

Extreme Optimization 1

Optimizing Radiation Therapy

AM121/ES121 — Fall 2016

Due 5:00 PM, Friday, October 7, 2016

Announcements

- Some ground rules. You are to work with members of your group but not across groups. You are allowed to use any course material that has been provided to you (textbook, lecture notes, section notes, etc), and any other textbooks and general purpose discussions on linear programming. But you may not use the internet in an attempt to look up solutions. If you have any questions or require any clarifications please contact a member of the course staff.
- Each group need only turn in one electronic copy. Please ensure that your write-up is coherent and complete. **Be sure to record the names of all team members and state what each team member contributed.**
- Your team will give a short presentation in class on Wednesday, October 12, 2016.
- Work as a team. Have fun!

Goals

This assignment gives you an opportunity to tackle a real-world problem using the techniques you have learned in the course. It will also give you a chance to exercise your creativity and insightfulness.

Contents

1 Saving lives: Radiation Therapy	2
1.1 An overview	2
1.2 The problem at hand	2
1.3 The mathematical formulation	4
1.4 Implementing in AMPL	5
2 Turning in extreme optimization	6
3 The presentation	6

1 Saving lives: Radiation Therapy

1.1 An overview

High doses of radiation can kill cells or prevent them from growing and dividing. This is true for cancer cells and normal cells, but radiation therapy is attractive as a cancer treatment because the repair mechanisms for cancer cells are less efficient than for normal cells. This observation inspires radiation therapy, where high energy radiation is beamed at a patient. Recent advances in imaging and radiation delivery have enabled radiation therapy as an effective form of cancer treatment. On the imaging front, advances in computed tomography (CT) and magnetic resonance imaging (MRI) provide a means for accurately visualizing the details of tumorous regions. On the delivery front, intensity modulated radiation therapy and tomotherapy enable a machine to beam large doses of radiation from various angles into the target.

As a beam of radiation enters a target area, the dosage of radiation dissipates as it passes through the target region, such that the largest dose is delivered where the beam enters and the smallest dose delivered when the beam leaves. Figure 1a shows the relative intensity of a beam delivered from the upper left side of the target region whose dosage has decreased somewhat when the beam hits the tumor. By using multiple beams aimed from different directions, a large dose can be delivered to the tumor as multiple beams intersect over the tumorous region. Figure 1b offers a simple example demonstrating the significance of having multiple beams, wherein the largest dose is now at the center of the target region where the tumor resides.

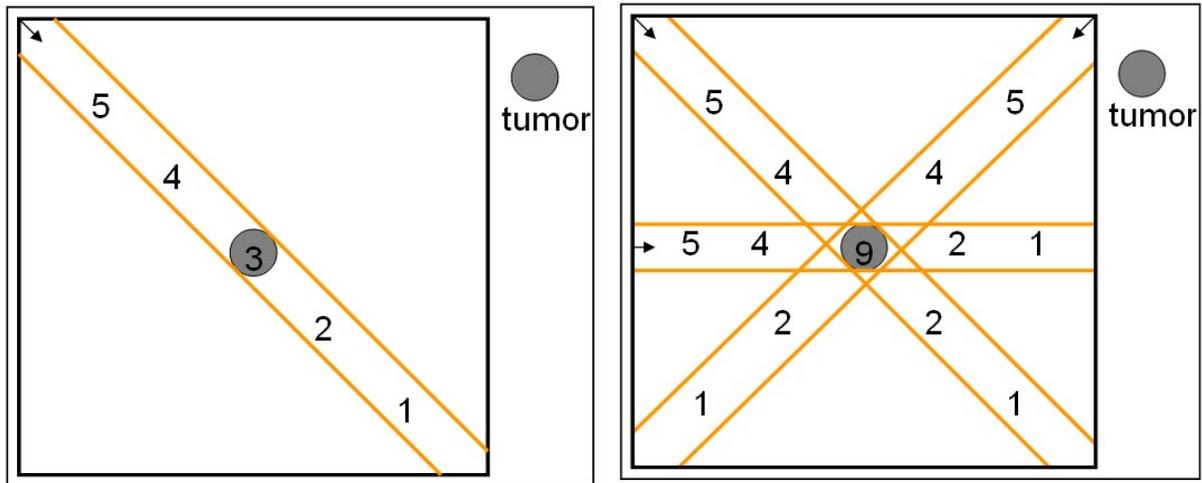
In conventional radiotherapy, only a small number of beamlets are used. The radiation oncologist and physicist work together to determine the set of beam angles and beam intensities, usually by a manual trial and error process. Determining the optimal beamlets is extremely difficult to solve by hand, as there are thousands of possible choices. Furthermore, a target area (see Figure 1c) is likely to contain not only tumorous areas but also critical areas where only a small amount of radiation (or none at all) is allowed. With a small number of beamlets, many beams may need to pass through the critical area to reach the tumorous region, making it difficult to deliver a high dosage over the tumorous region and low dosages over critical areas. This problem is particularly severe in the brain, where every non-cancerous cell is critical.

