DOUGLAS LANMAN: So today's talk is on computational imaging, which as you'll see, there's a lot of inspiration you can take from this work. And I'd just like to mention a lot of these slides came from Marc Levoy's SIGGRAPH course. So if you want any more detail, go to his website, and you can actually see the original course material for this.

So the basic idea is that scientific science is really driven by our technology. This is not surprising anyone in the room. The classic example Marc Levoy put here is that if you have Leeuwenhoek, plus microscope equals microbiology. And so really, in some sense in the sciences, you're waiting around for new technology to make new discoveries. We're waiting to put satellites up. We're waiting for birds to stop ruining the Large Hadron Collider. Technology is what we're waiting for.

And so the question Marc Levoy asks is, what's the most important instrument in the last century. Arguably, in terms of science, what would some people put up? If you could only choose one thing to take with on a desert island where they do science, what would you take?

RAMESH RASKAR: Microwave.

AUDIENCE: Computer.

DOUGLAS LANMAN: Exactly. This is not news to anyone that the computer is really arguably the greatest scientific device. And so what can we really do with this?

RAMESH RASKAR: You know, they asked people in the general population, which is the most important technology that has changed their life? And they said remote control and the microwave.

DOUGLAS LANMAN: The microwave, or a remote control microwave is the greatest product. So Professor Horn has said-- and this is, of course, is very similar to the ideas of this class-- that computational imaging is simply imaging methods in which computation is inherent in image formation. And so when we look at medical imaging and scientific imaging, really, this is the key story in the 20th century. And so in some of this talk, I'll show you what the modern way is to take some of image using computed tomography. And I'll show you how they did that 100 years ago, before they had computing power, how they did that using mechanical computing. And so I think that's an interesting process, to see how computation has really revolutionized how we take pictures in medicine and science.

And so if we look out well beyond the box we're in, which is optics-- most of what we talk about in computational photography is this wavefront coding, light field photography, holography, things you've looked at in your homework. But if you open the box up a little bit, these same ideas really come from a much broader field. And so the first part I'll talk about is medical imaging. I think the key story in medical imaging is tomography.

Really, that's the central problem. How do you take a cross-sectional image of an object non-invasively? And you can do that in two ways. You can use transmission or reflection tomography. We'll talk about both.

And then, that story is repeated throughout the sciences. Once they understood tomography, they said, OK, we can solve all of these problems. We can look at geophysics. The exact same methodologies can be applied there, which is really what's happening in optics now.

In the late '90s, they found this link between tomography and optics. And then all of a sudden, you have light field photography. Geophysics did that. About in the '80s, they discovered that link and they invented all of these methods for unobtrusive measurement underground.

And then in Applied Physics, we'll talk about when the problems become much more difficult. For light field photography, we look at ray optics. Once we start considering coherent optics or scattering tomography, that's when the problems really get interesting. And I think for those of you working on final projects, really looking at scattering tomography and its link to light field photography is where there's some exciting ideas. And hopefully, some of you are already thinking about that for your projects.

And then, we'll jump away from tomography. Again, I'm going to talk about all these things. But then, we'll move to something else, which is in biology, the story isn't tomography, necessarily. It's again, how can you take a cross-sectional picture? But the tricks are very different. So we'll talk about confocal microscopy, deconvolution, and deep learning, which some of you have probably seen already.

And the final area I'll end up on is astronomy, which is really where most of my work comes from. Many of these ideas, especially in biological sciences, involve refractive optics. And then, we'll look at how you can image without refraction. And so again, for X-rays, gamma rays. But you can't build refractive optics. And so then, you get into the ideas of code aperture imaging and interferometric imaging.

And then, the rest of this is stuff you've already seen. How do you how do you build panoramas? How do you take light field photographs?

And so this is a great example of medical imaging and the idea of tomography. And so the problem statement is very simple. You want to non-invasively image inside of a living object. And so of course, if any of you have been to the CT scanner at the hospital, you might have had them produce images like this. So the question is, how can you do this without simply slicing open a person and taking a photograph? And so in general, this involves basically taking the last 100 years of mathematics and solving an inverse problem. We'll talk about what that inverse problem is.

Then, there's other ways. Instead of inverting a system, we can of course, use endoscopes, thermal imaging microscopy to just simply directly measure inside the body. And I think the key thing here is to notice that the mathematics behind tomography, all of these devices are essentially the same. CT scanners, MRI machines, PET scanners, ultrasound, they all use the same underlying mathematics. The only difference is the wavelength of radiation to zeroth order. And of course, the other main areas are EEG and MEG.

And so if you think about designing the non-invasive medical device, the first problem you have is ionizing radiation. CT scanners use X-rays, whereas MRIs use radial waves. So that's the first major benefit. But really, the other problems to think about are again, minimizing the invasiveness, not inserting devices or probes, and also improving temporal spatial resolution, and finally, making these things very inexpensive for the developing world.

So now, let's just jump into to what tomography is. Well, it's a Greek word, Greek origin. And it's simply the word tomos, or "a section," or "cutting" plus photography, graphic, taking a picture. So it's a cross-sectional picture of an object.

So if you have some volumetric function in 3D tomography, we'll take a slice of that 3D function, and return a 2D image. That's our goal. The question is, how exactly can we do that non-invasively. And it's a very simple idea. The idea is-- this is what I'm describing-- known as transmission tomography. There's also reflection tomography, but we'll get to that later.

So the idea of transmission tomography is we have some density function, let's say our human body. And it absorbs electromagnetic radiation. And so we're going to put a single-point source in the scene, like an X-ray source, have it emit isotopically into the world, and put a film plane sensitive to X-rays on one side of the object. And our hope is that some wavelength of radiation that we can find will transmit ballistically through the object. So there'll be a little scattering, but absorption.

So along a given ray from the source to our film plane, the electromagnetic radiation-- the photons, if you will-- will simply travel through the object. And the only thing that happens is that they're attenuated somewhat by the object. That's our assumption.

Now of course, in the visible wavelength, that's not true for human bodies. You can't do this. And so that's why they chose X-rays, because that's easy to generate, easy to build detectors. And for the most part, human bodies satisfy the scattering assumptions.

And so if you look back at CT scanner history, if we go back 30 to 40 years, really, you have two technologies competing against each other. But the basic idea is the same. We're going to have some volumetric function, like our body. And we're going to take many projections of it. Because we're trying to estimate a three-dimensional function, we need to make a well-posed system in inverse. So we need to somehow sample the variation.

So a single picture will only take a two-dimensional projection. But then, if we move the source to a new position, we'll get a new projection. If we take a full set of projections, we hope that there's enough data to invert that system. And we'll see how you do that.

But really, before they invented the simple fan-beam tomography, there was actually a competing idea. The reason it became popular is that mathematical inversion is actually easier. And the idea is that instead of creating an X-ray point source, we have some cathode tube that generates X-rays. And then we collimate them, say with a series of lead fans. And then, we get a set of parallel beams.

So this is the parallel beam tomography problem, which is really the easiest one to solve mathematically. But now that we have plenty of computing power, it's very easy to solve this system, which is easier to build mechanically, which is to simply have point X-ray sources rotating around the patient. So that's our data set. It's simply a set of projections.

So now, we have a mathematical problem. It's very simple. There's some underlying density function. And so again, a ray is going to pass through this density function. You evaluate the line integral. And that will be the absorption. And our task is to invert that, those set of line integrals, to recover the density function.

So if you just take the case of the simple fan-beam projection, we have an X-ray source that's culminated, a film plate on the opposite side, and it rotates 180 degrees around the patient. And that gives you what's known as a radon transform. Maybe this is sounding familiar to a lot of those in the room. So this radon transform, again, is simply the projection of the density function along a given direction.

So this is the data set that a CT scanner will give you, for the most part. And then if you compare that to fan-beam projection, where you can imagine then putting a point source in the corner and then rotating it, for the most part, it looks almost identical. It's just this slight nonlinear transformation of the parallel beam projection. So now that we have ample computing power, it's very easy to resample this data set to be equivalent, and then use the inverse radon transform to solve.

But we still don't know how to invert this problem. If I give you this image, what algorithm can you apply to recover the original image? So that's really where the insight was. And so I recommend--

RAMESH RASKAR: Let's make sure everybody's up to us on the forward process. Is everybody clear how you go from the first image to the second image?

AUDIENCE: Wait. This one to that one?

DOUGLAS LANMAN: Yes.

RAMESH RASKAR: First on to second, that's the easiest one.

DOUGLAS LANMAN: Is that clear? I mean, just mechanically, it's very easy to simulate. Imagine this line is our collimated X-ray source. So it's just emitting parallel rays along the y-axis.

And this is our film plane. And so if I just simulate the light going through here along a given ray, I perform the line integral of this function. That gives me the absorption along that ray. So a given ray would be some rotation angle and position.

RAMESH RASKAR: So the first column of that middle image--

DOUGLAS LANMAN: Is the--

RAMESH RASKAR: --is the first image.

DOUGLAS LANMAN: --parallel-beam projection. And then we'd rotate both of these slightly, say in 90 degrees. That would bring us already here. We take another projection. Now, we just get the series of projections as a function of angle.

RAMESH RASKAR: Do you have a movie, by any chance?

DOUGLAS LANMAN: I don't have a movie in this one, no. Yeah, we'll get to it at the end. I probably should put that earlier.

And so then, the fan-beam projection is very simple to simulate, as well. Instead of a line source, we have a point source in this corner, or in the center here. And we have a set of fan beams going out. And then, we rotate those too, which is why you get this nonlinear transformation.

So that's our input data set. Now we're trying to invert the system and go from one of the right images back to the original image.

