slide1:

Good afternoon.Today, we will explain the system aspects of computed tomography. Since the CT data acquisition is a scanning process, a CT system is typically called a CT scanner.

slide2:

As you can see, the topic of this lecture is called CT scanner.Next time, we will have a MATLAB session.During that, we'll talk about filtered backprojection and related concepts.You'll have the opportunity to review MATLAB commands and functions so you can perform filtered back projection.

slide3:

So, let's look at the outline for today's lecture.

First, we'll talk about projection data truncation—the issue when the X-ray projections don't fully cover the object, which can cause artifacts.

Next, we'll discuss different scanning modes used in CT technology.

Then, we will move on to common image artifacts that can degrade image quality.

Finally, we'll cover X-ray radiation dose, which is important because it can be harmful to patients if not properly managed.

slide4:

As we discussed in the previous lecture, when you collect X-ray projection data from an unimpeded perspective, you aim to reconstruct the image layer by layer, starting from the outermost shell.

It's like peeling an onion, gradually revealing each layer beneath the surface.

To achieve this, you need to have X-ray measurements along every line passing through the region of interest.

For each X-ray, you want to record the line integral—the total attenuation along that path.

This situation pertains to 2D image reconstruction.

Always keep in mind the principle of unimpeded data acquisition.

We say there must be at least one point on the X-ray source trajectory for every line passing through the object.

In other words, you need a measured projection corresponding to every line through the cross-section.

This ensures you have complete data for accurate reconstruction.

.

slide5:

Now, in the 3D CT case, the X-ray source remains a point source. But instead of using a linear detector array, you use a planar or two-dimensional detector array.

This configuration results in a cone beam geometry originating from the point source.

The X-ray source emits radiation in a cone-shaped beam, which should be wide enough to cover the entire object.

In other words, the object must be fully contained within the cone beam to ensure complete data acquisition.

This concept is similar to the 2D case but extended to volumetric imaging.

Once you have all the rays, as shown here, you can rearrange, reorganize, or re-bin the X-rays into parallel beam geometry.

In parallel beam geometry, you can apply Fourier transform-based imaging methods.

The Fourier slice theorem allows the use of the one-dimensional Fourier transform on projections to recover the two-dimensional Fourier spectrum.

Each projection recovers a profile in Fourier space, and by changing the viewing angle, you sample different radial lines sweeping over the entire Fourier domain.

Likewise, in 3D imaging, an extended version of the Fourier slice theorem is required.

This involves applying Fourier transforms in higher dimensions to recover the 3D spatial frequency information of the object.

slide6:

The Fourier transform, as we learned, is the foundation of classic CT theory. It requires complete coverage of an entire cross-section or the whole object. When you have projection data, you can use the Fourier slice theorem in 2D or 3D, and then perform the inverse Fourier transform also in 2D or 3D. This classical framework is called the central dogma.

The whole cross-section or entire object must be fully covered by the X-ray beam. No matter the shape of the beam, you can always re-bin the data into parallel beam geometry. Then, the collected data can be inversely transformed back to the original image domain. That's the fundamental way image reconstruction is performed.

As we learned, filtered back projection is essentially just the inverse Fourier transform represented in polar coordinates. This core concept explains the process pretty well. If you follow this, you should grasp the essential idea. In my opinion, the key question is: what is the difference between classic CT theory and modern CT theory?

slide7:

This leads us to the issue of data truncation. If the data is incomplete, the Fourier slice theorem— which assumes a full projection profile for each one-dimensional Fourier transform—can no longer be applied straightforwardly.

What happens if part of the data is missing? If some portion of the projection profile is lost, the Fourier transformation cannot be performed fully. This is called the data truncation problem. Much of modern CT theory has been developed to address this data truncation problem, finding ways to reconstruct images even with incomplete data.

Before we delve into the types of data truncation, let me give you some context.

slide8:

First, let me show you the four generations of CT scanning geometry, often called generations or modes. This basically refers to how you systematically collect X-ray projection data.

In the first generation, you have a single point X-ray source and a single point detector. Each measurement records one line, one line integral. The setup moves by translation, producing one parallel beam projection in one direction. Then the entire apparatus rotates slightly to a new angular position, producing another parallel beam projection. This process depends on the viewing angle, often labeled theta. That's the first generation.

The second generation improves data acquisition efficiency. Instead of measuring one line integral at a time, you use a small detector array with a narrow fan beam. This detector array collects multiple lines simultaneously as it translates, covering the cross-section completely. Each translation collects multiple parallel beam projections, depending on how many detector shells you have. For example, if you have five detector shells, you get five parallel beam projections per translation. After a small rotation, you repeat this process.

The third generation further improves efficiency dramatically. Instead of a few detector elements, you have hundreds or thousands along a linear or arc-shaped wide detector array. This is a full fan beam imaging geometry. The detector array continuously rotates to collect data. This configuration satisfies the data sufficiency condition: performing a full 360-degree scan ensures every line through the patient's cross-section is measured somewhere, so the line integrals are known. This third-generation geometry is still very popular. Modern CT scanners mostly use this geometry but have expanded the one-dimensional detector arrays into multi-row or cone-beam configurations. We'll discuss those later.

The fourth generation differs by making the detector array a full stationary ring. Instead of rotating the detectors, only the X-ray source rotates around the patient. The detector ring stays fixed, continuously collecting signals. This fourth-generation geometry solves some mechanical problems associated with rotating the detector array.

slide9:

Next, we have electron beam CT, which uses a different approach. Instead of a mechanical X-ray tube rotation like earlier generations, this system uses an electron beam that's electromagnetically steered to hit a fixed target shaped as multiple metallic arcs.

Each arc produces X-rays upward, forming the scanning beam. The electron beam is steered by electromagnetic fields very rapidly without moving parts mechanically. This allows for much faster imaging speeds.

The driving force behind the evolution of scanning generations is increasing data acquisition efficiency: collecting the same amount of data in a shorter time to improve temporal resolution.

