define the common radiosurgery metrics:⁵²

$$\text{Coverage: } C = \frac{V(PIV \cap TV)}{V(TV)}, \quad (8)$$

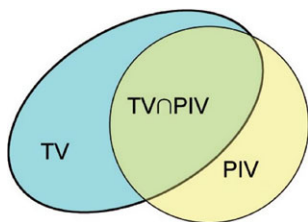


FIG. 1. The target volume (TV) and planning isodose volume (PIV). [Color figure can be viewed at wileyonlinelibrary.com]

$$\text{Selectivity: } S = \frac{V(\text{PIV} \cap \text{TV})}{V(\text{PIV})}, \quad (9)$$

$$\text{Gradient index: } GI = \frac{V(\text{PIV}_{ISO/2})}{V(\text{PIV}_{ISO})}. \quad (10)$$

In words, coverage is the fraction of the target volume that receives at least as high dose as prescribed, selectivity is the fraction of the planning isodose volume that encompasses the target volume, and gradient index describes the dose fall-off by the ratio of the volume that receives at least half the prescription dose to the planning isodose volume. In addition to these metrics, plan quality is often evaluated based on the Paddick conformity index (PCI),⁵³ which is defined as:

$$\text{Paddick conformity index: } PCI = C \cdot S. \quad (11)$$

3. RESULTS

We have evaluated both the overall performance of the proposed inverse planner as well as studied the novel components — the beam-on time penalization and the subsampling scheme — in isolation. We begin by reporting on the beam-on time penalization and the subsampling scheme. Then, we describe the overall performance on a set of 75 clinically acceptable reference plans. Finally, we consider one case in more detail. For this case, we illustrate the achievable trade-offs and show how to perform dose painting or create a homogeneous dose distribution.

3.A. Beam-on time penalization

To show the efficiency of the iBOT penalization, we compare with replacing the term $\Theta(t)$ in Eq. (4) with a simple summation over all times,

$$\Theta_{\text{simple}}(t) = \sum_{i=1}^{N_{\text{iso}}} \sum_{s=1}^8 \sum_{c=1}^3 t_{isc}, \quad (12)$$

which we refer to as simple BOT (sBOT). The rest of the objective function in Eq. (4) is left unchanged. We make the comparison on three different cases: one small acoustic neuroma, one medium-sized acoustic neuroma, and one irregular meningioma. By varying the weight of the (selectivity promoting) inner shell against the BOT penalization

while keeping the other weights fixed, we obtain plans with different plan quality and BOTs. To get adequate statistics, we ran 1000 optimizations for each case and choice of penalization. Plans with the same Paddick index and gradient index to within $\pm 1\%$ were tallied and the average BOT calculated for each bin. The bins were chosen to have Paddick and gradient indices within $\pm 5\%$ of the clinical plan. Table I summarizes the mean and standard deviation of the average BOT with the iBOT penalization and the average BOT with the sBOT penalization as well as the mean and standard deviation of the ratios computed for each bin.

Clearly, iBOT results in plans of equal quality but markedly shorter BOT than sBOT, in particular for the two larger targets. The optimization times are practically equal. Figure 2 gives an overview of all the simulations for the small acoustic neuroma, for which the BOT reduction was the smallest (but still almost a factor of two).

Figure 3 shows how the BOT, Paddick index, and gradient index depend on the iBOT weight. Even though the Paddick index may appear independent of the iBOT in the small acoustic neuroma case (middle left), recall from Eq. (11) that it is the product of coverage and selectivity. Separately these are not constant as a function of BOT; a shorter BOT is obtained by increasing the total dose rate, which for Gamma Knife is done by using larger collimators and more sectors simultaneously. This smoothens the dose distribution so that coverage increases at the expense of selectivity and gradient index.

3.B. Subsampling

Since we subsample stochastically, the resulting plan metrics also become stochastic. We do, however, want to choose the subsampling fractions such that the statistical fluctuations remain reasonably small. Here, we require the standard deviation of both coverage and selectivity to be below 1%. Furthermore, we want to ensure that our sampling strategy, which samples both interior and surface points, performs at least comparably to sampling only in the interior.

We compare the two strategies on the three cases studied in Section 3.A and an arteriovenous malformation case

TABLE I. Beam-on time (BOT) (min) at 3 Gy/min when using different penalization terms.

	Small acoustic neuroma	Medium acoustic neuroma	Irregular meningioma
iBOT	9.3 (2.6)	21.6 (1.5)	110.8 (3.1)
sBOT	16.8 (1.9)	75.2 (4.4)	303.2 (11.6)
iBOT/sBOT	0.55 (0.16)	0.29 (0.026)	0.37 (0.029)
Optimization times iBOT (s)	0.135	3.2	3.3
Optimization times sBOT (s)	0.125	3.4	3.5

(AVM). Before sampling, the number of target voxels are 6119 (small acoustic neuroma), 12 078 (medium acoustic neuroma), 27 765 (meningioma), and 3 0326 (AVM). We

adjust the voxel size for the interior point sampling so that the total number of points in the target is the same for both sampling methods. We performed 100 runs for each of seven subsampling fractions. Figure 4 presents the resulting standard deviation of coverage and selectivity as a function of the percentage of the total number of voxels, that is, the voxels belonging to all structures in the problem. For each sampling method, we used a fixed weight setting, $(w_T, w_S, w_G, w_{BOT}) = (1.0, 0.15, 0.15, 0.15)$ for our sampling method and $(w_T, w_S, w_G, w_{BOT}) = (1.0, 0.025, 0.025, 0.25)$ for interior sampling only, that rendered clinically acceptable plans with metrics within 1% of each other. From Fig. 4, we conclude that our sampling strategy performs comparably to the interior point sampling for both coverage and selectivity. Also, we see that the standard deviation of both metrics is below 1% when the subsampling fraction is at least 10%. The value of the metrics, and thus the standard deviation, is computed on the original grid. For small sampling fractions, the mean will deviate from the true value. However, for the sampling sizes, we are interested in, that is, about 10%; this is typically not an issue.

As illustrated in Fig. 5, subsampling shortens the optimization time. Evidently, the dependence between the optimization time and the subsampling fraction is approximately linear. In the cases considered above, a subsampling of 10% shortens the optimization time by a factor 8–22 compared to using all initial sampling points.