RAMESH RASKAR: And this is a common problem in a lot of virtual imaging situations, right? Because as Doug said, you don't want to slice the patient. So all your observations are from outside, but you want to infer what's inside. So from external observations, you want to infer what's internal to the patient, or to an object, or to an oil well.

DOUGLAS LANMAN: Right. So now, we'll go into the math. There's really one nice result. And again, if anyone's working on this, you should really look at Slaney and Kak. They have a really nice book that's online. I think it's out of print now, but the whole book is online and these figures are there.

But really, the core idea behind this-- and you've probably seen this in your own homework in light field photography. I think you did refocusing, I'm guessing?

RAMESH RASKAR: Mm-hmm.

DOUGLAS LANMAN: I don't if you did Fourier domain refocusing, or just--

RAMESH RASKAR: No, just [INAUDIBLE].

DOUGLAS LANMAN: OK. So you did the shift-and-sum refocusing. And it turns out, this is the first example. Again, Marc Levoy originally put these slides together. And he had a very nice paper-- I think it was 2005, the thesis.

But the idea was that for your lexical refocusing, you did shift-and-sum, which is computationally expensive. Using the Fourier projection-slice theorem, you can actually improve the running time of that algorithm. You'll probably see that once we get to this.

Let's just consider the tomography problem again. We have this density function in 2D. Let's just analyze this in 2D to begin with. So we have some density function of x and y. So this is our absorption function. And then, we're going to take a projection. Again, this is a parallel-beam projection along some angle. And so if we're just projecting along y, then the projection along y is simply the integral of this function over y. Does that make sense?

So we're projecting this function onto the x-axis. Right? That's our first. That's obvious.

Now, the question is, let's look at a Fourier domain analysis of this problem. So let's take the two-dimensional Fourier transform of our original density function. And so you can write that. It's very simple. Everyone remembers this. The Fourier transform is the function of the two frequency parameters. It's simply the original function times this complex exponential, which is the sum of the coefficients in each direction.

And here's the key insight. What we'd like to do is somehow relate the properties of this Fourier transform to a slice in our primary domain. And so this is really the neat trick behind the Fourier projection-slice theorem.

Let's just look at this 1D slice along the F-frequency axis in the Fourier domain. And so that that's just given by substituting 0 for the y into our previous expression. So essentially, this kx and ky can be thought of as a vector. And we're just selecting the one along x. And so here, you see that our y term drops out. And what this gives us is the Fourier transform along the kx-axis, just a 1D slice.

So the next step is really where all the magic happens. It's just regrouping terms in that integral. So if you look at this, you can see that it's separable. And so we can do the y integration, regroup it, and then do the Fourier transform. And so now, you can just substitute the first expression and you get this very nice result.

So can anyone interpret this? You see what this means? I can explain it to you, but it'd be nice to see if people are following.

RAMESH RASKAR: So remember, we started with the first row is p of x, which means it's the photo taken along the x direction. If we put a sensor, we get px photo.

And then, you go through the whole math. The second equation is Fourier transform. Third equation is Fourier transform along-- is it along the y-axis?

DOUGLAS LANMAN: Along, dropping it. So it's the Fourier transform along x.

RAMESH RASKAR: Yeah, along x. And then, what you get back at the bottom is the px term, which was the original photo that you started with.

DOUGLAS LANMAN: Right. So looking at this expression, there's now this relationship between the frequency domain, the Fourier transform, or a density function, and the original function. Do you see what that relationship is here? What does this saying.

AUDIENCE: The slice is just the Fourier transform of the original photo.

DOUGLAS LANMAN: Exactly. Very good. So the Fourier projection-slice theorem-- this is the key insight. It's that if I take 1D slice of my density function, take it's Fourier transform-- 1D Fourier transform-- that's equal to a slice of the 2D Fourier transform.

So we have this insight. So how could you use this to now solve our problem? Let's go back. Do you see how you could use this to now go from this image, which is a parallel-beam projection, back to the original density? Can anyone describe the algorithm roughly?

AUDIENCE: Just take the inverse?

DOUGLAS LANMAN: Take the-- what do you mean? Of?

AUDIENCE: --of all the p axis. And then reconstruct all the slices of--

DOUGLAS LANMAN: OK so-- so here's your image. What's the first step?

AUDIENCE: Take the line, one line.

RAMESH RASKAR: Every column here is a px.

DOUGLAS LANMAN: OK. Great.

AUDIENCE: And take all of them, take the inverses, inverse of the px, inverse for the transform.

DOUGLAS LANMAN: For the Fourier transform.

AUDIENCE: Yeah.

DOUGLAS LANMAN: OK.

AUDIENCE: And then, you will get kx.

DOUGLAS LANMAN: Yes.

AUDIENCE: And then, keep rotating all the kx's to generate [INAUDIBLE].

DOUGLAS LANMAN: That's absolutely correct. Great. And then the last step?

AUDIENCE: [INAUDIBLE].

DOUGLAS LANMAN: Exactly. So let me just summarize that up. That's great. That's exactly what they do.

So the algorithm is very simple right now. We're going to get a slice for each angle. So you simply take the projection, take its 1D Fourier transform, and initialize an accumulator that's all zeros. Assign the Fourier transform, then rotate your source and receiver. Take a new slice, 1D Fourier transform, assign it, and you're going to build up this whole space. After a full 180 degree rotation, you'll have populated this entire volume. And you're going to inverse Fourier transform and get everything back.

So I think I have a picture here of how that works. Yeah. So you just described this algorithm correctly.

So now you can see basically that procedure I outlined. At each step, we're going to populate our accumulator. Once it's full, we can inverse 2D Fourier transform, and we get our image.

So what are some problems you can see? This is the textbook-- and there's the textbook again-- solution to the inverse radon transform. So what are some limitations you can already see, practically? If you built a CT scanner that used this algorithm, what would be some of the limitations?

RAMESH RASKAR: The resolution of the angles at which you can take--

DOUGLAS LANMAN: Exactly. So you see that the problem here is we need to fully populate this transform. We need to fully populate our 2D Fourier transform before we inverse transform, or else we'll have missing data and artifacts. Does that make sense?

And so this insight is really at the core of tomography. You can see a lot of the limitations. So the first one you mentioned is resolution. You mean angular resolution, I think.

So if we built a CT scanner that took a projection every five degrees, or every 20 degrees, we'd have lots of zeros left in this data set.

RAMESH RASKAR: If we take every 20 degrees, you're filling up this image only at every 20 degrees. It's 20, 40, 60, 80, and so on.

DOUGLAS LANMAN: So you can imagine they'll be huge. Near the origin, you'll be OK. Your low-frequency reconstruction will be all right. But your high-frequency detail will be lost, or unknown. You'll have many zeros out here.

So that already implies that CT scanners need to have very dense set of rotations. Or, you need to have some prior about the scene. But that's one problem.

What's another problem you can see? So if I built a CT scanner, for instance, that could only scan a limited set of angles, not a full 180-degree rotation, then you'd also have a lot of missing data here. And so that's the second. Having sufficient angular sampling is one problem with CT scanning. The other is limited baseline. If I only rotate over a small set of angles, because maybe in whatever system I'm using, it's difficult to rotate over a full 180-degree rotation, then I won't be able to get a complete reconstruction.

So really, keep in mind those two limitations of CT scanning. You need to have sufficient angular density and sufficient baseline. And they actually map directly to problems with light field cameras.

But then, let's just step back again and think about the [INAUDIBLE]. So, again, our algorithm is we rotate our source and receiver around the object. If it's fan-beam or parallel-beam tomography, you end up with the same thing, where you take the Fourier, the 1D transforms, and you assign them to the frequency domain. And we may have some zeros.

But there's a third problem, which I didn't describe. And you can see it here. Again, the algorithm was we initialized and accumulated all zeros. We take the 1D Fourier transform after we accumulated, then rotate, add the next one, add the next one, add the next one. You can see it right here with the opacity. I made all of these equal, and you can see we're building up this really high density in the center. Does that make sense?

And so when I inverse Fourier transform, essentially, my low-frequency terms have been boosted. And so we'll lose edge contrast, if we don't do anything here. So the third problem we identified with the inverse radon transform is that you can't simply evaluate this accumulator and inverse transform.

So here's an example. So if you only take it every 30 degrees, you can see versus 5, versus 1 degree. There's a huge difference in the density of angular sampling. But then, come back.

We have this problem where we get this hot spot in the center. And so if you look at this, it goes in the inverse of the radial frequency. So as I go higher in radial frequency, I don't have this problem as much.

And so the solution-- it's up here, but does anyone want to tell me what the solution is to solve the hot spot?

AUDIENCE: [INAUDIBLE].

DOUGLAS LANMAN: Multiply by the inverse, right? So that inverse is interesting. So the direct inverse of this function is simply omega itself, absolute value of omega. So as a frequency domain filter, it's simply saying multiply by a gain, which is proportional to your radius, away from the center of the image.

So what kind of filter is this?

AUDIENCE: Bypass.

DOUGLAS LANMAN: Bypass filter, right? So the algorithm is telling us we're multiplying by this bypass filter, which means what we're doing in the primary domain is we're sharpening. Does that make sense? We're convolving with a sharpening filter.

And so that adds some other problems. Basically, this filter amplifies high-frequency noise. Right so in general, in CT scanning in the '80s and '70s, basically, they examined apodizing this filter, having it trail off at some frequency, so you don't arbitrarily amplify high frequencies. So again of this is the textbook solution. But it contains all of the ideas of CT. And certainly, how they apply to light fields, all of these ideas are there, as well.

In light field refocusing, you need to have these filters to prevent aliasing. In light field refocusing, you need a baseline to sample your light field, et cetera. So for those of you working on final projects, you should really try to think about the connection with tomography.

RAMESH RASKAR: Are you going to talk about the teeth?

DOUGLAS LANMAN: Oh, yeah. Sorry, I almost missed that.

