CARABEAMER: A Treatment Planner for a Robotic Radiosurgical System with General Kinematics

Rhea Z. Tombropoulos* John R. Adler† Jean-Claude Latombe‡

Stanford University, Stanford, CA 94305, U.S.A.

Abstract: Stereotactic radiosurgery is a minimally invasive procedure that uses a focused beam of radiation as an ablative instrument to destroy brain tumors. To deposit a high dose of radiation in a tumor, while reducing the dose to healthy tissue, a large number of beams are crossfired at the tumor from multiple directions. The treatment planning problem (also called the inverse dosimetry problem) is to compute a set of beams that produces a desired dose distribution. So far, its investigation has focused on the generation of isocenter-based treatments in which the beam axes intersect at a common point, the isocenter. However, this restriction limits the applicability of the treatments to tumors having simple shapes. This paper describes CARABEAMER, a new treatment planner for a radiosurgical system in which the radiation source can be arbitrarily positioned and oriented by a six-degree-of-freedom manipulator. This planner uses randomized techniques to guess a promising initial set of beams. It then applies space partitioning and linear programming techniques to compute the energy to be delivered along each beam. Finally, it exploits the results of the linear program to iteratively adapt and improve the beam set. Experimental results obtained with CARABEAMER on both patient and synthetic cases are presented and discussed in this paper. These results demonstrate that a radiosurgical system with general kinematics can deliver treatments in which the region receiving high dose closely matches the shape of the tumor, even in complicated cases. They also suggest new research directions which are discussed at the end of the paper.

1 Introduction

Stereotactic radiosurgery is a minimally invasive procedure that uses a focused beam of radiation as an ablative instrument to destroy brain tumors [42]. The beam is produced by a linear accelerator that is moved by a mechanical gantry. To deposit a high dose of radiation in a tumor, while reducing the dose in healthy tissues, a large number of beams are crossfired at the tumor from multiple directions (see Figure 1). The treatment planning problem considered in this paper (also called the inverse dosimetry problem) is the following: Compute a set of beams that will result in a given desired dose distribution in the tumor and surrounding tissues.

This problem has already received considerable attention in the context of conventional radiosurgical systems (for instance, see [6, 13, 15]). These systems, such as the LINAC [19] and the Gamma

*Current address: Silicon Graphics Inc., 2011 N. Shoreline Blvd., ms. 984, Mountain View, CA 94043-1389. E-mail: rzt@engr.sgi.com

†Neurosurgery Department. E-mail: adler@robotics.stanford.edu

‡Department of Computer Science. E-mail: latombe@cs.stanford.edu

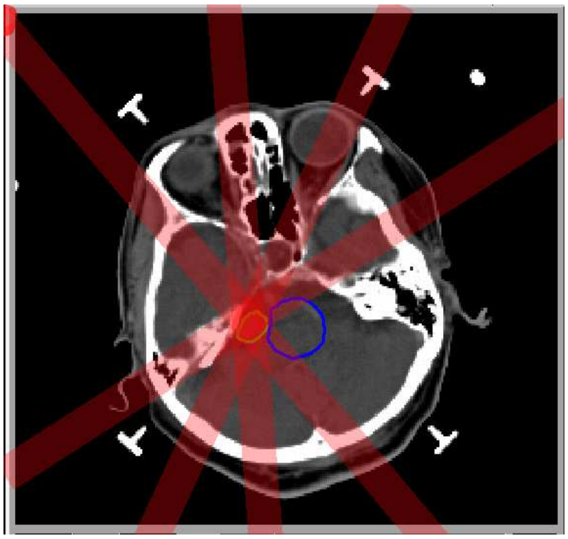

Figure 1: Crossfiring beams at a tumor

Knife [11], require that the beam axes intersect at a common point, called the *isocenter* (as is the case in Figure 1). The resulting isocenter-based treatments are limited to tumors having relatively simple shapes. In the ideal case where the tumor is roughly spherical, the isocenter is chosen at the centerpoint of the tumor, and the collimator diameter is set to the radius of the tumor. If a tumor is non-spherical, it may be approximated by multiple spheres [11]. However, switching between isocenters requires moving the patient relative to the gantry, which increases the total duration of the treatment. Typically, no more than five spheres are used to approximate the tumor shape; thus, the resulting treatment often yields a dose distribution with undesirable features. For instance, the spheres may not fully cover the tumor, so that some tumor areas will receive insufficient dose. Moreover, some spheres may overlap healthy tissue, which may create excessively high dose outside the tumor. Finally, some spheres may intersect in the tumor, resulting in a non-homogeneous dose distribution inside the tumor.

In reality, tumors vary widely in size and shape. Moreover, they can be located close to radiation-sensitive and/or critical structures (e.g., brain stem, optic nerves) which should receive particularly low doses. An important goal for radiosurgery is thus to generate conformal treatments, in which the region of high dose closely matches the 3-D shape of the tumor. To achieve that goal better than with conventional systems, a new radiosurgical system, called the CyberKnife, has been developed at the Stanford Medical Center by Accuray, Inc.[1]. In this system, the linear accelerator is moved by a general six-degree-of-freedom manipulator arm. This system, depicted in Figure 2, imposes almost no constraints on the beam configurations. In particular, the beam axes do not have to intersect at an isocenter. But, as a consequence, treatment planning for the CyberKnife is much more complex than for conventional systems, since it involves searching a much larger space of possible beam configurations, and subsequently of treatment plans.

This paper describes a planner, CARABEAMER, which we have developed to exploit effectively the kinematic generality of the CyberKnife system. This planner is not specific to the CyberKnife, however; it only assumes that the radiation source is moved by a gantry having general kinematics. The planning algorithm first narrows down the search space by selecting a set of promising beam configurations, and then it iteratively refines this initial set to produce a satisfactory solution. One input to CARABEAMER is a 3-D geometric model of the tissues of interest; these tissues are interactively delineated on computer

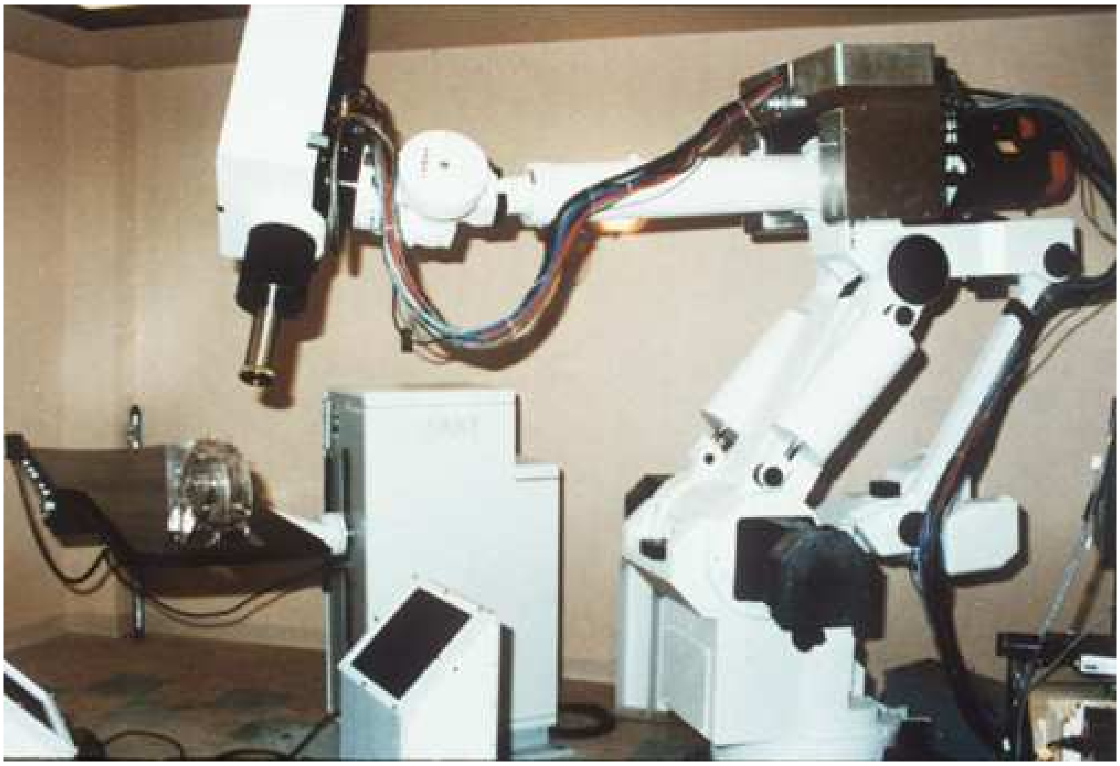

Figure 2: The CyberKnife system

tomography (CT) or magnetic resonance (MR) scans of the patient's head. Another input is a set of inequality constraints that bound the dose received by each tissue of interest. Additional inputs include the maximum number of beams allowed (N_{max}) and the diameter of the collimator. The maximum number of beams, usually on the order of two to four hundreds, roughly determines the duration of the treatment. If the planner can solve the problem, it outputs a set of weighted beam configurations; the *weight* of each configuration represents the energy emitted by the radiation source for this configuration. If no solution is found, the planner outputs a set of weighted configurations that results in a dose distribution that minimizes the violation of the input constraints. CARABEAMER can also be requested to select collimator diameters automatically, in order to deal with the case where the radiation source is equipped with an adaptable circular collimator, and to minimize the number of beams used to achieve the input dose constraints, in order to shorten the treatment.

The paper is organized as follows. Section 2 presents the general background of our research and gives an overview of CARABEAMER. It also relates our planner to previous work in treatment planning and defines terms used throughout the paper. Section 3 describes how regions of interest are defined and represented, and how dose constraints are specified. Section 4 presents in detail the core version of the planner, which we call the *basic planner*. Next, Section 5 discusses several extensions of the basic planner that have all been implemented in CARABEAMER. In Section 6 we discuss how CARABEAMER can be evaluated, and we give experimental results obtained with patient cases, as well as with synthetic cases. We used the patient cases to show that CARABEAMER can meet the requirements input by a radiation therapist, and the synthetic cases to explore the limits of CARABEAMER (and, simultaneously, of the CyberKnife). Section 7 concludes the paper and suggests areas for future research.

The planning techniques implemented in CARABEAMER are adaptations and extensions of techniques that we have previously presented in [2, 3, 29, 30, 31, 35, 36]. An extensive description of CARABEAMER and the experimental results obtained with it can be found in [34].

Throughout this paper, we will refer to the physician responsible for developing a treatement plan (hence, to the user of CARABEAMER) as the *neurosurgeon*. However, the neurosurgeon works closely with a radiation therapist throughout the treatment process. For regulatory reasons, the radiation therapist must approve all the treatment plans prior to their execution.

2 Background and Overview

The CyberKnife system consists of a six-degree-of-freedom manipulator with a 6-MeV X-ray linear accelerator (or linac) attached to its tip. The general kinematics of the manipulator make it possible, within a certain volume reachable by the tip of the manipulator, to achieve any orientation of the linac.

When activated, the linac produces a beam of radiation. A *beam configuration* refers to the volume in 3-D space that is traversed by the radiation beam for a given position and orientation of the linac. Hence, in the usual case where the linac is equipped with a circular collimator, one convenient way to parameterize a beam configuration is to use six parameters: three are the coordinates of the focal point of the linac optic in some world coordinate frame; two are angles defining the orientation of the radiation source relative to this frame's axes; the sixth one is the radius of the collimator. In the rest of the paper, we will simply use the word "beam" to designate a beam configuration. Hence, when we talk about multiple beams, we mean multiple beam configurations produced by the same radiation source at different positions and orientations.

When a point p in some tissue is traversed by a beam, the tissue at p's neighborhood absorbs some radiation energy, which is called the *dose* deposited at p. This dose depends on many parameters including the power of the linac, the time during which the beam is activated, p's distance from the source, the amount of tissue already traversed, p's distance from the central beam axis, and the absorption properties of the tissue at p. It is measured in *rads*, with one rad corresponding to the absorption of 100 ergs per gram of tissue. The doses deposited by several beams at the same point p are additive. This means that if the same point p is traversed by multiple beams, the total dose deposited at p is the sum of the doses delivered by each beam. This property is central to radiosurgery since a beam must travel through healthy tissue in order to access to the tumor and exit from it. It makes it possible to deliver a necrotic dose to the tumor by crossfiring many beams at it, while distributing the dose absorbed by healthy tissue over a greater volume (Figure 1).

A treatment plan is defined as a set of weighted beams. Given a 3-D geometric model of a tissue area and a treatment plan, a *direct dosimetry* program must be available to simulate the treatment and estimate the dose that will be delivered throughout the brain [1]. This program is the reference used by the neurosurgeon to assess treatment plans. It incorporates precise models of the shape of a beam, of the intensity of a beam as it travels from the source, and of the absorption properties of the tissue. These models are based on numerous data collected by radiation physicists using various experimental techniques (e.g., photographic phantoms). Several methods are used to visualize a dose distribution. One is to display $x\%$ isodose surfaces for several values of x (e.g., $x = 50, 80, 90$); the $x\%$ isodose surface bounds the set of points that receive at least $x\%$ of the maximum dose. Another way is to plot dose-volume histograms.

The *treatment planning* problem is to compute a set of beam configurations that will result in a dose distribution that satisfies input constraints. Typically, the constraints require that the dose anywhere in the tumor be greater than some given values and that the dose in critical/sensitive regions be lower than some other specified values. They may also specify the maximum variation of the dose within the tumor (i.e., dose homogeneity) and the dose falloff around the tumor. The approach to treatment planning that prevails today is to guess an initial plan, and then to iteratively run the direct dosimetry program and modify the current plan to deal with constraint violations. Such a trial-and-error approach is possible with conventional radiosurgical systems, because their restricted kinematics limits the size of the solution space and neurosurgeons have developed enough intuition to produce "good" initial plans in that case. But it is not applicable to the CyberKnife system, for which the solution space is much too large.

The majority of research in automatic radiosurgical planning has focused on isocenter-based treatments, in which the tumor region is first approximated by a collection of spheres. Although the inverse dosimetry problem has been mathematically solved in 2-D for a single isocenter [6, 7, 10, 16], the problem becomes substantially more complex in 3-D [13], especially in the absence of major symmetries. In that case, most proposed systems require the user to specify a set of beams or, in some instances, the locations of a few isocenters from which beams can be derived. They then use optimization techniques, such as linear programming [4, 17, 24, 27], quadratic programming [21, 22, 33], or simulated annealing [8, 23, 41], to determine the radiation energy to send along each beam to approximate the desired dose distribution as closely as possible. A geometric approach presented in [28] sees the critical structures as if they were obstacles obstructing access to and exit from the tumor. As noted in the introduction, isocenter-based planning methods are well adapted to the kinematics of conventional radiosurgical systems, but they have serious drawbacks. They typically do not allow the satisfactory treatment of complex-shaped tumors located close to critical structures. Multi-leaf collimators only partially eliminate this drawback.

