In this lecture, we'll discuss how linear optimization is used to design radiation therapy treatments for cancer patients.

Cancer is the second leading cause of death in the United States, second only to heart disease.

There were an estimated 570,000 deaths in 2013 due to cancer.

Additionally, over 1.6 million new cases of cancer will be diagnosed in the United States in 2013.

These trends are also seen throughout the world.

Worldwide, cancer is a leading cause of death, with 8.2 million deaths in 2012.

Cancer can be treated using radiation therapy, where beams of high-energy photons are fired into the patient to kill cancerous cells.

This is a very common form of treatment for many types of cancers, and in the United States, about half of all cancer patients undergo some form of radiation therapy.

Radiation therapy has a history going back to the late 1800s.

X-rays were discovered by Wilhelm Rontgen in 1895, who was later awarded the first Nobel Prize in Physics.

Shortly after the discovery, X-rays started being used to treat skin cancers.

This image shows an x-ray of Rontgen's wife's hand.

You can see the bones in her hand as well as her wedding ring on her finger.

A few years later in 1898, radium was discovered by Marie and Pierre Curie.

They won the Nobel Prize for this discovery in 1911.

Radium started being used to treat cancer, as well as other diseases.

Later in the 1900s, the first radiation delivery machines, called linear accelerators, were developed.

Then computed tomography, or CT scans, were invented in 1971.

These discoveries led to the invention of Intensity Modulated Radiation Therapy, or IMRT, in the early 1980s.

The invention of IMRT significantly improved the ability of radiation therapy to target cancerous cells.

To reach the tumor, radiation passes through healthy tissue and therefore damages both healthy and cancerous tissue.

This damage to healthy tissue can lead to undesirable side effects that reduce post-treatment quality of life.

For example, blistering and burning of skin can occur because of the damage to healthy skin cells.

For this reason, we want the dose to fit the tumor as closely as possible to reduce the dose to healthy tissue.

This became possible with the invention of IMRT.

In IMRT, the intensity profile of each beam is non-uniform.

Before IMRT, each beam had to have the same intensity, so the tumor could not be targeted very well.

But by using non-uniform intensity profiles, the radiation dose can better fit the tumor.

Let's see what this looks like.

In this image, we have a person's body outlined in black, and then the target, or tumor, and two critical structures also outlined.

We would like to maximize the radiation to the target, while minimizing the dose to healthy tissue, and especially to the critical structures.

Using traditional radiation therapy, each of the three beams has the same intensity throughout the beam.

So to deliver enough radiation to the tumor, we also have to deliver a significant amount of radiation to the critical structures and to other healthy tissue.

But by using IMRT, we can change the intensity throughout each beam to make it non-uniform.

Some pieces of the beam will have a higher intensity than others.

This allows us to deliver the necessary amount of radiation to the tumor, while minimizing the total radiation to healthy tissue, and thus, the critical structures get significantly less radiation.

Each of the pieces of the beam is referred to as a beamlet.

So in IMRT, we decide the intensity of each of the beamlets so that we can target the tumor with radiation while minimizing the radiation to healthy tissue.

So in designing an IMRT treatment, the fundamental problem is -- how should the beamlet intensities be selected to deliver a therapeutic dose to the tumor and to minimize damage to healthy tissue?

This is the problem that we'll address in this lecture, using linear optimization.