So the whole assumption that we made at the very beginning is that we could find some wavelength of light, or electromagnetic radiation, that travels on a ballistic trajectory through our object. Right? So our object becomes semi-transparent and simply absorbs. But that's not true for the human body, for the most part. Inside bones, skin, and in various cells, that's for the most part true that X-rays travel on a ballistic trajectory. There is some scattering for bone you need to handle.

But can anyone guess what these are here?

AUDIENCE: Those are fillings.

DOUGLAS LANMAN: Yes. So the problem with fillings is they're metal. And so they, for the most part, scatter and absorb X-rays. And so you'll see we get these artifacts where essentially we have missing data. And so when we reconstruct, we get these halos around the objects.

Yeah.

AUDIENCE: I think they're called starburst in the [INAUDIBLE].

DOUGLAS LANMAN: OK. So you get these starbursts, which have to do with your number of samples you took. Because you can see the starburst pattern, like in the 30 degree. If you have a limited angle, you can actually get starburst patterns. You can see it also in that one on the left.

And so then the question is, how you handle these occlusion functions? And there are some simple and more complicated methods to handle this. But modern CT scanners, for the most part, can now scan fillings using relatively simple algorithms.

And so now, putting all this together, we've solved the problem in flatland. So now, we can build a CT scanner to do the volumetric scanning task. It's very simple. We're just going to solve the problem repeatedly in two dimensions.

And so how many here have had a CT scan? Anyone? No? Right. The older ones of us have had CT scans, I guess.

So I think this is great.

RAMESH RASKAR: The volume? We have the volume for this.

DOUGLAS LANMAN: So if you go to YouTube. This is a CT scanner opened up. So maybe I should stop it and go back to the beginning, just to explain what you're seeing.

RAMESH RASKAR: Put the audio here.

DOUGLAS LANMAN: Oh. Let's see. We'll see.

RAMESH RASKAR: The audio is very important.

DOUGLAS LANMAN: So the radiologists at the hospital have to take apart the CT scanner periodically and inspect it for damage, because it's essentially a high-quality motor, like a jet engine. So you'll see in a minute why General Electric and other industrial jet engine manufacturers make CT scanners.

But this is the CT scanner with the cover removed. And so basically, the patient lays on this gurney. And they're slowly translated through this doughnut. and So you'll see, this is the x. It's a little blurry in the still, but see this is the X-ray source. It's a point source, so this is a fan-beam projection. And over here, with a bunch of fans on top of it, are the X-ray detectors because they generate a lot of heat.

And so basically, this thing will rotate around the patient as the patient translates slowly. So in the frame of the patient, the X-ray source will move on this helix. And at each orientation, it will take a fan-beam projection. And you can resample that data, if you have a dense enough helix, to be just a set of parallel-beam projections. Each of those independently apply the inverse radon transform.

That's not exactly what they do. But if you were to implement this in MATLAB, that would be basically good enough. And so then, you'll get this series of cross-sectional images. Then, you can simply do segmentation and thresholding, for instance, high density regions correspond to bone, other properties correspond to blood vessel networks.

So that's the basics of the CT scanner. So the radiologists and the techs often have to take this apart to inspect it. And so you'll notice this thing is quite scary with the cover off. And you can listen to what the radiologists are saying as they're inspecting this.

So again, it spins up.

[SPINNING SOUND]

TECH ON VIDEO: Shit.

[LAUGHTER]

DOUGLAS LANMAN: So if you ever get a CT scan, that's what's happening. But of course, the safest place to be is inside the doughnut. Outside's the problem.

AUDIENCE: Outside the room.

DOUGLAS LANMAN: Outside the room is a better place. You can see that there's a lot of problems they have to solve, mechanically. So they have-- it's a lot like a hard drive. They have optical encoders here, which track the rotation velocity. And then, they have what's known as slip rings. Because basically, one of the evolutions, believe it or not, was just getting the data off these detectors. And so actually, the bearings are used as a circuit to just transmit data across the bearings. And believe it or not, that was a big insight and a huge patent.

Before, they actually could only rotate once, and the wires would get tangled up. So having the slip ring bearing transfer mechanism was the big evolution that made million-dollar CT scanners possible. And so most of these cost $1,000,000.

And as you're watching this, just to point out, again, the story I'm trying to give you is that this very basic idea of tomography and the inverse radon transform gives you almost all the medical scanning devices that we were so proud of in the 20th century. You know, of course, X-ray has given you CT. If you move to gamma rays, you get SPECT. If you look at positron emission tomography, it's just a different wavelength of light. And again, if we move to ultrasound-- so we become acoustic-- that's really when you get into ultrasonography, which is a reflection mode version of this.

You can reformulate all of these equations, instead of to be transmission, to be a reflection mode. And that's [INAUDIBLE] but I actually won't go into it.

So the takeaway message is choose your wavelength, apply inverse radon transform, and you too can image non-invasively.

RAMESH RASKAR: So just as a thought experiment, can you see a relationship between this crazy jet engine and an electron camera? We'll come back to it, but--

AUDIENCE: Sort of.

RAMESH RASKAR: --just keep that in the back of your mind.

AUDIENCE: Just take the Fourier transform.

[LAUGHTER]

RAMESH RASKAR: That's it. That's all that is.

DOUGLAS LANMAN: That is it.

RAMESH RASKAR: And imagine now, the light field camera is compact and something you can carry in your pocket. The question is, can you carry the CAT scan machine in the future, in your pocket?

DOUGLAS LANMAN: Right. So we'll come back to Ramesh's puzzle at the end and see if anyone's solved it.

[LAUGHTER]

So before I go on, there's a slide later that I need to explain this for. I only explained the frequency domain version of this. And so for various reasons, if you don't have a fast FFT-- which you do, of course, now-- there's a purely spatial domain algorithm. So remember, our algorithm was to take all of these projections, take the 1D Fourier transforms, and populate our Fourier transform in inverse. But you can actually do this without ever taking a Fourier transform.

Does anyone see how you'd do that? So it's a purely spatial domain algorithm. You get the set of projections and you directly reconstruct the density function.

AUDIENCE: You just accumulate the projections?

DOUGLAS LANMAN: How so? Let's go back to our image.

So if you can tell, I overuse the Socratic method. But I like guys to figure things out, because I already know the answers.

So you give this data. There's no Fourier transforms now. How are you going to get back to this?

RAMESH RASKAR: Without Fourier transform, this time.

DOUGLAS LANMAN: Right? This is problem B on the final. Part A was Fourier transform. Now, you're not allowed.

RAMESH RASKAR: On The midterm, you mean.

DOUGLAS LANMAN: The midterm, OK.

RAMESH RASKAR: Which is next week-- there will be a question on tomography.

DOUGLAS LANMAN: Yeah.

AUDIENCE: I'm just shooting in the dark.

DOUGLAS LANMAN: That's great.

AUDIENCE: But I guess, I'd take the first [INAUDIBLE]. And since that-- no, wait. Actually, rotation [INAUDIBLE]. So I'll take the first row like this.

RAMESH RASKAR: First column.

DOUGLAS LANMAN: Row or column? What do you want?

RAMESH RASKAR: First column is the first photo.

DOUGLAS LANMAN: First column is first projection.

AUDIENCE: Oh. OK, first column.

DOUGLAS LANMAN: OK, you got first column.

AUDIENCE: And then, I'll basically create rays that are representative of those individual vector densities and just shoot them back out.

DOUGLAS LANMAN: OK.

AUDIENCE: Then, I'll take the next column. And then, I guess it's a slightly different angle. So I'll--

DOUGLAS LANMAN: Just shoot them back. So you're going to start with an accumulator, again, that's all zeros.

AUDIENCE: Yes, it's all zeros.

DOUGLAS LANMAN: That's your initial estimate.

AUDIENCE: And I'm going to take one line of that and basically just fill the entire thing with that.

DOUGLAS LANMAN: So you're going to take that value and just replicate it.

AUDIENCE: Basically.

DOUGLAS LANMAN: OK.

AUDIENCE: And then, I'll do that for the second one, but at an angle.

DOUGLAS LANMAN: OK?

AUDIENCE: And I'll normalize it.

DOUGLAS LANMAN: Yep.

AUDIENCE: And hopefully, I'll get something that looks like the skull.

DOUGLAS LANMAN: OK, that is the algorithm. It's called filtered back-projection. Can you prove why it's correct?

[LAUGHTER]

Does everyone understand why that would work, even intuitively? Intuition works pretty well here. Imagine I just had a big circle here. It's absorbing 50%. It's just a circle, right? So what I'll see here will just be a big rectangle, right?

RAMESH RASKAR: Big cylinder.

DOUGLAS LANMAN: Right. At some value, it'll be 50% everywhere inside and 0 outside. Right? So if I use Kevin's algorithm, first, my accumulator will be all-- everything in here will be 50%. Then, I'll turn, and the center will start getting higher. So I'll get that starburst pattern we saw earlier.

So intuitively, it works. But mathematically, why does it work? And I think, again, now you can use the Fourier transform, it's easy to see. For those of you familiar with the Fourier transform, is it's pretty direct. Let's go back to-- so any time you have a Fourier domain problem, you always use the same tricks. Take slices, or use Fourier transform pairs.

And so we know that the slice here is equivalent to the projection. So if I had some slice, some function-- again, the slice being this value-- and I take its inverse 2D transform of just a slice, what would I get?

AUDIENCE: [INAUDIBLE].

DOUGLAS LANMAN: Yeah. It's not easy to see. So give me some-- so if I just start and I have my accumulator-- I hope you can see that. I have kx . I have ky. And I have some slice, we'll say, along x.

RAMESH RASKAR: There's a switch right next to you on the right side.