In the context of the CyberKnife system, computing a set of beams directly from the geometric model of the tissue and the dose constraints is extremely difficult, if not impossible. Therefore our planner, CARABEAMER, internally uses an iterative trial-and-error approach. But it relies on several algorithmic components to make it practically efficient: First, it selects a promising initial beam set. Second, it incorporates a simplified direct dosimetry model to quickly evaluate this set and compute the weight of each beam. Third, it uses the results of this evaluation to improve the beam set. All these components and extensions will be described in detail in Sections 4 and 5. Note that CARABEAMER computes a treatment plan that achieves the input dose constraints according to the dosimetry model used in the planner. Thus, this plan still needs to be evaluated by the neurosurgeon using the reference direct dosimetry program. If it turns out not entirely satisfactory, the neurosurgeon must tune accordingly the description of the tissue regions and/or the dose constraints associated with them. Such modification could be done automatically by the planner itself; however, this is not the case in the current implementation of CARABEAMER.

CARABEAMER also attempts to minimize treatment duration. The energy delivered along a beam is proportional to the time during which the beam is activated. However, because the linac is a relatively heavy instrument (approximately 300 lbs), the motion of the manipulator is relatively slow and a rest of a few seconds at each position/orientation is needed before activating the beam to ensure that the linac is perfectly immobile. Hence, the number of beams roughly determines treatment duration. If a solution has been found with the maximum number of beams N_{max} allowed by the neurosurgeon, CARABEAMER tries to minimize this beam set by eliminating low-weight beams and adjusting the weights of the remaining beams. CARABEAMER can also be used to select automatically collimator diameters when the linac is equipped with a collimator of adaptable diameter or when it is possible to switch collimator. So far, the implemented planner accepts only circular collimators, although preliminary experiments with collimators of adaptable shapes have been presented in [30].

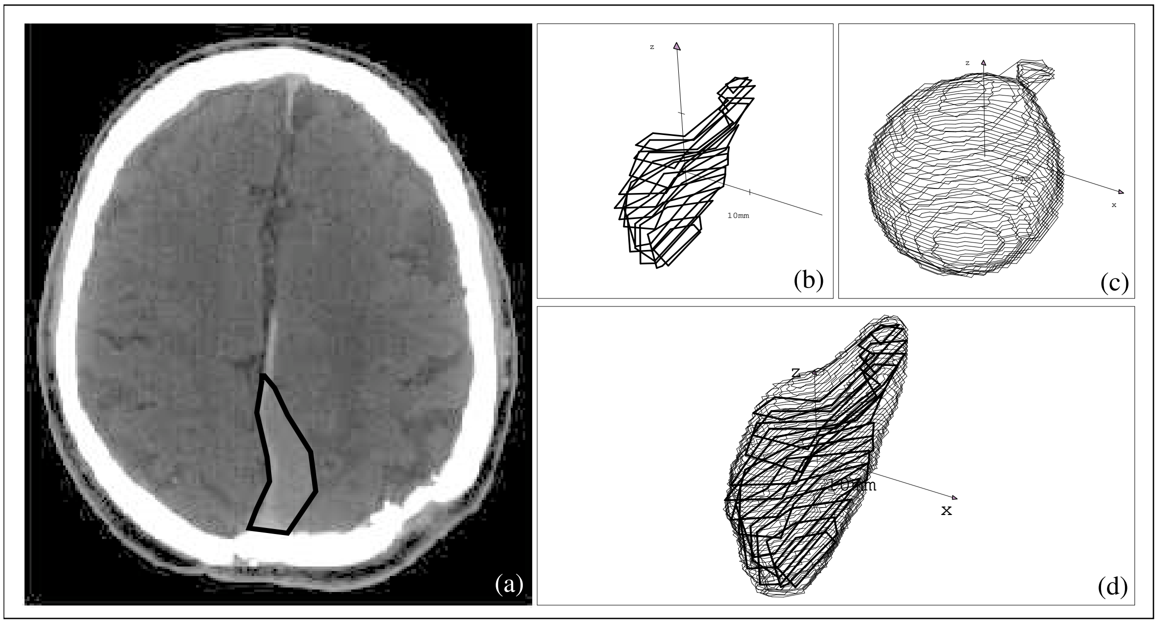

Figure 3: Comparison of an isocenter-based treatment with a CARABEAMER treatment

Figure 3 briefly illustrates the capabilities of CARABEAMER on one example. Figure 3(a) displays an elongated tumor in a CT scan. Figure 3(b) depicts the 3-D reconstruction of the tumor's geometry generated from multiple CT scans. Figure 3(c) shows the 80% isodose surface obtained with an isocenter-based treatment in which the tumor was approximated by two spheres of different sizes. Figure 3(d) shows the 80% isodose surface with the treatment planned by CARABEAMER. The surface of Figure 3(d) matches the tumor volume much more closely than that of Figure 3(c).

3 Tissue Representation and Dose Constraints

The neurosurgeon inputs the shapes and locations of *regions of interest* (e.g., the tumor and the critical structures). For each of those regions he also specifies *dose constraints*.

3.1 Tissue Representation

The neurosurgeon interactively defines the regions of interest by delineating them on a graphical display of the successive CT or MR scans of the patient's head. Each region is thus input as a stack of polygonal contours defining the region's shape and location in 3-D space. Figure 4 illustrates this construction. Figure 4(a) displays one scan in which the tumor and the brain stem have been delineated. Figures 4(b) and 4(c) show the stacks of contours describing those two regions (side and top views, respectively).

It is entirely up to the neurosurgeon to decide what tissue subsets are regions of interest; this choice is important since it will only be possible to assign dose constraints to these regions (see Subsection 3.2). The regions of interest usually coincide with the tumor and the critical structures, but this is not a

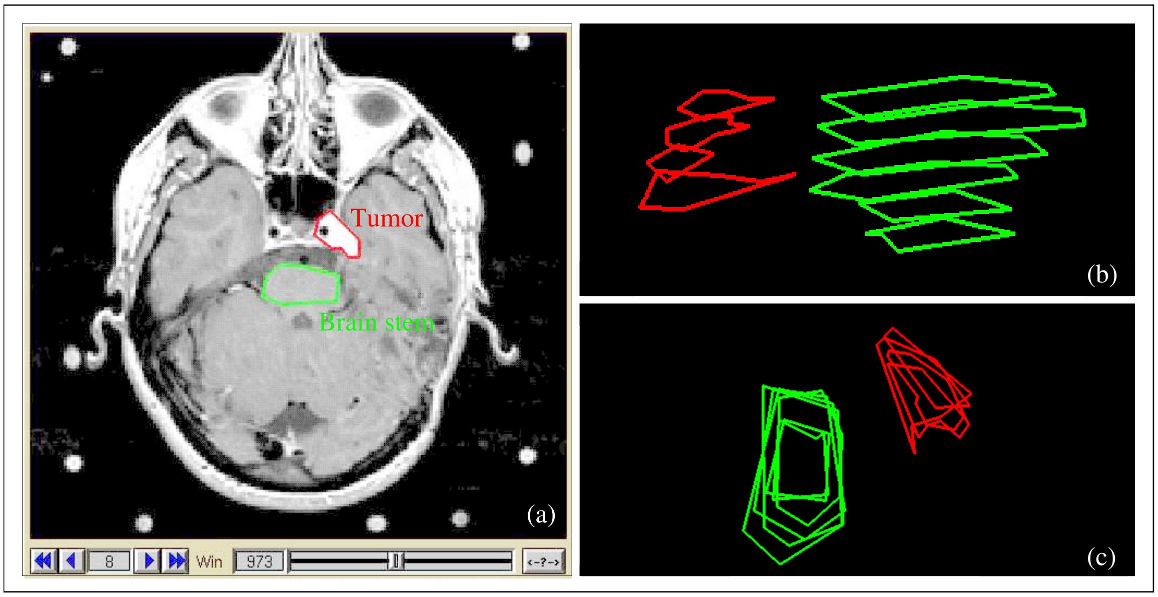

Figure 4: Delineation of tissues

requirement. A non-critical portion of healthy tissue can be input as a region of interest, as well. In fact, at one extreme, the entire healthy tissue can be partitioned into regions of interest; the only requirement is that no two regions overlap. The same connected structure (e.g., the tumor) can also be broken into several regions, which can then be assigned different dose constraints.

CARABEAMER converts the contours of the regions of interest into representations that are more directly exploitable by the planner. It computes a volume representation of each region in the form of a connected set of voxels in a 3-D array. It also constructs a surface representation of the tumor in the form of a triangular mesh. While the surface representation of the tumor will be used by the planner to select target points for the beams (Subsection 4.2), the volume representation of all regions (including the tumor) will be used to map the input dose constraints into linear inequalities binding the weights of the radiation beams (Subsection 4.3).

Volume Representation: A 3-D array of voxels that encompasses all the delineated regions is created. For simplification, let us assume that each 2-D slice of this array is aligned with a CT or MR scan. Each voxel is marked according to the region to which it belongs. To determine whether a voxel v is inside a delineated region R, CARABEAMER erects an arbitrary ray from the center of v in the 2-D slice containing v, and computes the intersections of this ray with the polygonal contour of R in the same slice. The voxel is marked as belonging to R if the number of intersection points is odd. It is easy to extend this computation to the more typical case where the resolution of the 3-D array is finer than that of the scans; CARABEAMER does not require that each slice of the voxel array be aligned with an input scan [34].

Surface Representation: Let the tumor be described by a series of r polygons $T_1, \ldots, T_r$. For every $i = 1, 2, \ldots, r-1$, CARABEAMER creates a triangulation between T_i and T_{i+1}. It then triangulates the two polygons T_1 and T_r. The triangulation between T_i and T_{i+1} is done using the algorithm proposed in [12]. The triangulations of T_1 and T_r are done using the algorithm presented in [32]. But other techniques (producing different triangulations) could have been used as well.

As mentioned above, CARABEAMER allows the user to define several disjoint regions to represent the tumor. (In fact, this must be the case when there are multiple tumors.) A separate surface representation is then computed for each region. However, for simplification, we will assume in the following that there is a single tumor region denoted T. We will refer to all delineated regions other than the tumor as the *critical regions*. We will denote them by C_i, $i = 1, \ldots, q$.

3.2 Dose Constraints

Two types of dose constraints can be input by the user: region and falloff constraints.

Region constraints: Such constraints specify the range of radiation energy that should be delivered to a given region. For the tumor T, the constraint is of the form:

$$\forall p \in T : D_T^{min} \leq d(p) \leq D_T^{max},$$

where $d(p)$ denotes the dose absorbed at point p. D_T^{min} specifies the minimum dose that any point in the tumor should absorb (to guarantee that it will be destroyed). D_T^{max} specifies the maximum dose allowed at any point in the tumor; it allows the user to control dose homogeneity inside the tumor.

For a critical region C_i, the constraint is of the form:

$$\forall p \in C_i : d(p) \leq D_{C_i}^{max},$$

where $D_{C_i}^{max}$ is the maximum dose that any point in C_i should receive. The minimum dose outside the tumor is always 0.

Falloff constraints: Each such constraint is defined by three parameters: a maximal dose D^{max}, a falloff rate ρ, and a falloff distance δ, where ρ is the amount by which the dose should decrease for every millimeter from the tumor and δ is the thickness of the region around the tumor to which the constraint applies. At any point p within δ from the tumor, the constraint implies that:

$$d(p) \leq D(p),$$

where $D(p) = D^{max} - \rho \times \ell$, where ℓ is the distance between p and the tumor.

If the tumor is close to critical regions, the user may input a different falloff constraint within each of these regions, as well as within the region around the tumor that has not been defined explicitly as critical. CARABEAMER automatically computes the region where falloff constraints apply. It partitions them into subregions and assigns a constant maximum dose to each subregion. Thus, for the planner, falloff constraints only result in increasing the number of distinct critical regions.

In the case where several tumor regions have been defined, these regions can be assigned different region and falloff constraints.

Unlike in some other treatment planners, no optimization criterion is input to CARABEAMER. It is thus implicitly assumed that any dose distribution satisfying the above constraints is satisfactory. However, if the user inputs impossible constraints, the planner will minimize their total violation (see Subsection 4.3).

4 Basic Planner

To facilitate the description of CARABEAMER, we first consider its core algorithm, which we call the basic planner. Next, we present several additions to this algorithm, which form the extended planner.

BASIC-PLANNER:

1. $B \leftarrow$ GENERATE-BEAM-SET(N).
2. **repeat** max **times** :
   (a) **if** ASSIGN-BEAM-WEIGHTS(B) succeeds in satisfying the dose constraints **then exit** with weighted beams.
   (b) **else** $B \leftarrow$ MODIFY-BEAM-SET(B).
3. **exit** with failure.

Figure 5: Outline of the basic planner

In this section we describe the basic planner. It should be clear, however, that the current version of CARABEAMER is the extended planner.

4.1 Outline

Directly inferring an appropriate set of weighted beams from tissue geometry and dose constraints seems at best extremely difficult. The geometric planner proposed in [28] can do so by mapping the critical regions into forbidden subsets of the beam configuration space and selecting beams outside those subsets. However, assumptions in this planner (e.g., the fact that treatments are using a single isocenter) restrict the configuration space to be a 2-D sphere. Moreover, completely avoiding critical regions may result in poor dose distributions, especially when the tumor is located close to critical regions. Mathematical methods have also been proposed to compute beam directions, but they have too many limitations. Here, the beam configuration space is a 6-D space. This led us to design CARABEAMER around the adaptive algorithm of Figure 5.

CARABEAMER first selects an initial beam set. In the basic planner, the size of this set is the input maximum number of beams that the user allows in a treatment plan (typically, several hundred). However, we will see that in the extended planner this size can be significantly greater than the maximum. The initial beam set is selected based only on the input 3-D tumor geometry. Ideally, if the chosen beams were all given the same weight, they would produce a region of high dose that conforms to the tumor volume. However, because initial beam selection does not consider the dose constraints, critical regions could also receive prohibitively high doses; furthermore, the dose might not be homogeneous throughout the tumor.

The planner then attempts to satisfy the dose constraints (in the context of the selected beam set) by assigning a distinct weight to each beam. In this second step, beams crossing critical structures tend to be given lower weights than other beams; but the complex shape of the tumor and its proximity to critical regions may result in non-intuitive weight assignments. If the planner succeeds in satisfying the constraints, it returns the computed set of weighted beams. If it fails, it replaces some of the selected beams by new beams and repeats the weight-assignment step.

The algorithm outlined in Figure 5 indicates that the basic planner invokes three main procedures,

GENERATE-BEAM-SET(N):

1. Pick N points evenly spaced on the tumor surface, and call these *target points*.
2. **for** each target point p, pick 3-D orientation θ at random.
3. **return** the set of selected beams.

