Treatment Planning for Image-Guided Robotic Radiosurgery

Rhea Tombropoulos[1,2], Achim Schweikard[2,3], Jean-Claude Latombe[2], and John R. Adler[3]

[1] Section on Medical Informatics
[2] Department of Computer Science Robotics Laboratory
[3] Department of Neurosurgery
Stanford University, Stanford CA 94305, USA, rhea@flamingo.stanford.edu

Abstract. Radiosurgery is a non-invasive procedure that uses focused beams of radiation to destroy brain tumors. Treatment planning for radiosurgery involves determining a series of beam configurations that will destroy the tumor without damaging healthy tissue in the brain, particularly critical structures. A new image-guided robotic radiosurgical system has been developed at Stanford University in a joint project with Accuray, Inc. It has been in clinical use at Stanford since July, 1994, and thus far three patients have been treated with it. This system provides much more flexibility for treatment planning than do traditional radiosurgical systems. In order to take full advantage of this added flexibility, we have developed automatic methods for treatment planning. Our planner enables a surgeon to specify constraints interactively on the distribution of dose delivered and then to find a set of beam configurations that will satisfy these constraints. We provide a detailed description of our treatment planning algorithms and summarize our first experiences using the system in a clinical setting.

1 Introduction

In radiosurgery, a moving beam of radiation is used as an ablative surgical instrument to destroy brain tumors [2, 3]. The goal of radiosurgery is to deliver a necrotic dose of radiation to the tumor while minimizing the amount of radiation to healthy brain tissues, particularly critical or dose-sensitive tissues (such as the brain stem or optic nerves). This goal is usually accomplished by *cross-firing* at the tumor: rather than aiming one powerful beam of radiation directly at the tumor, which would destroy not only the tumor but everything on its path, several weaker beams are aimed from different directions. Though none of these weaker beams can destroy tissue by itself, the sum of the radiation in the regions where they intersect is severely lethal to tissues.

Treatment planning involves determining a series of beam configurations (position and orientation) such that the beams will intersect to form a region of high dose at the tumor. The dose distribution produced by the beams should match the shape of the tumor, with the entire tumor region receiving at least 80% of the maximum dose delivered to any one location. The fall-off of the dose around the tumor should be rapid, to minimize radiation to healthy tissues, and there should be little or no dose within the critical structures.

Recently, a new system capable of image-guided robotic radiosurgery was developed at Stanford University Medical Center in a joint project with Accuray, Inc. This system was developed to overcome many of the limitations of conventional radiosurgery [1, 7]. It uses a six-degree-of-freedom GMF robotic manipulator arm to position the radiation source and a real-time imaging system to monitor the patient's motion continuously. The system has been described previously [1]. In [7] we propose a new method for treatment planning that takes advantage of the flexibility offered by this new system and enables us to create irregularly shaped dose distributions that closely match the tumor shape.

We have continued to refine the treatment planning methods described in [7]. Our new planner gives the surgeon more power to define the final dose distribution: it allows the surgeon to specify particular regions of interest, such as the tumor and critical structures, and to specify the range of dose that each of these structures can receive. Our system uses linear programming to optimize the plans and satisfy the constraints specified by the surgeon whenever feasible. A fully integrated version of the system is now installed at Stanford University Medical Center and has been in use since July 1994. Three patients have been treated to date. This paper describes our treatment planning methods and our first experiences using the system in a clinical setting.

2 Treatment Planning

The treatment planning process is divided into several steps. First, the surgeon specifies regions of interest (e.g., the tumor and critical structures) and imposes constraints on the amount of radiation that these regions should receive. The system then uses this data to construct a three-dimensional representation of the necessary geometry and to select beams that will best approximate the tumor shape. Next, linear programming is used to refine plan in terms of the amount of radiation sent along each beam and the width of each beam such that the dose distribution produced will meet the specifications of the surgeon. Finally, the dose distribution is calculated by the system and evaluated by the surgeon. If the distribution is satisfactory, treatment is delivered. Otherwise, the planning system can backtrack to several different points in the process in order to improve the plan.