DOUGLAS LANMAN: And this function along-- if I was just to plot it-- I'd have kx versus what we call s of kx. We'll say it's a continuous function. So I'd have some function here.

And now, what I'm asking is if I take the inverse Fourier transform, what does this function look like in the primary domain, which is x and y?

AUDIENCE: It's constant between one of these other dimensions [INAUDIBLE].

DOUGLAS LANMAN: Right. Because this guy, you can write as some function F of kx times the delta function of what? Of y, right?

AUDIENCE: Yeah.

DOUGLAS LANMAN: And we know that the Fourier transform pair of a delta function is--

AUDIENCE: It's all constant.

DOUGLAS LANMAN: It's all constant, absolutely. Great. So we get the 1D Fourier transform is this thing. And then, the values are simply the inverse Fourier transform of the 1D function, here. So we get some different values here.

RAMESH RASKAR: Bright in the middle and not so bright in the center.

DOUGLAS LANMAN: So this gets to why Kevin said he may have already known the algorithm. But this is mathematically why we smear. And now if we rotate, we have this one, which now comes in as a new smear along a new angle that you had the wrong way probably, but-- and so it's linearity. You can write the total Fourier transform, F of kx, ky is equal to the sum over all of your angles of this function F of, we'll just call it the angle of delta of another function of the angle.

And then, by linearity of the Fourier transform, we come over here. We get a sum over angles of the inverse Fourier transform of that thing. Linearity of the Fourier transform, that gives you-- does that roughly make sense to everyone?

So just by applying properties you already know of the Fourier transform, you can argue why it's correct using the Fourier projection-slice theorem. But if the end result is you don't need to take any Fourier transforms.

You still need to have that sharpening filter, though. Because you can see that you're going to build up too much weight in the center. But other than that pre-filtering step, that sharpening step, filtered back-projection will solve the problem.

So this algorithm I just described, rather than using the inverse 2D Fourier transform, it's called filtered back-projection, because the first step is you high pass filter all of the projections. The second step is back-project them. You smear them through, as Kevin described. So there's two ways to solve the inverse radon transform. You can take the 2D Fourier transform, or you can use the filtered back-projection

So now, we'll move on to so some my own work I'm just going to plug. So the idea Marc Levoy had in presenting this at SIGGRAPH is to take inspiration from medical and scientific imaging, so we can make better cameras. If we're computational photography through optics people, how can we prove this idea?

And so one of the limitations Ramesh identified with computed tomography, why we can't put it in our pocket, is that we have a lot of moving parts. We have these X-ray sources moving around the object. So the question is, how can we remove the moving X-ray sources? And so just as a thought experiment, I'll quickly explain what we did.

Basically, this is another version of CT you often find, is where instead of having a rotating point source, we simply have a linear array of sources. And then, we have a film plane [INAUDIBLE] X-ray detector on one side of an object. And then, we simply switch these sources at very high speed. Because only one source can be on at a time, or else we'll have to projections overlap.

And so there's a limitation with this design, as you have limited angle. If you think about this as a CT scanner, we've only gone over some small baseline. But using modern tomography, that's not such a big deal. You can use limited baseline reconstruction.

And so the main limitation with this type of device is how fast you can switch the sources on and off. And it turns out it's very difficult and expensive to do that. So what we'd like is to have an algorithm where we simply have an array of point sources on all the time. But what you'll have is on the detector, you'll have a linear superposition. You'll have more than one projection overlapping. So now, you have an interesting mathematical problem, which is how you invert those series of projections to get back to your tomographic data set. And as the number of point sources increase, that inversion becomes ill-posed.

And so what we did is we inserted a mask here. And so this is basically a lead panel with holes drilled in it. And so now, you can actually optimize this pattern of holes so that the inverse problem is well-posed. So again, our goal is to get the set of projections, but they're all overlapping. So now, we're going to design a mask to invert that system.

And so if you want to look up this, we have a SIGGRAPH paper on this called*Shield Fields.* But that's the basic idea. By putting this mask here, you can avert the system.

And if you think about it very simply, it's like a 3D TV. You could cope what's called a parallax barrier here. So you just have a series of pinholes. And then, if each pinhole is separated by the size of the image of this array, then you just get an array of pinhole cameras and you could invert that system. But of course, it doesn't pass much light. So you can optimize it and find better masks.

So because we don't like working with X-rays, because I didn't want cancer, we built this invisible wavelength. And so what you can see is this is like a CT scanner built for $100. So we have an array of LEDs, which serve as our X-ray sources. We have an opaque object, a wooden mannequin. And then, we have what simulates a large format sensor. So this is a trick-- maybe some of you are using this in your homework, as well.

To make a really large detector, you can use a camera and just a sheet of paper held between glass. So we take a photo from behind. If I was to turn one LED on, I'd see a shadow. To get a picture of a shadow, it's as if I had a giant sensor. So that's a cheap way of making huge sensors. And then, we put a thin sheet of glass here with our high-frequency mask.

And so now, we get to the computer vision problem. So that's what it looks like when the lights are on. So you can see the multiplex shadows. And then, this is the magic pattern we actually used, which you can see why, if you read the paper.

But again, if you just think about this as a mathematical problem, we have a set of projections that are all superimposed. We're trying to invert and get the individual shadows. And it turns out by putting this high-frequency pattern in, that inversion becomes well-posed. So I'll leave it to you to think why.

RAMESH RASKAR: You can also-- if you go back to the previous image. You can also think of this as if you're standing in front of an X-ray machine. And every one of those shadows would be the shadow you would get in front of the X-ray machine. And as you move the X-ray source, you will get one of those shots. They will move.

Except here, all the sources are on at the same time. And so you are getting this simultaneous projection of the X-ray. So the question is from this one photo, how can I resolve all the original shadows?

DOUGLAS LANMAN: And so here, you can see this is what the data set looks like. So you don't really see anything there. But then, just by inverting that system of equations, you can actually pull out, from that single image, all the individual shadows. So it's as if you have 36 projections all in one image.

So the key advantage here is we never strobe the light. So we can record this as fast as the camera can record images. So for really high speed tomography, instead of rotating or strobing the light sources, we can just have a light sources on all the time. We're only limited by the frame rate of the camera.

And of course, from these projections, you can use the visual hole algorithm and reconstruct. Again, this isn't exactly tomography, because it's an opaque object. But it's essentially the same idea. We're doing that filtered back-projection. And so you get this reconstruction.

You can see, looking at the mannequin, there's this starburst effect you described earlier, where you have this phantom because of the limited baseline. And so again, that just has to do with the algorithm used to reconstruct. If you had some prior on the smoothness of the object, you could take that out.

But again, usually, you'd have to strobe those lights on and off. So you'd have to take 36 pictures. Here, we just take one. So that was the inspiration we took from the medical imaging world. And it turns out all of this can be applied to take light field photographs, which is what Ramesh did in 2007, by putting that high frequency mask inside the camera. So you can see all of these ideas link back together. And they all start with tomography.

And so now, let's take a quick plug. We have a recent paper where we took this idea in yet another direction. And so the idea is that we're going to build this device. I'm going to step back--

RAMESH RASKAR: Doug, did Matt do a presentation?

DOUGLAS LANMAN: Have you already seen this?

RAMESH RASKAR: Yeah. I think he showed it. So you can go through it very quickly.

DOUGLAS LANMAN: OK. So the basic idea--

RAMESH RASKAR: You can describe the connection between the two.

DOUGLAS LANMAN: Yeah. So the basic idea is that we're going to take this and build an LCD screen that can sense depth. And so the trick is pretty simple. Ignore the LCD screen. We're just going to build the exact device I just described.

So we have an LCD panel. Instead of displaying an image, we're just going to use it to display [INAUDIBLE]. No big surprise there. Instead of printing the mask, we use the LCD display mask.

And then, behind the liquid crystal panel, we'll have a diffuser. Inside LCDs, you already have diffusers behind the liquid crystal. And so again, we'll have that giant sensor. We'll have a diffuser, our LCD panel, and some cameras behind it. And then, we'll be photographing the world.

And so it actually turns out that you can capture the light field with that setup. So by just building that setup, using the LCD as a mask, you can capture the light field. And then, you can just time multiplex.

On the first frame, we turned the back light on and we use the LCD to display a picture. On the next frame, we turned the back light off and used the LCD to display a mask and get the light filled and repeat at very high frequency, so that your eye can't perceive it. And so now, you're getting an LCD panel that can do multi-touch, like normally see. And then, as the user moves the hand off the table, you can sense depth. And so now, instead of having normal multi-touch, you can also have a z-axis. You can pull things off the table.

It was a very simple proof-of-concept demo. But it shows that with just that basic idea, now, we've gone into the HDR direction. Again, Matt Hirsch did all this work. So you should really talk to him to see the demos.

And so we're presenting this at E-Tech at SIGGRAPH Asia. So if you have any good ideas on demos you could use with a depth-sensing screen, please email me or Matt. And if you want to implement them, all the better. We'll give you credit when the reporters ask, who was the clever guy who programmed that?

So here, you can see just a very simple CAD Explorer. You can choose your model, which is a touch interaction. You can move your hand around and it just controls the rotation, translation and scale matrix being applied to that model.

But again, what's neat here is you don't see any cameras. There's nothing hidden in the bezel. So it's to the user, it's surprising that you're getting the light field. And the demo I'm not showing is that you can do light field transfer with this, if you know what that is.

Are we doing OK on time?

RAMESH RASKAR: I think so.

DOUGLAS LANMAN: I can move faster. I think Ankit already presented this project, right?

RAMESH RASKAR: Yeah, he presented this one. Yeah.

DOUGLAS LANMAN: OK, so I'll just go through this really fast, then. So again, I mentioned that once the computer was invented, tomography became easy. Because they knew the radon transform and they could invert it. But if you don't have a computer, how can you take a cross-sectional image of a patient?