Cardiac imaging is the most challenging application for CT, since the heart beats fast, anywhere between 60 and 100 beats per minute. Tomographic reconstruction assumes the object is stationary, but the heart is constantly moving. So, we need to shorten the data acquisition time drastically.

The electron beam CT scanner was designed specifically for cardiac imaging due to its high speed. However, it is costly, and image noise is often higher. Additionally, the fourth-generation geometry faces problems with scattered signals affecting other detectors, which causes degradation in image quality.

Because of these limitations, the third-generation geometry remains more widely used today, and electron beam CT scanners are now mostly considered outdated.

slide10:

Defining CT scanning modes, or how many generations to count, varies depending on the source. Different textbooks categorize them differently. Usually, the first three generations are well established: the first generation, the second generation, and the third generation.

However, beyond these, the classification becomes a matter of opinion and varies by source.

slide11:

Let's look at scanning modes. Here, you can visualize a diamond that helps you remember the progression of CT generations. The first generation uses a parallel beam, then a narrow fan beam in the second generation, followed by a wide fan beam with rotation in the third generation. The third generation features fan beam geometry with rotating detectors. Beyond that, we have the fourth and fifth generations. The next stages are called the sixth or seventh generation, which involve helical scanning.

Helical scanning is quite an innovative, out-of-the-box idea. Traditionally, with one-dimensional detector arrays, imaging is done by rotating the source around the patient—often a half circle or full circle—to collect sufficient data for reconstructing one cross-section using filtered back projection or iterative reconstruction algorithms. After reconstructing one slice, the patient table is translated slightly to image the next cross-section. This process is known as the “step and shoot” mode.

You scan one cross-section, then move the patient, accelerate or decelerate the table, perform another full rotation, and repeat. This approach is time-consuming, especially because anatomical and pathological features exist in three dimensions, and patients can only hold still for limited times. Scanning slice by slice leads to low throughput and motion artifacts.

Spiral scanning solves many of these issues. You ask the patient to remain still, often holding their breath. Then, the table moves continuously while the X-ray source rotates, and data acquisition happens simultaneously and continuously. From the patient’s perspective, the source follows a helical trajectory, which is why this mode is called helical imaging.

However, helical scanning introduces data truncation issues. In second-generation scanning, the data may be transversely truncated due to the small detector size. In helical scanning, data is longitudinally truncated because the source follows a helical path within any transverse plane, meaning a full two-dimensional data set for image reconstruction is absent.

slide12:

This slide illustrates a spiral, single-slice CT scan, showing the imaging geometry. You see the helical path around the patient clearly.

Now, the question is: how do we perform image reconstruction effectively despite this geometry?

How do we handle the challenges posed by the helical scan?

slide13:

With helical scanning, when you want to reconstruct an XY cross-section, ideally, you'd have rays confined entirely within the XY plane. But the actual green helical scanning path deviates from this ideal.

You measure line integrals only along rays that coincide with this helical path. Immediately, the X-ray source moves out of the plane, producing rays above and below it. To reconstruct a cross-section, say in the XY plane, you need every line integral through this section, including along particular rays such as Ray A.

In practice, you can't measure the line integral exactly along Ray A because the source never stays on that plane long enough. You only have rays above and below it. That's the central data truncation problem in helical CT.

Engineers found a practical solution. Helical scanning benefits from continuous table motion without stops or accelerations. The patient lies flat on the table, which moves steadily through the scanner. The patient simply stays still until done, making the process fast and smooth.

To compensate for missing data, we perform linear interpolation between the upper and lower rays, which are available. The assumption is that the data varies roughly linearly between these rays. This allows synthesis of an in-plane data set, which then can be used with filtered back-projection algorithms. This was the first trial approach.

Later improvements involved helical scanning which uses opposite rays, not just the nearest rays in the same direction. Before reaching a certain position, you could use complementary rays shot from an opposite angle to improve interpolation.

slide14:

Using opposite rays that are closer to ray A reduces the distance between the measured and estimated line integrals. The narrower this gap, the better the accuracy you can expect.

This "half-scan" interpolation enhances longitudinal resolution along the Z-axis compared to linear interpolation. Longitudinal resolution along the Z direction is often lower than in-plane resolution.

Back when I just graduated and started work at Washington University Medical School in St. Louis, helical CT was introduced. Radiologists appreciated the improved temporal resolution it offered.

However, they noticed motion blurring along the Z-axis caused by the interpolation process. Since the patient is continuously moved through the scanner, motion blur appears in the longitudinal direction, similar to a photo taken while running.

As they say, there's no free lunch—you gain temporal resolution but lose some longitudinal resolution due to motion blur.

slide15:

We analyzed helical CT scanning modes further. Traditional step-and-shoot CT uses multiple circular scans. For each scan, you reconstruct one CT slice, repeating many times for full coverage.

In helical CT, instead of discrete circular scans, you have one continuous helical acquisition. For each helical rotation, interpolation creates synthetic data points.

You can reconstruct images at many arbitrary longitudinal positions along the patient, generating a large number of slices computationally. This technique is called retrospective reconstruction.

Retrospective reconstruction has a significant clinical advantage. For example, in tumor detection, a lesion may fall between two slices in traditional scanning and go undetected.

With retrospective reconstruction, densely sampled slices mean one slice will include the tumor, improving detection confidence and diagnostic accuracy.

slide16:

Next, we have partial volume averaging artifacts. For each slice, if a tumor only partially lies within that slice, you might not see it clearly. Small tumors can be especially difficult to detect when they occupy only half of a slice, and image noise further complicates clear visualization.

However, with spiral CT, you can select longitudinal positions as densely as you want. This gives you a real opportunity to reconstruct images densely enough. One of the reconstructed slices is likely to contain the tumor fully, providing much better contrast resolution and visibility.

slide17:

Building on this heuristic idea, we performed a rigorous linear system analysis and published a journal paper in Medical Physics. We performed modeling the signal response of the CT system to compare the two scanning modes as I mentioned earlier in the lecture; that is, multiple scanning circles versus a spiral or helical scanning trajectory.