Figure 6: Beam selection algorithm

GENERATE-BEAM-SET, ASSIGN-BEAM-WEIGHTS, and MODIFY-BEAM-SET. These procedures are respectively described in Subsections 4.2, 4.3, and 4.4. The algorithm iteratively calls ASSIGN-BEAM-WEIGHTS and MODIFY-BEAM-SET until it finds a solution. If no solution is found after a maximum of *max* iterations, the planner exits with failure. Clearly, the planner is not complete in the strongest sense: it may not return a solution, even if one exists. However, we will argue in Subsection 4.4 that it is probabilistically complete.

The procedure ASSIGN-BEAM-WEIGHTS in the basic planner embeds a simplified dosimetry model in which each beam has perfect cylindrical shape, has constant intensity over any cross-section, and does not attenuate in intensity. We also assume that the dose absorbed at a point p traversed by a beam is a constant fraction of the beam intensity at p, independent of the type of the tissue at p. We consider that all beams have the same given radius, which corresponds to the case when a circular collimator of fixed radius is used. Most of these assumptions will be withdrawn in the extended planner described in Section 5.

4.2 Generation of Initial Beam Set

In the case where the tumor's shape is roughly spherical, a classical approach is to select beams whose axes pass through the center of the tumor along many evenly distributed directions. This approach assumes that the radius of each beam is set equal to the radius of the tumor. Our early algorithm for selecting beams extended this approach by approximating the tumor by a number of spheres. For each sphere, it computed a set of beams targeting the sphere's center point along multiple directions. In one variant (emulating isocenter-based treatments), the algorithm computed a small number of spheres and produced many beams per sphere with uniformly distributed orientations; in another variant, it computed many spheres and few beams per sphere. However, our experiments with multiple variants of the algorithm showed that this approach gives poor results for many tumor shapes and placements. Even after beam weighting, the obtained dose distribution tended to be highly irregular within the tumor volume and had insufficient falloff around the tumor.

Several experiments led us to the algorithm now used by CARABEAMER. This algorithm is outlined in Figure 6, where N is set to the specified maximum size N_{max} of the beam set. It first picks N points evenly spaced in the tumor surface. Next, for each point, it creates one beam targeting this point; the orientation of this beam is chosen at random with uniform distribution.

To pick target points uniformly on the tumor surface, GENERATE-BEAM-SET uses the triangulated surface representation of the tumor and performs the following three steps. First, it assigns a number of points to each triangle, proportional to the triangle's surface area relative to the total area of the triangulation. Next, in each triangle, it picks the points at random with uniform distribution [38]. Finally, it uses the point-repulsion technique proposed in [39] to move the points and uniformly space them over the entire tumor surface. The points obtained after those three steps are the *target points*.

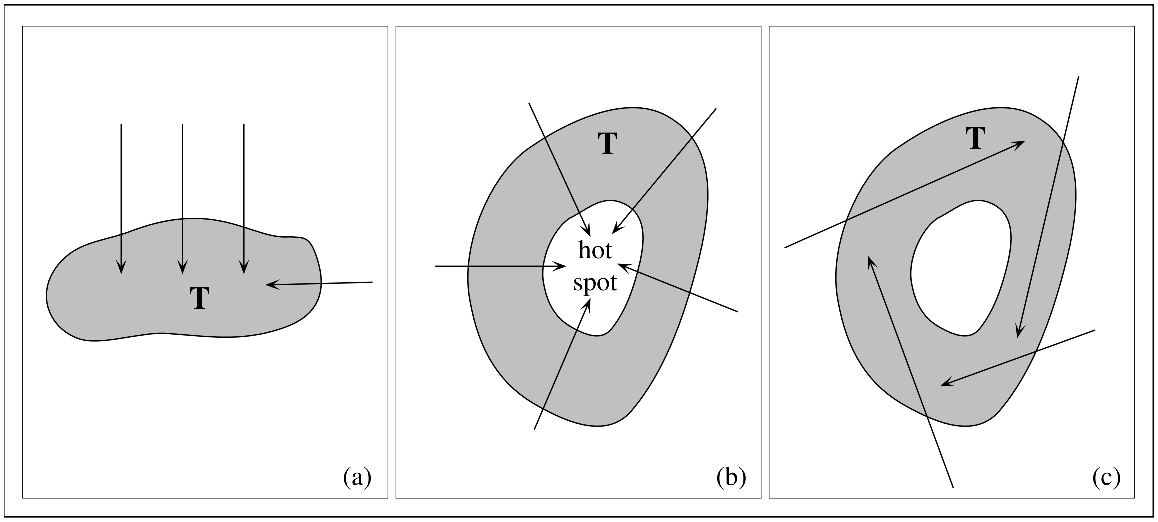

Figure 7: Strategies for choosing beam directions

The choice of beam orientations deserves some comments. Indeed, it would seem more natural to choose orientations perpendicular to the tumor surface at the target points. In a case like the one shown in Figure 7(a), this choice would work well. But in the case of Figure 7(b), it would result in a high-dose region outside the tumor. In the latter case, it is more appropriate to choose beams that are almost tangent to the tumor surface, as in Figure 7(c). In fact, no single deterministic technique seems to handle all cases well. One could try to classify tumor shapes into categories, and associate a distinct technique with each category. But this turns out to be more complicated than might be expected. Choosing beam orientations uniformly at random is attractive because it does not seem to have major pathological cases, as long as N is large enough. Indeed, in this case, each point in the tumor volume has a high probability of being contained in several beams, each with a different orientation.

Figure 8 shows the results of these computations on an example of triangulated surface. The target points chosen in the surface before and after applying the point-repulsion technique are displayed in Figure 8(a) and Figure 8(b), respectively. The selected beam orientations are shown in Figure 8(c). The 50% and 80% isodose surfaces when all the beams are assigned equal weight are overlaid on the tumor geometry in Figure 8(d). (Recall that the $x\%$ isodose surface is defined as the boundary of the set of all points receiving at least x% of the dose received by the maximum-dose point.) These results were obtained with 600 beams of radius 12.5mm, the "tumor" being contained in a box of size $21 \times 18 \times 9\text{mm}^3$. The tumor dose-volume histogram is displayed in Figure 8(e). This histogram should be read as follows: at any given abscissa x, the ordinate y indicates that a subset of the tumor of volume y receives at least $x\%$ of the maximum dose. Thus, Figure 8(e) indicates that the entire tumor volume (resp., half of it) receives more than 40% (resp., 70%) of the maximum dose. The results shown in Figures 8(d-e) were obtained with the simplified dosimetry model embedded in Assign-Beam-Weights.

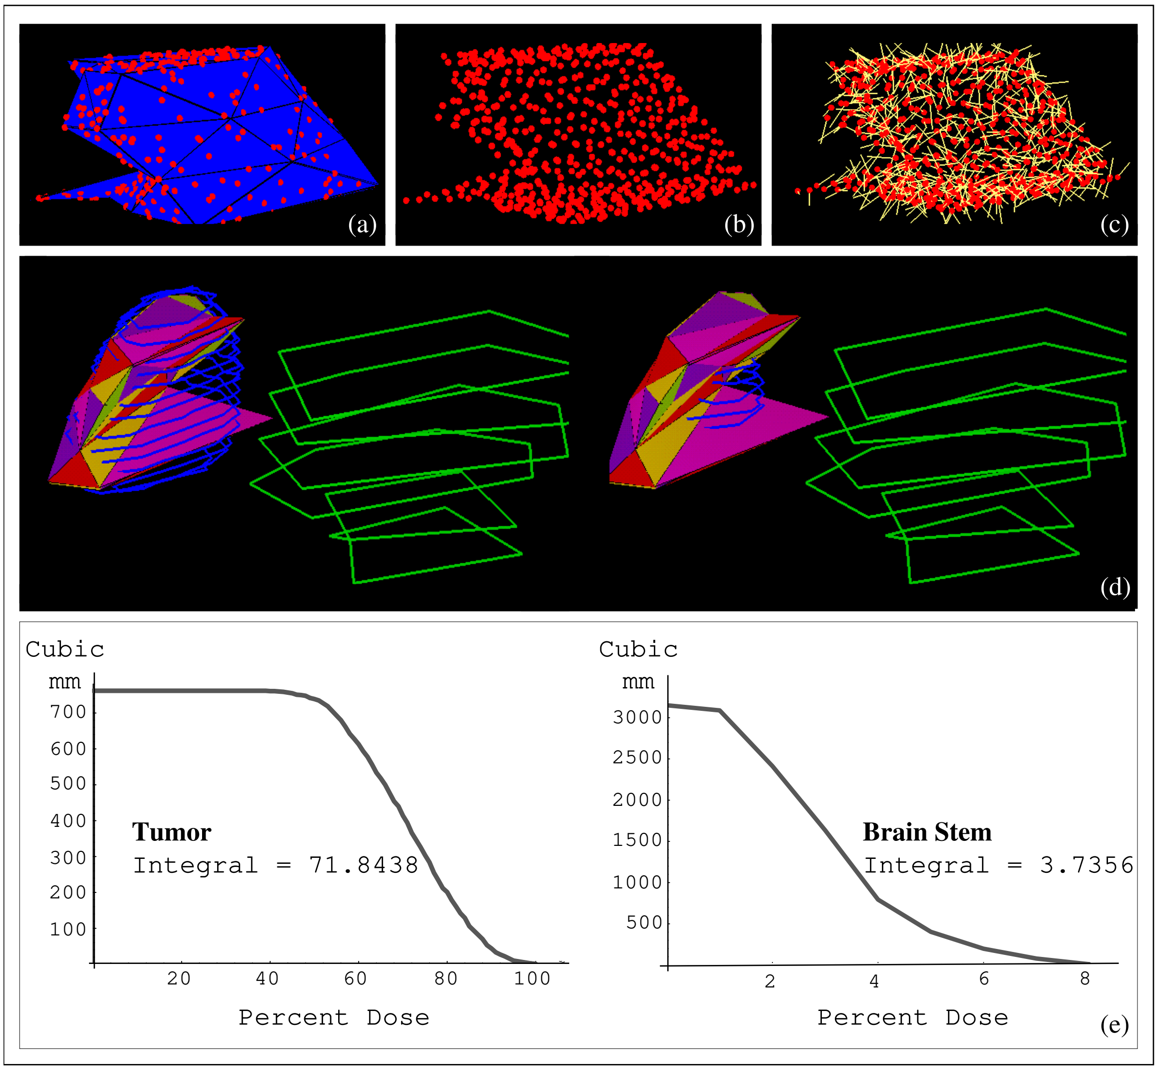

Figure 8: Example of beam set computed by Generate-Beam-Set

4.3 Weight Assignment

The procedure Assign-Beam-Weights attempts to achieve the input dose constraints by assigning an appropriate weight to each beam. Using the simplified dosimetry model, it first transforms the dose constraints into a set of linear inequalities binding the beam weights. It then tries to solve these constraints using a linear program. Figure 9 outlines the algorithm of Assign-Beam-Weights.

As indicated earlier, Assign-Beam-Weights represents every beam as a cylinder of constant intensity. Let $B_1, \ldots, B_N$ denote the cylinders representing the N beams. Recall that T designates the tumor region, in which the dose should be between D_T^{min} and D_T^{max}, and C_i, $i = 1, \ldots, q$ denote critical regions in which the dose should be lower than $D_{C_i}^{max}$. The surfaces of $B_1, \ldots, B_N, T, C_1, \ldots, C_q$ decompose space into an arrangement of cells, where each cell is a maximal 3-D region that does not intersect any of the surfaces of $B_1, \ldots, B_N, T, C_1, \ldots, C_q$. Let κ be a cell of this arrangement and $k_1, \ldots, k_r$ ($k_r \in [1, N]$ and $r \geq 0$) the indices of the beams containing this cell. If κ is contained in T, Assign-Beam-Weights establishes the following two inequalities:

ASSIGN-BEAM-WEIGHTS(B):

1. Compute the 3-D cell arrangement A of beams and tissue regions.
2. **for each** cell of A establish linear constraints on beam weights, and call L the set of all obtained constraints.
3. Remove redundant constraints from L.
4. Invoke linear program to solve L.
5. **if** a solution has been obtained **then exit** with beam weights; **else exit** with beam weights minimizing dose constraints.

Figure 9: Weight assignment algorithm

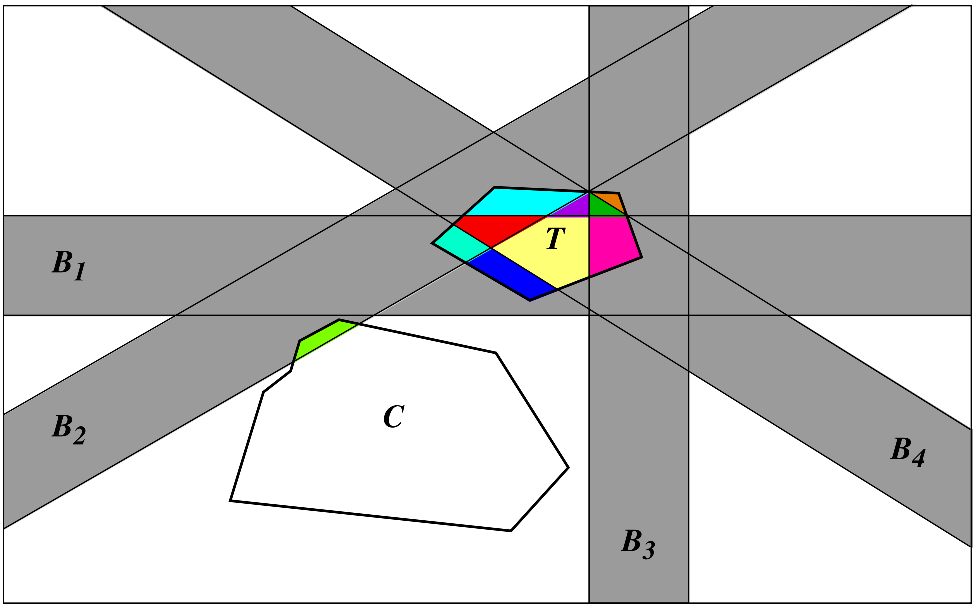

Figure 10: A simple arrangement of cells

$$D_T^{min} \leq w_{k_1} + w_{k_2} + \ldots + w_{k_r} \leq D_T^{max}.$$

Actually, since the weight w of a beam refers to the energy that is emitted by the radiation source during the time it is activated, only a fraction γ of this weight should be used in the inequalities. Because γ is a constant number in the simplified dosimetry model, we do not make it appear in the inequalities.

Similarly, if κ is in the region C_i, the following inequality is produced:

$$w_{k_1} + w_{k_2} + \ldots + w_{k_r} \leq D_{C_i}^{max}.$$

If κ is not contained in either T or a region C_i, no constraint is generated. If the number of beams is large enough, a cell κ contained in T has very low probability of being contained in no beams; should that happen, additional beams are selected to target points within κ.