So if I just have an X-ray source and an X-ray film, and I had the patient sitting on the gurney, I could turn the source on. I'd get a projection. But the main problem is I'm really only interested in seeing a cross-sectional image so I can identify a tumor in the brain, or something else. But everything else out of that plane, I don't really care about. And the problem is if I just have this perspective projection, everything out of focus will also be imaged on the sensor. so

The question is, how can I get an image just along a slice through my volume without using computation? And so I think Ankit explained this. But 100 years ago, they did this and it's called laminography. And it gives you almost identical results, purely with mechanical means.

So the idea is simple. You simply take the X-ray source and you mechanically translate it from left to right. And then, at the same time, you mechanically translate your film from right to left. And so it picks some point in the patient and you're stationary. You'll have some ray going through that point. And then, if I turn on the next source, we'll have some other ray coming in, as long as I move the sensor so that same pixel is illuminated. Then, that point will be focused, but everything else won't.

So it's a clever trick to get a CT scanner without any computers. And so we used this to publish a paper at ICCP.

RAMESH RASKAR: Actually, let's go back to that one. There's a minor difference, though, between doing this, achieving a cross section image-- this method, versus doing the whole spinning and multiple projection thing. What's the difference between the two?

AUDIENCE: You don't get as much depth.

RAMESH RASKAR: Sorry?

AUDIENCE: You don't get as much depth emitted by what--

RAMESH RASKAR: The plane of focus-- we can control that based on how close the X-ray source and the sensor is. So you can get a pretty narrow depth of field. But what happens to stuff that's outside the depth of field?

AUDIENCE: [INAUDIBLE].

RAMESH RASKAR: It's still going to be blurred. And it will be still part of the image. So you don't really eliminate what's above and below your plane of focus. It's just out of focus, but it's still there.

DOUGLAS LANMAN: Right. You can make it strongly out of focus, but your contrast will be lower in the CT.

So Ankit Mohan had a great idea to apply this to photography. And so her already explained this to you, so I'll go quickly through this. But the basic idea is if I have my normal thin lens equation, I have some plane in the world on the left, and my sensor on the right, the lens maps those. And as I stop down the aperture, the blur circle for the autofocus point decreases. And in the limit, if I just have a pinhole, then I get everything in the world in focus.

And so if you look at your iPhone camera, your cell phone camera, the apertures are so small that you're essentially taking pinhole images for everything. And what we'd like to do is take images with a cell phone camera that are as if they have a large lens. Because really, when you pay a professional photographer to do your wedding photos, it's really the blur. The blur is one of the big tricks they have.

RAMESH RASKAR: You're paying for the blur.

DOUGLAS LANMAN: You're paying for the blur, in a way you're. Paying for that big piece of glass and a bit of the talent to focus on the foreground objects so that the background has this really beautiful blur. That's one of the things that you'll immediately notice when you look at a wedding album. It's like, oh, gee. They just blurred the background. Great.

So the question is, can we get nice defocus, nice blur, with a cell phone camera? So in a way, we're going to make our camera worse. Right? A cell phone camera takes everything in focus, which maybe that's a good thing. But it doesn't have that aesthetic feel of having a nice defocused background.

And so the main trick-- and again, I'll credit Ankit with this-- is a nice insight. Well first, let's consider moving the pinhole. Because we just talked about laminography, and that's like moving the X-ray source. So our pinhole here is like our X-ray source. It's our center of projection.

And so if you translate a pinhole-- and again, we have two points in the world. They're mapped onto our sensor. Nothing's moving, other than the pinhole. Then, you see that without moving the sensor, you get these two blur circles. But one of them is slightly larger.

And so again, in the interest of saving time, you can go through a mathematical analysis. But what's really important here is it's the same idea as laminography applied to X-rays. Basically, this is like our X-ray source and this is like our film plane. And if we translate them at a certain velocity ratio, we can actually put certain planes in the world into focus. It's as if our patient has moved to the left side instead of the center.

Now, we have the X-ray source in the film. And the velocity now becomes in the same direction, instead of opposite directions. But if you think about this, it's identical. It's identical to the laminography we saw earlier.

And so then, what we can do is we can just choose this velocity ratio. And again, if you say, look at the blue ray here, it'll always move to the same pixel on the sensor. Because we're moving the sensor at just the right speed so that pixel always stays on the blue ray. But the red one is getting blurred. And by changing that velocity ratio, we can then focus at a different plane, a very simple idea.

And so if you look in the paper, what we've done is just recreated what a lens does, but over time. And so everyone knows this equation, the Gaussian thin-lens equation. The distance to the image plane and the object plane are at inverse reciprocals equal to the focal length. Hopefully, everyone knows that.

So normally, this focal length is selected by your chunk of glass. You choose some refractive index and some curvature, and that gives you the focal length. Here, you have a virtual focal length. It's very simple. It' the register distance, which is the distance to your pinhole to the sensor, times the velocity ratio. So as we change this velocity ratio, we can make arbitrary lenses. And then, it gives us a virtual F number.

And so then, just to blow through this, this is the prototype we built. So again, you can see the story here. We knew about laminography, which is a medical imaging topic. We said, oh, we want to publish something in computational photography. Let's build an optical version. And here you go.

So you have two translation stages. You can see those on the bottom. These are from [INAUDIBLE]. And then, we have lens and sensor. And then, you can see, this is the type of image. If you stopped down a lens, this is what your cell phone would give you. And then by adjusting that velocity ratio and making it small, we can focus on the foreground, the middle, or the background.

And so we're only getting blur in 1D, so there's some more complexity. But that's the basic idea. So you see that path. We start with medical imaging, we published a paper in-- not SIGGRAPH, but a graphics journal. It's a pretty straightforward algorithm.

And so now, just to finish up tomography, applying it to optics, like we did, is obvious. And it's also obvious how you would apply this to other applications. So this is an acoustic version of tomography, because we're just changing the wavelength, in a way.

And so the idea is that, say we have some underground region we're looking for oil deposits, or more noble things, like looking for the skeletons of dinosaurs. And so to find the density function, it's very easy. You just set up a set of explosive charges to generate pressure waves. And they're going to travel, hopefully, if we do things correctly and we model it. If we select things right, then they mostly travel on ballistic paths to a series of microphones. And then, again, we can just treat this as a limited baseline tomography and invert that system. And here, at each pixel or Voxel, we're reconstructing the velocity. Because the velocity is proportional to the density of a material. And so then, we can reconstruct underground objects using tomography, using explosives and microphones, which is a cool idea.

RAMESH RASKAR: So you can do this for anything. You can do this, not just underground, but this is how ultrasound works, as well, ultrasound scanners.

DOUGLAS LANMAN: Right. And so then, Marc Levoy proposed this in the course, which I always liked, which was, if you could convince the Italians that you can set up some explosive charges in their subway-- which they might have concerns with-- and use some microphones, you might be able to reconstruct underground Rome just from that algorithm, at least a rough estimate. And so it's not bad. I mean, if you go out to these archaeological sites, they're huge. And they often discover new sites underground. And so this is another way to compete with other scanning technologies. It's very inexpensive.

And so then, as long as I'm doing OK on time, I'll quickly blow through.

RAMESH RASKAR: Yeah. It's fine. Go ahead.

DOUGLAS LANMAN: So we started with tomography, which assumes everything travels on a ballistic path. And so now, let's start loosening our assumption on that ballistic path. So the first assumption is, let's allow things to refract. So instead of having an object which just attenuates rays, let's assume we have a lens here that's going to bend rays and diffracted them slightly. So we have weakly-refractive media. So we've loosened the prior model of the object.

So if you try applying the inverse radon transform to the data set you gather, it turns out we won't be able to reconstruct this. That assumption that things move on ballistic paths is essential. There's some interesting work recently at SIGGRAPH on Schlieren tomography, where they do manage to do that with certain assumptions.

But anyways, if we loosen this to allow refraction, then the algorithm changes. And so this is very clever work. Again, this is in Slaney and Kak, if you're interested. But the idea is, let's illuminate the object. Let's make a parallel beam set up again.

So we'll start with some monochromatic plane-wave. So we have some laser that's generating one wavelength of light. It's coherent, traveling towards our object. And then, it arrives on some detector. And so we get our projection, again. Again, what we're trying to reconstruct now is the index of refraction, not the absorption. OK?

So we receive the scattered wave. And it turns out we also need to measure the phase. So you'll have a second reference beam. So we'll essentially be taking a hologram of the object, for those of you familiar. But you can ignore that. Basically, you'll have a plane-wave projecting through the object, creating some scattered field.

And so if you run through the math, you get a new result, which is very interesting, before your projection-slice theorem changes. It turns out in your frequency domain, again, we take the 2D Fourier transform, our index of refraction. And in the transform domain, a plane projection, a parallel-beam projection mask to a curve, an arc in the frequency domain.

So now, to do tomography, we could do what we did before. We can rotate the source, or the emitter and detector, around the object. And then, we just have to know what this curve trajectory, which is just dependent on the wavelength of light. And then, we'll populate our accumulator, as before. Inverse transform done. Does that makes sense, hopefully, to everyone?

But there's a really clever trick. Again, no moving parts is a usual theme in these things. To reconstruct topographically this object, we have to rotate something to populate this accumulator. But it turns out the clever trick here is that this arc is dependent on the frequency of illumination. So does anyone see, without looking at the thing on the right, what you do?

AUDIENCE: Yeah. You change the frequency.

DOUGLAS LANMAN: Change the frequency. Great. That was a very good insight.

So if you just stuck through the frequency slowly, you'll stick this out. And you'll get at least half of the transform. But what if I told you that was a real-valued function? Do you know what the property is?

AUDIENCE: Symmetric.

