submitted: IEEE Trans. Rob. Autom.

Planning for Camera-Guided Robotic Radiosurgery

Achim Schweikard [⋆†] Mohan Bodduluri [**] Rhea Tombropoulos [⋆ ‡] John R. Adler [†]
Jean-Claude Latombe [⋆]

[⋆] Robotics Laboratory, Department of Computer Science
[†] Department of Neurosurgery
[‡] Section on Medical Informatics, School of Medicine
Stanford University, Stanford, CA. 94305,
[**] Accuray Inc., Santa Clara, CA. 95054
{adler | as}@flamingo.stanford.edu

Abstract:
Radiosurgery is a treatment for brain tumors (and other lesions) which employs a moving beam of photon radiation. Classical radiosurgical systems rely on rigid skeletal fixation of the anatomic region to be treated. This fixation procedure is very painful for the patient and limits radiosurgical procedures to brain lesions. Furthermore, due to the necessity of rigid fixation, radiosurgical treatment with classical systems cannot be fractionated. A new camera-guided system capable of tracking patient motion during treatment has been built to overcome these problems. The radiation source is moved by a six degree-of-freedom robotic arm. In addition to offering a more cost effective, less invasive, and less painful treatment, the robotic gantry allows for arbitrary spatial motion of the radiation source. The system introduces a new class of radiosurgical procedures, termed non-stereotactic, *or* camera-guided *radiosurgery. At the heart of these procedures are algorithms for planning both a treatment and the corresponding beam motion, given the geometric description of the tumor shape and relative locations in the particular case. Results obtained from phantom studies and clinical use suggest that we can treat non-spherical lesions with accuracies unachievable with classical radiosurgical systems.*

Keywords: Medical Robotics, Geometric Reasoning, Motion Planning, Collision Avoidance, Image-Guided Surgery, Treatment Planning, Radiosurgery, Radiation Therapy, Medical Physics

1 Introduction

In radiosurgery, brain tumors (and other brain lesions) are treated by moving an intense beam of radiation [3, 12, 15]. A necrotic dose is delivered to a tumor by cross-firing from multiple directions, to reduce the amount of energy deposited in healthy brain. The radiosurgical treatment consists of several phases. First, a precise 3D map of the anatomical structures in the area of interest (for instance, the brain) is constructed using computed tomography (CT) or magnetic resonance (MR). Next a motion path for the radiation beam is determined to deliver an appropriate dose distribution (taking into account a variety of medical constraints). Finally, a jointed mechanism moves the radiation source according to this path.

A new radiosurgical system has been built to overcome limitations of earlier systems. The new system is based on an on-line x-ray vision system and a six degree-of-freedom manipulator moving a compact and light-weight linear accelerator.

1.1 System Overview

The system consists of the following components (fig. 1):

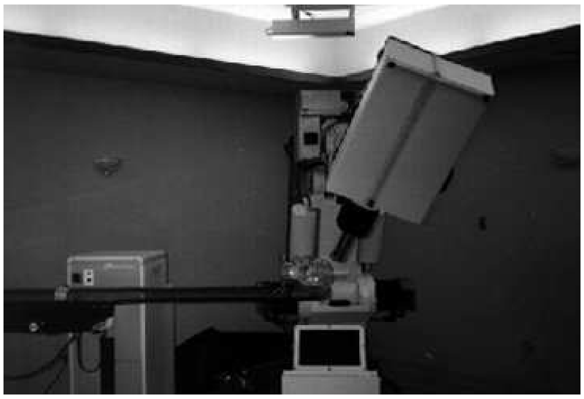

Figure 1: Robotic radiosurgical system, patient couch, x-ray camera system

Figure 2: For conventional radiosurgery, the patient's head must be immobilized with a frame via holes drilled directly into the skull. Image-guided radiosurgery obviates the need for rigid head fixation.

- *Radiation source:* A novel compact and light-weight linear accelerator running at high frequency generates a 6 MV (cylindrical) photon beam.
- *Mechanical gantry:* A GM-Fanuc 420 manipulator with six degrees-of-freedom moves the linear accelerator during treatment.
- *Static and Dynamic Imaging:* CT and MR scanners acquire pre-operative images of the brain. An online vision system with two orthogonal x-ray cameras takes images of the patient's skull at fixed time intervals during treatment.
- *Treatment planning:* Geometric planning algorithms compute a beam path given the constraints placed by the desired dose distribution. The beam path must observe additional constraints, i.e. the path must be collision-free and must not obstruct surveillance and on-line vision subsystems.

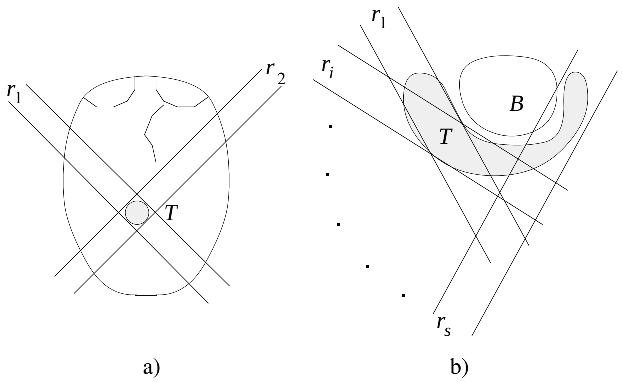

Figure 3: Cross-firing at a tumor

The treatment delivery proceeds as follows. Prior to treatment, a *collimator* (lead inset for beam shaping) of appropriate size is inserted into the radiation source. In our case the collimators are conical, i.e. the beam is a (nearly cylindrical) cone with very small conicity angle. The beam diameter (typically between 5mm and 40mm) remains constant throughout a given treatment. During treatment, the linear accelerator is moved in point-to-point mode through a series of configurations. At each configuration, the beam is activated for a small time interval, while the arm is held still. The duration of activation can vary from one configuration to the next. Radiographs are taken within the time intervals between beam activation. The vision system then computes the patient's current position by correlating the images to precomputed radiographs. Small movements of the patient are compensated for by a corresponding motion of the manipulator arm. In appropriate cases, the treatment is interrupted to change the collimator.