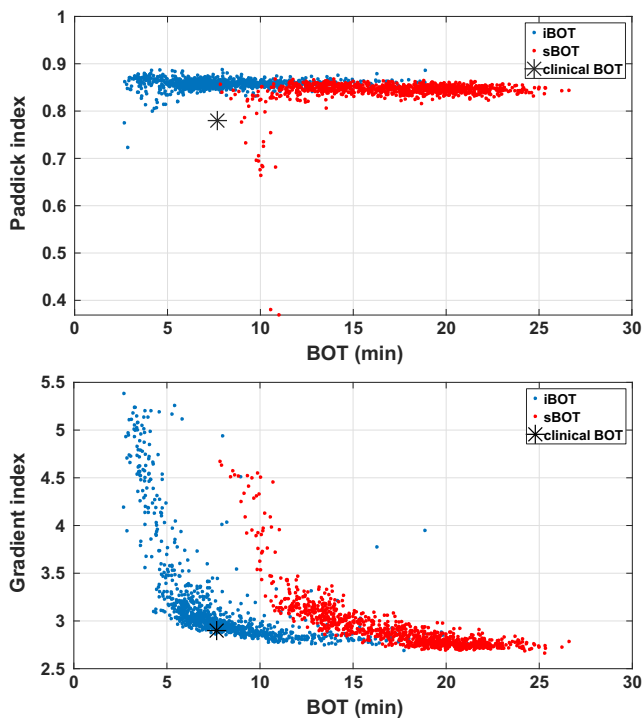


FIG. 2. Comparison of the two beam-on time (BOT) penalization terms for the small acoustic neuroma case. [Color figure can be viewed at wileyonlinelibrary.com]

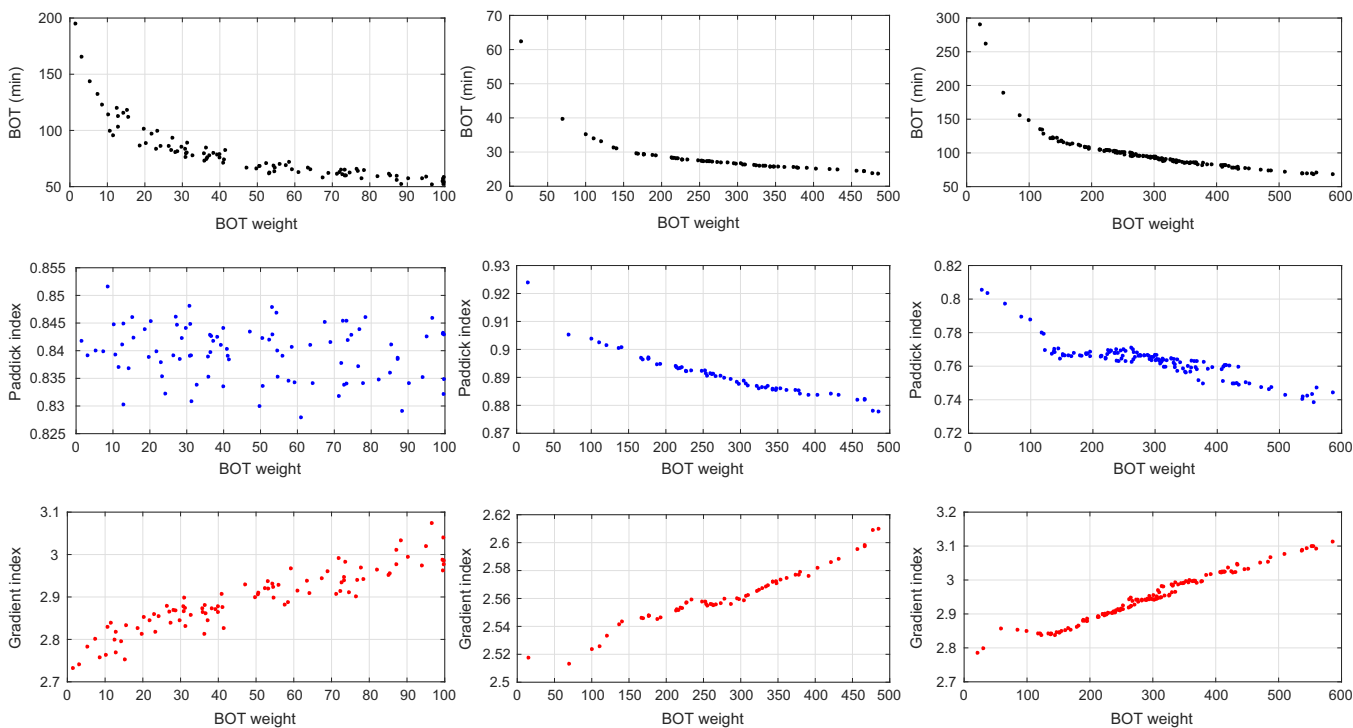


FIG. 3. Dependence of beam-on time (BOT) (top), Paddick index (middle), and gradient index (bottom) on the iBOT penalization. The columns correspond to different cases: small acoustic neuroma (left), medium acoustic neuroma (center), and irregular meningioma (right). [Color figure can be viewed at wileyonlinelibrary.com]

