Trinity Leave Proposal
Leave Dates: Spring 2007
Allen Holder - Mathematics

Introduction

Cancer is one of the leading causes of death among all demographic groups and is typically treated with a combination of surgical procedures, chemotherapy, and radiotherapy. Radiotherapy is a non-invasive treatment that focuses high-energy beams of ionizing radiation on targeted tissues, with the goal being to deposit enough radiation to kill cancerous cells while at the same time sparing surrounding organs. The treatment is based on the fact that cancerous cells are in a heightened state of reproduction, a state that does not allow them to recover from damaged DNA. Through direct and secondary interactions (due to the formation of free radicals) a cell's DNA is damaged as radiation passes through it. If the damage is not sever, a healthy cell can continue its normal cycle whereas cancerous cells can not. The intent of treatment design is to decide the angles at which radiation will travel through the anatomy and the exposure time for each of these angles.

Radiotherapy is delivered by immobilizing a patient on a couch (a movable horizontal table) and adjusting the position of the couch and gantry. The gantry focuses radiation and is capable of rotating around the patient in a great circle. Once the patient and gantry are properly aligned, the patient is exposed to the beam of radiation. The beam is focused with sub-millimeter precision, making the largest unknown factor the random adjustment in patient position due to such anatomical processes as breathing. There are several types of radiotherapy, with one of the more prominent techniques being Intensity Modulated Radiotherapy (IMRT). In this paradigm, the radiation beam is modulated by a multi-leaf collimator in the head of the gantry. The modulation allows the large rectangular beam to be divided into smaller sub-beams, often called pencils, and a patient can have different exposure times for each of these sub-beams. The exposure pattern over these sub-beams is called a fluency pattern, and Bahr [**?**] suggested in 1968 the use of an automated optimization process to decide fluency patterns (this initial research did not consider modulation because it had not yet been invented). Since each sub-beam in IMRT has a variable exposure time, the number of variables to decide during the design process is huge, and the complication of selecting such a large number of quality exposures makes it imperative to have an automated process to 'optimize' the fluency pattern. Today all commercial software packages use an optimization technique to decide a fluency.

Promising angles are pre-selected through the use of visualization software that allows a treatment planner to visualize a beam's path. Modern technology permits a treatment to contain from 7 to 14 angles. In general, treatment quality increases as the number of angles grows, but the number of angles is limited to reduce treatment time so that patient movement is restricted. If a patient's position is altered during treatment, the probability of success is likely to decrease significantly. Once the angles are selected, an optimization routine calculates a fluency and the anticipated dose delivered to the patient is estimated. The treatment is then judged by several metrics and if accepted, the planning is complete. The preponderance of initial treatments is unsatisfactory, forcing the planner to update the angle collection and re-calculate a fluency. This iterative process repeats until an acceptable treatment is designed, a process that typically requires 4 hours per patient.

As stated earlier, the idea of fluency optimization was introduced in 1968. Medical physicists have since developed 'optimization processes' that produce fluencies. However, all of these methods are variants of a stochastic process called simulated annealing, which is a mathematical algorithm based on the annealing process. Without going into details, simulated annealing depends on a cooling structure that mathematically controls how the search space is sampled. These mathematically elegant algorithms have a striking result, namely that there is a cooling structure that allows this method to solve any optimization problem provided that sampling is allowed to continue forever. Moreover, the algorithm is simple to implement, making it a natural fit for difficult real-world problems. However, the cooling structure is problem dependent and impossible to predict. So, while these techniques are mathematically capable of solving any optimization problem, they fall short of implementations that guarantee convergence to an optimal, or even a good, solution.

Applied mathematicians studying optimization became interested in the fluency problem in the late 1990s, bringing with them sophisticated modeling and solution techniques. In particular, the mathematicians were able to calculate optimal fluencies with algorithms that terminated in finite time with a guaranteed optimal solution. The guarantee of optimality was one of the significant contributions, since the medical physicist no longer had to wonder if there was a better treatment. Such certificates of optimality require a specific form (such as linear or quadratic), but this was not a hindrance because the typical models that were already being 'solved' by simulated annealing had the required form. So, the first mathematical advances were algorithms that, unlike simulated annealing, converged to a true optimal fluency in finite time.

The mathematicians are now beginning to investigate the plethora of problems outside of fluency optimization. Some have looked at how to best adjust the leaves of the collimator to achieve the desired fluency. The basic question is how do we deliver a treatment to achieve the highest degree of success? Others have looked at how to optimize the pre-selection of beams to alleviate planners from the tedious design process. Such questions are not incorporated in the current optimization models that decide fluencies, and this separation is not natural. Incorporating these concerns so that the entire design process is optimizated is a future goal. However, while it is simple mathematically to express an encompassing model, the result is far beyond are current ability to optimize. The problems are simply too large.

There are now numerous models and solution procedures suggested in the literature, and head-to-head comparisons are now needed to see which methods provide the most clinical benefit. It is important to stress **clinical** benefit, as the majority of mathematical work has been accomplished in academia in a collaborative effort with a clinic. The problems solved in academia are based on clinical data but are often simplified to accommodate the academic setting. Moreover, whatever numerical experiments are undertaken are not reproducible since the problems and implementations are developed within a research group and are not publicly available. The research field has two obstacles, 1) the lack of an academic treatment system, a system that allows users to alter the solution methodology and model parameters, and 2) the lack of a problem library accessible to the research community.