The new system offers the following advantages over earlier radiosurgical and radiotherapeutical systems (see e.g. [12, 15, 23]):

- The tissue does not have to be fixed in space. In earlier radiosurgical systems this was done with stereotaxic frames (fig. 2). Besides being very painful for the patient, this fixation was only possible for the lesions in the head, so the use of radiosurgical methods was limited.
- Treatment can be fractionated, i.e. the total dose can be delivered in a series of 2-30 treatments, where only a small dose is delivered during each treatment. Fractionation has turned out to be highly effective in radiation therapy, but could not be used in radiosurgery due to the necessity of rigidly fixing the tissue in space.
- Based on geometric planning algorithms the radiation dose can be focused within the lesion with high accuracy so that healthy tissue can be protected, and side effects of radiation can be reduced dramatically. In particular, our algorithms allow for generating dose distributions with prescribed non-spherical isodose surfaces and high homogeneity.

1.2 Geometric Constraints for Treatment Planning

We will first describe the main constraints arising in geometric treatment planning. The objective of treatment planning is to find a motion suitable for treating the particular condition presented by a patient. This planning problem is illustrated in Fig. 3. Fig. 3-a schematically depicts an axial cross-section of the brain with a circular tumor T. If T is irradiated from only one direction r_1 (with a cylinder beam), the healthy tissue along the beam absorbs approximately the same dose as the tumor. If, instead, we use two beam directions, r_1 and r_2, the dose deposited in the tumor is approximately twice the dose in healthy tissue. Using more beam directions can lead to further improvements of the dose distribution, and a very sharp drop-off of the dose in the tissue surrounding the (spherical) tumor region can be achieved. In this type of treatment, the axes of all beams cross a single point in space, called the *isocenter*. Notice that the region receiving high dose is spherical if the axes of all beams cross a single isocenter. However, for tumors of non-spherical shape this single-isocenter treatment procedure is problematic for the following reason: The (non-spherical) tumor region T in fig. 3-b surrounds a *critical* or particularly radiation-sensitive healthy region B. The dose in B should remain as small as possible. A single-isocenter treatment will generate a spherical treated volume. Covering T with a spherical region will lead to an over-radiation of B. The beams shown in fig. 3-b do not cross a single isocenter. However, the beam directions shown in the figure equally may not be adequate for this case, since this type of beam motion will accumulate dose along the upper (concave) boundary of T, leading to an undesirable inhomogeneity of dose inside T.

An appropriate dose distribution for a given tumor should satisfy the following constraints:

1. Critical and/or radiation-sensitive healthy structures near the tumor (i.e. nerves, blood vessels, sensitive brain structures) should receive very low dose.
2. For certain cases, the dose inside the tumor should be highly uniform, i.e. the dose in the tumor should not only be lower-bounded but also upper-bounded. Inhomogeneities, even inside the tumor can cause severe complications.
3. The *treated volume* (i.e. the region receiving $\geq 80\%$ of the maximal dose) should cover the entire tumor and should match the tumor shape.
4. The drop-off around the treated volume should be as sharp as possible, i.e. the (generalized) distance between the 80% isodose surface and isodose surfaces of lower values should be small.

For current radiosurgical systems a treatment plan is determined manually by the surgeon for each individual patient. In many cases, appropriate dose distributions require a large set of beam directions and/or complex motions of the radiation source.

2 Related Work

The most widely used hardware systems for radiosurgery are LINAC-based gantry systems (LINAC-systems, [3, 12, 15]) and Gamma-knife systems ([1]). For LINAC-based radiosurgery the beam is moved along circular arcs in space. In contrast, Gamma-knifes do not use beam motion. Instead, a large number of static beams are generated by appropriate radioactive sources. The spatial arrangement of the sources is such that all beam axes cross a single point in space.

The *direct dosimetry problem*, or forward planning problem is to compute the dose distribution in a tissue given a treatment plan. It relies on a model of fluency in tissue structures of different types, measured in phantom studies (see e.g. [15, 17]) The *inverse dosimetry problem*, also called inverse planning problem is to find a treatment plan whose execution will achieve a desired dose distribution.

The inverse problem in radiosurgery is related to the problem of finding treatment plans in radiation therapy. However, the technical constraints for the new robotic system differ substantially from those arising in both conventional radiation therapy and radiosurgery. Only simple motions of the radiation source can be executed with LINAC-systems. For conventional systems, the patient's head is immobilized by a stereotactic frame, and no patient motion is allowed during treatment. Thus pre-computed motion paths are not changed during treatment.

The new system has six degrees of freedom, and thus offers full kinematic flexibility for placing and directing the radiation beam in three dimensions. Furthermore, the new system is camera-guided. Thus patient motion is tracked by the radiation source, so that the pre-computed motion path of the robot may have to be adapted *during* treatment execution. The planned motion must be collision-free, and collision-avoidance must be performed dynamically.

The inverse planning problem in conventional radiosurgery and radiation therapy has been studied extensively in the literature (see e.g. [2, 5, 8, 16, 23]). Inverse planning can be regarded as an optimization problem, with constraints stemming from medical requirements, workspace geometry and gantry kinematics. In our case, fully computerized optimization of distributions is problematic and is not used for the following reasons:

- The constraints describing an adequate dose distribution are conflicting. As an example, in many cases it is difficult to achieve accurate shape matching (of tumor and treated volume) *and* sharp dose drop-off around the treated volume.
- The described planning problem is non-linear. In general, optimization methods will only yield a local optimum, with no information about global optima.
- The results of an optimization depend on settings of internal system parameters. Small changes in these values can yield largely diverging results. It is difficult to make the influence of such parameter settings and trade-offs transparent to the system user.

Instead, treatment planning is performed interactively by the surgeon, where each of the following three steps is carried out separately.

1. Selecting beam configurations
2. Computing beam weights, i.e. duration of activation at each configuration
3. Kinematics and path planning

The first planning step involves finding appropriate beam configurations, given the shape and the location of the tumor. New paradigms for beam selection have been developed to allow for full use of the new system's kinematic flexibility while addressing safety and verifiability requirements of the application.

The second step assigns weights to each of the selected beam configurations. Linear programming methods for computing such weights in radiation therapy have been proposed by several authors (see e.g. [16, 18]). We describe a framework for the application of linear programming in robotic radiosurgery, which relies on a standardized description of the constraints for a desired distribution. A particularly simple beam model allows for an efficient *ad hoc* reduction of the number of constraints.

In the third step, the selected beams are connected by a collision-free motion of the radiation source. When planning this motion, we must ensure that patient motion can be tracked during treatment.

Beam selection methods are described in the following section. Section 4 derives a constraint set from a standardized description of the input planning problem, given previously selected beam configurations. An algorithm for computing appropriate beam weights is then derived. Sections 5 and 6 describe methods for path planning, calibration and accuracy enhancement. The implementation of these methods as well as results obtained from phantom studies and clinical use are summarized in section 7.