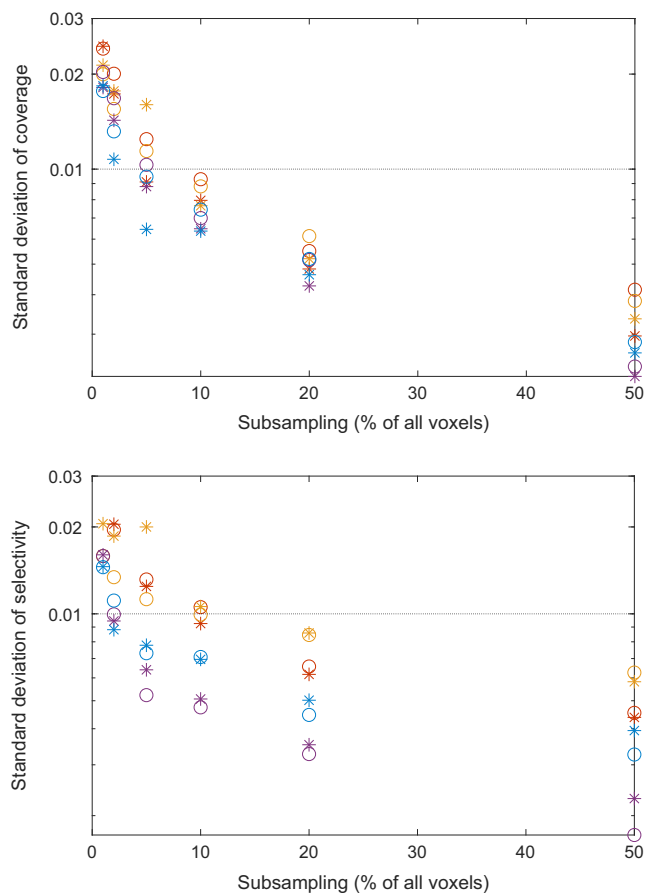


FIG. 4. The standard deviation of coverage (top) and selectivity (bottom) for the two sampling methods, interior and surface (*) and interior only (e) for the small acoustic neuroma (red), the medium acoustic neuroma (purple), the irregular meningioma (yellow), and the arteriovenous malformation (blue). [Color figure can be viewed at wileyonlinelibrary.com]

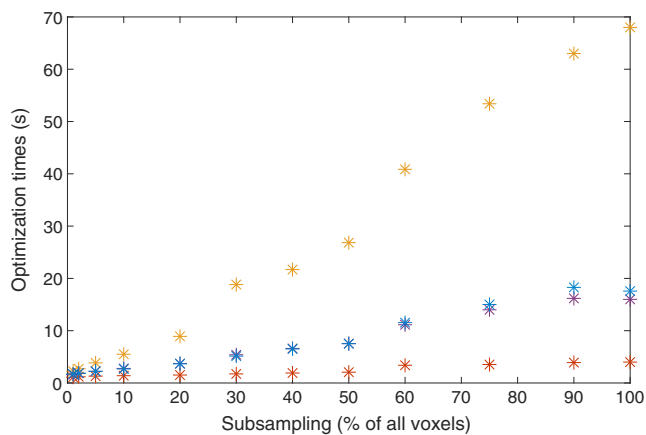


FIG. 5. The optimization time as a function of subsampling fraction for small acoustic neuroma (red), the medium acoustic neuroma (purple), the irregular meningioma (yellow), and the arteriovenous malformation (blue). [Color figure can be viewed at wileyonlinelibrary.com]

3.C. Overall performance

We evaluate the overall performance by comparing the plan metrics for optimized plans with manual forward

plans for 75 clinical cases (the majority of Gamma Knife users still do forward planning). The clinical cases are all single target cases: single metastases, acoustic neuromas, and meningiomas. The range in tumor size is 0.6–11.7 cm³ and the planning dose range is 12–24 Gy. In 18 of these cases, there are organs at risk (e.g., brainstem, cochlea, optic nerves) so, to make a fair comparison, the maximum doses to the OARs in the manual plans are used as hard constraints in the optimization. Furthermore, all the manual plans have 98%–100% coverage, so the weights were chosen to always satisfy this criteria. To compare the manual and optimized plans, we have to find weight settings for each optimized plan that give a similar trade-off between clinical objectives as the corresponding manual plan. For simplicity, we employed the following strategy: first, we created a range of plans with different characteristics by performing 101 optimizations for each case, with randomly sampled weights for the inner ring and BOT penalization; then, we selected, for each case, the optimized plan with the best gradient index among those with both higher selectivity and shorter BOT than the corresponding manual plan. The resulting plan metrics are shown in Fig. 6. In summary, for all of the 75 cases, we could find optimized plans with simultaneously higher selectivity and shorter BOT than the manual plans, and in 44 cases, these plans also had better gradient index than the manual plans. In other words, we found plans that dominate the manual ones in almost 60% of the cases.

To solve the optimization problems, we used the open-source solver Glop⁵⁴ with default settings. The optimization times ranged from 2.3 to 26 s with a median time of 5.7 s on a standard GammaPlan workstation (a HP Z640 with 32 GB RAM and 12 cores running at 2.9 MHz). Figure 7 shows how the optimization time depends on the number of nonzero elements in the constraint matrix B (cf. Appendix B). We conclude that there is a linear trend, but that the variation appears larger for small cases.

3.D. Detailed case study

As an example, we consider a large left-sided cavernous sinus meningioma, with a volume of 16.2 cm³ and with three adjacent OARs: the chiasm, the left optic nerve, and the left optic tract. The prescription dose to target is 15 Gy, the dose in the inner shell encompassing the target is penalized if the dose exceeds 15 Gy. In the outer shell, we penalize doses exceeding 7.5 Gy. The maximal allowed dose to all three OARs is 8 Gy, which is enforced by hard constraints. In Table II, we present two optimized plans, one promoting selectivity and one promoting short BOT, together with the clinical plan for comparison. Figure 8 shows snapshots and dose–volume histograms (DVHs) from Leksell GammaPlan[®] for the three different plans. In the plan where short BOT is promoted, the BOT is almost halved at the expense of a moderate decrease in