DOUGLAS LANMAN: What kind of symmetry?

AUDIENCE: [INAUDIBLE].

DOUGLAS LANMAN: Conjugate symmetric, right?

AUDIENCE: Yeah.

DOUGLAS LANMAN: So we know that this function is conjugate symmetric because it's a real-valued function. So getting only half of it's enough. Because then, we can replicate it. So again, tomography is basically remembering signal processing. If you remember all those transform pairs, you'll have a good luck in your final project, if you decide to use it.

And so now, we can take a sequence of images where we just vary the frequency parameter. We'll populate this guy and use conjugate symmetric symmetry and inverse transform. And so then, you get a new result, which is almost as fundamental in this field as the Fourier protection-slice theorem, which is that a white light hologram basically can reconstruct the index refraction of an object. Because we can use superposition, again. We can illuminate with multiple frequencies at the same time, as long as we can resolve them. Then, we'll get all of the data at once.

So you basically have a broadband clean, coherent way of traveling to the object. And from that, you instantaneously get the index refraction.

RAMESH RASKAR: So the refraction and diffraction are being used interchangeably here.

DOUGLAS LANMAN: Yes.

RAMESH RASKAR: Do you know why?

DOUGLAS LANMAN: Why refraction and diffraction?

RAMESH RASKAR: Yeah. I mean, you're trying to reconstruct something in presence of refraction using a diffraction theorem.

DOUGLAS LANMAN: Right.

RAMESH RASKAR: Do you know what's the reason for that?

DOUGLAS LANMAN: I don't think I understand this at a deep enough level to--

RAMESH RASKAR: OK.

DOUGLAS LANMAN: I mean, fundamentally--

RAMESH RASKAR: I just realized why people [INAUDIBLE] confusing the two.

DOUGLAS LANMAN: Yeah. I think-- Yeah, I'll have to get back to you on that. This is getting beyond my knowledge.

So this is what a reconstruction looks like of a singular object. So now, our artifacts have this arc trajectory. But again, what is interesting is if you look at that Schlieren tomography paper, it means you could get real-time reconstruction without strobing any lights by just doing broadband holographic imaging of your object.

And then, this is why I mentioned the filtered back-projection earlier. That was the frequency domain reconstruction transform. So it turns out there's a purely spatial domain transform for this one, as well. And so this one, I don't know. But the 2D Fourier transform of an arc maps to this strange depth-dependent function you see here.

So that's what filtered back-projection becomes. It's smearing that along these paths that become wider, as you go in depth. Because that's what the inverse Fourier transform is of this arc pattern. So you can do filter back-projection, but it turns out it's much more computationally expensive. So they generally do this using frequency domain, I think.

So that was weakly-refracting. But now, if we move to putting LEDs against your skin, now, we're talking about things that are strongly-scattered. So we've moved from complete, just absorbing along a ballistic trajectory, towards refracting a little bit. And now, we're just scattering completely.

RAMESH RASKAR: So imagine putting an LED on your finger and you want to see the bone inside your finger with just visible light.

DOUGLAS LANMAN: So this is an example you can see. This is a cross-section to our finger, we'll say. And we have photodiodes all around our finger. And then, we put this little fiber optic right against our finger and illuminate it. And so, you'll scatter light all through that volume and you'll measure its intensity, at some points, coming out on the surface.

So this is known as diffuse optical tomography, because diffusion is the key process we're trying to invert. And so in all of these cases, we're just creating this inverse system. We're creating something that we then apply the inverse Fourier transform to invert, or what have you.

So in strongly-refractive or scattering media, you end up with a very difficult model to invert. It's ill-posed and nonlinear, which means it's very hard. So what they generally do is they use some tricks to get to bootstrap inversion. So if you can start with a good initial guess of what the-- here, we're looking, maybe, for optical density. We're looking at how dense the material is. That's the 2D function we're trying to reconstruct in a cross section.

And if we had a good initial guess for that, then we can use a forward modeling process, where we perturb that density function and see how well it predicts the values we found. And you put that in an optimization framework. So we do some gradient descent and we optimize our reconstruction of the density function, so that the observations match.

The predictions match the observations. That's the general inversion framework you have, some of nonlinear gradient decent algorithm. But the question is, how do you get that initial guess Of the density function?

And so you see, it says it right there. You can use some other process that is ballistic that's correlated with your density function. So what they generally do here is they use time-of-flight. They strobe this, record the time it takes to travel through-- maybe not in this specific example, but you can imagine doing this. You turn the light on very quickly, and look at the time delay, and that will give you an initial rough estimate of your density function.

So generally, borehole tomography also does this. It uses time-of-flight to invert the process, but also to get the initial guess. And then, this is used a lot in non-invasive medical imaging. So if you don't want to do a CT scan because the patient has had too much dosage of X-rays, you can, for instance, put electrodes on their body, like you see here, and at least get a rough something not as high quality as a CT scan, but sufficient to make the diagnosis.

So here's an example of how you applied diffused optical tomography for diagnostic purposes. So here, we have two twins, two babies, the left one versus the right one. Again, they're twins. And then, one of them had a-- this is a specific thing-- left intraventricular hemorrhage. So they had a blood vessel rupture on the left hemisphere.

And you can see here. So they attached all these electrodes. And they generate these time-of-flight images for the initial guess. And then, they look at the conductivity, and they reconstruct this density function. Actually, they get two density functions. One is blood volume and oxygen saturation. And so here, you can clearly see the hemorrhage, because this is the twin that has the hemorrhage. This is the one that doesn't.

You can see there's a lot of blood volume. That's a good indication of a hemorrhage, but not a great one, because your prefrontal cortex has a lot of blood flowing to begin with. But then, you also can see low oxygenation, which tells you that blood hasn't been refreshed in a while. So it's a big, big pile of blood that's not oxygenated, which means hemorrhage.

So this is just an example where a CT scan really wouldn't show this, because the density would be the same. And so by using a diffuse optical tomography, you can get a good reconstruction. But because it's diffused, you're only going to get low-frequency detail.

And then, here's another example where you can see the non-invasive part is starting to be weakened. This looks pretty invasive to me, because I don't want to take those electrodes off. So here, you can see this is off Wikipedia. But this is just an example I found, I thought was interesting.

Say we want to take CT live images of the heart to study some sort of function. We could just coat the patient in electrodes, wire them up, strobe these electrodes on and off. And the density function we're reconstructing here is conductivity and resistivity, which we can then use to look for various problems-- watch the heartbeat, et cetera. And so this problem is actually very difficult. It's called the Calderón problem. And inverting this system is very challenging, which is why the images are so low-quality.

But again, you can see this theme is just being extended and extended. We're taking tomographic-like projections of the data set. But our reconstruction equations are not just inverse Fourier transform because of the scatter.

And so then, you might think about, well, if we're now going to make a computational photography project out of this, what if I have scattered light and a light field photograph, then I end up with this problem. So there's your final project-- a plus b equals project.

So now moving. We're almost through the talk now, but a lot of this was on tomography. But the general idea of getting a cross-sectional image reappears throughout the sciences. And so in biology, rather than using X-rays, they do this optically, and they just use 3D deconvolution.

So this is a very simple idea. Probably, all of you understand it already. But the idea is, say I have some specimen in a microscope that has some depth to it, and I want to reconstruct it in 3D. So how do I do that?

Well, the simple observation is to model the image formation process. So what if I just had a single beam in space and I focused on it? Then, I see this image. That would be like the impulse response function for focusing on this plane. And then, if I focus slightly below our point source, I'll see a defocused blur. And as I defocus more and more, I see a larger blur.

So this is known as a focal stack. I take my microscope and just translate the specimen slide without changing the optics any. And I create a focal stack, this three dimensional impulse response to a point source.

So if I look at this as a function of depth, if I just take this slice through the point spread function, and then I look as a function of depth into the object, I have this three-dimensional blur, where you can see it goes out with these two inverted cones. And so now, we can model the image formation process. So you assume that the specimen you're observing does not scatter significantly, then you can use a linear-image-formation model, which is very simple.

You have some three-dimensional source function, which is just decomposing your object into a series of point sources. Each point source generates a PSF weighted by whatever the intensity of that point is. And that generates the focal stack I see. Does that make sense to everyone? And again, this is only true under the assumption that scattering is not significant, or else this model is not correct.

So generally, what we're doing is we're taking a three-dimensional function convolving with another three-dimensional function. And again, you always use the same trick in medical imaging. The first thing you do is take the Fourier transform.

So if we take the 3D Fourier transform of this, then everyone knows the convolution theorem. To simulate this in the frequency domain, we take the 3D Fourier transform of the object, the 3D Fourier transform of the PSF, multiply the two, and that gives us the Fourier transform or the focal stack. Everyone follow?

So now, if I want to get rid of the blur-- because remember, we talked about laminography. Things out of the plane will be blurred. So I want to a photograph with a microscope, but I want it to be like a pinhole. I want every single thing to be in focus, so I can look at structures that have some depth to them.

But the problem with that is we have the blur. And so then, we just invert this. So we take the 3D Fourier transform of the focal stack, and divide by the 3D transform of the PSF, and we end up with our object, which we then inverse 3D Fourier transform. And that gives us the focal stack all in focus.

So probably, most of you are familiar. Who's familiar and have done deconvolution before, understand it? Really? Deconvolution is a new concept to everyone? Interesting.

OK. Well, I probably should've explained this in 2D, but hopefully, you followed what I was explaining. So if this was 2D, all of these transform become two-dimensional. That's the basic idea.

RAMESH RASKAR: Yeah. So deconvolution is very similar to what you did for your light field assignment, which is shift and add. Shift and add is the [INAUDIBLE]. And imagine if given all those refocused images, you wanted to go back and construct the light field. That would be a form of deconvolution.