Our key theoretical conclusion is that, for a given X-ray dose, meaning a fixed number of total incident photons, you can choose to scan in either the conventional step-and-shoot mode or helical mode to have multiple scanning circles or a single helical trajectory. Given the same X-ray dose, helical CT offers substantially better longitudinal resolution than conventional step-and-shoot.

This improvement comes from the inherent capability of helical CT to perform retrospective reconstruction—that is, reconstructed overlapping slices along the longitudinal axis.

In our papers, we specifically recommended reconstructing roughly three to four slices per helical turn for the best results.

Thus, helical CT not only improves temporal resolution but also enhances longitudinal resolution, so there's no loss in image quality. This important finding helps explain why helical scanning is so beneficial.

Some believe that you gain temporal resolution with helical scanning but lose longitudinal resolution. However, with proper overlapping reconstruction, both temporal and longitudinal resolution can be improved. This is a key insight from our earlier work and kind of a free-lunch in clinical practice.

slide18:

Helical or spiral CT began with single-slice scanning using a linear detector array, capturing rays in two dimensions.

Expanding this, cone beam helical spiral CT was developed. For cone beam spiral CT, the issue of data truncation is addressed more intricately.

If you slice through the patient in any plane, the actual X-ray beam forms a fan beam in that plane.

This fan beam is longitudinally truncated because the patient's body is quite long, and the detector size limits the coverage. You cannot cover the whole body with a single cone beam at once.

Often, you focus on smaller regions of interest like the heart or inner ear, so using a huge cone beam isn't practical.

Consequently, data truncation naturally occurs in helical cone beam scanning, making it an important consideration in advanced CT imaging.

slide19:

In the early 1990s, I wrote a paper called the General Cone Beam Reconstruction Algorithm.

It explains that for cone beam helical scanning—where the beam geometry is divergent—you cannot simply apply linear interpolation to reconstruct images.

The mathematics involved is more complicated, requiring specialized algorithms to handle data truncation and beam geometry correctly.

slide20:

This particular paper became the first seminal work on helical cone beam scanning and has been widely cited.

Today, multi-slice cone beam helical scanning is widely used, with between 100 and 200 million such CT scans performed worldwide each year. This number highlights how prevalent and important this technology has become.

Reflecting on that time, we didn’t file any patents on this work, which turned out to be a major missed opportunity. Had we filed for intellectual property, it could have been extremely valuable.

Helical scanning combines multiple technologies, including slip ring transmission, X-ray sources, and sophisticated reconstruction algorithms. The software scanning model itself is crucial to performance.

To put it in perspective: even if each scan were worth only one cent, with a million scans annually, that’s a million dollars a year in potential revenue.

So, not securing patent protection was probably the biggest mistake we made. Nonetheless, our original paper still stands as the prototype of this important process and algorithm.

slide21:

Now, let’s discuss longitudinal data truncation from a single slice. When you perform one helical turn with a single slice, you have a parallel beam and adjacent rays, allowing you to perform interpolation.

However, when you move to a wider angle cone beam, simple interpolation is no longer sufficient, and more advanced algorithms are needed. The first paper I wrote on cone beam helical scanning proposed an approximate algorithm for this challenge.

slide22:

It then took over a decade for the field to develop what is called the exact helical cone beam algorithm. This algorithm still uses the filtered backprojection framework but is mathematically exact.

The author, Dr. Katsevich, is a mathematician, so the paper is quite dense and challenging for most engineering students. Nonetheless, it demonstrates that modern CT theory can effectively handle longitudinally truncated data.

Earlier, I mentioned that in the second generation, data truncation happens transversely, while helical CT mainly deals with longitudinal data truncation, which can be seen as a symmetric problem.

slide23:

There is also an interesting transverse data truncation problem.

Imagine your cross-section represents the green base, while the blue area is a smaller region of interest, or ROI.

Perhaps you’re only interested in the heart, or in the ear if planning a cochlear implant.

In this case, the ROI is small. You might use a narrow fan beam targeted only through the ROI, then rotate as in third-generation fan beam geometry. This means the projection data is transversely truncated on both sides.

With truncation on both edges, the Fourier slice theorem no longer applies, making reconstruction very challenging.

Mathematically, this “interior problem” has no unique solution, meaning multiple structures could produce identical data, and you wouldn’t know which is correct.

Several years ago, my group revisited and solved this problem, which we felt very proud of.

While the classical interior problem’s non-uniqueness holds in unrestricted function spaces, we showed that introducing prior knowledge allows exact reconstruction.

For example, if you know a small sub-region of the ROI, such as air in the airway or blood in the aorta, you can use that known X-ray attenuation coefficient as a reference. This prior information enables theoretical, exact reconstruction within the ROI from truncated data, a method called interior tomography. Our work was well received, although reviewers raised important questions.

One student tested whether blood density varies among different races or sexes, finding it consistent across populations.

This allowed us to use blood density as a stable reference in cardiac imaging, so you only need to scan the heart, reducing radiation dose.

Later, reviewers pointed out that contrast agents—like iodine injected during CT angiography—change blood density, so you can’t always rely on a fixed reference.

This challenge led us to develop a new solution.

slide24:

In the new approach, we assume there is no known sub-region.

Instead, we model the region of interest as composed of finitely many sub-regions, each modeled by piecewise polynomials.

The boundaries of these sub-regions need not be known in advance. Each sub-region can be approximated by constant, quadratic, or more general multivariable polynomials.

This assumption matches the idea of band-limited signals from sampling theory.

With this sparsity-based model, where the ROI is decomposed into a small number of smooth regions, the interior reconstruction problem becomes solvable.

For example, in a chest CT scan, you can reconstruct the whole chest from all data or just from truncated data restricted to an ROI.

Here is a real dataset scanned on a Siemens scanner demonstrating this.

Interior tomography is no longer a mathematical curiosity but delivers practically good reconstruction results.

slide25:

Let me add some comparison between classic CT and interior tomography.

According to classical theory, which we learned and which I explained in textbooks, you need to collect all data to satisfy the data sufficiency condition—usually a fan beam or cone beam that covers the entire cross-section or object.