Current Research Directions

I began applying optimization techniques to radiotherapy design in 1997-98, and my first major contribution was a model and solution procedure for fluency optimization that guaranteed an adherence to a physician's prescription. Prior to treatment design, a physician delineates structures on a series of patient images and limits the radiation to be delivered to the identified regions. The prescribed limits are upper bounds on critical structures and lower bounds on cancerous tissues. These bounds constrain the optimization process and are often not simultaneously achievable. The model and solution procedure developed in [**?**] shows that it is possible to find an optimal fluency that satisfies the constraints in a pre-ordained manner. For example, it is possible to guarantee the lower bound placed on the target before satisfying restrictions on the healthy organs. The important contribution is that we can do this in one solve instead of iteratively, saving time and improving numerical stability. This research was awarded the 2000 Pierskalla prize as the best annual research contribution in health care. Following the award, the investigator was invited to author two tutorials on optimizing radiotherapy treatments.

In collaboration with a Trinity undergraduate, he adapted his original work to photodynamic therapy, a treatment that destroys tissue through a chemical reaction that occurs in the presence of inferred light. This research addresses the question of whether or not it is possible to treat a deep tissue tumor with a current photosensitizer (the chemical compounds that reacts to inferred light). This numerical work showed that even in the best of situations the chemical attributes of an FDA approved photosensitizer are incapable of treating targets more than a few millimeters from the surface of the anatomy. Photodynamic therapy will not be a general treatment option until the biochemists improve photosensitizers.

The majority of the my current research is divided among the following,

Beam Selection: How do we automate the pre-selection of angles so that patient quality adheres to a minimum standard?

Academic Treatment System: I have directed 7 Trinity students over the last two summers as part of the Summer Undergraduate Research Fellowship (SURF) program. The primary goal was to build an academic treatment system that was accessible to the research community through a web interface.

A Problem Library: The treatment system above interacts with a problem library that is open to all, allowing for head-to-head comparisons.

The complete details of this work are beyond the scope of this proposal, but the fact that much of this work stemmed from undergraduates is worth mentioning. I attempted to solve the Beam Selection problem for 4 years with only the most limited of successes (several colleagues were also unsuccessful). The problem appeared to be intractable, one that many in the research community, both mathematicians and medical physicists, believed would not be adequately solved. I began describing this work to a few students in 2000-01 and was approached in 2003 by Josh Reese, who suggested that beam selection was similar to image compression. I did not give this much thought at the time but a year latter decided to investigate the possible link. After a few months of research, I discovered that Josh was correct and that beam selection could be modeled as a data compression technique called vector quantization. The initial numerical work was promising, and I am convinced that vector quantization will successfully solve the beam selection problem. Josh did not stop with this discovery, and his senior project shows how beam selection and vector quantizer are related to a well studied problem in graph theory called the p-median problem. The importance of this relationship is that it brings a sizeable literature from the p-median problem to beam selection. We are now in a position to identify which of the numerous heuristics from graph theory are best suited to select quality beams. Our combined efforts will generate 3 papers over the next year on these topics.

The other 6 SURF students primarily worked on the software engineering aspects of developing an academic treatment system and the related problem library. This has been a monumental task, one that I could not possibly understand at the outset 2 years ago. Every day of the program, every single day, there have been computational and modeling hurdles to overcome. These students have not only overcome these obstacles but have developed elegant solutions that streamline the numerical procedures. We are collectively working on a paper that discloses our techniques, making our system the first that is reproducible (and open source).

Goals and Timeline

There are two primary goals of the leave. First, there is insufficient crossover between the mathematical and medical physics communities. Small steps have been accomplished with both communities beginning to cite each other. However, neither group is publishing in the others journals. I became an adjunct member of the Department of Radiology at the University of Texas Health Science Center at San Antonio in 2003, and I plan on spending the semester at the Cancer Therapy and Research Center (CTRC). I already have office space and computer resources, and my colleague, Dr. Bill Salter - Director of Physics, is looking forward to assisting us in transforming the academic work from Trinity into clinical efforts that are appropriate for publication in medical physics journals. This is an important step, for publishing in these journals will highlight the advances permitted by the mathematical efforts to the community that these efforts were directed. In particular, our academic software will be linked to patient data, allowing us to experiment with parameters unavailable in commercial systems.

The second goal of the leave is to cement the relationship between mathematical optimization and radiotherapy design by co-authoring a monograph on the topic. Dr. Salter and I have signed a contract with Wiley to publish a text directed at graduate students and researchers who are interested in the field. We estimate that it will take 3 to 4 months to complete the first draft, and a significant portion of the leave will be directed at completing this draft. This text will be the first on the topic and will be used in a graduate course at the Health Science Center.

Trinity students are already benefiting from this research, and this leave will enhance my understanding of the clinic. The more I learn about the clinic, the more questions there are to investigate.

Many of these questions are appropriate for undergraduate research, translating a positive leave experience into continued opportunities for students.

I have specifically asked for the leave to be granted in the Spring of 2007. The spring semester is suggested for two reasons. First, a colleague from New Zealand who works on the same problems may be able to visit during the spring, and it is to our advantage to have uninterrupted research time. Second, the educational load is not as significant in the spring semester, making it easier for the department to cover courses.