And instead of doing it in frequency domain for the assignment, you did it in the primary domain. You just shifted and added [INAUDIBLE]. Shifting and adding in primary domain is same as projection of convolution. Or, convolution is basically [INAUDIBLE].

DOUGLAS LANMAN: So hopefully everyone understood that to deblur, to put everything sharply in focus-- so if I apply this algorithm to the impulse response itself, what would you expect to see? What would these images become? Maybe that's a check to see if you understood this.

AUDIENCE: Won't they all just become the point source?

DOUGLAS LANMAN: All become the point source? Close-- 50%.

AUDIENCE: [INAUDIBLE] point source [INAUDIBLE].

DOUGLAS LANMAN: So I'm trying to remove the blur. If I had a point source-- again, these images are at different depths. Right?

AUDIENCE: Mm-hmm.

DOUGLAS LANMAN: But by applying the deconvolution algorithm, I expect to see certainly a sharp point right here. Right?

AUDIENCE: Yeah.

DOUGLAS LANMAN: But what about if I move just a little in depth?

AUDIENCE: [INAUDIBLE] the other way.

RAMESH RASKAR: Remember, there's a blur here, right?

DOUGLAS LANMAN: Yeah.

AUDIENCE: Right.

RAMESH RASKAR: So convolution means blur. And in most of the cases, deconvolution means removing to blur.

AUDIENCE: So you wouldn't see anything.

AUDIENCE: No, you should see the other image in focus.

DOUGLAS LANMAN: Which would you see?

AUDIENCE: Why wouldn't you see the image on the other depth in focus?

DOUGLAS LANMAN: Because there's just a point in the world.

AUDIENCE: Oh, OK. So you don't see anything.

DOUGLAS LANMAN: You don't see anything, right. So you get the gold star. And so that checks. Hopefully, if you understand why all of these images look completely dark, except for the center one, then that's deconvolution in general.

So then, just to point out-- again, when we're doing this inversion, we're dividing by the Fourier transform of our blur kernel. And so the problem here, if you remember earlier, when we were doing the inverse for transforming and the frequency domain, you really don't want zeros, right? So if I divide by a 0 I'm going to have problems.

But if you look at our focal stack and look at that the blur kernel, we have lots of zeros. And these zeros actually come in due to the numerical aperture of the lens. If the lens sees at a very sharp angle, we can make this have few zeros.

So since you haven't seen convolution before, I think that's probably enough, just to observe that when you're deblurring, you really need this function to not have zeros. So if you're going to apply the concepts we saw earlier and the computational photography idea, the idea would be somehow to modify the optics, add apertures codes, to make this blur kernel not have zeros. So when you do the inversion, you don't amplify high frequencies [INAUDIBLE].

AUDIENCE: If I understand what you're saying correctly, you're taking multiple images at different depths, right?

DOUGLAS LANMAN: Yes.

AUDIENCE: So why can't you focus at different depths? Why do you have--

DOUGLAS LANMAN: That's what we're doing.

AUDIENCE: So then, if you are taking multiple images, why don't you shift your focal plane to different depths?

DOUGLAS LANMAN: That's exactly why we're doing. But it's just like laminography from before. It's a thick specimen. So let's go to a picture. Because pictures are always worth a lot more.

So we have this mandible of some insect. We're trying to get a cross-sectional image. But if we don't do anything, if we just focus at some depth in the specimen and it's back-illuminated, you get this halo. Because everything out of the focal plane is blurred.

AUDIENCE: Yes.

DOUGLAS LANMAN: Right?

AUDIENCE: Yes.

DOUGLAS LANMAN: So now, I can focus at another depth and generate a new image, which is exactly what a focal stack is. But I'll still have that blur presence. So what I'm trying to do is collect that focal stack and then invert that whole system, so that I get, really, what is a cross-sectional image, so that nothing is blurred in it, which is what laminography does, which is what tomography does, in general.

So you can see that theme. Something has to be done to remove the blur. And so in this case, what's done to remove the blur is that division by our impulse response Fourier transform, this deblurring algorithm, rather than using a tomographic-type algorithm.

Just to show you the gist, I think, since it seems that deconvolution is a new concept for everyone, definitely go and read the Wiki entry. Because I think that idea has been beaten to death in computational photography field. So you can definitely get some final projects out of it. You can see how a team can use that.

So here's some nice results. This is from Marc Levoy's work. This is actually a light field microscopy image, but the ideas are the same. You can do deconvolution. You can see it goes from something that a typical microscope would produce. We have this thick specimen, but it's optically transparent. We produce a nice, tomographic image, which we can then volume render. And so if you go to some really expensive commercial packages, this is what you'll see for deconvolution. So you'll go from something on the left, which has all the out-of-focus blur, just some nice, sharp, all in-focus image of an object with a lot of depth here so you can study the fine structure.

And so then, I think we have a little time left. So let's look at some other tricks they use in biology. So again, the theme in all of these-- laminography, tomography-- is always to remove the out-of-focus blur so we get a nice, cross-sectional image, so we can see the tumors, we can see structures clearly.

And so this idea is really nice. And again, I don't have time to show you, but there have been two or three computational photography SIGGRAPH papers, again, from Marc Levoy using this idea. So this theme just keeps coming up. You go to the medical literature, look what they were doing in the '60s and '70s, figure out how exactly it maps onto a camera, and it. You're published. That's basically the algorithm.

[LAUGHTER]

Although, it's becoming harder to do that, because too many of us use that algorithm for papers. So the idea behind confocal microscopy, which you can be proud of. Marvin Minsky is arguably credited with the invention of it here at the Media Lab as well.

So the basic idea is I have some plane in my specimen that I want to get a nice cross-sectional image of. And I want everything that's out of focus to not contribute anything to the final image. So I can do that computationally by doing this deconvolution. And that's like CT scanning. Or, I can do it mechanically, which is like laminography.

There'll be no computation here at all. I'm just going to make sure that the point spread function-- so if I go back to the point spread function, this thing-- if this was just a point, then if you look at this math, if you convolve the point with the object, you just get the object again. So just like laminography, solve the problem mechanically. If we can make this a point, our impulse response is a point, then we don't need to do deconvolution. And that's the main trick.

And so confocal microscopy does this in a very clever way. We have some secondary light source. Put a pinhole in front of it. Now, it goes through my thin lens. And it goes, gets focused based on the focal length, down onto some point in the plane we care about.

So if I was to now take a picture somehow of this plane, then all I'd see would be things that get illuminated by the beam of light. Everything else is not illuminated. So we're already doing OK, because things way over here aren't going to contribute to the blur at all. So that's the purpose of the pinhole. Let's make this clear.

Again, let's think about the mandible we saw earlier. If we have some point that's higher in depth than the plane we're trying to focus on, then the light will be spread over a disk. Does that make sense? We're going to take this cut through our cone of light. So as a function of radius, or depth, away from the plane we care about, it's going to fall like what, the brightness of that point?

AUDIENCE: Square.

DOUGLAS LANMAN: Yeah, [INAUDIBLE]. Great.

So already, if you think about our impulse response we had earlier, it follows 1 over arc squared, now. Right? That's not bad, but can we make it even sharper? Any ideas?

AUDIENCE: Higher numerical aperture.

DOUGLAS LANMAN: Higher numerical aperture-- that will spread the energy out. It'll make it somewhat sharper. That's right.

So the key trick-- oh, you just saw it. A preview, if anyone saw it.

So does anyone think-- so we're only doing things on the lighting side. We haven't considered how we take the picture yet. It turns out, you can use the exact same trick.

So to take the picture of this one point on the plane of focus we care about that, we're just going to put a big photo cell here, a big photodiode. So if I didn't do anything and I put a beam splitter here, this light would just fall on the photo cell over a big region, so you'd have low SNR, if you didn't have a lens here. You'd need the lens, of course.

So if we put a lens, that focuses on this plane as well. So that collects all the light scattered by this point back onto our photodiode. So if we put our photodiode at this point, we get a nice image. But then, if we took a focal stack, then the points out of focus would have intensity 1 over r squared.

But we can do better than that. Because now, we can put a pinhole in front of our photo cell. So if you think about this, our image-- say you take this blur disk and you image it. It Images to a disk, but then we're just going to take that little part in the center of the disk. So it turns out the imaging side, by putting a pinhole on your photodiode, also gives you a function of 1 over r squared. So it actually goes 1 over r to the fourth now, which is quite strong. So that deconvolution isn't really needed anymore.

Does that make sense? And so that's why it's called confocal. You have two pinholes, two focal systems. And they're aligned with each other. But then of course, how are you going to get a picture? You're going have to scan this thing, which can take a lot of time. So you're going to have to move the pinholes on the light source and the detector, in 2D raster scan, the whole object.

So that's really the limitation of this technique. By doing deconvolution, by doing tricks to make the PSF convertible, we can do it all in one shot. But with confocal microscopy, before we had computers, we could do that without needing inversion. So you see the trick, right? And it should be the same trick as laminography and tomography.

And so then, I should finish up, right?

RAMESH RASKAR: Yeah, that's fine.

DOUGLAS LANMAN: So we're getting right near the end, now. So this is used in practice, not quite the way I described. It took a decade or two to commercialize Marvin Minsky's idea. And this is how it's used in biology in practice now.

It's called laser scanning confocal microscopy. The idea is basically the same. You have a laser source, a pinhole aperture on both the source and the detector and a photodiode, and a special means splitter. We'll get to that. What we do is-- really, it's just like CT scans. You don't really just care about the absolute density function. Often, you add some contrast agent.

So if you're looking for a pulmonary edema, you inject iodine into the veins. Iodine absorbs X-rays. And so then, you can look at the vein structures and how they evolve over time.