The data must comprehensively cover the object for accurate reconstruction.

Interior tomography, however, only requires partial data limited to the region of interest.

Thus, with filtered backprojection, classic CT reconstructs a global image, like a wholesale business covering the entire area.

Interior tomography reconstructs a smaller, targeted area—like retail shopping.

This opens up new opportunities, providing flexibility and precision.

You only need to shoot X-rays through the important region, reducing dose and focusing on clinical needs.

slide26:

I have given multiple presentations on interior tomography. One key idea is that interior tomography uses less data—less is more in many ways. Less data means a deeper understanding of imaging principles. We can handle larger objects, use less radiation dose, and achieve faster scanning.

Why is it faster? All scanning geometries, from first to helical generations, aim to speed up scanning. A cone beam is better than a fan beam because it allows parallel data acquisition. Interior tomography takes this even further by enabling higher-level parallelism with multiple X-ray sources and smaller detector arrays.

In the extreme, you could have many X-ray focal spots, possibly based on emerging technologies like carbon nanotube cold emission X-ray sources. These enable distributed X-ray cells with small detector elements arranged similarly to the fourth-generation geometry.

At any instant, such a system could collect multiple truncated projections. Handling data truncation like this allows reconstruction focused on small regions of interest, achieving the fastest possible tomographic imaging speed. You cannot go faster than the speed of light, which governs X-ray photon travel.

We filed for intellectual property on this concept at Virginia Tech. Later, I will discuss the potential of interior tomography in cardiac imaging.

slide27:

Now, let me explain CT scanner architecture. To reconstruct images, we rely on line integrals from X-ray projections, using filtered backprojection and Fourier transform-based reconstruction, or interior tomography for region-of-interest reconstruction.

To acquire this data, you need hardware: the CT scanner itself. This is a donut-shaped structure called the gantry, which is the framework holding all components.

Inside, you find X-ray sources and collimators that shape the beam and filter out low-energy X-rays. Low-energy photons won’t penetrate the patient and only add an unwanted dose.

Collimators also reject scattered radiation to ensure the measured signal corresponds closely to the direct line integrals.

The scanner is a sophisticated device: you place the patient in the gantry, and from that, you acquire many images that resolve internal structures clearly.

slide28:

Let me show you a typical CT scanner layout. You have the X-ray tube, which is quite heavy—for example, our lab's donated GE tube is heavy.

Powering the tube requires a high-voltage supply and a cooling system; scanning consumes energy and heats the tube.

Cooling is done with water or air to keep components at safe operating temperatures.

This entire setup is called the gantry. It includes the patient table where the patient lies, the detector array opposite the source, and collimators in front of the detectors.

In conventional CT, the source, detectors, and associated hardware are mounted on the rotating gantry.

During scanning, the gantry physically rotates, typically half or a full circle, then reverses direction. This back-and-forth rotation is limited by cabling, which can tangle with repeated turns.

For helical scanning, slip ring technology is employed. The slip ring lets you rotate continuously without cables twisting or breaking.

It consists of rings with carefully designed contacts to transmit power to the X-ray source and data from the detectors.

Slip ring technology is critical for continuous and fast scanning. It enables the transfer of data to the scanner's computer or local storage during rotation.

When designing a CT scanner, the goal is faster and faster imaging speed, with the most demanding application being cardiac imaging—the holy grail of CT.

slide29:

At RPI, we have collaborated with GE Global Research Center and with Bruno Diem and others to develop optimized architectures for dedicated cardiac CT scanners.

We published a paper last year on this topic.

Designing a CT scanner is much like designing a building or room—you consider architecture and optimal layout.

Our team is currently preparing a grant proposal to advance this research. The proposal is due soon and covers aspects not addressed in previous reviews.

slide30:

I am especially excited about this project, collaborating with GE.

Currently, cardiac CT scanners mostly use third-generation geometry, but with cone beam helical scanning mode, which some call improved third generation or helical cone beam generation.

This system is already very fast, capable of rotating three to four times per second.

The mechanical acceleration involved is astonishing—much higher than what's involved in launching a space shuttle.

The current best temporal and spatial resolution for cardiac CT is about 500 microns.

But according to cardiac surgeons I know, 100 micron resolution would solve many clinical challenges.

Improving cardiac CT resolution requires higher imaging speed to freeze the heart's rapid motion on this fine scale.

However, as you increase rotation speed, the X-ray tube has less time to emit photons, limiting achievable image quality.

Our out-of-the-box solution uses interior tomography. We designed a stationary ring scanner to boost X-ray flux by orders of magnitude.

Imagine a cardiac patient lying still while being scanned. You get immediate images.

When higher resolution is needed, the patient or ring can be translated to focus precisely on the region of interest.

This stationary ring can move up, down, left, or right to align perfectly with conventional cardiac CT images.

I believe this design is a promising way forward—an out-of-the-box solution for high-speed, high-resolution cardiac imaging.

slide31:

Now, let's finish with the last two sections. First, image artifacts. I've mentioned image artifacts several times throughout this course. Tomographic images are meant to reveal the true underlying anatomical structure. However, for various reasons, the images sometimes show structures that aren't actually present. These false structures are what we call image artifacts.

If there's no real structure present, what you see might just be noise. Sometimes, artifacts show up as streaks or stripes in the image, but obviously, there are no stripes inside the human body. They are just artifacts. We'll discuss CT image quality considerations next, which include how artifacts impact the overall image quality.

slide32:

Starting with data sampling: In modern CT scanners, such as those from GE or Siemens, a typical rotation collects about 1,000 projections—sometimes slightly fewer, depending on the manufacturer, like 800 or 790.

At each projection, the detector array works like a comb beam or multi-row detector, capturing thousands of data points along each ray path. For example, along one ray direction, you might have about 1,000 detector elements.

All these detector readings combine to define the scanner's spatial resolution, often achieving sub-millimeter precision.

In the future, we are aiming to improve this further, reaching resolutions on the order of 100 microns. This is challenging because every wire and detector element has limits, so current strategies often focus on targeting smaller regions of interest.