The imaging and radiation delivery technologies are now at a point where a much larger number of beamlets can be used and radiation can be delivered accurately. This calls for an automated process exists to determine how to deliver radiation to maximize dosage to the tumor area and minimize dosage to the critical area.

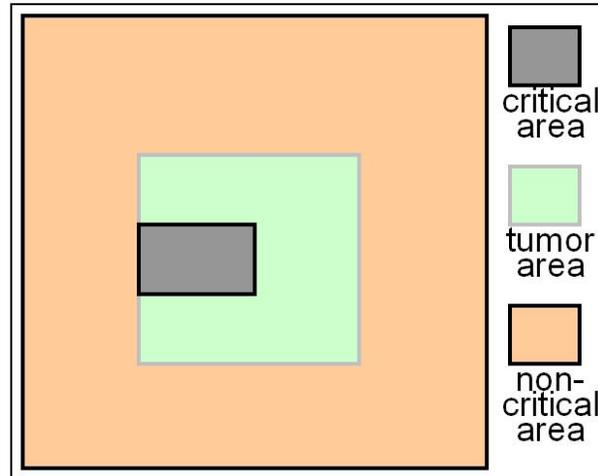
1.2 The problem at hand

The radiotherapy oncologists have provided us with sample 2D images of the tumorous and critical areas, along with a set of possible beamlet origins and angles. The course staff have transformed the data to a more friendly format for you to allow for uninterrupted extreme optimization. Each provided image has been transformed into a binary matrix where each element of the matrix corresponds to a pixel in the original image, where a 1 entry indicates the presence of the attribute shown in the image. For example, the binary matrix that corresponds to the tumor area will have a 1 entry if a tumor was present at the corresponding pixel and 0 otherwise. Note that the tumorous and critical areas need not be rectangular, and multiple tumorous regions may exist within the considered target region.

The beamlet data has been converted to a set of relative intensity matrices, where each matrix corresponds to the relative intensity delivered by a particular beam, normalized between 0 and 1 inclusively. This is the data that captures the variation in intensity of radiation as the beam passes through a region and also the imprecision in a particular beam. Each entry of a matrix corresponds to a pixel in the image of the target area, such that an entry in the matrix is nonzero only if the beam passes through the corresponding pixel. Note that each beam matrix corresponds to a beam origin-angle combination; it is possible for there to be multiple beams in the set of possible beamlets with the same origin but aimed at different angles.



(a) Relative intensity of radiation delivered from one beam. (b) Relative intensity of radiation from three beams. The beams intersect to deliver a concentrated dose to the tumor.



(c) The important areas of the treatment region.

Figure 1: Illustration of beams, intensity and critical, tumor and non-critical areas.

The oncologists have provided two sets of data, one corresponding to a small example for testing purposes and another to a case of a patient with a brain tumor. We have transformed the data and provided them to you in two separate folders (`smallexample` and `actualexample`). Each folder contains three files that contain the tumor area binary matrix, critical area binary matrix, and beamlet data matrices respectively, as well as another file with specifications relevant to the problem. This information includes the number of beams considered (the number of beamlet data matrices), the number of rows and columns in the provided images (the number of rows and columns in the matrices), and lower and upper limits on the amount of radiation delivery required over the tumorous areas and allowed over the critical areas, respectively. The oncologist wants us to determine the intensity levels to set for the beams in the provided examples, but even more so, to provide general mathematical formulations and corresponding AMPL code that can then be used on other cases as well. The oncologist needs this from us by 5:00 PM, Friday, October 7, 2016.