So they want to look at specific structures inside cells. They don't just want a gross picture of a cell. They want to enhance certain details. So what they use is fluorescent beads, usually with some antigen tags for the biologists. Anyways, I'll skip that.

They basically highlight. They put fluorescent dyes somehow attached to structures they care about, for instance, mitochondria, cell structures, DNA, what have you. And that'll enhance the contrast in the final image. And then, what happens is the [INAUDIBLE] designed to fluoresce at some wavelength which is different than your source wavelength.

And then, you have a beam splitter again, but you have a filter on this beam splitter that only reflects the fluorescent wavelength, not the stimulating wavelength. So that improves your contrast even more. So that's the basic trick behind laser scanning confocal microscopy.

But again, you have to scan mechanically both the light source aperture and the detector aperture. So this is a slow process. But cellular division is also a slow process. So we can create videos of that. For instance, here, you can see typical cellular division process, where I think, they probably accentuated the telomeres. So you can see them separate, I'm guessing. So again, cellular division occurs on a long enough time scale that we can raster scan that without any difficulties and create videos.

And then, just to show you how good the cross-sectional image is, I think this is pollen grain. Here, you can see that this is, again, without any deconvolution, no computation. You're creating this function. And then, you can take a slice of the function. And you can see it evolve on the right, here. These cross-sectional images are very sharp. And again, by using the fluorescent beads, you can enhance the contrast as well. So is this idea taken to its natural conclusion. This is the commercial product you get at the end.

And so now, again, to plug some of my own work-- well, this isn't my own work. But I want to talk just briefly about coded aperture imaging. So we talked about biology, medical imaging, a little bit about geology, and so now, let's not leave the astronomers out.

And so what ideas can we mine from the field of astronomical imaging? I think the main takeaway message that's impacting computational photography is the idea of coded apertures. So if I'm imaging in X-rays, I'm looking for supernova bursts, what have you, it's impractical to build refractive optics. Really, all you can do is attenuate X-rays. You can build lead sheets, like we did in our project, and block them in various ways.

And so then, the question is, how can you image non refractive wavelengths? So any ideas? So the first idea, of course, is build pinhole apertures. So say I want to take an X-ray image of the world, but I can't build a thin lens that refracts X-rays. So I have this X-ray, gumballs. And I just take a lead sheet, drill a hole in it. I have an X-ray film. No problem, I get an image, but I have exposure issues, as always.

So any guesses about how to solve the exposure issues, as engineers? What would be some tricks? Any ideas?

AUDIENCE: Since it's a pinhole, you want something else.

DOUGLAS LANMAN: Yeah, Kevin?

AUDIENCE: Coded aperture?

DOUGLAS LANMAN: Yeah, it's in the title.

[LAUGHTER]

So what would that mean? What would that mean?

AUDIENCE: Choose some pattern that lets in more light, and then deconvolve--

DOUGLAS LANMAN: Nice.

AUDIENCE: --the pattern.

DOUGLAS LANMAN: Exactly. So if we didn't know about deconvolution yet-- which, it seems, the wiki will tell you later-- the first step to getting more light is just make a larger pinhole. So if you don't really care about the resolution of your scene, you can afford to blur. So if we're just looking at stars, we can blur them. We can make bigger pinholes and then deconvolve that, maybe, even though it's difficult to deconvolve a circular aperture.

So that's not a great solution, but we let in more light by a function of the radius squared. But the real solution here is, again, called coded aperture. And the idea is to drill many holes, as Kevin said, in a way that all of the images that overlap are somehow invertable.

And you saw this. Hopefully, you remember back to when I was talking about the [INAUDIBLE] project. We put all these holes in front of the plane so that you could invert that system. And so we weren't the first to discover this, at least for imaging of point sources. And the basic idea is that as long as we design this mask appropriately, so that system of equations is well-posed, we can invert it.

And so here, I'll show you something that's not well-posed. This is like the [INAUDIBLE] setup we saw earlier, where we had 3 X-ray sources on all the time. And those three images were overlapping on the sensor. So if we do this in the X-ray domain, we have three images coming from three centers of projection. They're all just shifted from one another.

So you can imagine if I just had two pinholes and some prior on the scene, it wouldn't be too difficult to separate two images. So that would double my light and I'd probably be able to invert that. As Kevin said, that would be easy enough to deconvolve, probably, although it wouldn't be a linear process.

So then, as we add more and more pinholes, we let in more light. But that inversion, that system of equations, becomes more ill-posed. It's condition number is much worse.

So then, that's where we end up in this field of optimizing coded apertures, which is what our work was on. And so if any of you are considering using masks for light field capture, or other things, you'll very quickly arrive at similar results. And the idea here is to use-- in this case, they generally use something called a MURA code. And it's just designed-- I'll tell you the main fact about the MURA code is that its autocorrelation function is equal to a delta function.

And so if you think about deconvolution, for those of you [INAUDIBLE] that aren't familiar with it, that should tell you why you can do this. But the basic idea is that the image formation process we can now model as linear. We have a sequence of pinholes of varying sizes and distributions to simulate the image we receive. We can evolve our aperture, scaled aperture with our function. This is the image we get on the sensor.

And then, believe it or not, to deconvolve, you simply convolve again. Convolve this blurred image, the superimposed image, with the aperture function. And in the absence of noise, you get back exactly the image. And really, this is the key problem of coded aperture, and something you can investigate. There have been a couple of papers from [INAUDIBLE] and others in recent years, where they applied this.

Again, astronomy applied it to computational photography. And they said, OK, we're going to put coded apertures in cameras. But the key trick is what's the aperture going to be so that the inversion is well-posed in the presence of noise-- paper. So that should inspire you, hopefully, for your final projects.

And then, this has also been applied for tomography. And so you can do exactly what our goal in our project was, which is to have no-moving-part, instantaneous tomography using some of these concepts. The main challenge is doing it when objects are close to the detector. And so that's really the problem that we solved in our work.

So that concludes my talk. So thank you for your attention. If you have any questions, I'll be glad to take them.

RAMESH RASKAR: [INAUDIBLE].

[APPLAUSE]

All right? So light field and tomography-- one in the same thing. Is that clear to everybody? All right let's draw on the board.

DOUGLAS LANMAN: I think that link is the one you're interested in.

RAMESH RASKAR: And so similar to [INAUDIBLE] which was all about shifting and adding, [INAUDIBLE] focusing. So you had a bunch of cameras. And you have this [INAUDIBLE] here. And when you want it to refocus on a particular plane, if you just sum up all these images, then you're focusing at infinity. But if you wanted to focus on the close up, then you would appropriately shift these images in the pattern.

But what is this particular image? It's basically an image of a scene where all these rays are-- if you take a very simplified pinhole [INAUDIBLE] model here, the projection of the scene onto the sensor. When you move it over here, it's almost the same world being projected from a slightly different [INAUDIBLE].

And then using tomography, it's almost the same thing as when you have [INAUDIBLE] sensors. And you have an X-ray source. And you have an object here. And you have a projection, this object, onto the sensor. And when it's shifted around and put extra source here, you are projecting this from a slightly different viewpoint.

So in this case, it's as if the object was outside and you're projecting to this pinhole and taking an image on the detector. And by moving this pinhole, in this case here, you're taking different projections on the scene. And over here, you're taking different projections of what's inside. Here, what's outside. And here, what's inside. But the basic principle is the same, which is you are changing the viewpoint and taking the projection of what's [INAUDIBLE].

And as that explained, this data set is sufficient. You can explicit as a radon transform. And you can invert that to figure of what's inside. So if this was some material that's simply attenuating and not scattering, then you can figure out the density function.

In case of refocusing, what you were doing was you were shifting and adding to focus on a particular plane. And that's far more for tomography, or at least, a form of laminography, as you were looking at earlier. Because you're just looking at a slice, which is exactly what the word tomography means, recording of a slice.

And the [INAUDIBLE] is identical because, at least to the first order, because the projections of your 3D world on a 2D sensor eventually allow you to compute something about the 3D world. From a set of 2D images, you can say something about the 3D world.

In light field tomography, you mostly care about creating refocused images one layer at a time. But if this was also some kind of an object which had some density of pure attenuation, then again, from all these sequences of images, you can reconstruct what the volumetric representation of this object is.

There's one major difference, though, between a CAT scan here versus an electric camera that has an area of light field that's made up of [INAUDIBLE]. What's the major difference?

AUDIENCE: [INAUDIBLE].

RAMESH RASKAR: Far field beam-- that's one example. In fact, if you have an area of pinholes, it's very similar to the parallel-beam tomography. There are two or three variations that are mentioned here.

AUDIENCE: [INAUDIBLE].

RAMESH RASKAR: Go ahead.

AUDIENCE: When you have the angular [INAUDIBLE] your [INAUDIBLE] you're losing data?

RAMESH RASKAR: You're losing a lot of data, in terms of angle, as well as resolution. Here, you will take thousands of positions of extra source to construct this volume. But here, you may have just a few tens of cameras. And also, here, you may have a full 180-degree rotation of your source and detectors. But here, you're only within-- you're limited almost by the field of view and the arrangement of this camera, which is what you see.

So with respect a given point, you might only span, say, 30 degrees, or 40, degrees or so, depending on the field of view. And this is [INAUDIBLE]. So you could have that missing component [INAUDIBLE].

So those are, of course, just limitations, in terms of the type of reconstruction we can achieve from light fields. But again, it [INAUDIBLE] how you set up the system [INAUDIBLE].

So that's why you can construct a tomography machine using the light field idea, which we saw as [INAUDIBLE]. And hopefully, we can convert many of these complex concepts of tomography and deconvolution and confocal imaging and all that, and achieve them with principles that we're unfamiliar with, in the visible spectrum, with optics, or without optics, and make them available on possibly [INAUDIBLE] for cameras.