To resolve fine details, the angular sampling intervals and detector element sizes must be small enough, consistent with Fourier analysis and sampling theory.

If the sampling rate is too low or the detector elements are too big, small features won't be resolved clearly. For instance, with about 512 rays, you get fairly clear images. With fewer rays, small details become difficult to detect.

Angular sampling is also critical. If you have nearly 1,000 views around the object, image quality is good. Using fewer views leads to aliasing artifacts and lower image quality.

slide33:

CT reconstruction relies on CT numbers, which are derived from Beer's law. Beer's law relates the incoming and attenuated X-ray intensities via a parameter called the linear attenuation coefficient, mu.

Data processing converts raw measurements into line integrals, which are then used to reconstruct mu values throughout the image.

On the CT monitor, you don't see mu values directly but relative values compared to water, since the human body is mostly water.

Thus, CT numbers indicate how much more or less dense a region is than water.

For example, bone has a CT number over 1000, air has about -1000, and water is defined as zero. The range is often magnified by 1000 times for visualization.

This scale is sometimes called Hounsfield units, named after Sir Godfrey Hounsfield, the British engineer who developed the first CT scanner, which earned him a Nobel Prize for its revolutionary capability to reveal internal structures.

slide34:

Different body regions use different CT number windows for display. These windows are defined by minimum and maximum CT numbers, depending on the manufacturer or clinical application.

For instance, the full range of CT numbers might be mapped to 256 grayscale levels on a display, from white to black.

The human eye can typically distinguish about 50 shades within a window.

If too many shades are presented, subtle differences become difficult to see.

For lung imaging, for example, you might set the window maximum to about 250 Hounsfield units and the minimum to the air value.

Within this window, 256 gray levels help render tissues clearly. Anything outside the window appears clipped and is not visible.

This windowing is critical to tailoring visualization to different clinical needs.

slide35:

There are multiple resolution measures in CT. High-contrast resolution detects small objects adjacently positioned; for example, cardiac CT currently resolves about 500 microns.

If you want to detect even smaller details, down to 100 microns, you may need advanced methods like interior tomography that we've discussed extensively.

Low-contrast resolution differentiates subtle intensity differences. This is important because tumors may look very similar to surrounding tissue.

Narrow window settings, chosen carefully with a mean and range, help highlight these subtle differences.

If the CT number difference is very small—say 5 to 10 Hounsfield units—and noise is comparable, detecting such differences requires boosting contrast effectively.

Hence, low contrast resolution depends not just on hardware but also on image noise characteristics.

Temporal resolution is crucial for dynamic organs like the heart. To get high-resolution cardiac images, both spatial and temporal resolution must be adequate.

If your detector resolves fine details but scanning speed is too slow, motion blurring will reduce effective spatial resolution.

So, achieving both fast scanning and high spatial resolution is essential.

Photon-counting detectors with elements as small as 55 microns are available, enabling very high-resolution imaging.

Combining this with interior tomography concepts and distributed thousands of photon-counting detector elements, it's feasible to build a "magic ring" scanner that focuses on small regions of interest with resolution down to 100 microns.

This is an area of active research that we continue to pursue.

slide36:

The noise you see in CT images is due to quantum effects and other random variations. Several factors influence this noise level, such as the tube current, often referred to as MAS. When the tube current is high, the X-ray tube emits more photons, so the statistical fluctuations decrease, which means less noise.

Image resolution also plays a key role. Think of pixel or voxel size in a simple way. If you have a large pixel, it collects many photons, resulting in a lower noise level. But if you target a very small, high-resolution pixel or voxel, fewer photons pass through it, so you need even more photons overall to maintain good image quality.

This is why interior tomography is advantageous. Since it directs photons only through the small region of interest, it reduces unnecessary exposure while maintaining image quality for the target structure.

slide37:

Low contrast image resolution can be measured using resolution phantoms containing small test objects with subtle differences in density. These phantoms help to assess if the scanner and reconstruction algorithm can visualize objects with minor contrast differences.

When contrast is poor, like between blood vessels and surrounding tissue, contrast agents such as iodine are injected. These agents highlight the vessels, making blood, tumors, and related structures more visible.

This approach significantly improves low-contrast resolution, enabling better diagnosis and visualization of soft tissues.

slide38:

Temporal resolution, so you see conventional CP scanning. You do one slice at a time. Then you

move the patient into the country by a small increment. Do again, patient, keep pressing.

So from slice to slice, you see some discontinuity. So this is not due to some pathological feature.

It's really due to motion artifacts that show the temporal resolution agent is good enough. But

With helical scanning, it's smoothed out, but you have some longitudinal blurring.

Here is the earlier algorithm. Later, I mentioned overlapping reconstruction. You really can

recover an image longitudinally and a better resolution than conventional scanning. Anyway,

so we wouldn't have time to go into details.

slide39:

Scattering artifacts occur due to scattered X-ray photons. Let me explain this with an example. Imagine sending an X-ray beam from the south direction through the body. Follow my cursor as I illustrate. The primary X-ray beam travels straight through a data element, creating the expected line integral for reconstruction.

However, during measurement, some X-rays scatter in various directions. You might see scattering toward the right, left, or even backward directions. These scattered photons add a background signal that does not correspond to the direct transmission along the beam path.

Because of this scattered radiation, the detector picks up additional signals on top of the real primary X-rays. This extra background leads to artifacts in the image—usually appearing as low-frequency streaking or haze that reduces contrast and clarity.

If you correct for scattering components by estimating and subtracting them, you can restore the image quality close to the phantom reconstruction without scatter.

One common method to reduce scatter is using beam blockers placed strategically. These beam blockers function to block certain parts of the primary radiation, allowing measurement of scatter separately and improved image correction.

slide40:

Beam hardening artifacts occur when X-ray beams pass through an object, resulting in changes in beam energy distribution that affect image quality. Let me explain this in some detail.

Imagine an X-ray beam going through the body in a certain direction, forming a parallel beam projection. This is a real and common setup.