1.3 The mathematical formulation

You should always start with the verbal and mathematical formulation of the optimization problem before implementing the formulation into a model in AMPL. In writing down your model, make sure the formulations are linear programs. Unless we specify otherwise, do not use integer or binary variables! **For each formulation, be sure to describe the new elements of your model.**

Task 1

1. Formulate a mathematical model (not using AMPL at this point) for the problem specified to determine the intensity of radiation from each beamlet. To effectively kill tumor cells while protecting the critical areas, you must ensure that the dosage over any portion of the tumor area satisfies the specified lower limit and dosage over any portion of the critical area satisfies the specified upper limit. Note that these upper and lower limits are pre-set by the oncologist and apply uniformly over the target area (that is, all parts of the critical areas have the same upper limit and all parts of the tumor areas have the same lower limit).
2. Naturally, an oncologist will try to set the lower limit for tumor areas as high as possible while setting the upper limit for critical areas as low as possible. Even with lots of experience, an oncologist has a tendency to set the limits so that there are no feasible solutions. When this occurs, the oncologist would keep adjusting the limits until a feasible solution exists. To say the least, this is time consuming. To help with this process, you are to formulate a model that allows for a slight violation of the limits where it is necessary to ensure feasibility. This will allow the oncologist to set the limits as desired and nevertheless always obtain a solution. As with the previous subtask, you should still follow the principle of killing tumor cells while protecting critical areas. Your solution should try to get as close to the dosage limits as possible.
3. While imaging and radiation delivery technologies are quite advanced, they are not perfect. Images are limited in resolution, and minor inaccuracies can appear in the delivery process. To address these concerns, formulate a variation of your model from the previous subtask that slightly penalizes radiation delivery to parts of the non-critical area that border a critical area.
4. While you have done all that's asked for, you decide to be heroes and consider possible additions to your model that may enhance the efficacy of radiotherapy. Write down three or more possible enhancements to your models. For now don't provide the mathematical models but briefly provide details in English about the data, constraints, and objective necessary to accomplish each of the goals that the additions aim to address. It is okay if your enhancements may require additional data from the oncologist, but do comment briefly on whether any additional required data is likely to be available. Make sure your additions are well-motivated. Push your creativity and imagination to the extreme.
5. Choose two of the enhancements you have written down and provide mathematical formulations that will enable these enhancements. If your formulations operate on additional data, be sure to introduce any variables or constants that correspond to the required data. It is up to you whether to present one model that contains both enhancements or to create two separate models (this probably depends on the nature of the enhancement). You also need to make sure your formulations are indeed linear programs and thus can actually be effectively solved once implemented. If you do have a great idea that would require binary or integer variables (but is otherwise linear), let us know and we may okay it. Do note that an idea is useful in application only when you can scale up to actual data!

End Task 1

1.4 Implementing in AMPL

With mathematical formulations at hand, it is time to begin implementing your models in AMPL. Three things should strike you. First, given that the provided matrices raw data files are quite large, it seems undesirable to include their contents directly in an AMPL data file. Second, once you do obtain a solution, it would be nice to be able to visualize it to see how radiation is being delivered through the target area. We've got you covered on both these points.

Take a look in the folder `max_matrix`, which contains an AMPL program that takes as input matrices of the same size and outputs a solution matrix whose entries are the corresponding maximal entries across all input matrices. This program has nothing to do with radiotherapy, but showcases the `read` command for reading data from a file and the `print` command for outputting the solution matrix. Read through and understand all the files provided in this folder. In writing your data files, make use of the `read` command for the beam, critical area, and tumor area data files provided (you can hard-code details from the specification file directly into the data files). In writing your output script, make use of the `print` command to generate a matrix containing the dosage delivered to each pixel, and save this matrix to a file (similarly to `max_matrix`).