Figure 10 illustrates this decomposition with a simple arrangement of cells in a 2-D space. In this example there are only four beams and $q = 1$ (with the corresponding region denoted by C). There are a total of nine cells contained in T and one in C yielding a total of nineteen linear inequalities on w_1, w_2, w_3, w_4:

$$\begin{array}{lclcl}
D_T^{min} & \leq & w_2 + w_4 & \leq & D_T^{max}, \\
D_T^{min} & \leq & w_4 & \leq & D_T^{max}, \\
D_T^{min} & \leq & w_3 + w_4 & \leq & D_T^{max}, \\
D_T^{min} & \leq & w_3 & \leq & D_T^{max}, \\
D_T^{min} & \leq & w_1 + w_3 + w_4 & \leq & D_T^{max}, \\
D_T^{min} & \leq & w_1 + w_4 & \leq & D_T^{max}, \\
D_T^{min} & \leq & w_1 + w_2 + w_4 & \leq & D_T^{max}, \\
D_T^{min} & \leq & w_1 & \leq & D_T^{max}, \\
D_T^{min} & \leq & w_1 + w_2 & \leq & D_T^{max}, \\
 & & w_2 & \leq & D_C^{max}.
\end{array}$$

Since the weight of a beam cannot be negative, several linear inequalities may be redundant. For example, consider the two inequalities $D_T^{min} \leq w_2 + w_4$ and $D_T^{min} \leq w_4$. The satisfaction of the second entails that of the first, regardless of the value of w_2. So, we can remove the inequality $D_T^{min} \leq w_2 + w_4$ from the set of inequalities. More generally consider two cells κ and κ' in the same region of interest. Let $K = \{k_1, \ldots, k_r\}$ and $K' = \{k'_1, \ldots, k'_{r'}\}$ denote the sets of indices of the beams containing κ and κ', respectively. If $K \subset K'$, then the inequality for κ bounding the dose from below is subsumed by the corresponding inequality for κ'. Similarly, if $K \supset K'$, then the inequality for κ bounding the dose from above is subsumed by the corresponding inequality for κ'. In the example of Figure 10, the set of inequalities is thus reduced to the six following inequalities:

$$\begin{array}{lclcl}
D_T^{min} & \leq & w_1, & & \\
D_T^{min} & \leq & w_3, & & \\
D_T^{min} & \leq & w_4, & & \\
 & & w_1 + w_2 + w_4 & \leq & D_T^{max}, \\
 & & w_1 + w_3 + w_4 & \leq & D_T^{max}, \\
 & & w_2 & \leq & D_C^{max}.
\end{array}$$

Interestingly, the same concept of minimally and maximally covered cells has also been used in [14] to design an efficient assembly planning technique.

Solving the obtained set of inequalities for $w_1, \ldots, w_N$ is a linear programming problem. We solve this problem using the optimization software MINOS[25] which uses a modified version of the simplex algorithm. To facilitate the work of MODIFY-BEAM-SET, in the case where no solution is found, we introduce one variable per inequality, which measures the amount by which the inequality is violated and we ask the linear program to minimize the sum of all these variables. In this way, if the linear program does not find a solution for the beam weights, it still returns a set of weights that minimizes the total cummulative violation of the inequalities. Obviously, we may weight differently the violation of each inequality to prioritize the constraints.

The arrangement of N cylinders contains $O(N^3)$ cells and takes $O(N^3)$ time to compute [9]. Its exact computation is also delicate to implement. Instead, we lay a thin grid of specified resolution over the

3-D space of interest. The voxel matrices of the volume representation described in Subsection 3.1 are exactly aligned with this grid. For each sample point in the grid contained either in T or a region C_i we establish inequalities as above. Elimination of redundant inequalities occurs concurrently. Hence, if several grid points are in the same cell, the inequalities are generated for only one of these points. The apparent drawback of using a grid is that small cells may be missed. But the input geometric representation of the tissue, the positioning of the linac relative to the patient's head, and the dosimetry model used are not totally accurate either. Therefore, cells small enough to fall between the points of an appropriately thin grid are irrelevant in practice. The resolution of the grid is set by the user. In our experiments we usually set the distance grid points along every coordinate axis to 1mm, which corresponds to the resolution of the CT image.

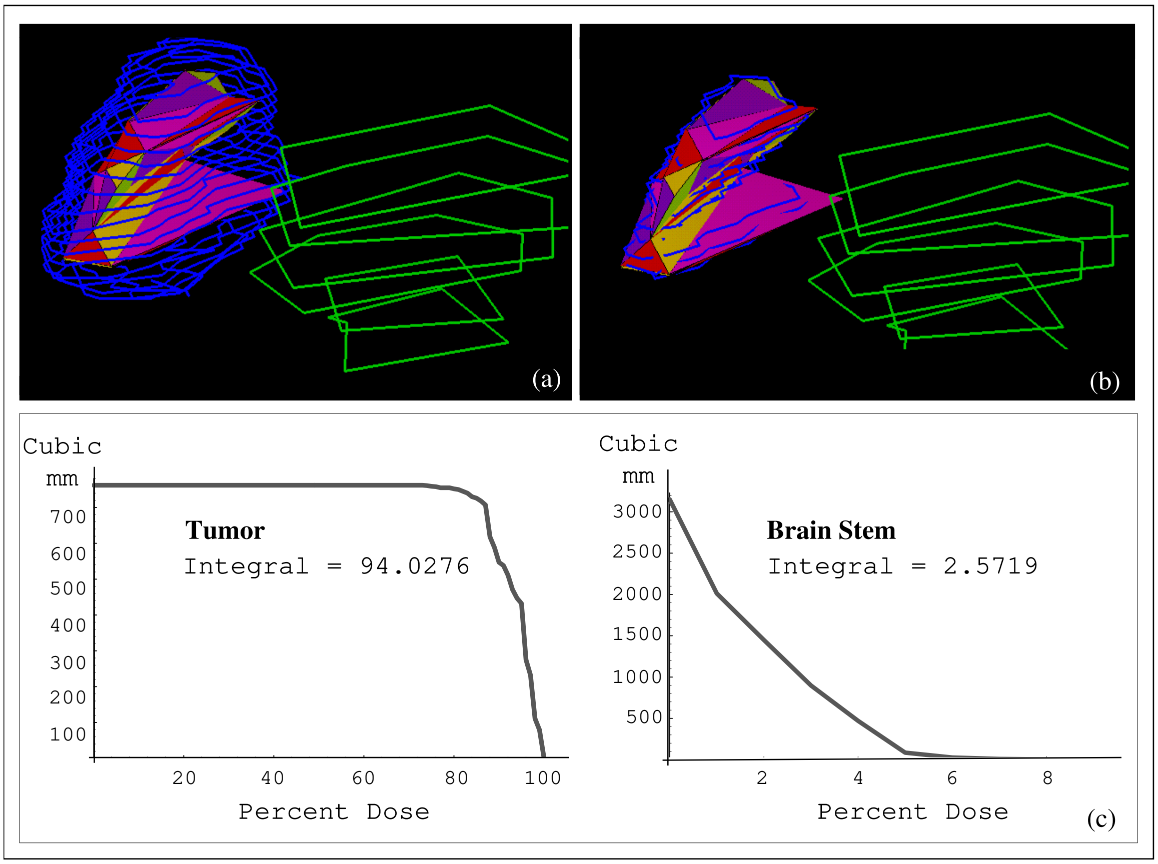

Figure 11: Example of dose distribution after beam weighting

Figure 11 shows the effect of running Assign-Beam-Weights on the beams displayed in Figure 8(c). The dose constraints were $D_T^{min} = 2000$ and $D_T^{max} = 2400$ in the tumor, and $D_C^{max} = 500$ in the critical region (brain stem); a falloff constraint of 100/mm was set around the tumor. The linear program succeeded in finding a solution. The 50% and 80% isodose surfaces are shown in Figure 11(a) and Figure 11(b), respectively. Note that the region of high dose conforms to the tumor shape much more closely than in Figure 8(d). The tumor dose-volume histogram is displayed in Figure 11(c). The entire tumor volume now receives more than 70% of the maximum dose, which means that the dose distribution in the tumor is also more homogeneous than in Figure 8. Again, these results have been obtained with the dosimetry model embedded in Assign-Beam-Weights.

MODIFY-BEAM-SET(B):

1. Remove every beam whose weight is less than w_{min} from B. Let k be the number of removed beams.
2. Generate k new beams and insert them in B.

Figure 12: Beam modification algorithm

4.4 Modification of Beam Set

In the case where the procedure ASSIGN-BEAM-WEIGHTS fails to satisfy the dose constraints, one possibility for the basic planner would be to give up, despite the fact that a more fortunate choice of beams could have led to a solution. Another possibility would be to run GENERATE-BEAM-SET again, that is, to pick an entirely new set of beams and to submit them to ASSIGN-BEAM-WEIGHTS. However, the fact that the first trial failed provides some evidence that the problem at hand either has no solution, or is somewhat tricky, suggesting that a more involved planning technique be applied. Moreover, by minimizing a measure of the dose constraint violation, ASSIGN-BEAM-WEIGHTS provides information about which beams are relevant and which ones are not. The procedure MODIFY-BEAM-SET makes use of this information.

The algorithm of MODIFY-BEAM-SET is outlined in Figure 12. First, using the weights computed by ASSIGN-BEAM-WEIGHTS, it treats the beams whose weights are below a given threshold w_{min} as irrelevant to a potential solution. It removes them from the beam set and replaces them by new beams. It considers the beams with the highest weights as promising ones, and computes a fixed fraction (say, half) of the new beams by randomly perturbing the target points and orientations of these existing beams. To deal with the fact that there may exist other relevant beams, MODIFY-BEAM-SET produces the other new beams by picking new target points evenly distributed on the tumor surface and selecting a 3-D orientation uniformly at random for each target point (i.e., by using the same technique as for the selection of the initial beam set).

Both GENERATE-BEAM-SET and MODIFY-BEAM-SET make several choices at random. Such randomization has the advantage of being robust and somewhat free of pathological cases. However, as a result, CARABEAMER is obviously not complete. It may not find a solution, even when a solution does exist. Nevertheless, if the size N of the beam set is large enough, every point in the tumor is likely to be contained in multiple beams. Moreover, since the beam set is iteratively modified by MODIFY-BEAM-SET and a fraction of the new beams at each iteration are picked at random, with high probability every point in the tumor will eventually be traversed by a beam having a desirable orientation. Hence, we believe that the planner is probabilistically complete. This means that if we let it run long enough (by setting the maximum number of iterations max to infinity), it will eventually produce a solution whenever one exists. Due to the complex behavior of CARABEAMER, it seems hard to formally prove this claim. One can nevertheless notice that similar randomized techniques have also been successfully used to plan collision-free robot motions in high-dimensional configuration spaces [5]. In that case, because obstacles divide the robot's configuration space into only two types of regions (collision and non-collision), it was possible to prove probabilistic completeness, as well as to relate the planner's running time to its probability of failure.

EXTENDED-PLANNER:

1. $B \leftarrow$ GENERATE-BEAM-SET(N).
2. **repeat** max **times** :
   (a) **if** ASSIGN-BEAM-WEIGHTS(B) succeeds in satisfying the dose constraints **then:**
      i. REDUCE-BEAM-SET(B)
      ii. **exit** with reduced set of weighted beams.
   (b) **else** $B \leftarrow$ MODIFY-BEAM-SET(B).
3. **exit** with failure.

Figure 13: Algorithm of the extended planner

5 Extended Planner

In this section we describe several implemented additions to the basic planner.

5.1 Reduction of Treatment Time

Reducing treatment time is an important goal, since shorter treatments both make it possible to treat more patients in a single day and reduce the psychological and physical stress on the patients.

Short treatment plans also make it possible to fractionate a treatment into several short sessions (each of 20 minutes or less), which has proven medical advantages. Indeed, a simple way to produce k sessions is to divide the beam weights by k, and apply the same treatment k times. However, this is feasible only if the treatment plan is short enough.

As mentioned in Section 2, the major determinant of the treatment time is the number of beams. Reducing treatment time can thus be achieved by reducing the number of beams needed to satisfy the input dose constraints. On the other hand, the randomized selection of beam done by CARABEAMER supposes that the size of the initial beam set be large enough to ensure that all points in a tumor are contained in several beams with high probability. To deal with these conflicting requirements, the extended planner (see Figure 13) runs the algorithm of the basic planner with a large number of beams, N. Once it has found N weighted beams that achieve the dose constraints, it attempts to reduce the number of beams by invoking REDUCE-BEAM-SET.

REDUCE-BEAM-SET iteratively replaces the α beams with the lowest weights by a smaller number, β, of new beams, where α and β are two constants. It runs ASSIGN-BEAM-WEIGHTS on each new beam set, until this procedure fails to satisfy the dose constraints. When failure occurs, the extended planner exits with the last set of weighted beams that satisfies the dose constraints. In reality, CARABEAMER is a bit more tenacious; when failure occurs, it tries multiple beam substitutions before giving up [34].

This extension allows the planner to set the size N of the initial beam set larger than the maximum number of beams, N_{max}, allowed by the user in a treatment plan. A larger N usually increases the

REDUCE-BEAM-SET(B):

1. $B_{old} \leftarrow B$.
2. **repeat** max' **times** :
   (a) Remove the α beams from B with the smallest weights.
   (b) Add β new beams to B.
   (c) **if** ASSIGN-BEAM-WEIGHTS(B) fails
   then exit with the previous weighted beam set B_{old}.

Figure 14: Beam set reduction

running time of the linear program, since it yields both a larger number of weights to compute and a larger number of linear inequations to deal with. On the other hand, it also increases the chance of finding a solution with a smaller number of iterations through ASSIGN-BEAM-WEIGHTS and MODIFY-BEAM-SET, which in turn reduces the number of invocations if the linear program. Our experience is that, in most cases, a good compromise is to set N to 2 or 3 times N_{max}.

5.2 Selection of Beams

Our experiments have revealed that the target-selection scheme described in Subsection 4.2 has two drawbacks.

First, when the tumor is very close to a critical region or when the difference $D_T^{max} - D_T^{min}$ is small (i.e., the requirement for dose homogeneity in the tumor is stringent), achieving the constraints either requires a large initial number of beams, and/or more iterations through ASSIGN-BEAM-WEIGHTS and MODIFY-BEAM-SET. A better way is then to select the target points on a surface at a distance ε from the tumor surface. If ε is negative, the target points are inside the tumor; then it is easier to achieve a steep falloff of the dose around the tumor, hence to satisfy dose constraints in surrounding critical regions. Instead, if ε is positive, the target points are outside the tumor; then it is easier to deal with the case where the difference $D_T^{max} - D_T^{min}$ is small.