When using a medical X-ray tube and current detector technology, both measure the attenuated signals, resulting in what we call the line integral. However, if you reconstruct the image based on the uncorrected projection profile, the reconstruction may look distorted, often represented by a dotted projection line in diagrams.

Why does this happen?

It's due to beam hardening. X-ray beams are polychromatic, meaning they contain photons of various energies, forming a spectrum. Initially, the beam may have many lower-energy photons that get attenuated more as the beam penetrates the object.

Lower-energy photons have higher attenuation coefficients, so they are absorbed more readily in tissue.

As the beam progresses into the body, the relative amount of higher-energy photons increases because they penetrate deeper. This process "hardens" the beam—increasing its average energy.

Since higher-energy photons are less easily blocked, the beam becomes more penetrating along longer paths.

Despite this, many reconstruction algorithms assume a monochromatic beam, leading to errors.

This results in an underestimation of the linear attenuation coefficient mu because the beam appears to pass more easily than expected.

Consequently, the computed line integrals and preprocessed data are based on a single energy assumption, but the reality is a changing, polychromatic beam.

This discrepancy causes characteristic artifacts in images, such as streaks and cupping.

Proper correction methods and models that account for beam hardening are necessary to reduce these artifacts and improve image quality.

slide41:

This is an example of CT images affected by beam hardening without any correction. You can see a darker strip in the image, which obscures details. After applying correction techniques, you can visualize structures like gray matter much more clearly.

Beam hardening artifacts originate from the X-ray source and the polychromatic nature of the X-ray beam. When you use a photon-counting detector, it can measure X-ray signals at different energy levels separately.

By distinguishing between energies, photon-counting detectors effectively eliminate beam hardening artifacts. This technology significantly improves image quality by reducing these artifacts that cause distortions in conventional CT images.

slide42:

Now, let’s discuss volume averaging and various related artifacts. These artifacts are not difficult to understand conceptually, but effectively removing them can be challenging.

In CT imaging, volume averaging occurs because of finite detector element size and limitations in the reconstruction matrix’s resolution. These hardware elements are not small enough to capture features smaller than the detector size or the X-ray focal spot.

Imagine you have a fine detail—a small yellow cross in an image. But your image pixels are arranged in a grid, and this cross might fall entirely within a single pixel.

That one pixel will then represent an average of all the structures within it, so the small cross may simply blend into a larger, less distinct area.

For example, if the cross is in one pixel and that pixel appears as a light yellow block, you cannot resolve the cross’s true shape.

This phenomenon is called volume averaging.

In CT images, various factors cause blurring and averaging effects, causing images to look less sharp.

Detector pixels have a limited size, and the X-ray source’s focal spot is not a perfect point but has an elliptical shape, which contributes to image blur.

To partially counteract this, image processing techniques like deblurring and deconvolution are applied, recovering some lost resolution.

This relates to our earlier study of linear systems and convolution in imaging.

So, volume averaging is a fundamental problem intrinsic to CT imaging, limiting the spatial resolution achievable.

slide43:

Next, we have what are called metal artifacts. Inside the human body, you may have metal implants like pacemakers or dental fillings. Metal is much denser than bone, so it can block X-rays effectively.

Because of this, X-rays cannot penetrate metal implants well. The attenuated signal behind the metal may be very noisy or even just noise without a clear signal.

When you perform filtered backprojection reconstruction in such cases, the results are often poor, showing streaking artifacts.

This happens because the line integral passing through the metal is very long, with many photons blocked or scattered. Beam hardening and other effects further degrade the signal.

As a result, the detected signals are compromised with a low signal-to-noise ratio. When estimating line integrals, the data is inconsistent and inaccurate.

Filtered backprojection on such corrupted data generates stripe-like artifacts in the images, which is a typical manifestation of metal artifacts in CT.

Reducing metal artifacts is an important research topic, and multiple algorithms have been developed, though it remains challenging.

slide44:

The last topic I'm going to explain is motion artifacts. Imagine you have a phantom object, like this one, viewed at angle zero. The phantom has a taiji pattern, a yin-yang shape.

As you slowly increase the projection angle, the X-ray source rotates around the phantom.

Over time, the data you collect changes. When the angle reaches 90 degrees, the phantom appears compressed, shrunk to a minimum in the projection.

As the source continues rotating, at 180 degrees, the phantom appears expanded again.

This back-and-forth pattern creates a sinusoidal oscillation in the data.

This example illustrates what happens when the object isn't stationary, but moves periodically.

You gather projection data while the object moves according to this periodic motion.

If you then apply filtered backprojection to reconstruct images from this data, you get motion artifacts—blurring or distortion caused by the object's movement during acquisition.

However, if you know the object's motion trajectory, you can incorporate that information into iterative reconstruction algorithms.

With this additional information, it is possible to recover accurate images despite motion.

Tomographic reconstruction remains feasible, provided you have either explicit knowledge of motion or use algorithms that infer motion implicitly.

Due to time constraints in this lecture, we won't delve deeply into motion artifact correction algorithms.

For cardiac imaging, one way to reduce motion artifacts is through hardware: performing ultra-fast scanning so the heart doesn't have significant motion during image acquisition.

This instantaneous imaging "freezes" the cardiac structure, minimizing motion blur.

Another way is through smart computational algorithms that predict motion and correct for it during reconstruction.

Often, cardiac CT patients are given beta blockers to slow the heart rate, reducing motion during scanning.

Additionally, CT acquisition can be synchronized with the ECG signal, coupling image data with specific cardiac phases.

The ECG indicates when the heart is largest or smallest. You can select data segments corresponding to the heart's maximum volume and reconstruct images at that phase.

Similarly, you regroup data segments for the minimum volume phase and reconstruct images accordingly.

This approach allows reconstruction of a beating heart in different cardiac phases.

This gives a rough idea of how motion artifacts are managed in dynamic organ imaging.

slide45:

Let's finish with some important points about X-ray radiation. X-rays are much more energetic than visible light. We experience visible light all the time—sunshine, for example—which is generally considered healthy.