We have also provided some basic MATLAB code for visualizing the output in the `visualizer` directory. First open the `visualizer` folder on your computer. Then open MATLAB (if you don't have it installed on your computer, any lab computer should work). Drag the `ExOpt1Visual.m` file from the `visualizer` folder into the left pane of the MATLAB window titled "Current folder". You should see it appear in the list. Next, drag the `ExOpt1Visual.m` file from the left pane into the "Command Window" in the center. Alternatively you can simply type `ExOpt1Visual` in the Command Window. The visualization script will run and ask you to select `beam.out` (your dosage output from AMPL), `tumor_raw.txt`, and `critical_raw.txt`. It will then show you the image of your radiation treatment with critical areas in orange and the beams in green. In this figure window you can go to `File-->Save as` and save it as a jpeg, tiff, pdf, etc. If you have any questions regarding the visualizer, please feel free to ask us! ¹

Task 2

Make sure to follow similar naming conventions for all your implemented models. The AMPL code, like the write-up itself, should be coherent and complete. Please be sure to comment your AMPL code, but to avoid redundancy, you do not need to comment things that should be clear, given either the names of the variables or from the comments provided for the mathematical model. Since the mathematical model from the previous task should already be well-explained, just add a brief comment wherever you were not able to directly map something into AMPL and there might be some confusion.

1. Implement your basic mathematical model in AMPL. Create a model file, two data files (one for the small example and one for the actual example), a run script, and output files containing the dosages delivered. Visualize obtained solutions using MATLAB. Describe your findings.
2. Implement your relaxed threshold mathematical model in AMPL. Create a model file, two data files (one for the small example and one for the actual example), a run script, and output files containing the dosages delivered. Visualize obtained solutions using MATLAB – this should take the place of printing out the dosage delivered at each point! Describe your findings, focusing mainly on the actual example and drawing from both the solution itself and the visualization.
3. Implement your variation of the relaxed threshold mathematical model in AMPL that penalizes radiation delivery to regions close to critical areas. Create a model file, two data files (one for the small example and one for the actual example), a run script, and output files containing the dosages delivered. Visualize obtained solutions using MATLAB – this should take the place of printing out the dosage delivered at each point! Describe your findings (focus mainly on the actual example). Did the solution change? Offer an interpretation for your answer.

¹You are free to use a different software to perform the visualization as well. For example, Mathematica.

4. Implement the enhanced model(s) you have come up with. Create a model file, data files, a run script, and output files containing the dosages delivered. Presumably, your data files will still be based on the two examples we have been provided, but you may have added something extra in there. When doing so, feel free to create (that is, make up ☺) any data your models depend on. Visualize obtained solutions using MATLAB – this should take the place of printing out the dosage delivered at each point! Describe your findings.

End Task 2

2 Turning in extreme optimization

You must submit an electronic submission before 5:00 PM, Friday, October 7, 2016. Gather all AMPL model and data files you have created and scripts containing the AMPL commands you used to solve the problems. Create a README file that will let us know which files refer to which problems. Put all these files in a zip archive, and submit the zipfile to the Canvas dropbox under the “Extreme 1” folder before 5:00 PM, Friday, October 7, 2016 (No late days!). Make sure you note in your zipfile, in the README, or some other obvious place, the names of your group members and a team name. **Remember to include a statement of what each team member contributed to your project.** You can choose any team name you’d like as long as it isn’t lame.

3 The presentation

You will be giving a **short** presentation (around 5 minutes) in class on Wednesday, October 12, 2016. The purpose of the presentation is to convince your fellow classmates that the enhancements you have introduced to the model can do great things for cancer treatment via radiotherapy. Choose your favorite or most exciting enhancement to discuss; be sure to motivate and describe the enhancement, offer the corresponding mathematical formulation, and present your findings on test data. **You will have limited time, so focus on the enhancement you select and not the basic formulation and its variations!**

Notice that the goal is as much to convince as it is to describe: your best shot of convincing audience is by doing a good job of motivating your work.

Submit your slides by 5:00 pm Tuesday, October 11, 2016 to Google slides. We’ll make a link available.

Congratulations on completing your first AM121 extreme optimization assignment!