Second, the portions of the tumor surface that have high curvature also tend to require more beams and/or more iterations to find a satisfactory set of beams. This observation led us to modify GENERATE-BEAM-SET to produce a greater density of target points where surface curvature is higher. Our planner uses the algorithm in [40] to compute the curvature of a triangulated surface at every vertex. The number of target points picked at random in each triangle then grows with the area of the triangle, as well as with the average curvature of its vertices. The point-repulsion algorithm is modified so that a greater number of target points remain in the high-curvature areas. See [34] for more detail.

It is likely that the choice of beam directions could be improved as well. Rather than picking directions at random uniformly, we could introduce a bias in the probabilistic distribution such that those directions which define beams traversing regions with very low dose requirements have lower probability of being picked than other directions. But we have not yet implemented this variant.

5.3 Improved Dosimetry Model

The dosimetry model used in the basic planner makes simplifying assumptions (Subsection 4.1) that one would like to remove. Indeed, it may yield treatment plans that are not found satisfactory when simulated using the reference direct dosimetry program. The extended planner embeds a more realistic model in which the dose deposited by a beam of weight w at a tissue point p is $\gamma(p) \times w$, where $\gamma(p)$ is a function of the distance between p and the radiation source, the amount of tissue traversed to reach p, the distance between p and the beam's central axis, and the absorption coefficient of the tissue at p. These parameters make it possible to take into account that a beam widens as the radiation gets away from the source, that in any cross-section radiation intensity is greater near the central axis than on the periphery, and that beam intensity attenuates as it travels through tissue.

Implementing such an improved model is straightforward due to the fact that ASSIGN-BEAM-WEIGHTS discretizes space into a grid of points at which linear inequalities on beam weights are established. However, the downside of this extension is that there is less redundancy in the inequalities, so that more inequalities must eventually be input to the linear program. A possible strategy for the user is to first run CARABEAMER with the simplified model, get a solution for that model, and then switch to the improved model with the previous solution as a starting point.

5.4 Collimator Selection

An input of the basic planner is the collimator diameter. In the extended planner, the user can enter a set of several diameters and let CARABEAMER select the diameter of each beam from that set. Each beam is first represented by several concentric cylinders. The innermost cylinder corresponds to the smallest collimator, and each surrounding cylinder corresponds to a successively wider collimator. The procedure ASSIGN-BEAM-WEIGHTS assigns s weights $w_{i_1}, \ldots, w_{i_s}$ to each beam Bi, where s is the number of collimator diameters considered. The weight w_{i_1} corresponds to the innermost cylinder, while w_{i_s} corresponds to the outermost cylinder. If a grid point lies in several cylinders representing the same beam B_i, the inequalities established at this point only involve the weight of the outermost cylinders to which the point belong. In addition, since each outer cylinder represents a wider version of the beam represented by an inner cylinder, the planner also imposes the following constraint for every beam B_i:

$$\forall k \in \{1, \ldots, s-1\} : w_{i_{k+1}} \leq w_{i_k}.$$

After computing the weights for each cylinder of every beam, ASSIGN-BEAM-WEIGHTS analyzes these weights in order to select a single diameter per beam. For each beam B_i, it starts from the innermost cylinder and moves outward. If it notes a substantial decrease between the weights w_{i_k} and $w_{i_{k+1}}$ assigned to two successive cylinders, it selects the cylinder k and assigns the weight w_{i_k} to the beam. If no substantial decrease is noted, it selects the outermost cylinder and assigns the weight w_{i_s} to the beam.

The set of weighted beams obtained after removing cylinders may no longer satisfy the dose constraints. For that reason, ASSIGN-BEAM-WEIGHTS recomputes the weights of the beams (now modelled, each, as a single cylinder) using the previously computed weight as a starting point.

5.5 Subsampling

When the initial beam set is large and the improved dosimetry model is used, the number of linear inequalities submitted to the linear program is also large (several thousands or more). This can seriously increase the run time of the planner. A simple technique, which we found very effective to reduce this time, is to first pick a limited number of grid points and to consider only the inequalities at these points. The solution for this simplified problem is then used as a starting point from which the planner can more efficiently find a solution to the fully sampled problem.

6 Experimental Results

Evaluating a planner like CARABEAMER is a difficult task. Ideally, we would like to compare its performance to an "optimal" planner, but such a planner does not exist. Moreover, there is no consensus on what makes one treatment better than another. In our case, we could say that a planner is optimal if it always returns a solution with minimal number of beams in minimal amount of time, whenever a solution exists. There seems to be no formal or empirical way of measuring the distance that separates CARABEAMER from such a planner.

CARABEAMER could be compared to already existing planners. But, as we mentioned before, existing planners compute isocenter-based treatments. Comparing their outputs to those of CARABEAMER would be unfair, since CARABEAMER's plans benefit from the general kinematics of the CyberKnife manipulator. Inversely, we could use CARABEAMER to generate isocenter-based treatments. But this would at least require us to change the way CARABEAMER selects beams. It would not be fair for our planner, since the smaller size of the beam configuration space makes it possible to use more specific techniques.

Until now, the FDA has only permitted the CyberKnife system to be used with isocenter-based treatment plans. So, a clinical evaluation of the patients after applying treatment plans produced by CARABEAMER has not been possible yet. In addition, such evaluation would better qualify the inputs provided by the user than the planner itself.

These considerations have led us to conduct two series of experiments, each with limited goals: experiments on patient cases (cases which were treated with a conventional LINAC system at Stanford Medical Center) and experiments on synthesized cases. The first series of experiments allowed us to verify that CARABEAMER can produce plans that meet the requirements taken into consideration when defining the treatments that were actually delivered. The second series was intended to explore the performance of CARABEAMER on particularly difficult cases.

A number of other experiments have been done with previous versions of the planner. See [2, 3, 29, 30, 31, 35, 36]. Additional experiments with CARABEAMER are also reported in [34].

6.1 Experiments on Patient Cases

The test cases used in this series of experiments correspond to 12 patients treated for brain tumors with the LINAC system at Stanford Medical Center during the Summer of 1996. Because several cases involve multiple lesions, 20 tumors were considered in total: 3 acoustic neuromas, 2 cavernous meningiomas, 4 hemangioblastomas, 1 hemangiopericytoma, 2 metastases (primary breast cancer), 2 metastases (primary ovarian cancer), 3 metastases (adenocarcinoma of lung), 1 oat cell carcinoma, and 2 pituitary adenomas at the right cavernous sinus. The critical structures noted in these cases were the brain stem, the right and left optic nerves, the optic chiasm, and the cornea.

The same neurosurgeon who developed the plans used in treatment defined the inputs for CARABEAMER. For each case, he specified the geometry of the regions of interest and the dose constraints in these regions. This required the neurosurgeon to specify the problems in a more explicit manner than that to which he is accustomed. Indeed, to define the LINAC treatments, he manually iterated with the direct dosimetry program in a trial-and-error manner; he usually did not know the precise objectives prior to planning, but, instead, let them evolve through experimentation. He could have invoked CARABEAMER multiple times as well, with different dose constraints, but he did not do so. In any case, it is clear that two so different planning environments would not lead the neurosurgeon to eventually consider the same constraints.

CARABEAMER was run on each of the 12 test cases with the direct dosimetry extension described in Subsection 5.3, which evaluates the dose delivered at a tissue point p as a function of the distance between p and the beam's central axis and the distance between p and the radiation source. All dosimetry results presented below have been obtained using this same model. CARABEAMER used a grid resolution of 1 cubic millimeter to formulate linear inequalities and calculate the dose distribution. The grid size varies with each case, from 64^3 to 128^3. We recorded the maximum and minimum doses to each region of interest, as well as the maximum and minimum dose to the "external" tissue (defined as the tissue outside any region of interest in the volume covered by the grid).

Although a comparison with the plans given to the LINAC system is not really possible, we did simulate these plans as well to observe the differences in the dose distributions. However, for logistical reasons, we used the same dosimetry model as for our plans, which is somewhat different from the model in the direct dosimetry program used by the neurosurgeon when he planned the cases. For that reason, we avoid reporting detailed numerical data about the LINAC plans.

Tables 1 and 2 summarize the results on the 12 cases. In each table, the first column lists the regions of interest and the external tissue defined for each test case (note that cases 1, 2, 6, 8, 9, 10, and 12 involve more than one tumor). The second column indicates the constraints that the neurosurgeon specified for each of the regions of interest. Column 3 gives the minimum and maximum dose computed by our dosimetry program for the plan generated by CARABEAMER. These results essentially show that CARABEAMER was able to compute plans that satisfy the constraints input by the surgeon. This does not mean, however, that the plans are entirely satisfatory. Indeed, for every test case, the maximum dose in the external tissue (for which no dose constraint was specified) is of the same order of magnitude as the maximum dose in the tumor. Points of high dose in the external tissue are close to the tumors and could have been limited by specifying falloff constraints around the tumor. But the surgeon made little use of such constraints. This probably reveals his bias toward isocenter-based treatments, where the highest dose necessarily occurs at the isocenter and never outside the tumor.

Figures 15 through 26 provide a visual comparison of the dose distributions of the plans produced by CARABEAMER and those produced manually for the LINAC system. The triangulated surfaces depict the tumors, while the stacks of green contours represent the critical regions. The stacks of blue contours show the 80% isodose surfaces, except in Figures 20 and 26 where they are used to show both the 50% and 80% isodose surfaces. Figures 15 through 26 do not allow us to say which of the two plans specifies a clinically better treatment for each given case. However, they do show that CARABEAMER's plans produce a better shape matching between the region of high dose and the tumor volumes. We believe that shape matching could have been further improved had the surgeon specified more stringent constraints.