But X-rays have much shorter wavelengths, roughly from 0.1 nanometers to 10 nanometers, and much higher energy.

Because X-ray photons carry so much energy, they can break molecular bonds in biological tissues. This damage can potentially cause genetic mutations, which are often harmful.

In the long term, such mutations may increase the risk of cancer or other genetic diseases.

There is some debate, though. Some argue that small doses of radiation might even be beneficial, similar to exercise, causing small stress that strengthens the body.

In this view, controlled, moderate radiation exposure acts like a vitamin, gently challenging the body to repair and strengthen itself.

While that perspective exists in some articles, the general consensus is caution.

As CT imaging, especially advanced methods like helical CT for lung cancer screening, becomes more widespread, concerns about radiation dose increase.

Patients receive higher cumulative doses, and this raises public health questions.

So, it is crucial to focus on how we can reduce radiation dose and how we measure it in the first place.

slide46:

Let me explain the basic concepts from your Green Book regarding radiation dose.

X-ray radiation is a form of electromagnetic radiation, and when it passes through your body, energy is deposited into your tissues. This energy deposited is called the absorbed dose.

The absorbed dose is a physical measurement of energy per unit mass of tissue, expressed in grays. One gray equals one joule per kilogram.

This is the fundamental concept of radiation dose—the energy your body absorbs from X-rays.

However, from a health perspective, the effective dose equivalent is more relevant.

To calculate this, the human body is divided into multiple organs or regions because different organs have different sensitivities to radiation.

Some tissues, like those involved in reproduction or rapidly regenerating cells, are more sensitive to radiation damage.

Therefore, younger women, children, and reproductive organs require special caution to avoid excessive exposure.

This sensitivity is accounted for by weighting factors, denoted by W.

The dose equivalent, represented by H, combines the absorbed dose with factors related to the type and energy of radiation.

Higher energy photons or particles generally cause more biological damage per unit of absorbed energy.

For example, X-rays, gamma rays, and alpha particles differ in their relative biological effectiveness.

To account for this, a quality factor, QF, is applied to quantify the biological impact of different radiation types.

For X-ray CT, the QF is normalized to one, but other radiation types have different values.

Thus, the effective dose equivalent is calculated using these weighted factors, resulting in a more accurate estimation of biological risk.

This calculation is summarized by a formula sometimes noted as equation 1.30 in radiation safety literature.

slide47:

Let me share with you some typical radiation dose values.

For example, in a breast exam, the physical energy absorbed is measured in grays. But when you take biological sensitivity into account and use standard formulas, the unit becomes the sievert, often abbreviated as Sv or sometimes called C-word.

To clarify, grays measure the pure physical energy deposited, while sieverts factor in biological effects, making them more meaningful for assessing health risk.

Because a sievert—or the C-word—is a relatively large unit, radiation dose is often expressed in millisieverts (mSv).

For a typical breast examination, the effective dose is quite low.

For high-dose CT scans, the effective dose can be in the range of three millisieverts.

For a full-body CT scan, it could be around ten millisieverts.

These values give you an idea of the radiation dose magnitude.

In CT scanning, another important metric is the CT Dose Index, or CTDI, which is used to estimate how much radiation dose is delivered.

CTDI is not organ-specific; it measures the average dose over one full rotation of the scanner around the patient, providing a rough estimate of radiation exposure.

Later, I will show you a figure to illustrate this concept visually.

slide48:

Before we dive deeper, let me give you a natural reference to help put radiation dose levels in perspective.

If you live in India, you receive an annual natural radiation dose of over 10 millisieverts—or 10 million microsieverts—which comes from natural sources like soil and cosmic rays.

In China, the natural background radiation exposure is roughly half that of India.

If you are concerned that radiation is harmful, then living in Toronto would be better, as the natural background radiation there is lower.

In the United States, the average natural background radiation dose is about three millisieverts per year.

These background levels are unavoidable since radiation comes from natural sources in the environment.

Understanding these reference levels helps put medical radiation doses into context.

slide49:

Now, let me give you some comparisons to help you understand radiation doses better.

For example, a chest X-ray typically delivers a radiation dose of about 0.1 millisieverts.

A whole-body CT scan, on the other hand, can deliver around 10 millisieverts.

Researchers and medical professionals are working hard to reduce these doses.

Our goal is to bring CT radiation doses down to about one millisievert or less per scan, which would be considered comparatively safe.

To put this in perspective, in your everyday life, you naturally receive about two millisieverts of radiation annually from environmental sources.

So, when you receive a CT scan with a dose around one millisievert, that’s roughly half of your average yearly natural exposure.

You can think of it like visiting a high natural radiation area, such as parts of India or other regions, where you might be exposed to somewhat higher background radiation if you stay longer.

There are other units you might encounter, such as M-R or similar, but the gray and sievert (or millisievert) remain the standard units for expressing radiation dose. This comparison helps warm you up and offers a bit of perspective on how these numbers relate to everyday life.

slide50:

Now, let me explain three key concepts related to radiation dose.

First, the absorbed dose. This represents the energy deposited by ionizing radiation into the body's tissues. It is a purely physical measure.

Second, the effective dose accounts for the sensitivity of different organs and tissues to radiation. This is where weighting factors come into play.

For example, certain organs involved in reproduction or rapid cell regeneration are much more sensitive to radiation. We want to protect children and younger women, especially from high doses. While adults may be less sensitive, it is still prudent to minimize exposure.

The dose equivalent combines dose with information about the type of radiation. Some radiation types cause more damage per unit energy than others.

However, when discussing CT scans, the radiation type factor is effectively one, so it generally does not alter dose estimates.

slide51:

Now, let me explain the concept of the radiation dose profile, sometimes called the releasing profile.

Imagine you perform a circular CT scan. This scanning method reconstructs images within a single plane, called the XY plane.

Ideally, the X-ray beam would be finely collimated so that all photons are confined within the slice thickness, represented as a neat rectangular region—say, highlighted in yellow.

However, in reality, the system is not perfect.

You have scattered photons, and the emission is not perfectly confined.