2.1 User Specifications and Geometry Reconstruction

The physician first views successive CT scans of the patient's brain. The tumor and any dose-sensitive critical regions are outlined on the images. Once the physician has outlined all of the regions of interest on the CTs, the system makes a three-dimensional reconstruction of the geometry. The surgeon then sets upper and lower bounds for these regions. For example, he might specify that the tumor must get a minimum 2000 rads but no more than 2200 rads, and that the brain stem should receive a maximum of 500 rads. These constraints are crucial in defining the distribution of the dose that will result, as they are later used to optimize the plan. For example, the narrower the range of possible values for the total dose in the tumor tissues, the more homogeneous the dose will be in that region.

2.2 Beam Selection

The beam selector must find a series of *beam configurations* such that the beams will intersect to form a region of high dose that closely matches the shape of the tumor. The configurations must be chosen and ordered in such a way that the robot can move from one to the next without colliding with objects in the workspace or obstructing the view of the imaging system. By beam configuration, we refer to the position and orientation of the radiation source producing the beam. Our system uses two points to define each beam configuration: the *source point*, where the beam originates, and the *target point*, at which the beam will be aimed. The source point corresponds to the position of the radiation source and, together with the target point, defines the orientation of the beam at that position. Our beam selector first finds a set of target points and then matches these with appropriate source points on a presynthesized path that the robot can traverse. The beam selector generates several hundred beams. However, not all of these beams will ultimately be used; during the plan-refinement step, non-useful beams are removed from the path.

Target Point Selection We have experimented with several different methods for selecting target points [1]. The method that has produced the best results involves selecting target points that are evenly spaced on the surface of the tumor. To approximate the tumor surface, we use the three-dimensional reconstruction generated from the surgeon's description of the tumor. Target points are evenly spaced on this polygonal surface using Turk's algorithm [8]: first points are randomly placed on the surface of the tumor; each point repels the other points with a force proportional to its distance from them; the points thus cause each other to spread apart; and, after a certain number of iterations, they stabilize at locations that are roughly equidistant from each other.

Source Point Selection Each target point must then be associated with a source point in order to define a beam configuration. Source points are selected from a presynthesized, well-tested template path. The path is comprised of several hundred points, or *nodes*. For each node, we precompute the range of possible joint angles for the robot. We select source points from the nodes in the following manner. For each target point, we define a beam passing through it at a random orientation. We then associate that beam with the nearest node in the presynthesized path, making sure the association is possible given joint constraints of the robot. That node becomes the source point for the given target point. We have experimented with more deterministic methods for orienting the beams at the target points [1] but have found that using random orientations leads to a more homogeneous dose distribution within the tumor.

Path Generation The final step in the beam selection process is to connect all of the beam configurations into a path that is feasible for the robot to traverse without colliding with objects in the workspace or obstructing the view of the imaging system. For safety reasons, however, we do no optimization in connecting the beam configurations. Instead, we simply use the template path used to determine the source points. The robot arm will move through each node in the path when executing the treatment, but will only stop and deliver radiation at the nodes where treatment beams have been defined.

2.3 Plan Refinement

The beam configurations are chosen by the beam selector such that they will intersect to form a region of high dose in the tumor. After beam selection, shape-matching is quite good: the 80% isodose surface produced by the selected beams closely matches the shape of the tumor. However, beam selection does not guarantee that all of the constraints specified by the surgeon will be satisfied: the beam selector does not consider the location of the critical tissues, nor does it guarantee a highly homogenous distribution of the dose within the tumor.

To satisfy these constraints, we refine the plan, using *linear programming* to adjust the weights and diameters of the individual beams once they have been selected. Linear programming is a mathematical technique for solving a set of linear equalities and inequalities subject to a linear objective function. It has been successfully applied to different aspects of the treatment planning problem in radiotherapy [4, 6]. We transform the constraints specified by the surgeon into linear inequalities involving beam weights. The inequalities are input into MINOS [5], a large-scale optimization program developed at Stanford University. If it is mathematically feasible, MINOS returns optimal weights for the beams such that all of the surgeon's constraints are satisfied. The weights are used to determine how much radiation to send along each beam and also what the optimal width for each beam is. This process is described below.