3 Selecting Beam Configurations

Spherical Tumors: First assume the tumor is spherical, with center p and radius r. We set the radius of all beams to r. The beam configurations are then selected in the following way. The axes of all beams cross the tumor center p. (p is the isocenter.) We consider beam directions determined by vertices of a regular polyhedron centered at p. The choice of orientation of this polyhedron is problematic if there are critical healthy structures in the tumor vicinity. Furthermore there is no such regular polyhedron with more than 20 vertices. A more even distribution is obtained using more directions. The methods in [22] provide point grids on spheres which are adequate for our purposes. Our planning method thus uses a large set of beams (typically $n \geq 100$) for a single sphere, and addresses the presence of critical regions not during the beam selection phase, but during the beam weighting phase.

Ellipsoidal Treated Volume: While the scheme in the previous paragraph gives both very sharp dose drop-off around a sphere-shaped lesion and sufficient homogeneity, generalizations of this scheme for treating non-spherical tumors are not obvious.

A first attempt is to use two spheres instead of a single sphere, where the two spheres are treated independently. However, it is generally difficult to cover an arbitrary tumor shape with spheres, such that the volume of overlap with the tumor environment is small, and the entire shape is covered. This is aggravated by the fact that the two

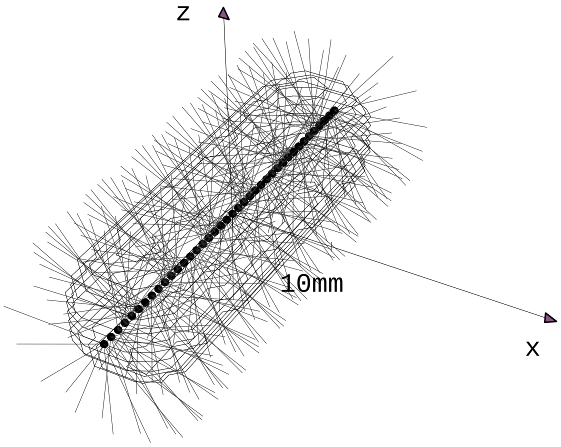

Figure 4: Isodose surface, beam directions. Region inside surface absorbs 50% or more of maximum dose

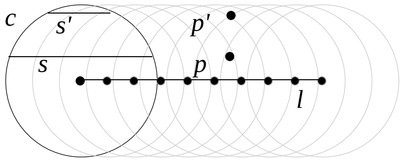

a)

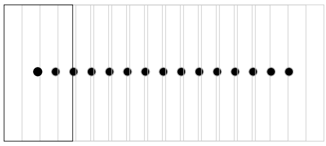

b)

Figure 5: Sweep-motion a) beam has circular cross-section b) Beam cross-section is rectangular

spheres should not intersect, since this would yield very inhomogeneous dose distributions, i.e. high dose (so-called hot spots) in the intersection region of the spheres and/or insufficient coverage of the tumor region. Studies on complication rates show that hot spots, even inside the tumor, should be avoided if at all possible.

However, a simple generalization of the above single sphere method is obtained in the following way. Instead of using a single isocenter, we place a series of isocenters evenly spaced along a line segment. Each such isocenter is treated in the same way as a single isocenter. Central to this approach is the fact that each sphere receives very low dose, and many spheres are overlaid. In this way we can avoid hot spots in the intersection of two spheres.

The dose distribution resulting from such a motion as well as isocenter points and beam directions for generating this distribution are shown in figure 4. The figure shows a 50% isodose surface, i.e. the region inside the surface receives 50% or more of the maximum dose.

There are several problems with this method:

1. If we use the same set of directions for all isocenters thus chosen, many of the beams will be parallel. Accumulated dose fall-off from such beams will cause dose inhomogeneity in the surrounding tissue.

2. Dose inhomogeneity in the tumor may occur. The dose in the tumor center will be much higher than the dose along the tumor boundary, as illustrated in example 1 below. This inhomogeneity is in large part due to the cylindrical beam shape.

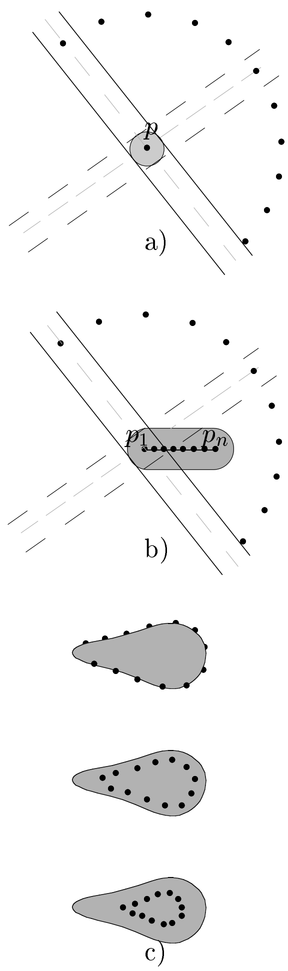

Figure 6: Beam selection paradigms. a) Beam axes cross a single (common) isocenter p, and treated volume is spherical. b) Isocenter points $p_1, \ldots, p_n$ are placed on a line segment, resulting in pseudo-ellipsoidal treated volume. c) General shape paradigm, based on a grid of isocenter points on the tumor surface.

To address the first problem, we choose rotated beam sets for each isocenter. In practice, rotating the entire beam set aimed at one isocenter by a fixed amount is equivalent to choosing each individual beam direction from a finer spherical point grid. Furthermore, instead of using several beams for each isocenter, we can use a larger set of isocenters, evenly spaced on a line segment, and a single beam (from a finer grid) for each isocenter. Thus we place a single beam through each isocenter point along the line segment. Each beam direction is then chosen (in random order) from a fine grid of points on the unit sphere.

The second problem can be addressed in several ways. The following example shows that the mentioned problem is related to the collimator shape.

Example 1: Fig. 5-a shows the cross-section of a cylindrical beam swept along a line segment l. The beam axis is orthogonal to the plane of cross-section. The beam cross-section is a circle c. The point p will absorb higher dose than point p' during this sweep-motion. The dose at these points will be proportional to the length of the segments s and s', thus causing substantial inhomogeneity.

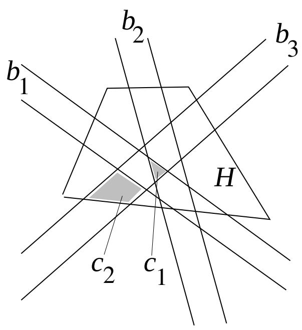

Figure 7: Arrangement with three beams, critical Region H. The beams determine a partitioning of space into cells. To upper-bound dose in the entire region H, we must only bound dose within the cell c_1.