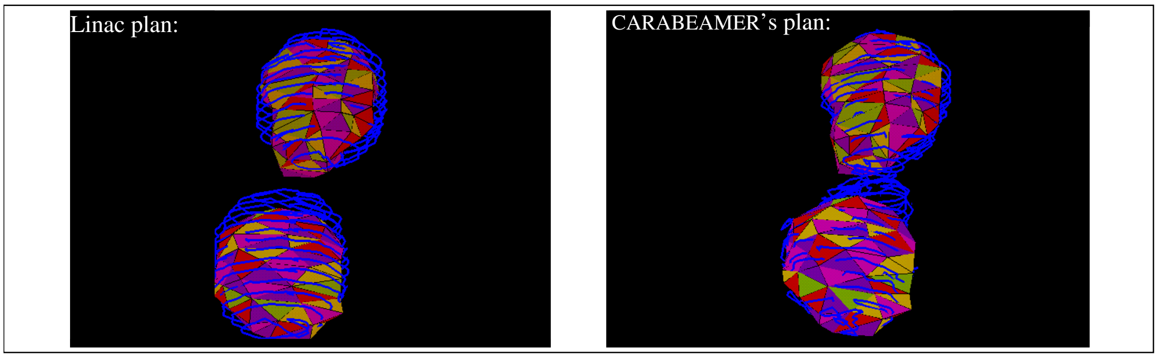

Figure 15: Patient case 1

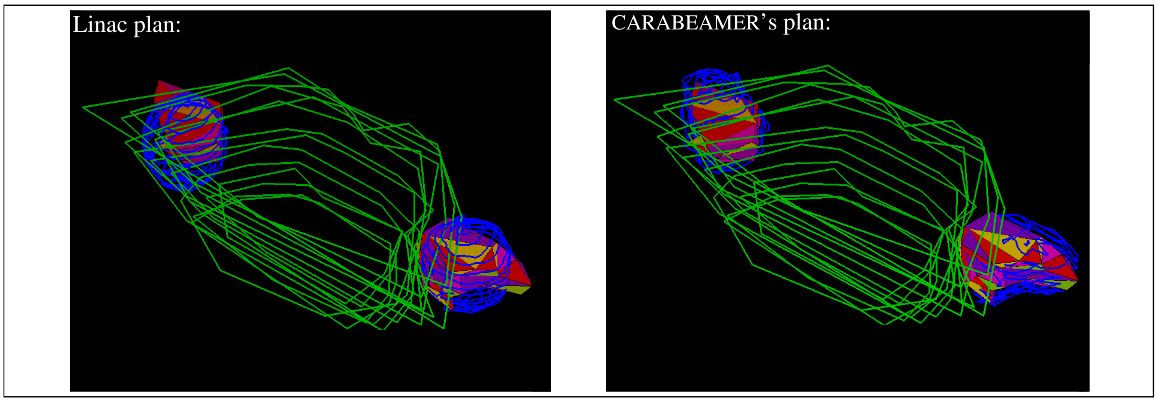

Figure 16: Patient case 2

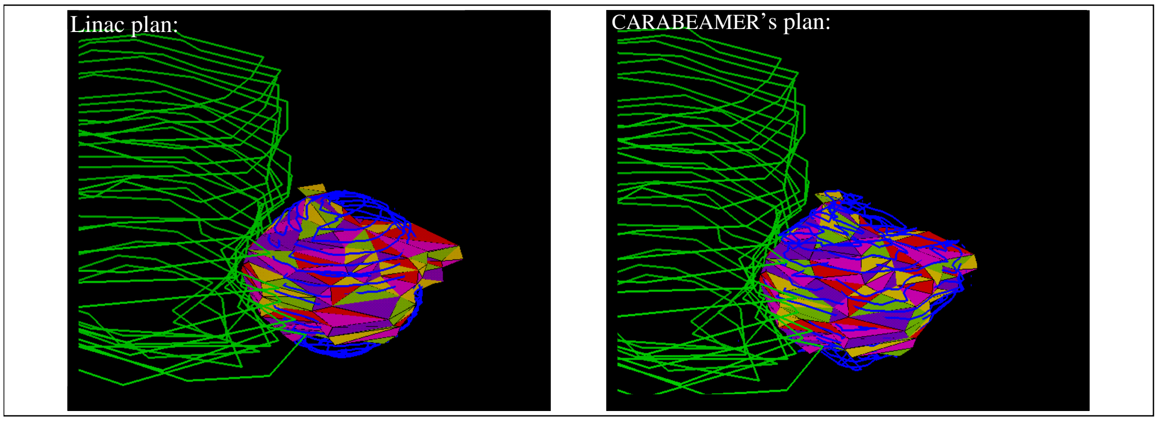

Figure 17: Patient case 3

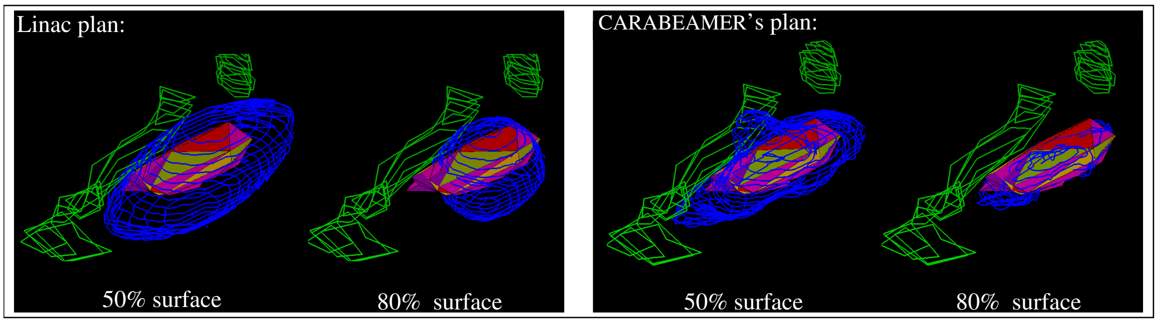

Figure 18: Patient case 4

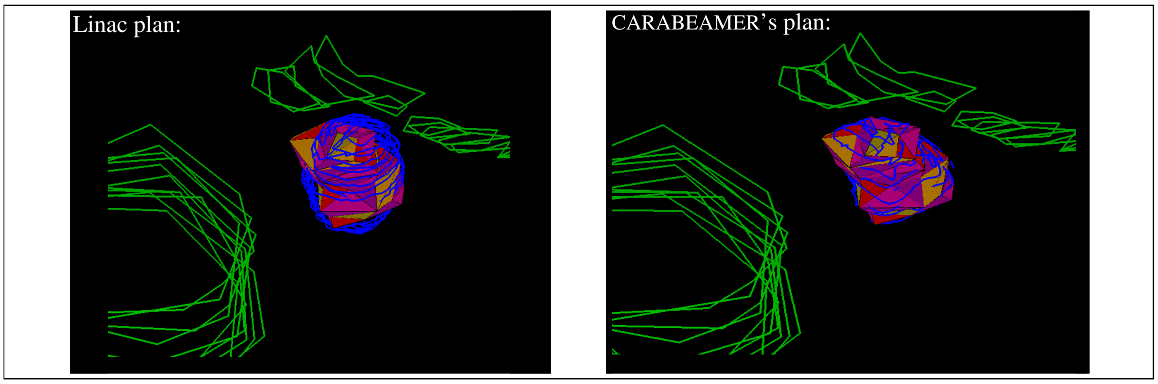

Figure 19: Patient case 5

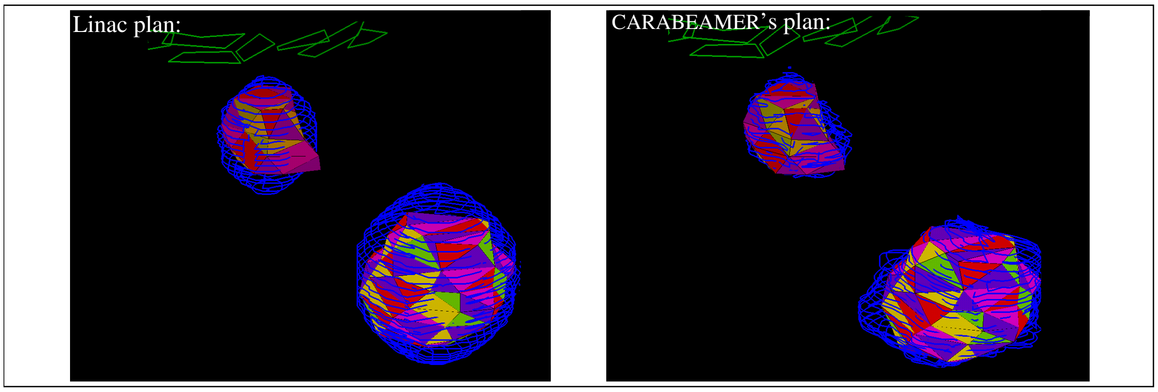

Figure 20: Patient case 6

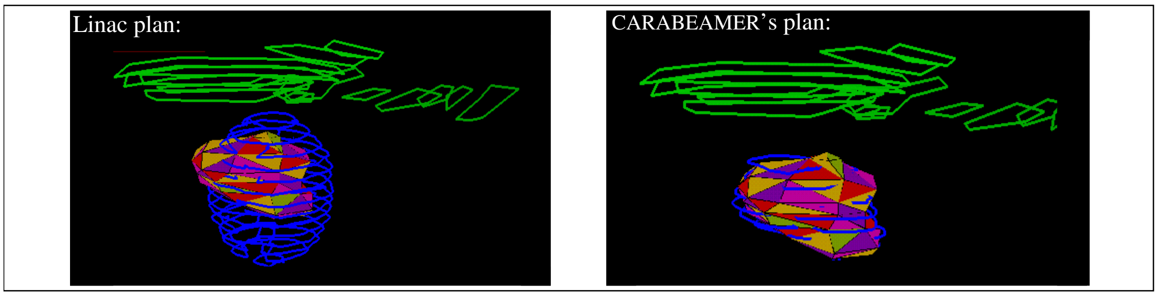

Figure 21: Patient case 7

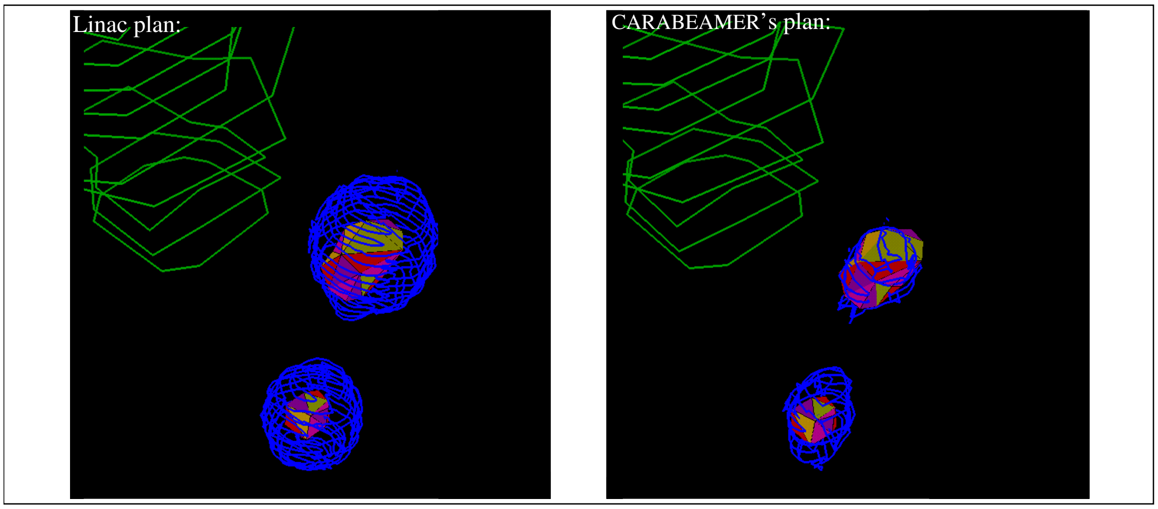

Figure 22: Patient case 8

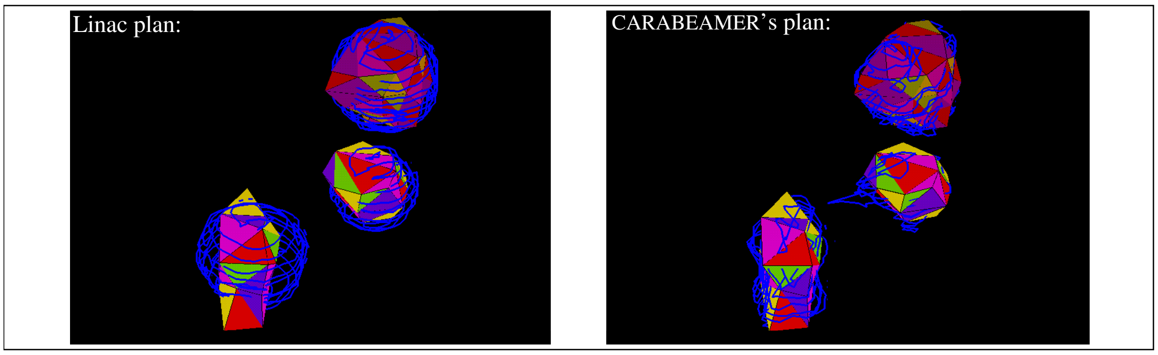

Figure 23: Patient case 9

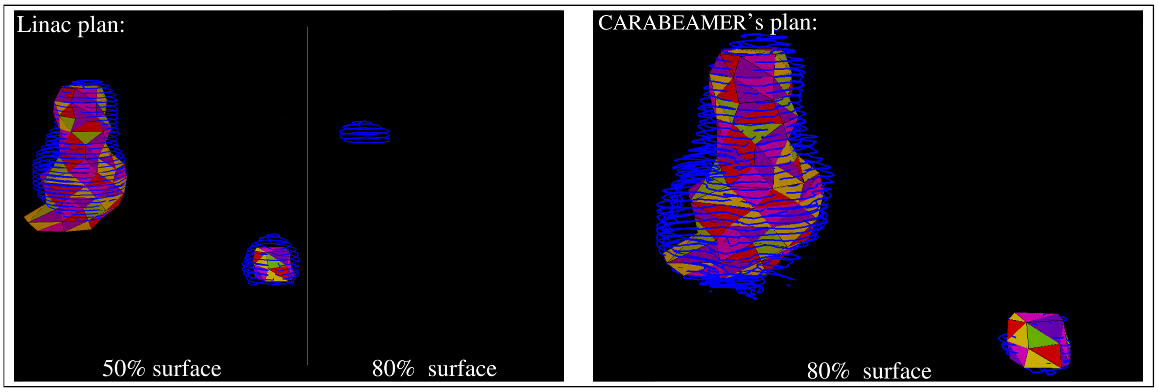

Figure 24: Patient case 10

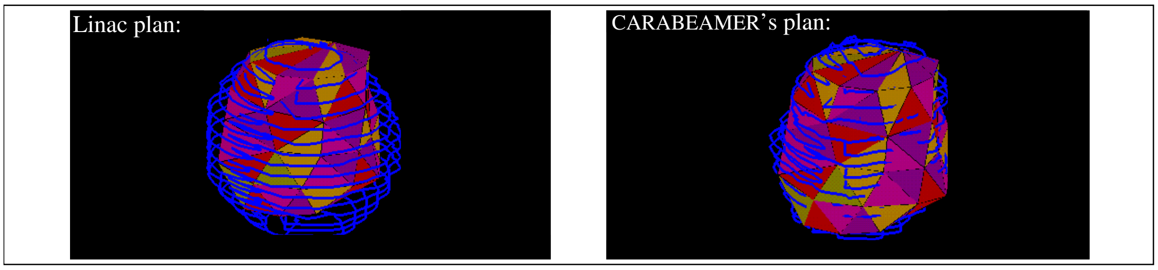

Figure 25: Patient case 11

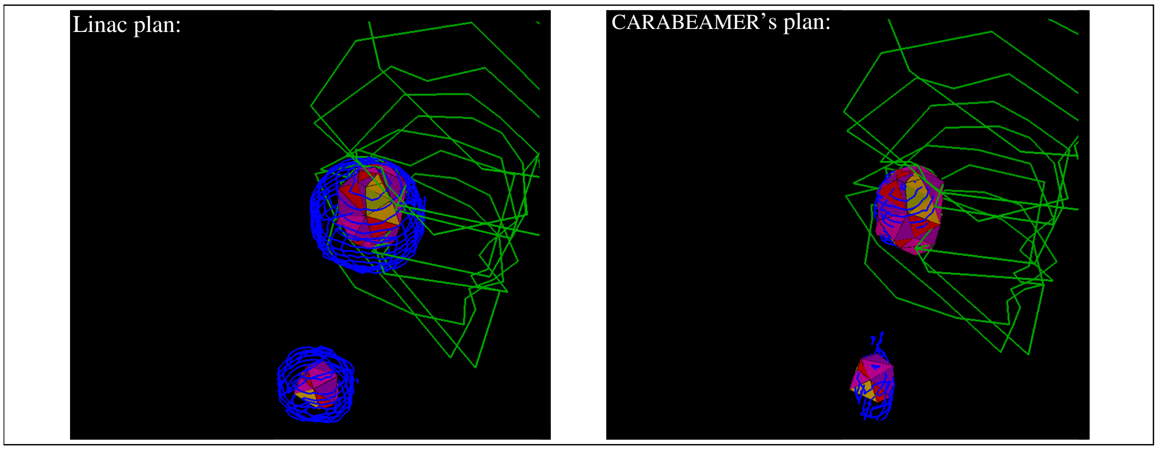

Figure 26: Patient case 12

Case Description		Neurosurgeon		CARABEAMER	
		min	max	min	max
Patient case 1:	lesion 1	1620	1980	1620	1980
	lesion 2	1620	1980	1620	1980
	external	–	–	0	2009
Patient case 2:	lesion 1	1890	2310	1890	2310
	lesion 2	1890	2310	1890	2310
	brain stem	0	800	0	800
	external	–	–	0	2310
Patient case 3:	lesion 1	1890	2310	1890	2310
	brain stem (within 2mm of tumor)	0	1890	530	1890
	brain stem	0	800	0	800
	external	–	–	0	2364
Patient case 4:	lesion 1	1620	1980	1620	1980
	right optic nerve	0	800	0	222
	left optic nerve	0	800	0	800
	optic chiasm	0	800	0	184
	cornea	0	800	0	203
	external	–	–	0	2204
Patient case 5:	lesion 1	1800	2200	1800	2200
	right optic nerve	0	800	0	800
	optic chiasm	0	800	155	800
	brain stem	0	800	0	314
	external	–	–	0	2218
Patient case 6:	lesion 1	1620	1980	1620	1980
	lesion 2	1620	1980	1620	1980
	right optic nerve	0	800	–	–
	left optic nerve	0	800	–	–
	optic chiasm	0	800	–	–
	external	–	–	0	2014

Table 1: Dosimetry results for radiosurgical-patient test cases

6.2 Experiments on Synthetic Data

In this series of experiments, we constructed a range of synthetic cases of varying difficulty, in order to explore the limits of the capabilities of CARABEAMER. We constructed 36 test cases of increasing difficulty relative to tissue geometry, dose constraints, and maximum number N_{max} of beams. In these experiments, CARABEAMER did not attempt to reduce the beam set further than N_{max}.

The 36 cases involve the four different tissue geometries shown in Figure 27, involving a bean-shaped tumor growing into a taurus-shaped one around a spherical critical region. Note that the radius of the critical regions is greater in Figures 27(c-d) than in Figures 27(a-b).