Beam Weighting We define our constraints in the following manner. In order to determine the weight of each beam, we divide the regions of interest into *cells* based on the arrangement of the beams inside them.

Cells are defined by the pattern of intersection of the beams in the regions of interest. Constraints are imposed on each cell such that the sum of the weights of the beams passing through each cell must be in the range specified by the surgeon for that region of interest. For example, a cell in the tumor intersected by beams 1, 3, and 4 would give rise to the following constraint:

$$TumorMin \leq w1 + w3 + w4 \leq TumorMax$$

where $TumorMin$ is the minimum amount of radiation the tumor should receive, $TumorMax$ is the maximum amount of radiation, and $w1 - w4$ are the weights of beams 1-4 respectively. For cells in the critical regions, constraints are usually one-sided, as there is no minimum. For example, the constraint introduced by a critical cell intersected by beams 2 and 3, is

$$w2 + w3 \leq CriticalMax$$

where $CriticalMax$ is the maximum amount of radiation that can be tolerated by this particular critical structure. It is important to note that all beam weights are constrained to be positive by the system, so negative beam weights (which are physically impossible) are never assigned.

Beams assigned zero weights are discarded. Because the linear programming assigns on average a weight of zero to 60% of the beams, this method provides an effective way of pruning our initial beam selection and identifying beams that are important for obtaining the desired dose distribution. By pruning the number of necessary treatment beams, we are able to decrease the time needed to deliver the radiation.

Collimator Selection We can augment this method to determine optimal diameters for each beam. As stated previously, the diameter of the beam depends on the collimator used. We can create beams with diameters ranging from 5 mm to 60 mm, along 2.5 mm increments. For any given treatment, however, we realistically cannot use more than four different beam diameters, as the collimator must be manually set in place. Hence, our goal is to find a maximum of four optimum collimator sizes to use for a given treatment plan.

Consider case where we must pick the optimal diameter for a beam from three possibilities: narrow, medium, and wide. Instead of representing the beam as a cylinder of fixed diameter, we represent it as a set of concentric cylinders. The innermost cylinder represents the smallest diameter, and each surrounding cylindrical tube corresponds to a wider diameter. Again, we analyze the arrangements of the cylinders in the regions of interest and form constraints based on these. We create additional constraints such that outer cylinders must receive weights that are less than or equal to those assigned to their respective inner cylinders.

Let ww correspond to the weight for the widest beam diameter, wm for the medium width, and wn for the narrowest. We have the constraint $ww \leq wm \leq wn$. If the linear program outputs that ww, wm, and wn are all equal, we use the widest diameter for that beam. If wn and wm are equal, but ww is substantially smaller, then we choose the medium width diameter. And if wn is substantially larger than both wm and ww, we choose the narrow diameter.

2.4 Dosimetry Calculations and Plan Evaluation

The dose distribution for a given set of beam configurations and weights is calculated and evaluated by the surgeon. The surgeon is presented with several different methods for visualizing this dose, including three-dimensional isodose surfaces, dose-volume histograms, color washes of the dose on the CT scans, and planar cross-sectional slices through the dose volume. If the surgeon is satisfied with the resulting dose distribution, treatment is delivered. If he is not satisfied, then planning can be restarted at several different points in the process. The surgeon can delineate new regions or alter constraints to specify a distribution that will be more satisfactory. Furthermore a new series of beam configurations can be generated.

One method of evaluating plans produced by our new planner is to compare them with the actual treatment plans used with the LINAC-System. Figure 1 shows such a comparison using dose-volume histograms. The black lines correspond to our new plans and the gray lines correspond to the old plans. In terms of both delivering a homogeneous dose to the tumor (Figure 1a) and minimizing the dose to the healthy tissues (Figure 1b), the planner for our new system performed much better.