A beam with rectangular cross-section avoids this problem. (figure 5 b).

Treated Volumes of General Shape: The above approach is restricted to the pseudo-cylindrical shapes shown in figure 4 and shapes obtained from bending such shapes. A direct generalization of this scheme is the following. With an algorithm in [22], we compute a grid of n points on the *tumor surface*. Each such point is now treated as an isocenter. As above, a single beam is placed through each isocenter. This scheme addresses both the first and the second of the mentioned problems: Accumulation along certain principal directions is avoided by using a fine grid. Accumulation of high dose in the tumor center is avoided by placing isocenters on the tumor surface. To achieve homogeneity, the point grid on the input surface can be expanded or retracted interactively. A path thus generated is refined by computing beam weights.

Fig. 6 summarizes the described beam selection paradigms.

4 Computing Beam Weights

During the second planning step, weights are assigned to each of the previously selected beams. The weight of each beam is its activation duration. Assume $b_1, \ldots, b_n$ are the beams selected in the first step. $b_1, \ldots, b_n$ are cylinders.

We must compute weights for each of the individual cylinders such that a set of global constraints on the distribution is satisfied.

Our method will assign a dose value $w_i \geq 0$ to every b_i. The values $w_1, \ldots, w_n$ determine a distribution D, defined as follows. If p is a point, and $b_{i_1}, \ldots, b_{i_k}$ are the cylinders containing p, then $D(p) = w_{i_1} + \cdots + w_{i_k}$.

We consider two disjoint regions T (for tumor) and H (for healthy tissue), with the following constraints: the dose delivered at each point in T must be larger than a fixed value α, while the dose at each point of H must be below β ($\beta < \alpha$). The n cylinders together with T and H define an arrangement of *cells* in space. Each cell is one region given by the intersection pattern of the beams and the boundaries of T and H. A cell is thus a maximally connected region not containing any points on the cylinder boundaries nor points on the boundaries of T and H. For each cell we compute a *label*. A cell c in cylinders $b_{i_1}, \ldots, b_{i_k}$ has label $l = \{i_1, \ldots, i_k\}$.

Example: Fig. 7 shows an arrangement with three beams b_1, b_2, b_3, and a critical region H. The partitioning induced by H and b_1, b_2, b_3 consists of several cells. Two cells are shown shaded. The labels of the cells c_1 and c_2 in this example are $\{1, 2, 3\}$ and $\{3\}$ respectively.

Assume, in the example in fig. 7, we wish to upper-bound the dose in H by a fixed value β. The cells c_1 and c_2 will give rise to inequalities

$$w_1 + w_2 + w_3 \leq \beta$$

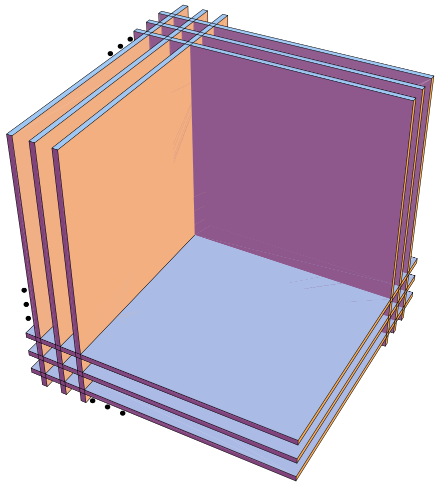

Figure 8: The number of maximal cells in an arrangement of n tiles (and cylinders) is $\Omega(n^3)$ in the worst case. The tiles are bounded by planes, with orientation pointing to the inside of the tiles. Maximal cells are cells in which three tiles intersect.

and

$$w_3 \leq \beta.$$

Similarly, we obtain an inequality from each of the remaining cells.

If we add a region T in the vicinity of H, then the resulting equation system will contain inequalities of the form

$$w_{i_1} + \cdots + w_{i_r} \leq \beta$$

and

$$w_{j_1} + \cdots + w_{j_s} \geq \alpha.$$

Specifically, the latter inequality corresponds to a cell in T with label $\{j_1, \ldots, j_s\}$.

The inequalities for all cells in T determine a convex polyhedral set P_α, in n dimensions. Similarly, the inequalities derived from the cells in H determine a polyhedron P_β. If P_α and P_β intersect, any point $(w_1, \ldots, w_n)$ in the intersection gives a dose distribution that satisfies the given constraints. Otherwise, the problem admits no solution. Thus we can use a linear programming feasibility test (simplex phase 1 computation) to decide whether the specified distribution can be achieved.

More generally, we can specify several healthy critical or non-critical regions $H_1, \ldots, H_q$ marked by distinct maximal doses $\beta_1, \ldots, \beta_q$. We then obtain polyhedra $P_\alpha, P_{\beta_1}, \ldots, P_{\beta_q}$. Any point in the intersection of these polyhedra determines a dose distribution that satisfies the input constraints.

Bounding the number of inequalities:

In the following, we will assume that the number of faces bounding T and H is small and fixed.

Fig. 8 shows that the number of inequalities in an arrangement of n cylinders in 3D is $\Omega(n^3)$ in the worst case. The figure shows n planar tiles, where n is divided into subsets of size $n/3$. The tiles in each subset are orthogonal to tiles in the other two sets. The resulting (plane) arrangement has $(\frac{2n}{3} + 1)^3$ cells. The example can be extended to cylinder arrangements by replacing the bounding planes of the tiles by cylinders with large diameter, so that the cylinder boundaries are almost planar in the area of interest. Conversely, an arrangement of hyperplanes in 3D has at most $O(n^3)$ cells (see e.g. [7]). The proof in [7] can be extended to the case of cylinders, and it can be shown that there are $\Theta(n^3)$ inequalities in the worst case.

However, we can reduce the number of inequalities by a large factor. Consider fig 7. In this case, we wish to upper-bound the dose in the region H. The selected beams induce a partitioning of H, where each cell in the partitioning gives rise to a constraint. It is easy to see that only the constraint in the cell c_1 must be considered. This is due to the fact that all other (upper-bounded) inequalities stemming from cells in H are implied by the constraint for this cell. Specifically, the constraint $w_3 \leq \beta$ is always satisfied if $w_1 + w_2 + w_3 \leq \beta$, since all weights w_i must be positive.

This simple observation will be generalized to reduce the number of constraints in the weight computation.

