

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high quality educational resources for free. To make a donation or to view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

INSTRUCTOR: He's the world's expert in understanding wavelengths and so on.

[LAUGHTER]

No, he's been a collaborator with our group, and he has done a lot of work in the space. He's built all kinds of interesting camera, also compressed sensing, and now he's working for an organization called MITRE, M-I-T-R-E.

MICHAEL Correct.

STENNER:

INSTRUCTOR: And he's going to tell us-- he did this really beautiful study of different ways of capturing multispectral images. So he's-- we're delighted to have you here. He's going to tell us in more detail.

MICHAEL OK, so by way of introduction, and I'll talk about this a moment, but it's an interesting thing. I think in this
STENNER: computational photography or computational imaging community, I think it's a new enough discipline that very, very few people actually study it.

Most of the people who have come to this community have come either from a computer vision, or graphics, or other computer background, and others have come from an optics-y, physics-y background, and I've come from the other side of it, which may shine through in some of the things that I say to you today.

So I'm going to talk, in general, about this spectral imaging. And this is a talk that I've put together not exactly for a group like this. So it's a little bit short on introductory slides. It kind of just dives right in. So I'm going to chat for a moment, and I'm also going to talk a little bit. These guys provided me some background slides, since you guys haven't talked about tomography yet, which is going to come up at some point in my talk. So this is going to be a little disjointed, but consider this background before I get into the body of the talk overall.

So some of the techniques that I'm going to talk about rely on tomography, which most likely, everybody has heard about. That's the T in a CAT scan or CT scan.

It's also a major part of how MRI scans work. And the basic idea behind tomography is that you take a measurement where you're integrating something along a path through some object, and then you measure, for example, the intensity of these things. So you end up with this, basically, line integrals through this thing of some property.

In this case, it's the, effectively, the density of an object, and then you measure that at the bottom on the other side. You do that for a bunch of different angles as you rotate your source and your detector around the object. And from that, as we'll talk about in a minute, you can reconstruct the object that you're interested in.

Now this slide is just demonstrating that there are a few different geometries for doing this. In one case, you have this parallel beam tomography. In other cases, you have this fan beam tomography. That is, to a large extent, a choice made purely on the convenience of your physical system.

So for example, here, if your beams are X-rays or something then generating a whole bunch of parallel X-rays is kind of a pain, whereas generating a-- creating a point source of X-rays that just go off radially, and then you detect, over here, that's quite a bit easier. So you just have some rearranging of rays to deal with on the [? editor, ?] but that's the basic idea. So then go two slides forward. Well, no that second one is fun. So let's go do that.

So here's an example of looking at the computed tomography for ahead apparently. So this is the density function. And if you look at the parallel beam projections, this is-- mathematically, this is called the radon transform. But if you look, this is the angle that the parallel beams are going through, the hedge, in this case. And this axis is which of those many parallel beams you're looking at.

So you see this-- it's really hard to describe without some familiarity with this problem. But this is the sort of pattern that you get. And basically, there is enough information here to reconstruct this.

INSTRUCTOR: So you can think of each vertical slice as one image taken from each direction--

MICHAEL
STENNER: As one-- yes, so for example, if you were to just take a standard X-ray, like you're looking at here, like your doctor puts up on the wall, if a vertical slice is a 1D version of that X-ray, right? So you take that same kind of X-ray, but looking through in a whole bunch of different directions, and this is what you get.

The same thing here, except that they're just kind of rearranged a little bit. Instead of a vertical slice, now it's a curved surface is that normal thing. Now as you get all of the same parallel rays, they just don't happen all at the same time anymore. Next one. The next one now.

So now, the question is, how do you turn this back into the image that you're looking for. And as it turns out, basically, the-- and then, it's been a while since I thought about it. So I might get the details wrong. Feel free to help me out here if I screw this up.

So if you are projecting through this way, then that vertical slice is what you get for a single angle, and that vertical slice basically is a line through the 2D Fourier transform of your image. OK? So you take all of those vertical slices, and instead of plotting them this nice rectangular thing, you take each one and lay it out at the angle it was taken.

And it's going to give you an estimate of the Fourier transform of your object. Except you're going to have samples along these radial lines based on how you scanned, all right.

INSTRUCTOR: Maybe you can draw it quickly?

MICHAEL
STENNER: That's a very good idea. I can see that. Yeah, I can do that, thank you.

So here's my Fourier space. Here's my physical space. Start out with my parallel rays. I measure the result with my detector over here. And that tells me the Fourier transform-- so this is a 2D Fourier transform-- and I now know the Fourier transform sampled along that line. OK?

So now I do it again in this direction. And now, I know the 2D Fourier transform here. All right, that's great. So one might only slightly naively think that, hey, if I know the Fourier transform of this thing, then I know the thing, itself. All I got to do is Fourier transform back. Right, that's almost completely true, except that you realize that the pattern you have, this thing Fourier transformed in, is not a rectangular grid.

So you can't just pop it into your standard 2D Fourier transform algorithm, and get the Fourier transform back, or the direct function back. There are lots of crazy interpolation problems, and it gets a little bit ugly.