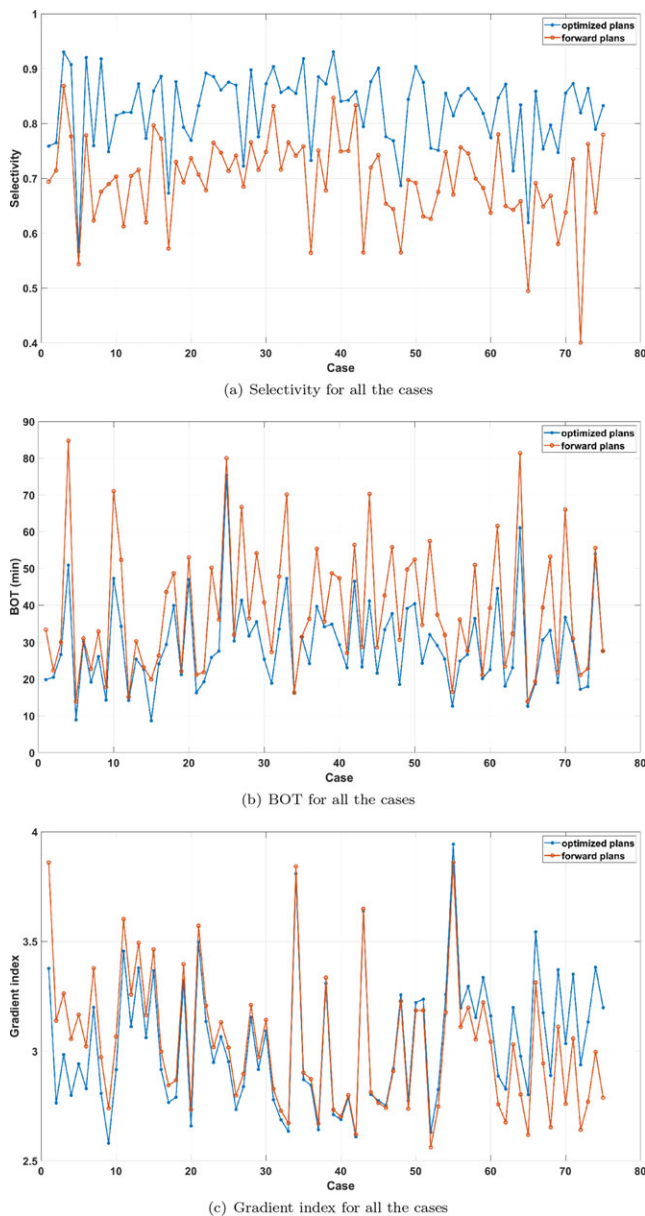


FIG. 6. Metrics for the optimized plans compared to the forward plans. The cases are sorted by the difference in gradient index. [Color figure can be viewed at wileyonlinelibrary.com]

selectivity compared to the second plan. However, we consider both plans clinically acceptable.

3.D.1. Dose painting and homogeneous plans

The flexibility of the proposed inverse planner opens up new possibilities for treatment planning. Here, we will exemplify both dose painting and how to create plans with a homogeneous dose distribution in the target.

Dose painting is achievable since both weights and prescription doses can be modified on a voxel-by-voxel basis. To illustrate this, we introduce a hotspot, with a volume of 1.2 cm³, in the center of the target. The dose prescribed to the hotspot is twice the dose prescribed to the rest of the

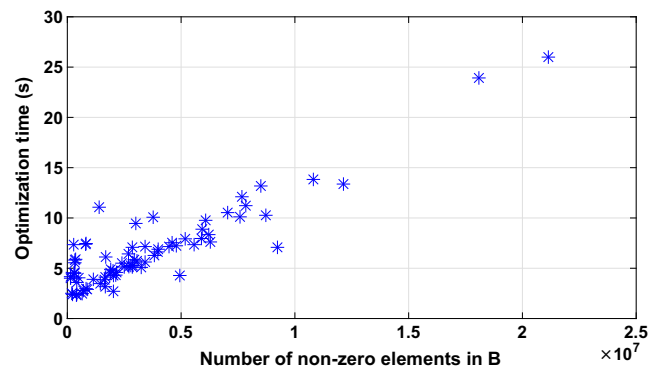


FIG. 7. Optimization time as a function of the number of nonzero elements in B for the 75 clinical cases. The times shown are averages over the 101 runs made for each case. [Color figure can be viewed at wileyonlinelibrary.com]

TABLE II. Plan metrics for the reference plan and two optimized plans.

	Reference plan	Promoting selectivity	Promoting BOT
Coverage	0.95	0.95	0.96
Selectivity	0.88	0.91	0.78
Gradient index	2.73	2.79	2.71
BOT (min) at 3 Gy/min	89	135	41
Max dose (left optic nerve)	8.95	8.0	7.2
Max dose (left optic tract)	7.3	8.0	7.2
Max dose (Chiasm)	7.3	8.0	7.2
Planning isodose	46	44	53

structure, that is, 30 Gy. In the present example, we get a hotspot coverage of 99%.

It turns out that the iBOT term will in many cases favor fairly homogeneous plans. However, such plans can be further promoted by penalizing overdosage of the target. We thus penalize doses exceeding 15/0.85 Gy to get a planning isodose of at least 85%, which is very difficult using forward planning. In Table III, we present an example of a plan for the hotspot case and one plan with a homogeneous dose distribution. Overall, we consider both plans acceptable, although naturally the additional requirements result in some form of trade-off. The hotspot case has similar plan quality but longer BOT than the clinically acceptable reference plan presented in Table II. The homogeneous plan trades off selectivity and we need to loosen the hard constraint on dose to OARs to achieve homogeneity. In Fig. 9, we present snapshots of one plan with a homogeneous dose distribution and one plan with a hotspot.

4. DISCUSSION

In this work, we have described a linear programming approach to Gamma Knife radiosurgery based on the division of the planning into three distinct phases: isocenter placement, optimization, and sequencing. Our main focus has, however, been the optimization phase. We have shown that