We consider inequalities from two cells c and c' in T. Let $L = \{i_1, \ldots, i_r\}$ and $L' = \{j_1, \ldots, j_s\}$ be the labels of c and c'. If $L \subset L'$, then the inequality given by c implies the inequality given by c', since all w_i are positive or null. A T-cell is called *minimal*, if its label is not a superset of any other T-label. Similarly, a cell in H_i is called *maximal*, if its label is not a subset of any other label in H_i. Thus, in the above feasibility test, we can remove non-minimal T-cells and non-maximal H-cells.

Fig. 8 shows that the number of maximal cells is not asymptotically lower than the total number of cells in the underlying arrangement. As noted above, there are $\Omega(n^3)$ cells in this arrangement in the worst case. A maximal cell in this case is an intersection of three tiles. The example shows that the number of maximal cells is equally $\Omega(n^3)$ in the worst case, and can be applied to cylinders as above.

Thus pre-computing minimal and maximal cells does not reduce the *asymptotic* number of equations in the worst case. However, in practice the reduction obtained by computing minimal and maximal cells is substantial. In the case of fig. 8, the number of inequalities is reduced by a factor of $(\frac{\frac{2n}{3}+1}{\frac{n}{3}})^3 \sim 8$.

Beam model:

In the above description it was assumed that each beam is a cylinder with constant fluency. A geometric beam model described in [19] allows a more accurate representation of the physical characteristics of photon beams. This model takes into account the decrease of dose with depth in tissue and distance from the beam's central axis, as well as the beam widening with distance from the source.

The above scheme for reducing the number of inequalities can be generalized to the refined beam model. However, experiments show that in our case the differences between the refined model and the simple cylinder model are small when compared with respect to the global distribution. This is partly due to the fact that a large number of beams is used. For the cylinder model, all coefficients in the inequalities are zero or one. Therefore we can compute a run-length encoding of labels. This encoding allows for rapidly detecting whether a label is a subset/superset of a label previously found. We thus use a cylinder beam model for the (inverse) weight computation, and verify the computed distribution with standard techniques for forward distribution calculation.

5 Kinematics and Collision Avoidance

The beam configurations selected by the planning system must be connected by a collision-free path of the linear accelerator. For safety reasons we use a hard-coded sequence of beam configurations, called *nodes*. A fixed motion, connecting the nodes, provides a path template. This path template is given in advance, and will not be changed during treatment planning. Beam configurations are always selected in such a way that their deviation from the hard-coded path template is small and upper-bounded. This is possible due to the following construction. For all (hard-coded) node configurations in the path template, the beam axis crosses the work space origin. A node is thus simply given by a the coordinates of the robot's tool center point. We use tool center points on a sphere of appropriate radius centered at the work space origin (fig. 9). Fig. 9-a shows the workspace with the base coordinate system at the upper end of the patient couch. To compute treatment beams given the particular tumor shape and location (section 3), we connect isocenter points to the tool center points of nodes. Thus each treatment beam is given by an isocenter point p and a tool center point v_i (fig. 10). Given appropriate bounds for the tumor size we can assume that all isocenter points p are close to the workspace origin. In contrast, the source-target

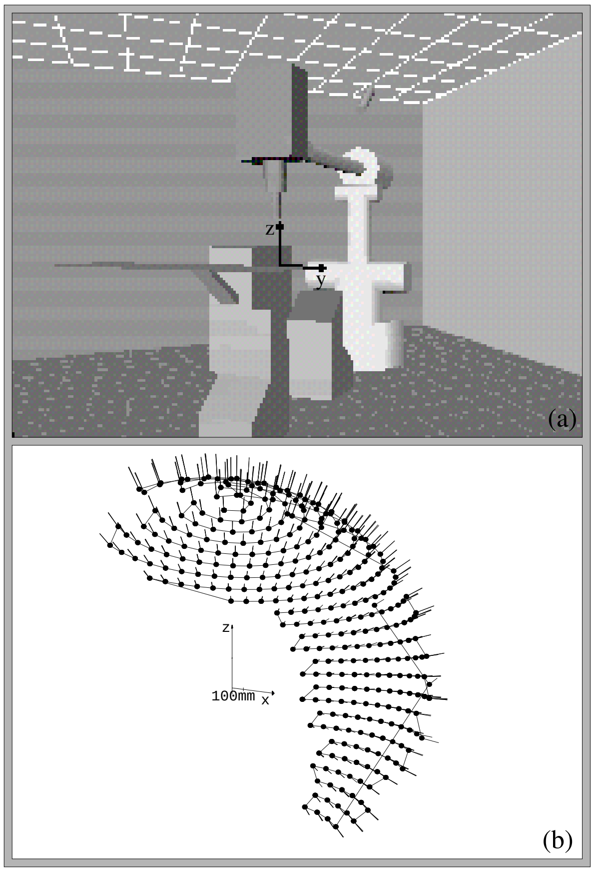

Figure 9: Path template connecting node configurations

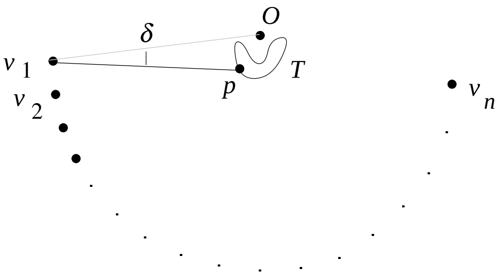

Figure 10: Origin O, target region T, tool-center points $v_1, \ldots, v_n$ for hard-coded nodes

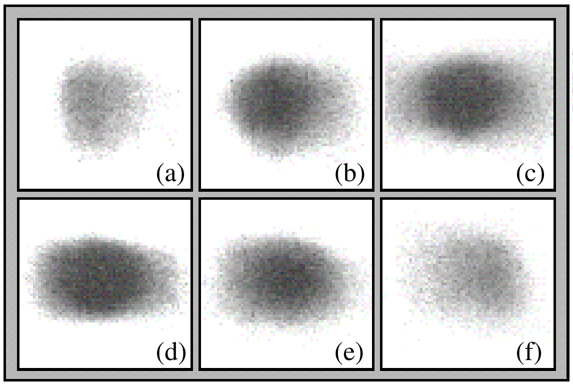

Figure 11: Photographic film phantom for spherical distribution (axial cross-sections)

distance for the beam (i.e. the distance between p and v_i) is large. Thus we can ensure that the angular deviation δ between a beam configuration and the corresponding node remains below a fixed threshold (figure 10). Given this threshold, we can compute a local joint space of appropriate size for each node configuration. Thus, angle limits within which the arm and the accelerator can be moved without colliding or getting into the line of sight of a camera are associated with each node. The angle limits are chosen such that small motions of the patient can be compensated without collisions.

