

Get with the (Sequentially Linear) Program: A Robust Approach to Zapping Cancer

By Barry A. Cipra

Each year in the U.S., according to the American Cancer Society, there are approximately 1.4 million new cases of cancer and about 560,000 deaths due to cancer. Given the current rate of 4 million births per year, these statistics suggest that an “average” American has a 1 in 3 chance of developing some form of cancer during his or her lifetime, and a 1 in 7 chance of dying from cancer. Alternatively, arguing from current mortality statistics—2.4 million deaths per year from all causes (including the ever-popular heart attack)—one might conclude that the latter ratio is closer to 1 in 4. Either way, cancer does a lot of damage.

You can improve your odds, of course, by taking up smoking or by handling certain forms of hazardous waste. (You could also alter the odds by choosing different parents.) It’s widely recognized that preventive measures—eating broccoli, say, or wearing sunscreen—would, if more widely adopted, dramatically reduce the incidence of cancer. Early detection is also seen as helpful (although the high incidence of false-positives carries its own risks). Finally, re-searchers in a gamut of disciplines are exploring new techniques for treating cancer in hopes of reining in its deadly toll.

Linear programming, one such discipline, is well known for saving money—many billions of dollars every year. Can it also help save lives?

Steve Wright thinks so. He and colleagues at the University of Wisconsin are studying the use of linear programming to optimize a new approach to cancer treatment known as intensity-modulated radiotherapy. The group has developed a robust formulation that accounts for uncertainties inherent in the procedure by introducing certain nonlinearities into the optimization; they have found an efficient way to solve the robust formulation that generalizes to a class of similar problems. Wright described this work in a minisymposium on complex processes at the 2006 SIAM Annual Meeting.

Sculpted X-ray Beams

The basic idea of intensity-modulated radiotherapy (IMRT) is to “sculpt” x-ray beams so that they kill cancer cells but not the surrounding healthy tissue. If, say, a tumor wraps around a critical structure, such as the spinal cord or esophagus, the desired outcome is a cancer survivor who can walk or chew gum. IMRT seeks such outcomes by irradiating the body from various angles with a beam composed of numerous “beamlets,” each with its own adjustable intensity. (It’s the adjustability—the “IM” of IMRT—that’s only about a decade old. Without it the technique is known as 3-D conformal radiotherapy, itself only a couple of decades old. “Conventional” radiotherapy, which is still the standard in much of the world, blasts away with a single beam.) Each beamlet delivers a dose of radiation to all the cells along its path, cancerous or not. The trick is to pick a multitude of paths that converge at the cancerous cells but are relatively sparse elsewhere.

The linear programming formulation starts by chopping the body into volume elements (voxels) labeled $i = 1, 2, \dots, m$. These voxels are partitioned into three sets: the target set T containing the cancer, a critical set C containing structures one can’t afford to sacrifice, and the remaining set of normal cells N . (The target set is usually further partitioned into a “clinical tumor volume,” CTV, surrounded by a “planning tumor volume,” PTV, which is intended to include regions likely to contain undetected cancer cells. The critical set is often further partitioned according to function.) Using $j = 1, 2, \dots, n$ to label the set of beamlets, the primary LP input is a “dose” matrix D , whose ij entry is the amount of radiation received by voxel i from one unit of radiation in beamlet j . The other inputs are various thresholds for the amounts of radiation to be received by voxels in the three main sets and penalties for exceeding those thresholds. The LP problem is to minimize the total (linear) penalty incurred by the dosage vector $x = Dw$, where w is the (non-negative) weight vector of beamlet intensities, usually subject to upper and lower bounds on the dosages received by the voxels in the target set T . (The lower bound ensures that the entire tumor is treated.)

Obviously, feasibility can be problematic—sometimes the treatment plan specified by the physician is impossible to deliver without modification. The calculation of D , the allocation of voxels to T , C , and N , and the determination of thresholds and penalties are all important operations in setting up the model. Once the numbers are in place, the linear programming problem is relatively straightforward. But what if the numbers are wrong?

Bring On the Cones

Errors and uncertainties can arise at numerous points as the model is formulated. The dose delivered to each voxel, calculated by simulation, is represented better as a distribution than as a single number. The locations of the voxels can change during the course of treatment, as the tumor shrinks or as the patient’s internal organs move. The tendency of living, breathing people to move while radiation is being delivered is particularly problematic in determining voxel locations.

Several groups have reformulated the linear programming problem to include aspects of uncertainty. Timothy Chan and John Tsitsiklis of MIT, together with Thomas Bortfeld and Alexei Trofimov of Harvard Medical School, have developed a linear programming model that incorporates probability density functions to describe breathing motions. Millie Chu, Yuriy Zinchenko, and Shane Henderson of Cornell University and Michael Sharpe of Princess Margaret Hospital in Toronto have also studied tumor-location uncertainties. Their robust reformulation turns the nominal LP problem into a convex nonlinear optimization known as a second-order cone programming (SOCP) problem by adding L_2 -norm terms to the linear terms in the objective and constraints. The robust version can be dispatched with off-the-shelf SOCP solvers based on interior-point algorithms. The downside is run-time: These solvers take considerably longer than LP solvers for the nominal problem.

Wright and one of his former graduate students, Arinbjörn Ólafsson, have taken the next step. Their reformulation, which also produces a second-order cone problem, encompasses dose-matrix calculation as well as positional uncertainties. Furthermore, they solve the SOCP with a sequential linear programming approach rather than an interior-point method, a strategy made possible by the fact that the objective function and constraints are smooth (and hence linearizable) at all points of interest. Wright believes that this property may be shared by other SOCP applications, making the simple linearization approach broadly applicable.

Wright and Ólafsson have tested their method on a data set from a clinical case of nasopharyngeal cancer. The nasopharynx is the upper part of the throat (pharynx) behind the nose. Their model has 24,000 voxels and 1989 beamlets (39 beamlets from each of 51 angles).

The clinical tumor is roughly rectangular in cross section, about 10×30 mm; it lies between the salivary glands and in front of the spinal cord (see Figure 1).

The optimizers ran the nominal and robust formulations of two models: one with the original, rather large planning volume for the data set, which translated into a linear programming problem with 16,489 constraints, and the other with a smaller PTV “collar” around the clinical volume, which reduced the number of constraints to 3751. The smaller planning volume, made possible by the inclusion of positional uncertainty in the robust formulation, reduces the run-time by an order of magnitude, from 27.3 to 3.8 minutes. (The corresponding nominal LP problems were solved in 3.5 minutes and 44 seconds, respectively.) The sequential LP algorithm took 15 iterations in the first case and 14 in the second.

The differences between the nominal and robust solutions are relatively small, which is encouraging because of the implication that the nominal approach already does a decent job of delivering radiation where it’s wanted and sparing locations where it’s not. The potential for improved clinical outcomes—much study remains to be done before any of these algorithms can be considered routinely reliable—can be likened to the extra sliver of profit an optimization analysis can add to a company’s bottom line. Linear programming and its robust, second-order-cone cousins are poised to save lives as well as dollars.

Barry A. Cipra is a mathematician and writer based in Northfield, Minnesota.




Figure 1. A nasopharyngeal cancer lies above the spinal cord and between the salivary glands. Properly optimized, intensity modulated radiotherapy (IMRT) can focus x-rays on the tumor, sparing the critical structures. (Figure courtesy of Steven Wright and Arinbjörn Ólafsson.)