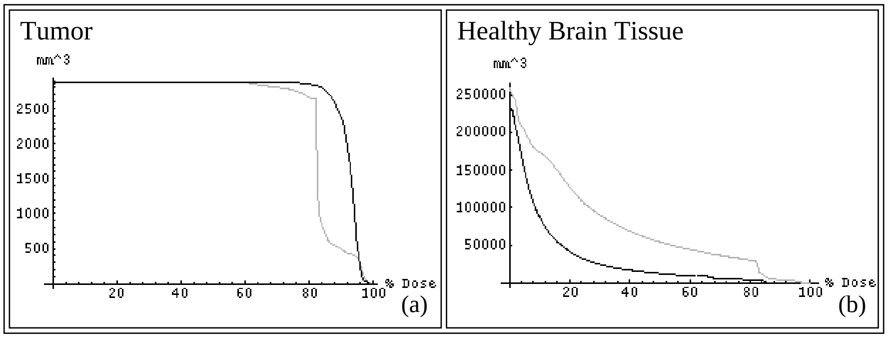

Fig. 1. Comparison of Dose-Volume Histograms

3 Clinical Use

Three patients have been treated with the system at Stanford Medical Center. In all cases, the tumor was reduced or was stabilized. Simple arc-shaped paths were used in planning for all of these cases so as to simulate treatments done with the LINAC-System and facilitate comparisons between the two. We hope to soon be able to use more complex paths in our planning system and make full use of its capabilities. All planning for these cases was done prior to the day of treatment. The robot was tested with each of the proposed paths before using it on the patient.

Treatment time for these first few cases averaged approximately 45 minutes. We expect this time to decrease as more patients are treated and we become more familiar with using the system. Even now, it is comparable to the duration of treatment with the LINAC-System, which usually takes between 30 and 45 minutes. Furthermore with this new system, the patient was able to walk into the clinic in the morning, receive treatment, and leave approximately 1 hour later. Fractionation of radiosurgical treatments will clearly be possible with this system.

4 Conclusion

A new system for frameless, robotic radiosurgery has been developed. This system is installed at Stanford University Medical Center and has been used to treat three patients. This system overcomes many of the limitations of traditional stereotaxic radiosurgery. By eliminating the need for a frame, it not only provides increased patient comfort, but also the ability to fractionate treatments and to treat extra-cranial tumors. By using a six-degree-of-freedom robotic arm to position the source of radiation, the system enables us to deliver highly accurate treatments and to create irregularly shaped dose distributions that conform closely to the shape of the tumor. In order to be able to fully utilize the advantages offered by this new system, we have implemented an automatic planner. This planner allows the surgeon to set constraints on the dose distribution produced and to find a set of beam configurations that will produce a dose distribution that satisfies these constraints. We hope the new features and abilities of our system will lead to new directions for this field and cures for previously untreatable tumors.

References

1. Adler, J. R., Schweikard, A., Tombropoulos, R., and Latombe, J. C., Image-Guided Robotic Radiosurgery. *Proc. First International Symposium on Medical Robotics and Computer Assisted Surgery*, 2:291-297, 1994.
2. Betty, O. O., Munari, C., and Rosler, R. Stereotactic Radiosurgery with the Linear Accelerator: Treatment of Arteriovenous Malformations. *Neurosurgery*, 24(3):311-321, 1989.
3. Larsson, B. et. al. The High Energy Proton Beam as a Neurosurgical Tool. *Nature*, 182:1222-1223, 1958.
4. Morrill, S. M., Lane R. G., Wong, J. A., and Rosen, I. I., Dose-Volume Considerations with Linear Programming Optimization. *Medical Physics*, 18(6):1201-1210, 1991.
5. Murtagh, B. A. and Saunders, M. A., *MINOS 5.4 User's Guide*, Technical Report SOL 83-20R, Dept. of Operations Research, Stanford University, March, 1993.
6. Rosen, Issac I., Lane, Richard G., Morrill, Steven M., and Belli, James A., Treatment Plan Optimization Using Linear Programming. *Medical Physics*, 18(2):141-152, 1991.
7. Schweikard, A., Tombropoulos, R., 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*, 1720-1727, 1994.
8. Turk, G., Re-Tiling Polygonal Surfaces, *Computer Graphics*, 26(3):55-64, 1992.

This article was processed using the LaTeX macro package with LLNCS style