Therefore, the actual radiation dose profile follows a curve, often shown as a blue curve.

This means that, even though you intend to irradiate only the slice thickness, some dose extends beyond this region.

This dose "tail" outside the intended slice area also contributes to the total patient radiation exposure.

slide52:

Now, let's talk about the CT dose index, or CTDI.

Suppose the slice thickness is represented by T.

You integrate the radiation dose measurements over the slice thickness, effectively summing the dose contributions from the slice and imaginary adjacent slices—say, seven slices above and seven slices below.

All these doses are added together, and the total is then averaged by normalizing with the slice thickness T.

This result is what we call the CT dose index, or CTDI.

It's a simple yet important concept, as it includes any additional leakage radiation in the dose computation.

The CTDI provides a standardized measure of radiation output from the CT scanner, reflecting dose distribution within the irradiated volume.

slide53:

Another way to think about radiation dose is through the dose profile measured over multiple circular scans.

When you perform one circular scan, you get a certain dose distribution.

If you then do a second circular scan, you'll get a similar dose profile again.

When multiple scans are performed, each scan adds its own dose profile.

These multiple dose profiles overlap, contributing to the total radiation dose the patient receives.

slide54:

This concept of multiple scan average dose, or MSAD, is defined by an integral calculation.

Radiation dose from neighboring scans contributes to the dose in the current scan.

The number of adjacent scans considered is represented by N.

You perform the integration over the radiation dose profile, then normalize it by the slice thickness.

This normalization accounts for the actual tissue volume receiving radiation.

MSAD is another classic measure of radiation dose in CT.

With the rise of helical CT, there are new adaptations of these dose measurements, including for helical scanning and combined bimodal CT systems.

However, the fundamental principle behind measuring radiation dose remains the same.

As I mentioned earlier, understanding how we measure the release dose is essential to managing patient exposure.

slide55:

To measure radiation dose, we use specialized instruments called dosimeters.

I personally had the opportunity to work on radiation dose measurements for CT scanners during my time at Washington University School of Medicine.

It was engaging and valuable research work, and I frequently used dosimetry equipment to evaluate radiation output.

These dosimeters help assess the radiation delivered by the scanner, ensuring safety and proper calibration.

slide56:

Now, let's talk about some general guidelines regarding CT dose.

Radiation dose measured during multiple scans is roughly proportional to the tube current, or mA. Higher tube current means more photons emitted, resulting in a higher radiation dose.

Longer scan times also increase dose, as the patient is exposed for a longer period.

Increasing the tube voltage, or kVp, raises photon energy, which also influences dose. Although higher energy photons are more penetrating, they tend to cause slightly higher patient dose.

When you reduce slice thickness—that is, make the X-ray beam narrower—you increase spatial resolution. But thinner slices require more photons since the beam is less focused, which can slightly increase the dose due to some photon wastage.

Many factors affect dose increase or decrease, including patient size or phantom size. For a standard head phantom, the dose at the center and periphery is roughly the same.

For larger phantoms, like the whole body, the central dose may be lower due to increased attenuation, with a higher dose in peripheral regions.

In extreme cases, like imagining irradiating the Earth, the very center would receive almost no photons, so the central dose would be very low, while peripheral areas get a higher dose.

slide57:

Let me share some evidence that supports this approach.

Several experts, including myself a few years ago, co-authored a landmark article.

This article mapped out strategies and technologies aimed at achieving routine sub-millisievert CT scanning.

It addresses how low-dose CT scans can be effectively performed while maintaining diagnostic image quality.

slide58:

We have discussed various technologies to reduce radiation dose in CT scanning, including interior tomography and photon counting detectors.

Additionally, iterative reconstruction algorithms play a key role. Unlike traditional analytic algorithms, which do not account for image noise or statistical fluctuations, iterative methods model these factors explicitly.

By doing so, they improve image quality while allowing dose reduction.

Those represent technological approaches to lowering patient radiation exposure.

Now, I'd like to share something not in the slides—a more biological approach.

Ionizing radiation causes real damage through DNA breaks and other effects.

Some medications or natural products contain antioxidants, which can help mitigate radiation-induced damage.

For example, orange juice and blueberries contain antioxidants.

If you search the literature, you'll find several studies suggesting that antioxidants may have protective effects against radiation damage.

While this is not part of any clinical protocol, and doctors don't routinely prescribe antioxidants before CT scans, I personally believe in their potential benefit.

Next time I have a CT scan, I plan to drink two bottles of orange juice an hour beforehand.

If you find this interesting, you might want to do your own research and consider consuming antioxidants before imaging.

One of my students is setting up experiments to scientifically evaluate the protective effect of antioxidants, like orange juice, against radiation damage during CT scans.

This is an exciting area of ongoing research, and we hope to quantify the benefits more precisely in the future.

slide59:

Now, let's talk about the last slide and some course housekeeping.

These three questions in the green book are quite typical and not very difficult.

Regarding homework, you have one week—seven working days—to complete it.

On the other hand, for a bit of fun and creativity, you can try designing a futuristic portable CT scanner in an autonomous vehicle—think of it like something fancy and innovative.

I believe future technology will enable very affordable, fully reconfigurable CT scanners.

Imagine robotic X-ray detectors integrated into cars—if you want a screening, you just call like ordering an Uber.

The robotic scanner arrives at your home, performs the scan, and you get your results without leaving the house.

I think that will be incredibly cool and transformative.

slide60:

Several years ago, we undertook an innovative project and received some news coverage about it.

We built a CT scanner that does not rotate the gantry as in conventional designs.

Instead, the patient lies on a rotating table—literally a sheet that turns—while the X-ray tube and detector remain fixed in place.

The gantry is simplified to just a tube and a detector, with no mechanical rotation during imaging.

We then perform the translation of the table and the scanning components.

Although this setup limits the viewing angle compared to traditional systems, it still produces good-quality images.

If you like, you can think creatively or artistically about how this concept could be extended to robotic scanners deployed in autonomous vehicles or portable scanners for versatile use.

This is all just for fun and imagination.

That’s all for today.

Thank you for your attention.