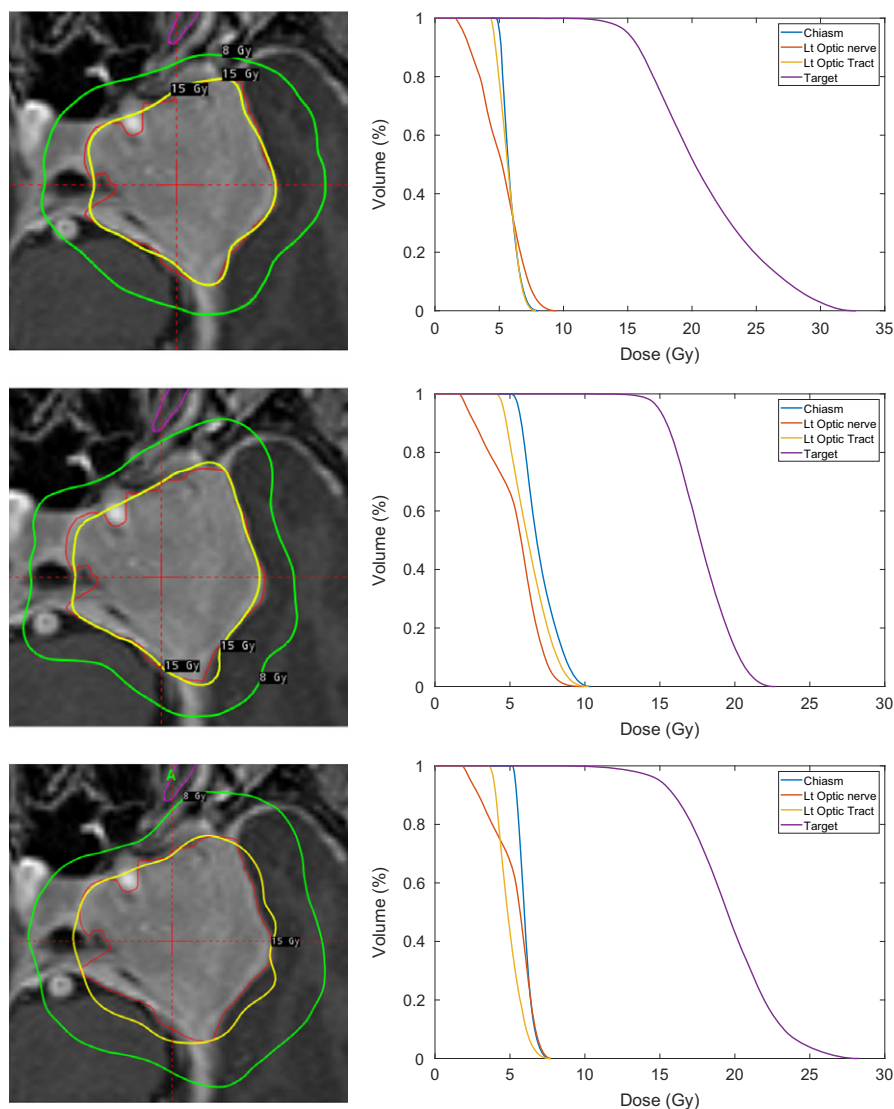


FIG. 8. Snapshots and dose–volume histograms (DVHs) of the reference plan (top), the plan promoting selectivity (middle), and the plan promoting beam-on time (bottom) for the meningioma with adjacent organs at risks. The DVH for the target is in red and the three OARs in magenta, light blue, and pink. [Color figure can be viewed at wileyonlinelibrary.com]

even without changing the isocenter positions, our optimization approach can find plans that dominate manual forward plans in almost 60% of the cases investigated. Reusing the isocenter positions in this way clarifies the improvement that results directly from the optimization, but it is — of course — not how our approach would eventually be deployed. A natural next step is thus to investigate methods for isocenter placement with an emphasis on the interplay with subsequent phases.

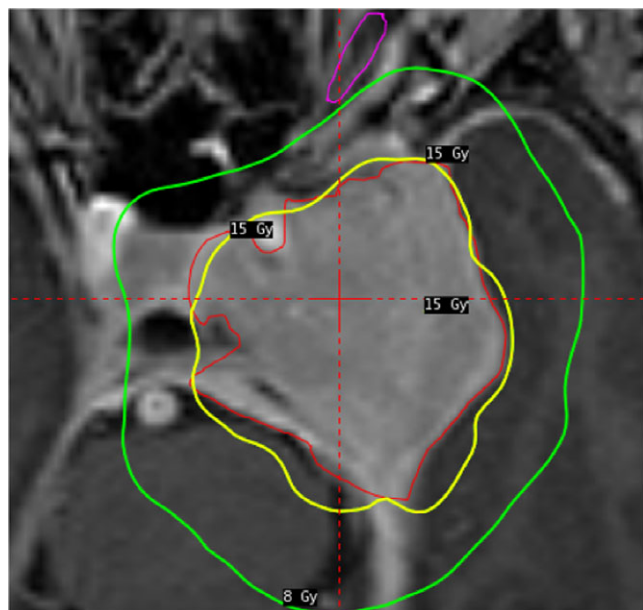
Our optimization problem, with a cost function that is a weighted sum of multiple competing objectives, is fundamentally a multicriteria optimization problem. A useful concept in multicriteria optimization is that of the Pareto surface. Simply put, the Pareto surface is the set of solutions in which you cannot improve one objective without impairing another. This means one should never be satisfied with a solution that is not on the Pareto surface (Pareto optimal) unless other considerations than explicitly expressed in the cost function is taken

into account. For example, our cost function only models gradient index and selectivity indirectly, which means that we cannot guarantee Pareto optimality with respect to these metrics. Our results show that manual forward plans are, in general, not Pareto optimal either and that the proposed inverse planner can often find plans that dominate them.

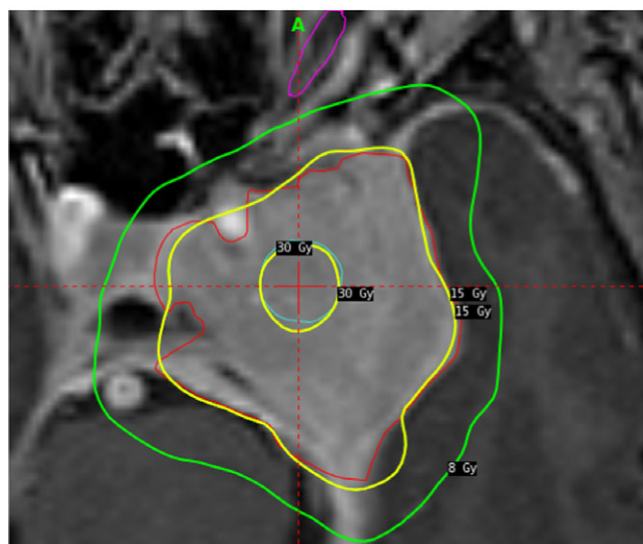
For convex multicriteria optimization problems, such as ours, every choice of weights corresponds to a Pareto optimal solution. By specifying the weights before optimization, we have treated this problem as a single-criteria optimization problem.⁵⁵ However, every Pareto optimal point corresponds to the solution of the optimization problem for a set weights on the unit simplex, that is, such that $w_i \geq 0$ and $\sum_{i=1}^n w_i = 1$. In other words, the formulation as a weighted sum does not restrict the set of optimal plans that we can obtain. In practice, our normalization is not ideal since the number of active points in each structure depends on a number of factors, including the volume and the sampling

TABLE III. One plan with a hotspot and one homogeneous plan.