The following three sets of dose constraints were specified, where D_T and D_C denote the doses in the

Case Description		Neurosurgeon		CARABEAMER	
		min	max	min	max
Patient case 7:	lesion 1	1800	2200	1800	2200
	optic nerve/chiasm	0	600	0	300
	external	–	–	0	2200
Patient case 8:	lesion 1	2250	2750	2250	2750
	lesion 2	1800	2200	1800	2200
	brain stem	0	800	0	131
	external	–	–	0	2327
Patient case 9:	lesion 1	1620	1980	1620	1980
	lesion 2	1620	1980	1620	1980
	lesion 3	1620	1980	1620	1980
	external	–	–	0	2040
Patient case 10:	lesion 1	1620	1980	1620	1980
	lesion 2	1620	1980	1620	1980
	external	–	–	0	2028
Patient case 11:	lesion 1	1620	1980	1620	1980
	external	–	–	0	1878
Patient case 12:	lesion 1	1800	2200	1800	2200
	lesion 2	1800	2200	1800	2200
	brain stem	0	800	0	800
	brain stem (within 2mm of tumor)	0	1800	0	1800
	external	–	–	0	2259

Table 2: Dosimetry results for radiosurgical-patient test cases

tumor and the critical region, respectively:

$$\begin{array}{lccc} (1) & 2000 \leq D_T \leq 2400 & \text{and} & D_C \leq 500, \\ (2) & 2000 \leq D_T \leq 2200 & \text{and} & D_C \leq 500, \\ (3) & 2000 \leq D_T \leq 2100 & \text{and} & D_C \leq 500. \end{array}$$

Finally, we considered 3 maximum sizes for the beam set: 500, 250, and 100. The 36 cases were obtained by forming all combinations of the 4 geometries, the 3 constraint sets, and the 3 maximum beam set sizes.

Table 3 summarizes the performance results obtained. Each of the 4 major columns refer to one of the 4 geometrical cases. Each major column is further divided into two columns, the one on the left indicating the run time to obtain a solution, the one on the right giving the number of iterations of the planner (number of times ASSIGN-BEAM-WEIGHTS is invoked). The decomposition of the cases by dose constraints and size of beam set is shown horizontally. To account for the randomization in beam selection, the run times and numbers of iterations reported in the table are the average numbers over 10 runs, each with different seeds of the random number generator. The two cases marked with asterixes were run only once. The blanks indicate that for one of these cases CARABEAMER could not find a solution in less than 25,000 iterations. All cases were run on a Silicon Graphics Challenge

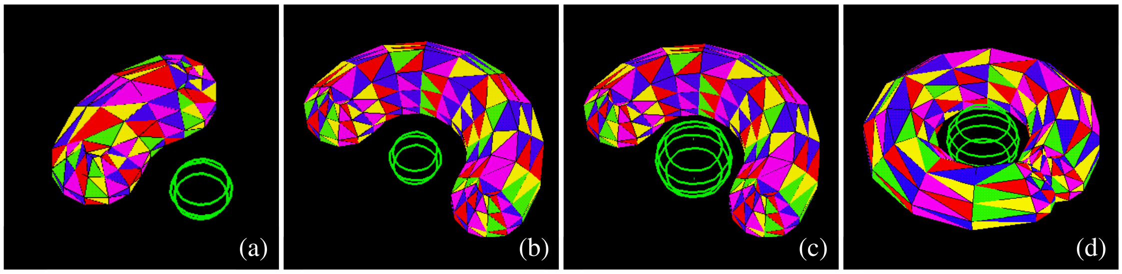

Figure 27: Tissue geometries for the synthetic test cases

Constraints	**Case 1**		**Case 2**		**Case 3**		**Case 4**	
	runtime	#iter.	runtime	#iter.	runtime	#iter.	runtime	#iter.
2000 – 2400								
N_{max} = 500	:41	1.4	:05:23	2.5	:06:45	2.7	1:40:55	3.9
N_{max} = 250	:40	1.3	:05:33	2.5	:07:19	3.0	1:44:18	4.0
N_{max} = 100	:51	2.1	:07:19	3.3	:07:06	3.1	1:41:19	40,9
2000 – 2200								
N_{max} = 500	:50	1.4	:08:37	2.7	:13:05	3.8	6:33:01	8.6
N_{max} = 250	:59	1.8	:08:42	2.8	:12:16	3.5	7:11:01	10.4
N_{max} = 100	:01:02	2.0	:10:43	7.4	:21:06	32.9	176:25:02*	14364.0*
2000 – 2100								
N_{max} = 500	:01:41	2.9	:27:39	3.4	1:03:06	12.9	44:11:04	34.0
N_{max} = 250	:01:31	2.6	:24:43	3.9	1:07:12	13.7	84:21:27*	426.1*
N_{max} = 100	:02:38	6.4	1:02:27	115.1	5:06:29	731.9	–*	–*

Table 3: Performance results for the 36 synthetic cases

XL, with 8 R10000 processors, 512MB of RAM, and 2MB of level 2 cache per processor. The run time was defined as the sum of the time that the CPU spends running instructions in CARABEAMER and the time that the operating system spends running on behalf of CARABEAMER. The results of Table 3 depend somewhat on the initial number of beams, but they are indicative of CARABEAMER's performance.

Figures 28 displays the 80% and 90% isodose surfaces for each of the tissue geometries when the tightest constraints on the dose distribution were used. In all cases, the dose distribution is skewed away from the critical region. That is, the falloff around the tumor is more rapid on the sides that are adjacent to the critical region. In all cases, the region of high dose matches the tumor volume extremely well.

7 Conclusion

We have described a new radiosurgical treatment planner, CARABEAMER, which makes it possible to exploit the kinematic generality of the CyberKnife system. Our experiments with this planner show that it generates satisfactory plans in complex cases. Obviously, longer clinical evaluation is needed

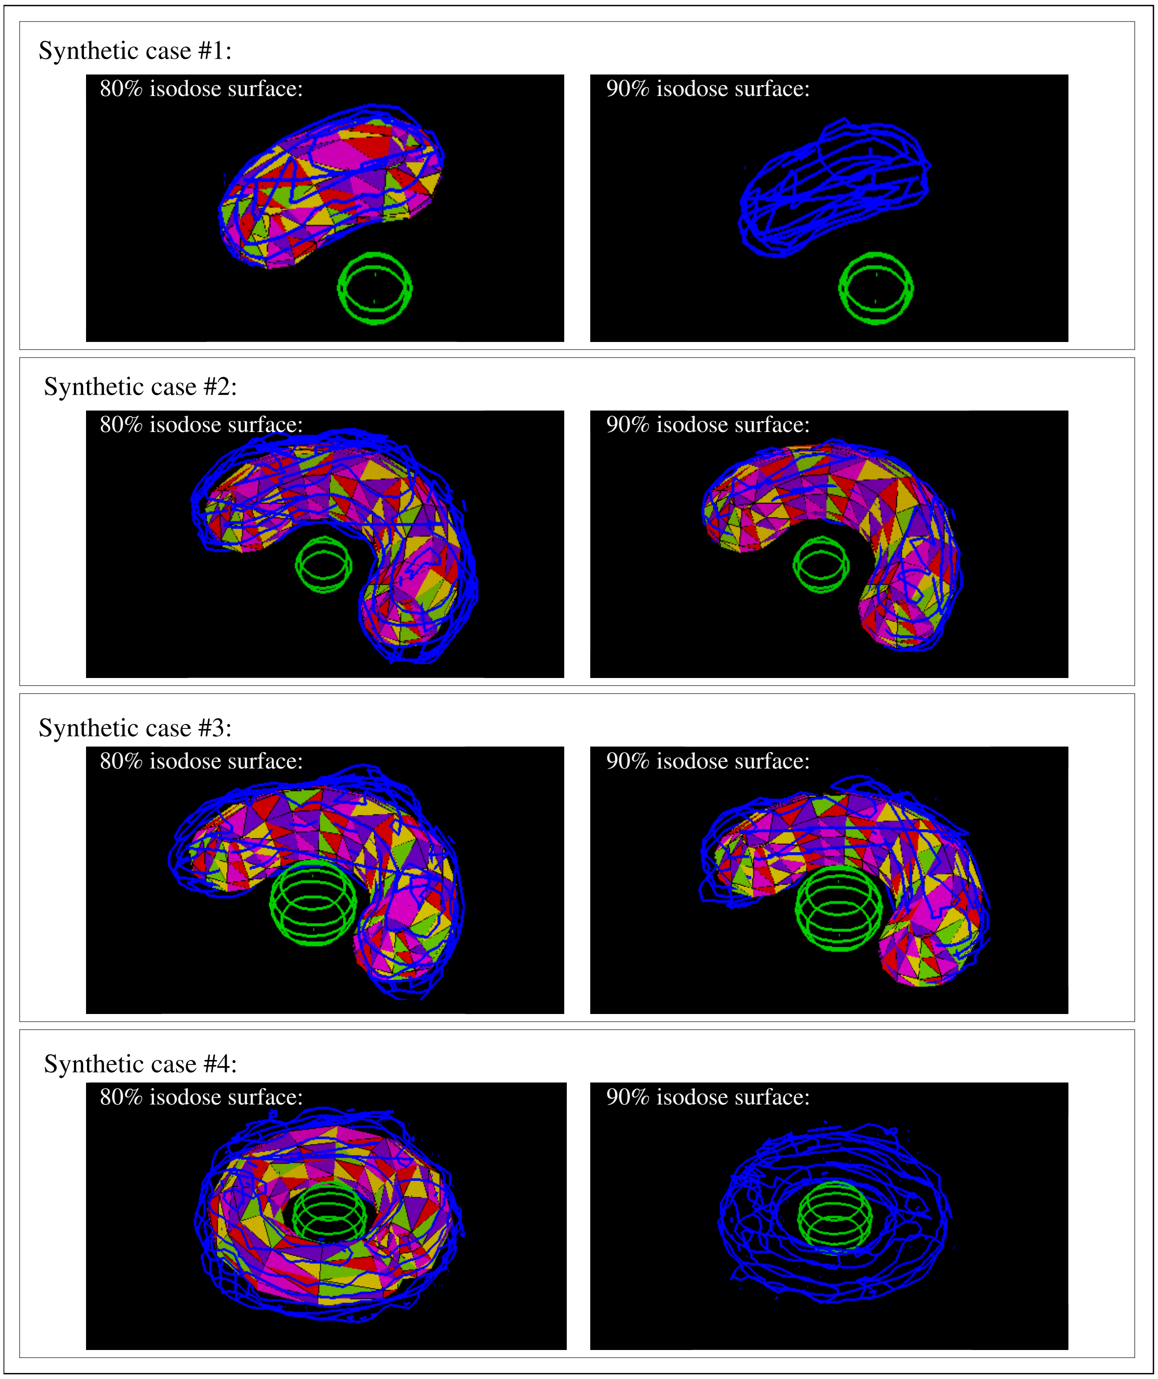

Figure 28: Dosimetry results for the synthetic test cases

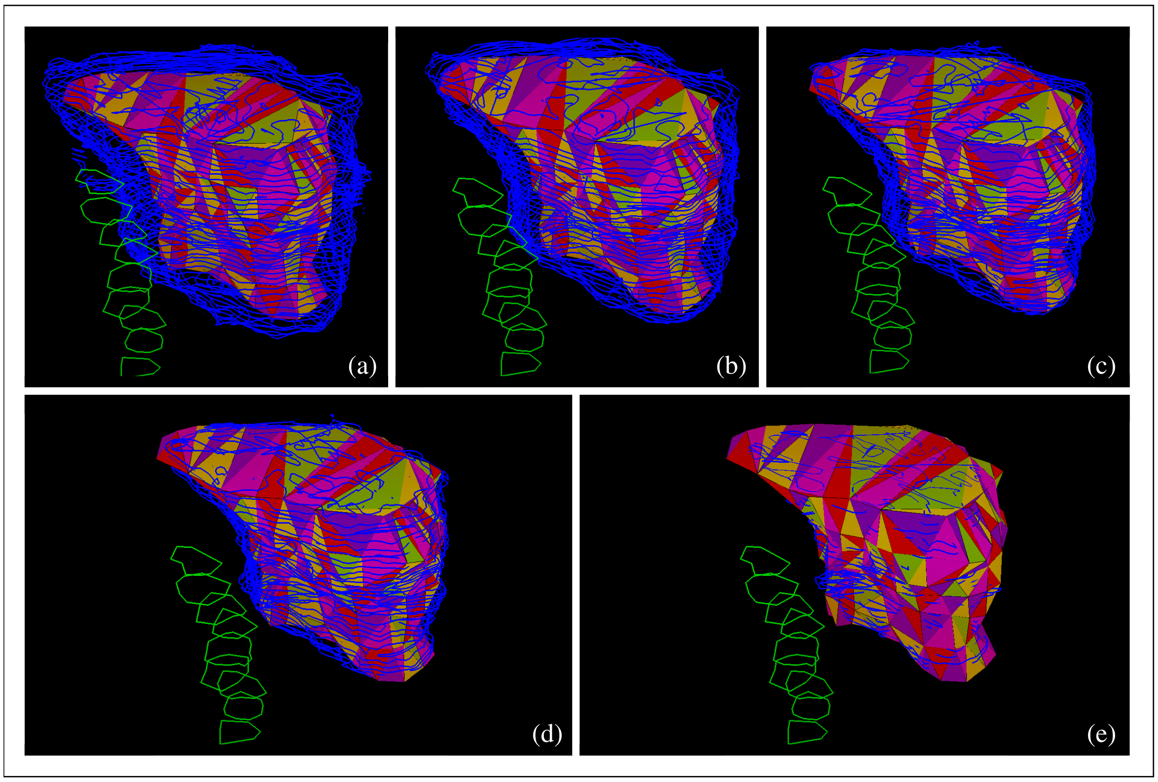

Figure 29: Dosimetry results for spine tumor case

to prove rigorously that the CyberKnife system, in conjunction with CARABEAMER, offers significant advantages over conventional systems, such as the LINAC. However, our results are extremely promising. In particular, they show that our planner can generate satisfactory plans for tumor geometries and placements that are clearly beyond what conventinal systems can treat.

Besides conducting further evaluation of the system, another goal of our current research is to explore the feasibility of applying CARABEAMER (and the CyberKnife system) to extracranial tumors. One major issue outside the brain is to reliably localize tissue structures. This localization problem can be dealt with for spine [18] and prostate tumors [37]. In those two cases, tissue of interest can be localized relative to nearby bone structures. This led us to experiment with CARABEAMER on several spine and prostate tumor cases. The challenge with spine tumors is that they tend to be relatively large; they also grow close to the spinal cord, an obvious critical structure. Figure 29 shows a patient spinal-tumor case and the 50%, 60%, 70%, 80%, and 90% isodose surfaces of the distribution produced by CARABEAMER. Falloff constraints were specified to ensure rapid falloff of the dose between the tumor and the spinal cord. The challenge with a prostate tumor is that it tends to grow between two large critical structures, the bladder and the rectum. Figure 30 shows a prostate case on which we run CARABEAMER. Notice how closely the tumor molds between the bladder (left) and the rectum (right). In that example, the neurosurgeon gave a higher priority to delivering a low dose in the critical regions than to achieving a homogeneous dose in the tumor. This led to placing the target points on a contracted representation of the tumor surface. The figure displays the 50%, 60%, 70%, and 80%. Figure 31 shows the dose-volume histograms in the three regions of interest. Such results are very promising and show that CARABEAMER and the CyberKnife may be effective tools to treat complicated spine and prostate tumors.