The placement of a cylinder beam in space is kinematically redundant. The beam can be translated along and rotated about its axis without changing spatial placement. Thus a cylinder has four degrees of freedom. In contrast, the robot has six degrees of freedom. The equation system for computing the robot joint angles from a given beam placement is thus under-determined, i.e. valid solutions form a two-dimensional subspace of the six-dimensional joint space. However, practical considerations provide additional constraints. Firstly, the radiation source should remain at a fixed and constant distance from the workspace origin. This is due to the fact that the beam widens with depth in tissue. By remaining at an even distance from the origin, we can ensure that the beam diameter remains (nearly) constant within the target region. Secondly, the robot's joint angle ranges are limited. The range of beam directions should be as large as possible to allow for homogeneous distributions. This requires complex motions of the arm. For ill-chosen angles, the limitations in joint ranges will require frequent posture changes (i.e. changes between elbow-up and elbow-down configurations of the arm, etc.). We thus resolve the remaining kinematic redundancy in such a way, that posture changes along the template path are entirely avoided. This is done based on methods in [20].

For verification, the template path in fig. 9-b was executed and a photographic film in a phantom skull (Plexiglas reconstruction of a skull, filled with water) was exposed to the generated radiation. An equal dose value of 20 cGy was delivered at each of 300 nodes, resulting in an execution time of one hour. For each node, the beam axis aimed at the work space origin. The set of reachable beams does not cover a full sphere (fig. 9-b). The restrictions in the workspace result in a pen-umbra along the horizontal direction in fig. 11-c, -d.

6 Calibration

The positioning accuracy of the mechanical gantry is much lower than the repeating accuracy. The radiation source must remain at sufficient distance from the tissue to be irradiated. Small inaccuracies in the positioning of the radiation source are thus amplified and may cause inaccuracies in target localization. Thus the hard-coded motion path was calibrated in the following way. An optical crystal was placed at the work space origin. A laser beam pointing along the beam central axis was placed inside the radiation source. The crystal is activated by the laser beam. Activation intensity is proportional the distance between the laser beam and the crystal. Thus a distance measure can be read from the crystal, giving a measurement for placement accuracy. The calibration process considers each node separately. The end effector is placed at grid points in a small local joint space surrounding the node. At each grid point the laser is activated and a measurement is taken from the crystal. For the given node we can thus find joint angle offsets placing the beam in such a way that the beam axis crosses the origin. The angular deviation between treatment points and hard-coded nodes is small, and we use the offsets thus computed as offsets for the treatment beams as well. A bound of 0.3 mm for inaccuracies was determined experimentally for isocenter points close to the origin.

7 Planning System

The tomographic images show the area of interest in a series of parallel 2D cross-sections, called slices. To specify a desired dose distribution, we proceed as follows. The slices are displayed of the control interface screen. The surgeon delineates regions of interest (tumor, critical regions) on the screen in each individual slice. Dose constraints are specified for each region. To prescribe a desired homogeneity, we specify both upper and lower dose constraints for this region. The desired drop-off of dose around the target volume may vary in certain areas. We thus allow for setting individual upper and lower constraints for edges of a surrounding polygon.

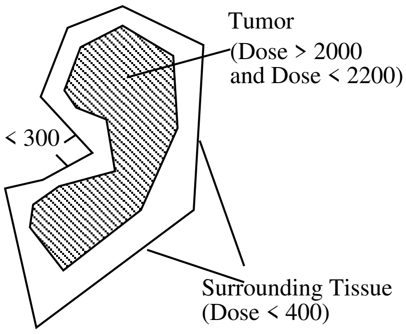

Figure 12: Specifying distribution characteristics. To specify desired homogeneity in the tumor, upper and lower dose bounds are set. To ensure shape matching and appropriate drop-off in surrounding tissue, dose constraints for edges of a surrounding polygon are entered.

Fig. 12 illustrates the specification of the constraints for the desired dose distribution. The tumor region (shown gray) is delineated by a polygon, and prescribed upper and lower dose bounds are entered. In addition, a surrounding polygon is entered, and upper dose constraints for the edges of the surrounding polygon are given.

The delineation is done manually, and results in a stack of polygons for each region. We convert each stack into a three-dimensional representation of this region, by thickening between slices. The radius r of the (cylindrical) beam and the number n of beam configurations are given as input. The beam selection method in section 3 is invoked to compute a set of beam directions corresponding to the input tumor shape and location.

The distribution is displayed and evaluated after completion of phase 1 (beam selection), where all beams are weighted evenly. We have found that the following criterion is particularly useful for deciding whether the selected set of beams is adequate. For evenly weighted beams, we test whether an appropriate isodose curve is contained entirely within the surrounding polygon specified by the user. Furthermore, the same isodose curve must entirely contain the target region. If this criterion fails in one of the slices, appropriate changes are suggested. Thus, the collimator size, the placement of isocenter points (fig. 6), or the shape of input polygons can be adapted interactively.

The second planning phase computes beam weights. We place a fine regular 3D grid of points across the brain map (here the grid resolution is 1mm). The label of a grid point is the set of cylinders containing this point. Labels are represented as bit strings. For each label, we compute the corresponding run-length encoding. This encoding allows for a rapid computation of subset and superset relations between labels. Prior to collecting the inequalities

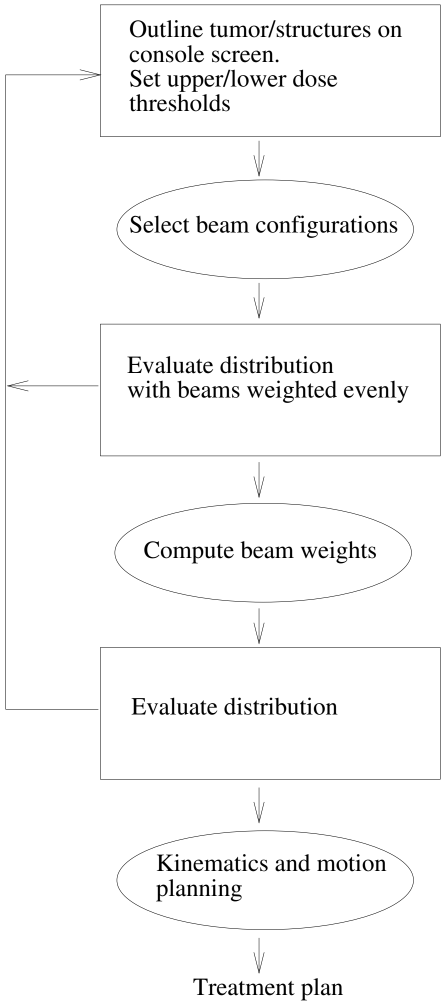

Figure 13: Planning steps

determined by the labels, all non-minimal labels for points in the tumor are removed. Similarly, labels of non-maximal cells in critical regions are removed. The resulting set of equations is then solved with the (simplex-based) MINOS linear programming system ([14]). Fig. 13 summarizes the planning steps.