	Hotspot	Homogeneous
Coverage	0.97	0.95
Selectivity	0.85	0.87
Gradient index	2.9	2.8
BOT (min) at 3 Gy/min	136	152
Max dose target (Gy)	38.5	17.7
Max dose left optic nerve (Gy)	8.0	8.0
Max dose left optic tract (Gy)	8.0	8.0
Max dose chiasm (Gy)	8.0	8.0
Planning isodose	40	85



(a) Homogeneous plan



(b) Plan with hotspot

FIG. 9. Snapshots from one plan with homogeneous dose distribution and one with a hotspot. [Color figure can be viewed at wileyonlinelibrary.com]

strategy. Consequently, the exact weight settings are unlikely to generalize from one case to another even though they may be quite similar. We thus expect that the planner will have to adjust the weights a few times before arriving at a satisfactory solution.

Although not explored here, we anticipate that further computational gains are attainable by cleverly manipulating the dose rate matrix $\Phi(r)$ and adapting the matrix algebra accordingly.^{56–58} We also expect that tuning or replacing the solver, or both, could shorten the optimization times further.⁵⁹

5. CONCLUSIONS

We present the first sector-duration optimization that uses linear programming. It uses a beam-on time (BOT) penalization tailored to the Gamma Knife, which reduces the BOT by a factor of 2–3 compared to the naïve alternative. In addition to using linear programming, we describe two techniques that reduce the problem size and thus further reduce the solution time: dualization and representative subsampling. Dualization leads to an equivalent problem that can be solved 5–20 times faster than the primal one. With representative subsampling, we refer to a stochastic sampling of positions, both in the interior and on the surface of relevant structures, to use in the optimization. We show that using 10% of the original number of voxels is enough to generate plans for which the statistical fluctuations in coverage and selectivity are below 1% while resulting in time-savings comparable to dualization.

The importance of different objectives, such as coverage, selectivity, or beam-on time (BOT), are controlled by adjusting the weights of the corresponding terms in the cost function. In a comparison with 75 clinical plans, we show that in 44 of these we can find plans that simultaneously have better selectivity, BOT, and gradient index (with coverage close to 100% in all cases) than the forward plans.

Treatment planning for Gamma Knife has always been highly interactive. Thanks to the combination of techniques to reduce the computational cost that we have presented, it becomes possible to fit sector-duration optimization — with all its benefits — into the clinical workflow. This is our main contribution.

ACKNOWLEDGMENTS

The authors thank Björn Somell and the anonymous reviewers for helpful comments and suggestions in the preparation of the manuscript. The Sharknado team — nuff said. The research was supported by the VINNOVA/ITEA3 projects BENEFIT (grant 2014-00593) and IMPACT (grant 2018-02230).

CONFLICTS OF INTEREST

All authors are employed at Elekta Instrument AB, which holds several patents related to this area. The work presented may become part of a future commercial product.

APPENDIX A: EXPLICIT FORMULATION OF THE PRIMAL PROBLEM

Any linear programming problem can be written in the form

$$\begin{aligned} &\text{minimize} && w^t x \\ &\text{subject to} && Ax = b \\ &&& x \geq 0, \end{aligned} \tag{A1}$$

where $x \in \mathbb{R}^n$ are the optimization variables and w, b , and $A \in \mathbb{R}^{p \times n}$ define the objective and constraint functions of the problem.

We introduce auxiliary variables to rewrite the hinge and iBOT functions using linear programming: if $y^+, y^- \geq 0$ and $D - \hat{D} = (y^+ - y^-)$, then $y^+ \geq (D - \hat{D})_+$ and $y^- \geq (\hat{D} - D)_+$; if $\tau_i \geq \sum_{c=1}^3 t_{isc}$ for $s = 1, \dots, 8$, then $\tau_i \geq \max \sum_{c=1}^3 t_{isc}$. We may then express the primal optimization problem (4) in standard form by making the following identifications:

$$A = \begin{pmatrix} \Phi_T & -I & I & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \Phi_S & 0 & 0 & -I & I & 0 & 0 & 0 & 0 & 0 \\ \Phi_G & 0 & 0 & 0 & 0 & -I & I & 0 & 0 & 0 \\ \Phi_O & 0 & 0 & 0 & 0 & 0 & 0 & I & 0 & 0 \\ C & 0 & 0 & 0 & 0 & 0 & 0 & 0 & I & -T \end{pmatrix},$$

$$x^t = (t, y_T^+, y_T^-, y_S^+, y_S^-, y_G^+, y_G^-, p, q, \tau),$$

$$w^t = \left(0, 0, \frac{w_T}{N_T D_T}, \frac{w_S}{N_S D_S}, 0, \frac{w_G}{N_G D_G}, 0, 0, 0, \frac{w_{BOT}}{D_T / \varphi_{cal}} \right),$$

$$b^t = (D_T, D_S, D_G, D_O, 0), \tag{A2}$$

where

$$C = \begin{pmatrix} 1 & 1 & 1 & 0 & 0 & 0 & \dots & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & \dots & 1 & 1 & 1 \end{pmatrix}$$

$$= I_{N_{iso}} \otimes I_8 \otimes (1, 1, 1) \in \mathbb{R}^{8N_{iso} \times 24N_{iso}}, \tag{A3}$$

$$T = I_{N_{iso}} \otimes 1_{8 \times 1} \in \mathbb{R}^{8N_{iso} \times N_{iso}}, \tag{A4}$$

and \otimes denotes the Kronecker product.

APPENDIX B: EXPLICIT FORMULATION OF THE DUAL PROBLEM

Most modern linear programming solvers begin with a presolve step that intends to reduce the problem size.⁶⁰ However, for completeness and because we have not encountered a solver that automatically detects the benefit of dualizing our problem, we will here carry out the dualization explicitly. The

dual problem corresponding to a linear programming problem on standard form, Equation (A1) is^{16,39,40}

$$\begin{aligned} &\text{minimize} && b^t v \\ &\text{subject to} && A^t v + w \geq 0, \end{aligned} \tag{B1}$$

where the dual variable v is the Lagrange multiplier for the linear constraints of the primal problem. By eliminating redundant variables and grouping constraints, our problem can be stated as

$$\begin{aligned} &\text{minimize} && \rho^t \tilde{v} \\ &\text{subject to} && B^t \tilde{v} \leq \sigma, \\ &&& \ell \leq \tilde{v} \leq \mu, \end{aligned} \tag{B2}$$

where we have introduced the following rescaled entities

$$\tilde{v} = \left(\frac{v_T}{D_T}, \frac{v_S}{D_S}, \frac{v_G}{D_G}, \frac{v_O}{D_O}, \frac{v_{iso}}{D_T / \varphi_{cal}} \right),$$

$$\varphi_T = \frac{\phi_T}{D_T}, \varphi_S = \frac{\phi_S}{D_S}, \varphi_G = \frac{\phi_G}{D_G}, \varphi_O = \frac{\phi_O}{D_O}, C_s = \frac{C}{D_T / \varphi_{cal}}. \tag{B3}$$