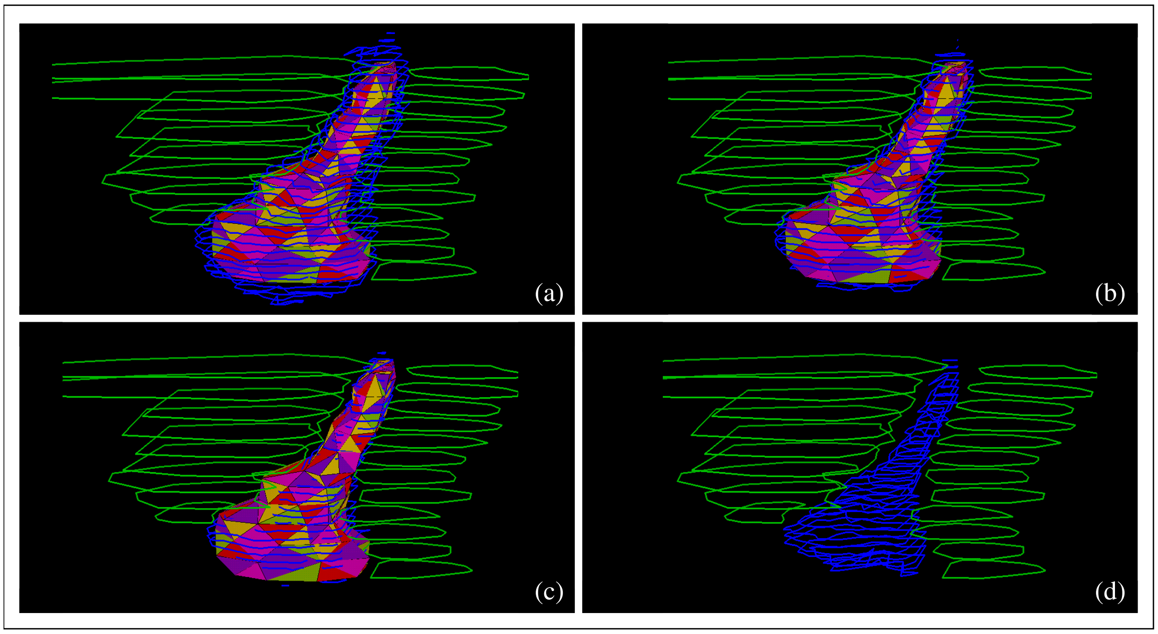

Figure 30: Dosimetry results for prostate tumor case

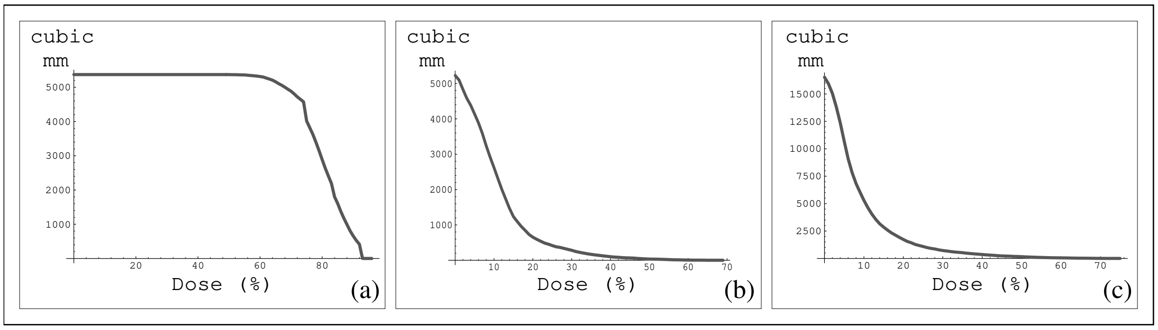

Figure 31: Dose histograms for prostate tumor case

CARABEAMER can still be improved in many ways. For example, the matrix of the linear inequalities generated by the procedure ASSIGN-BEAM-WEIGHTS is often abnormally dense in comparison to other real-world problems solved by linear programs. This increases the run time of the simplex method used by MINOS. It seems that the dual formulation of the linear-programming problem can often be solved more efficiently. Another improvement would be to allow collimators of adaptable shapes, e.g., multileaf collimators, to be used, rather than just circular collimators. By adjusting the shape of the collimator to the profile of the tumor along a beam direction, one could hope to produce better treatments with less beams. See [30] for some preliminary work on this topic.

Acknowledgments: R. Tombropoulos was funded by the NLM-NCI training program grant LM07033. J.C. Latombe was supported in part by MURI grant DAAH04-96-1-0007 (Army). This research was also supported in part by Alice & Stanley Schlickman Gift Fund and Arrillaga Family Gift Fund. We thank Achim Schweikard for his numerous contributions.

References

[1] J.R. Adler and R.S. Cox. Preliminary Clinical Experience with the CyberKnife: Image-Guided Stereotactic Radiosurgery. *Radiosurgery*, 316–326, 1995.

[2] J.R. Adler, A. Schweikard, R.Z. Tombropoulos, and J.C. Latombe. Image-Guided Robotic Radiosurgery. In *Proc. 1st Int. Symp. on Medical Robotics and Computer-Assisted Surgery*, 291–297, Pittsburgh, PA, Sept. 1994.

[3] R.Z. Tombropoulos, J.C. Latombe, and J.R. Adler. Inverse Treatment Planning for the Cyberknife. *Radiosurgery*, 2:236-250, Karger, Basel, 1998.

[4] G.K. Bahr, J.G. Kereiakes, H. Horwitz, R. Finney, J. Galvin, and K. Goode. The Method of Linear Programming Applied to Radiation Treatment Planning. *Radiology*, 91:686–693, 1968.

[5] J. Barraquand, L.E. Kavraki, J.C. Latombe, T.Y. Li, R. Motwani, and P. Raghavan "A Random Sampling Scheme for Path Planning," *The Int. J. of Robotics Research*, 16(6):628–649, Dec. 1997.

[6] N.H. Barth. An Inverse Problem in Radiation Therapy. *Int. J. Radiation Oncology Biol. Phys.*, 18(2):425–431, 1990.

[7] A. Brahme. Optimization of Stationary and Moving Beam Radiation Therapy Techniques. *Radiother. Oncol.*, 12:129–140, 1988.

[8] M.P. Carol. *Conformal Radiosurgery*, chapter 18, pages 249–265. Stereotactic Surgery and Radiosurgery. Medical Physics Publishing, Madison, Wisconsin, 1993.

[9] B. Chazelle, H. Edelsbrunner, L. Guibas, and M. Sharir. A singly-exponential stratification scheme for semi-algebraic varieties and its applications. *Lecture Notes in Computer Science*, Vol. 372, Springer, New York, NY, 1989.

[10] A.M. Cormack and R.A. Cormack. A Problem in Rotation Therapy with X-rays: Dose Distributions with an Axis of Symmetry. *Int. J. Radiation Oncology Biol. Phys.*, 12(12):1921–1925, 1987.

[11] J.C. Flickinger, L.D. Lunsford, A. Wu, A.H. Maitz, and A.M. Kalend. Treatment Planning for Gamma Knife Radiosurgery with Multiple Isocenters. *Int. J. Radiation Oncology Biol. Phys.*, 18(6):1495–1501, 1990.

[12] H. Fuchs, Z.M. Kedem, and S.P. Uselton. Optimal Surface Reconstruction from Planar Contours. *Comm. ACM*, 20(10):693–702, 1977.

[13] M. Goitein. The Inverse Problem. *Int. J. Radiation Oncology Biol. Phys.*, 18(2):489–491, 1990.

[14] L.J. Guibas, D. Halperin, H. Hirukawa, J.C. Latombe, and R.H. Wilson. Polyhedral Assembly Partitioning Using Maximally Covered Cells in Arrangements of Convex Polytopes. *Int. J. of Computational Geometry and its Application*, to appear.

[15] J. Jacky. 3-D Radiation Therapy Treatment Planning: Overview and Assessment. *Am. J. Clin. Oncol.*, 13(4):331–343. 1990.

[16] H.M. Kooy and N.H. Barth. The Verification of an Inverse Problem in Radation Therapy. *Int. J. Radiation Oncology Biol. Phys.*, 19(2):433–439, 1990.

[17] M. Langer and J. Leong. Optimization of Beam Weights under Dose-Volume Restrictions. *Int. J. Radiation Oncology Biol. Phys.*, 13(8):1255–1260, 1987.

[18] S. Lavallée, P. Sautot, J. Troccaz, P. Cinquin, and P. Merloz. Computer Assisted Spine Surgery: A Technique for Accurate transpedicular Screw Fixation Using CT Data and a 3-D Optical Localizer. In *Proc. 1st Int. Symp. Medical Robotics and Computer Assisted Surgery*, Vol. 2, 315–321, Pittsburgh, PA, Sept. 1994.

[19] W. Lutz, K.R. Winston, and N. Maleki. A System for Stereotactic Radiosurgery with a Linear Accelerator. *Int. J. Radiat. Oncol. Biol. Phys.*, 14:373–381, 1988.

[20] G. Luxton, G. Jozsef, and M.A. Astrahan. Algorithm for Dosimetry of Multiarc Linear-Accelerator Steretotactic Radiosurgery. *Med. Phys.*, 18(6):1211–1221, 1991.

[21] L. Ma, A. Chug-Bin, and R. Malyala. Dose Optimization in Conformation Therapy. *Med. Phys.*, 14(3):488, 1987.

[22] S.C. McDonald and P. Rubin. Optimization of External Beam Radiation Therapy. *Int. J. Radiation Oncology Biol. Phys.*, 2(3):307–317, 1977.

[23] R. Mohan, G.S. Mageras, B. Baldwin, L.J. Brewster, G.J. Kutcher, S. Leibel, C.M. Burman, C.C. Ling, and Z. Fuks. Clinically Relevant Optimization of 3-D Conformal Treatments. *Med. Phys.*, 19(4):933–944, 1992.

[24] S.M. Morrill, R.G. Lane, G. Jacobson, and I.I. Rosen. Treatment Planning Optimization using Constrained Simulated Annealing. *Phys. Med. Biol.*, 36(10):1341–1361, 1991.

[25] B.A. Murtagh and M.A. Saunders. *Minos 5.4 user's guide*, Tech. Rep. SOL 83-20R, Dept. of Operations Research, Stanford Univ., Stanford, CA, 1993.

[26] L.A. Nedzi, H. Kooy, E. Alexander, R.S. Gelman, and J.S. Loeffler. Variables Associated with the Development of Complications from Radiosurgery of Intracranial Tumors. *Int. J. Radiation Oncology Biol. Phys.*, 21(3):591–599, 1991.

[27] I.I. Rosen, R.G. Lane, S.M. Morrill, and J.A. Belli. Treament Plan Optimization using Linear Programming. *Med. Phys.*, 18(2):141–152, 1991.

[28] A. Schweikard, J.R. Adler, and J.C. Latombe. Motion Planning in Stereotaxic Radiosurgery. *IEEE Tr. on Robotics and Automation*, 9(6):764–774, 1993.

[29] A. Schweikard, M. Bodduluri, R.Z. Tombropoulos, and J.R. Adler. Planning, Calibration, and Collision Avoidance for Image-Guided Robotic Radiosurgery. In *Proc. Int. IEEE Int. Workshop on Intelligent Robots and Systems (IROS)*, 854–861, 1994.

[30] A. Schweikard, R.Z. Tombropoulos, and J.R. Adler. Robotic Radiosurgery with Beams of Adaptable Shapes. In *Proc. 1st Int. Conf. on Computer Vision, Virtual Reality and Robotics in Medicine (CVRMed'95)*, Lecture Notes in Computer Science No. 905, Springer, NY, 138–149, 1995.

[31] A. Schweikard, R.Z. Tombropoulos, L.E. Kavraki, J.R. Adler, and J.C. Latombe. Treatment Planning for a Radiosurgical System with General Kinematics. In *Proc. IEEE Int. Conf. on Robotics and Automation*, 1720–1727, 1994.

[32] S.W. Sloan and G.T. Houlsby. An Implementation of Watson's Algorithm for Computing 2-Dimensional Delaunay Triangulations. *Advanced Engineering Software*, 6(4):192–197, 1984.

[33] G. Starkschall. A Constrained Least-Squares Optimization Method for External Beam Radiation Therapy Treatment Planning. *Med. Phys.*, 11(5):659–665, 1984.

[34] R.Z. Tombropoulos. *Treatment Planning for Image-Guided Robotic Radiosurgery.* PhD Thesis, Dept. of Medical Informatics, Stanford Univ., Stanford, CA, March 1997.

[35] R.Z. Tombropoulos, J.C. Latombe, and J.R. Adler. A General Algorithm for Beam Selection in Radiosurgery. In *Preprints of the IARP Workshop on Medical Robotics*, 91–98, Vienna, Austria, 1996.

[36] R.Z. Tombropoulos, A. Schweikard, J.C. Latombe, and J.R. Adler. Treatment Planning for Image-Guided Robotics Radiosurgery. In *Proc. 1st Int. Conf. on Computer Vision, Virtual Reality and Robotics in Medicine (CVRMed'95)*, Lecture Notes in Computer Science No. 905, Springer, NY, 131–137, 1995.

[37] J. Troccaz. *Conformal External Radiotherapy of Prostatic Carcinoma: Requirements and Preliminary Results.* Tech. Rep. No. 9121, TIMC-IMAG, Faculté de Médecine, Grenoble France, 1993.

[38] G. Turk. Generating random points in triangles. In A.S. Glassner (ed.), *Graphic Gems*, pp. 24-28, Academic Press, San Diego, CA, 1990.

[39] G. Turk. Generating textures on arbitrary surfaces using reaction-diffusion. *Computer Graphics*, 25(4):289–298, 1991.

[40] G. Turk. Re-tiling polygonal surfaces. *Computer Graphics*, 26(2):55–64, 1992.

[41] S. Webb. Optimization by Simulated Annealing of three-dimensional Conformal Treatment Planning for Radiation Fields defined by a Multileaf Collimator. *Phys. Med. Biol.*, 36(9):1201–1226, 1991.

[42] K.R. Winston and W. Lutz. Linear Accelerator as a Neurosurgical Tool for Stereotactic Radiosurgery. *Neurosurgery*, 22(3):454–464, 1988.