Several paradigms for evaluating radiosurgical and radiotherapeutical dose distributions are in clinical use. The main paradigms are *dose-volume histograms* [11, 12, 15] and color representations of distributions in 2D overlaid upon tomographic images.

A dose-volume histogram (DVH) shows the volume (ordinate values) absorbing a dose x as a function of x (abscissa values). Dose-volume histograms are particularly useful for evaluating homogeneity of dose as well as matching between target region and treated volume. A separate DVH is computed for the tumor, critical tissue, and for tissue surrounding the tumor.

In addition, the system uses 3D isodose surfaces overlaid on reconstructions of anatomic regions, tissue-volume ratios [11], and integrals over dose-volume histograms. These paradigms refer to dose distributions computed in simulation. Calibrated photographic film phantoms (CFP) are used for verification. Film phantoms are processed electronically to validate distributions computed in simulation. Thus dose-volume histograms can be computed directly from actual radiation exposure and not from simulation.

In a test series, all cases treated during a comparison period with an earlier system for frame-based stereotactic

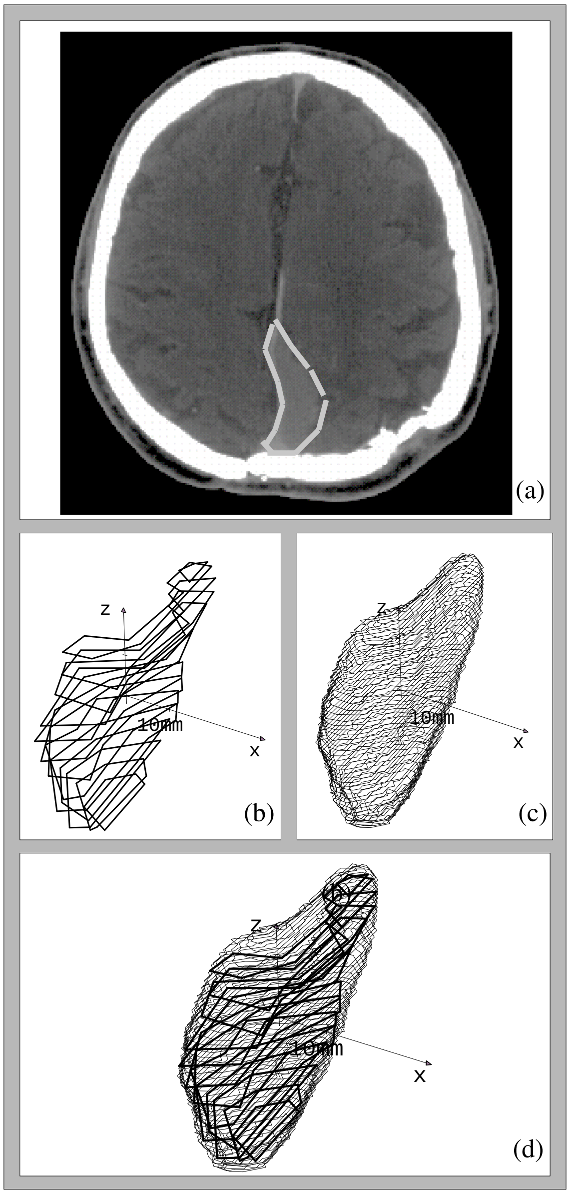

Figure 14: Sample case (Stanford Medical Center). (a) Axial CT with tumor delineated in black (b) 3D representation of tumor delineation (c) Isodose surface (d) Isodose surface overlaid on tumor reconstruction

radiosurgery (LINAC-system, [12]) were considered. In each case, data for the actual treatment were retrieved. Plans for the same cases were then computed for the new system. The comparison is based on dose-volume histograms, 2D and 3D dose visualization paradigms as well as photographic film phantom studies in water.

An example is shown in figures 14-16. Fig. 14-a shows an axial cross-section of the brain for a patient presenting a large non-spherical tumor of the falx cerebri, delineated in the cross-section. The case shown is one of the cases considered in the test series. Thus comparison data are available for the actual treatment performed with the LINAC-system. Fig. 14-b shows the 3D reconstruction of the tumor outline. A regular grid of 600 points was expanded on a surface reconstruction, giving isocenter points. A lower dose threshold of 2000 cGy and an upper threshold of 2200 cGy was prescribed for the tumor region. The total computing time for one planning cycle (selecting beam configurations and computing beam weights, but excluding final forward dose calculation) was 3

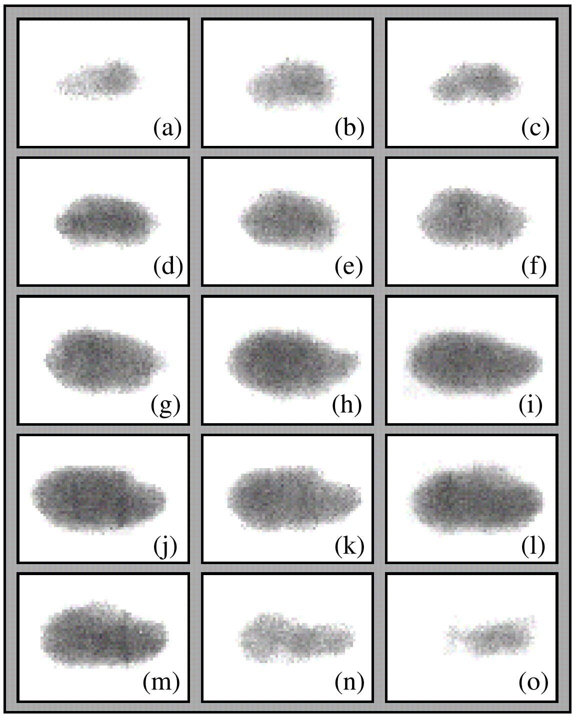

Figure 15: Film phantom for sample case in fig. 14, showing the shape of the region receiving high dose. Film slices in the phantom are arranged in the same spatial structure as the tomographic images.