We may then express the dual optimization problem corresponding to Eq. (A2) by making the following identifications

$$\begin{aligned} \rho^t &= (-1, 1, 1, 1, 0), \\ B^t &= \begin{pmatrix} \varphi_T^t & -\varphi_S^t & -\varphi_G^t & -\varphi_O^t & -C_s^t \\ 0 & 0 & 0 & 0 & T^t \end{pmatrix}, \\ \sigma^t &= (0, w_{BOT}), \\ \ell &= (0, 0, 0, 0, 0), \\ \mu &= \left(\frac{w_T}{N_T}, \frac{w_S}{N_S}, \frac{w_G}{N_G}, \infty, \infty \right). \end{aligned} \tag{B4}$$

The number of nonzero elements in B is $(24(N_T + N_S + N_G + N_O) + 32)N_{iso}$.

^{a)}Author to whom correspondence should be addressed. Electronic mail: jens.sjолund@elekta.com.

REFERENCES

1. Barnett GH, Linskey ME, Adler JR, et al. Stereotactic radiosurgery—an organized neurosurgery-sanctioned definition. *J Neurosurg.* 2007;106: 1–5.
2. Kondziolka D, Lunsford LD, Coffey RJ, Flickinger JC. Stereotactic radiosurgery of meningiomas. *J Neurosurg.* 1991;74:552–559.
3. Kondziolka D, Patel A, Lunsford LD, Kassam A, Flickinger JC. Stereotactic radiosurgery plus whole brain radiotherapy versus radiotherapy alone for patients with multiple brain metastases. *Int J Radiat Oncol Biol Phys.* 1999;45:427–434.
4. Lunsford LD, Niranjan A, Flickinger JC, Maitz A, Kondziolka D. Radiosurgery of vestibular schwannomas: summary of experience in 829 cases. *J Neurosurg.* 2005;102:195–199.
5. Sheehan JP, Niranjan A, Sheehan JM, et al. Stereotactic radiosurgery for pituitary adenomas: an intermediate review of its safety, efficacy, and role in the neurosurgical treatment armamentarium. *J Neurosurg.* 2005;102:678–691.