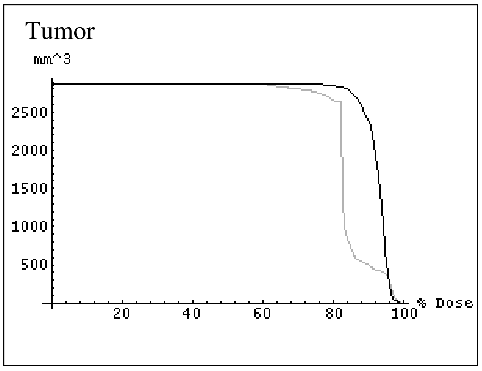

Figure 16: Comparison of dose-volume-histograms for sample case (tumor). Gray curve: conventional treatment; black curve: general shape planning paradigm.

minutes on the SGI-Indigo control interface workstation.

Fig. 14-c shows the 80% isodose surface resulting from the computed motion. Fig. 14-d overlays the input tumor shape with the generated isodose surface to illustrate the matching between input and output shapes. For verification, the motion was executed with the robotic system while exposing a photographic film phantom (figure 15). The phantom shows a close matching between the input tumor shape and the actual shape of the region receiving high radiation dose. The resulting distribution was then compared to the treatment actually performed in this case. Fig. 16 shows the dose-volume histograms for both cases. The black curve is the curve obtained for the new system, planned according to the general-shape beam-selection paradigm. For the conventional system, a plan with two isocenters was used, resulting in a significantly less homogeneous distribution (gray curve in fig. 16).

8 Conclusion

The described methods divide treatment planning into three steps: (1) selection of beam configurations. (2) computation of dose weights for these configurations. (3) kinematics and collision avoidance.

18 treatments have been performed thus far with the new system. These treatments were cases presenting tumors with spherical shape. Results obtained from the phantom studies suggest that the general shape planning paradigm is an alternative often preferable to conventional procedures. The new methods allow for applying radiosurgery to extra-cranial tumors, which is not possible with conventional radiosurgery, although highly desirable. It is likely that the new systems will allow for cost-effective radiosurgical treatment of other lesions. In particular, extensions are desirable for very radiation-resistant tumors and cases in which conventional treatment would cause too much damage to healthy tissue and is not applicable.

Acknowledgments: This research was funded in part by the Sheik Enany Fund, Lorraine Ulshafer Fund. The authors thank Joe Depp, Dan Halperin, Lydia Kavraki, Martin Murphy and Bo Preising for comments and suggestions.

References

[1] Becker, G.; Kortmann, R., Kaulich, T.W., Duffner, F., Bamberg, M. Gamma knife versus stereotactic linear accelerator. Utilization, clinical results and cost-benefit relations. *Radiologe* 36(4): 345-53, 1996.

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

[3] Betty, O. O., Munari, C., and Rosler, R. Stereotactic Radiosurgery with the Linear Accelerator: Treatment of Arteriovenous Malformations. *Neurosurgery*, 24(3):311–321, 1989.

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

[5] Brahme, A. Optimization of stationary and moving beam radiation therapy techniques. *Radiother. Oncol.*, 12:127-140, 1988.

[6] Chazelle, B., Edelsbrunner, H., Guibas, L., and Sharir. M. A singly-exponential stratification scheme for semi-algebraic varieties and its applications. In *Proc. 16th Int. Colloq. Automata Lang. Programm.*, Lecture Notes in Computer Science, Springer, 372:179–192, 1989.

[7] H. Edelsbrunner. *Algorithms in Combinatorial Geometry.* Springer, Heidelberg, 1987.

[8] Eklöf, A., Ahnesjö, A., Brahme, A. Photon beam energy deposition kernels for inverse radiotherapy planning. *Acta Oncologica*, 29:447-454, 1990.

[9] Halperin, D. and Yap, C.K. Combinatorial Complexity of Translating a Box in Polyhedral 3-Space *9th Annual ACM Symp. on Computational Geometry*, 29–37, 1993.

[10] ICRU Report 50: Prescribing, recording and reporting photon beam therapy. Technical report. Int. Comm. Radiation Units and Measurements, Bethesda, MD, 1992.

[11] Nedzi, L. A., Kooy, H. M., Alexander, E., et. al. Dynamic Field Shaping for Stereotactic Radiosurgery: A Modeling Study *Intern. J. Radiation Oncology Biol. Phys.*, 25:859–869, 1993.

[12] Lutz, W., Winston, K. R., and Maleki, N. A System for stereotactic radiosurgery with a linear accelerator. *Int. J. Radiation Oncology Biol. Phys.*, 14:373–381, 1988.

[13] J. J. More', S. J. Wright. Optimization Software Guide. *Frontiers in Applied Mathematics*, 14, SIAM, 1993. See also: "http://www.mcs.anl.gov/home/ otc/Guide/SoftwareGuide/".

[14] Murtagh, B. A., Saunders, M. A., Minos 5.4 user's guide. Technical Report SOL 83-20R, Dept. of Operations Research, Stanford University, 1983.

[15] Podgorsak, E. B., et al. Dynamic Stereotactic Radiosurgery. *Intern. J. Radiation Oncology Biol. Phys.*, 14:115–126, 1988.

[16] Renner, W. D., O'Connor, T. P., Bermudez, N. M. An algorithm for design of beam compensators. *Intern. J. Radiation Oncology Biol. Phys.*, 17: 227–34, 1989.

[17] Rice, R. K., et al. Measurements of Dose Distributions in Small Beams of 6 MeV X-Rays. *Phys. Med. Biol.*, 32:1087–1099, 1987.

[18] Rosen, I., Lane, R. G., Morrill, S. M., Belli, J. A. Treatment Plan Optimization using Linear Programming. *Medical Physics*, 18:141–152, 1991.

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

[20] Schweikard, A., Hommel, G., Berechnung erreichbarer Effektorstellungen für Handhabungsgeräte mit weniger als sechs Freiheitgraden. (Calculation of reachable effector configurations for manipulators with less than six degrees of freedom). *Robotersysteme*, 4, 177-182, 1988.

[21] Troccaz, J. et al. Conformal External Radiotherapy of Prostatic Carcinoma: Requirements and Preliminary Results, Rep. No. 912I, TIMC-IMAG, Faculté de Médecine, Grenoble, 1993.

[22] Turk, G. Re-Tiling Polygonal Surfaces. *Computer Graphics*, 26(3):55–64, 1992.

[23] Webb, S. The Physics of Three-dimensional Radiation Therapy, Institute of Physics Publishing, Bristol and Philadelphia: 1993.