6. Kondziolka D, Flickinger JC, Niranjan A, Lunsford LD. Trends and importance of radiosurgery for the development of functional neurosurgery. *Surg Neurol Int.* 2012;3:53.
7. Lunsford LD, Kondziolka D, Flickinger JC, et al. Stereotactic radiosurgery for arteriovenous malformations of the brain. *J Neurosurg.* 1991;75:512–524.
8. Maesawa S, Salame C, Flickinger JC, Pirris S, Kondziolka D, Lunsford LD. Clinical outcomes after stereotactic radiosurgery for idiopathic trigeminal neuralgia. *J Neurosurg.* 2001;94:14–20.
9. Régis J, Metellus P, Hayashi M, Roussel P, Donnet A, Bille-Turc F. Prospective controlled trial of gamma knife surgery for essential trigeminal neuralgia. *J Neurosurg.* 2006;104:913–924.
10. Ferris MC, Shepard DM. Optimization of gamma knife radiosurgery. *Discrete Math Probl Med Appl.* 2000;55:27–44.
11. Ferris MC, Lim J, Shepard DM. An optimization approach for radiosurgery treatment planning. *SIAM J Optim.* 2002;13:921–937.
12. Ferris MC, Lim J, Shepard DM. Radiosurgery treatment planning via nonlinear programming. *Ann Oper Res.* 2003;119:247–260.
13. Ghaffari HR, Aleman DM, Jaffray DA, Ruschin M. A tractable mixed-integer model to design stereotactic radiosurgery treatments. Technical Report, Department of Mechanical and Industrial Engineering, University of Toronto; 2012.
14. Lee EK, Fox T, Crocker I. Optimization of radiosurgery treatment planning via mixed integer programming. *Med Phys.* 2000;27:995–1004.
15. Vandewouw MM, Aleman DM, Jaffray DA. Robotic path-finding in inverse treatment planning for stereotactic radiosurgery with continuous dose delivery. *Med Phys.* 2016;43:4545–4557.
16. Boyd S, Vandenberghe L. *Convex Optimization.* Cambridge, UK: Cambridge University Press; 2004.
17. Ghobadi K, Ghaffari HR, Aleman DM, Jaffray DA, Ruschin M. Automated treatment planning for a dedicated multi-source intracranial radiosurgery treatment unit using projected gradient and grassfire algorithms. *Med Phys.* 2012;39:3134–3141.
18. Bortfeld T. Optimized planning using physical objectives and constraints. In: *Seminars in Radiation Oncology.* Vol. 9. Amsterdam, Netherlands: Elsevier; 1999:20–34.
19. Romeijn HE, Dempsey JF. Intensity modulated radiation therapy treatment plan optimization. *Top.* 2008;16:215.
20. Ghaffari HR. *Optimization models and techniques for radiation treatment planning applied to Leksell Gamma Knife® Perfixion™.* PhD thesis, University of Toronto; 2012.
21. Ghaffari HR, Aleman DM, Jaffray DA, Ruschin M. Incorporation of delivery times in stereotactic radiosurgery treatment optimization. *J Global Optim.* 2017;69:1–13.
22. Ghobadi K, Ghaffari HR, Aleman DM, Jaffray DA, Ruschin M. Automated treatment planning for a dedicated multi-source intra-cranial radiosurgery treatment unit accounting for overlapping structures and dose homogeneity. *Med Phys.* 2013;40:091715.
23. Ehr Gott M, Güler Ç, Hamacher HW, Shao L. *Mathematical optimization in intensity modulated radiation therapy.* 4OR. 2008;6:199–262.
24. Holder A, Salter B. A tutorial on radiation oncology and optimization. In: *Tutorials on Emerging Methodologies and Applications in Operations Research.* Berlin: Springer; 2005:4–1.
25. Reemtsen R, Alber M. Continuous optimization of beamlet intensities for intensity modulated photon and proton radiotherapy. In: *Handbook of Optimization in Medicine.* Berlin: Springer; 2009:1–40.
26. Rosen II, Lane RG, Morrill SM, Belli JA. Treatment plan optimization using linear programming. *Med Phys.* 1991;18:141–152.
27. Shepard DM, Ferris MC, Olivera GH, Mackie TR. Optimizing the delivery of radiation therapy to cancer patients. *Siam Rev.* 1999;41:721–744.
28. Bahr GK, Kereiakes JG, Horwitz H, Finney R, Galvin J, Goode K. The method of linear programming applied to radiation treatment planning. *Radiology.* 1968;91:686–693.
29. Lodwick WA, McCourt S, Newman F, Humphries S. Optimization methods for radiation therapy plans. In: *Computational Radiology and Imaging.* Berlin: Springer; 1999:229–249.
30. Holder A. Designing radiotherapy plans with elastic constraints and interior point methods. *Health Care Manag Sci.* 2003;6:5–16.
31. Wu QJ, Chankong V, Jitrapakulsarn S, et al. Real-time inverse planning for Gamma Knife radiosurgery. *Med Phys.* 2003;30:2988–2995.
32. Schweikard A, Tombropoulos R, Kavraki L, Adler JR, Latombe J-C. Treatment planning for a radiosurgical system with general kinematics. In: *IEEE International Conference on Robotics and Automation, 1994. Proceedings, 1994:* 1720–1727. IEEE; 1994.
33. Schweikard A, Bodduluri M, Adler JR. Planning for camera-guided robotic radiosurgery. *IEEE Trans Robot Autom.* 1998;14:951–962.
34. Kilby W, Dooley JR, Kuduvalli G, Sayeh S, Maurer Jr CR. The CyberKnife® robotic radiosurgery system in 2010. *Technol Cancer Res Treatm.* 2010;9:433–452.
35. Wagner TH, Yi T, Meeks SL, et al. A geometrically based method for automated radiosurgery planning. *Int J Radiat Oncol Biol Phys.* 2000;48:1599–1611.
36. St John TJ, Wagner TH, Bova FJ, Friedman WA, Meeks SL. A geometrically based method of step and shoot stereotactic radiosurgery with a miniature multileaf collimator. *Phys Med Biol.* 2005;50:3263.
37. Wu QJ, Bourland JD. A Study and automatic solution for multishot treatment planning for the γ -knife. *J Radiosurg.* 2000;3:77–84.
38. Doudareva E, Ghobadi K, Aleman DM, Ruschin M, Jaffray DA. Skeletonization for isocentre selection in Gamma knife® perfixion™. *Top.* 2015;23:369–385.
39. Bertsimas D, Tsitsiklis JN. *Introduction to Linear Optimization.* Vol. 6. Belmont, MA: Athena Scientific; 1997.
40. Nocedal J, Wright SJ. *Numerical Optimization.* New York, NY: Springer Science+Business Media; 2006.
41. Romeijn HE, Ahuja RK, Dempsey JF, Kumar A, Li JG. A novel linear programming approach to fluence map optimization for intensity modulated radiation therapy treatment planning. *Phys Med Biol.* 2003;48:3521.
42. Alber M, Paulsen F, Eschmann SM, Machulla HJ. On biologically conformal boost dose optimization. *Phys Med Biol.* 2003;48:N31.
43. Bentzen SM. Theragnostic imaging for radiation oncology: dose-painting by numbers. *Lancet Oncol.* 2005;6:112–117.
44. Ling CC, Humm J, Larson S, et al. Towards multidimensional radiotherapy (MD-CRT): biological imaging and biological conformality. *Int J Radiat Oncol Biol Phys.* 2000;47:551–560.
45. Svensson J. Multiobjective optimization in radiosurgery: How to approximate and navigate on the Pareto surface; 2014. urn:nbn:se:kth:diva-145133.
46. The International Commission on Radiation Units and Measurements. *J ICRU.* 2010;10. <https://doi.org/10.1093/jicru/ndq001>
47. Scherrer A, Küfer K-H, Bortfeld T, Monz M, Alonso F. IMRT planning on adaptive volume structures—a decisive reduction in computational complexity. *Phys Med Biol.* 2005;50:2033.
48. Ungun B, Xing L, Boyd S. Real-Time Radiation Treatment Planning with Optimality Guarantees via Cluster and Bound Methods Working paper.
49. Martin BC, Bortfeld TR, Castanon DA. Accelerating IMRT optimization by voxel sampling. *Phys Med Biol.* 2007;52:7211.
50. Wang M, Bertsekas DP. Stochastic first-order methods with random constraint projection. *SIAM J Optim.* 2016;26:681–717.
51. Nordström H, Johansson J. Sequencing sector fields 2015. US Patent 8,938,047.
52. Torrens M, Chung C, Chung H-T, et al. Standardization of terminology in stereotactic radiosurgery: report from the Standardization Committee of the International Leksell Gamma Knife Society. *J Neurosurg.* 2014;121:2–15.
53. Paddick I. A simple scoring ratio to index the conformity of radiosurgical treatment plans. *J Neurosurg.* 2000;93:219–222. <https://developers.google.com/optimization/lp/glop>.
54. Ehr Gott M. *Multicriteria Optimization.* Vol. 491. Berlin, Germany: Springer Science & Business Media; 2005.
55. Cho PS, Phillips MH. Reduction of computational dimensionality in inverse radiotherapy planning using sparse matrix operations. *Phys Med Biol.* 2001;46:N117.
56. Thieke C, Nill S, Oelfke U, Bortfeld T. Acceleration of intensity-modulated radiotherapy dose calculation by importance sampling of the calculation matrices. *Med Phys.* 2002;29:676–681.
57. Zakarian C, Deasy JO. Beamlet dose distribution compression and reconstruction using wavelets for intensity modulated treatment planning. *Medical physics* 2004;31:368–375.
58. Ólafsson A, Wright SJ. Linear programming formulations and algorithms for radiotherapy treatment planning. *Optim Methods Softw.* 2006;21:201–231.
59. Andersen ED, Andersen KD. Presolving in linear programming. *Math Programm.* 1995;71:221–245.