So there's a-- I don't if you have-- yeah, so there is-- I'm not going to talk about how it works-- but there is this other algorithm that's very popular for this of thing called filtered back projection. Really, it messes with where you do the interpolation. And in general, it gets past a lot of the artifacts that you would otherwise see.

The bottom line, the thing to take away from this little pre-talk talk is the basic concept of tomography, that is, you perform integrals through some object from a bunch of different angles. And then it gives you all you need to know in order to figure out the full structure of the object, itself.

INSTRUCTOR: Is that clear?

MICHAEL Any questions?

STENNER:

INSTRUCTOR: So X-ray tomography, you have to put X-ray with the sensor, take multiple projections, ' from there reconfigure what's inside. And we're going to use that same exact principle now, but think about wavelength.

MICHAEL Yeah, we're going to stop looking at heads and look at data cubes. Go ahead.

STENNER:

AUDIENCE: So you have some kind of object that's spherically symmetrical, but if you cut it, actually it has different layers, how do you get the layers out?

MICHAEL How do you mean, it's spherically symmetrical?

STENNER:

AUDIENCE: So let's say the density or whatever you're measuring is exactly the same for every direction, but internally, it's actually heterogeneous, instead of homogeneous?

INSTRUCTOR: Like the hair?

AUDIENCE: Well, that's not perfectly spherical.

MICHAEL So are you saying it's got layers?

STENNER:

AUDIENCE: Yeah, but like perfectly symmetrical layers.

MICHAEL OK--

STENNER:

[INTERPOSING VOICES]

MICHAEL So a spherical shell of glass, a spherical shell of metal, like that built up?

STENNER:

AUDIENCE: Yeah.

MICHAEL That's not a problem, because what you're telling me, I mean, what you're suggesting that might be a problem is
STENNER: that every measurement you take, every image that you take, is going to look exactly the same.

AUDIENCE: Right, I get it now.

MICHAEL But the object really does look exactly the same from every direction. So that's OK.

STENNER:

AUDIENCE: OK.

INSTRUCTOR: But that's a good way of thinking about this problem, because it's a [INAUDIBLE]-- if it looks the same from the other side, then in a traditional camera, it must be a sphere. But in the X-ray tomography, or tomography general, if it looks the same from every projection, then it must be spherically concentric.

MICHAEL I'm a big fan of symmetry. OK, so this talk is basically a comparison-- it was put together for a customer who
STENNER: wanted a survey of hyperspectral imaging techniques and the trade-offs between them. So there's not a whole lot of intro here, but so what we're trying to do-- I'm going to talk about a data cube, which is x, y, and lambda. And we've got some resolution in all three of those-- in general, I'm just going to call that and nx, ny, and L-- so the number of resolution elements in lambda.

What else was I going to say? OK, so that's a reasonable introduction. So these are all-- I don't think any of these-- some of these have become commercial products. But these are all basically research level devices, nominally looking at the same thing, except for the last couple. One of which is Ramesh's, and [INAUDIBLE]. And the other one is all-- well, maybe it doesn't matter where it's coming from, but it's another one that's not exactly a spectral imager.

One more thing, by way of terminology, spectral imager means you get some spectral channels-- a few, five, 10. Multispectral generally means you get a bunch, dozens. Hyperspectral generally means you get like 1,000.

So I tend to be very sloppy with those terms. But that's really all it means.

INSTRUCTOR: There are some people who make multispectral as visible and hyperspectral as [INAUDIBLE] surface.

MICHAEL That's true.

STENNER:

INSTRUCTOR: People use all kinds of different terminology.

MICHAEL Yeah, so basically, I'm just copping to the fact that I'm as sloppy with these terms as everybody else is. OK, so go
STENNER: ahead, we'll get started.

So the points of comparison I'm going to talk about here are the data volume of the scene that we're looking at-- for almost all of these is this going to be nx times ny times the number of spectral channels. The physical volume, just how big is the physical device, the architectural impact on acquisition time.

Some of these are devices that you point at something, push a button, and everything's acquired in an instant. Other things have to be scanned. Another one is the computational reconstruction and scaling. You guys are pretty comfy at this point that sometimes data has to be processed after you acquire it.

The photon efficiency-- this is a big one that Ramesh was alluding to earlier, that in a lot of these devices, you end up throwing away a lot of light that actually comes through your aperture but never gets used.

And then compression is an interesting one-- have you guys talked about compressive sensing at all?

INSTRUCTOR: No, not yet.

MICHAEL STENNER: OK, so compressive sensing is this-- basically the idea that if I'm trying to measure a data volume, n_x times n_y times L , I might take some number of measurements that is smaller than that. And based on some assumptions about the space, I might try to reconstruct that full data volume. There might be artifacts involved, but that's the general idea.

So a couple of caveats-- some of these quantities are basically rough. And I'm not talking about the data quality here, because it's very, very dependent on the specific device and the way that you operate.

All right, first and easiest, by way of introduction, is this-- what kind of a baseline camera, where you just have a scanning filter. And this is one that Anka talked about earlier, as well. You just have a lens, you're imaging a scene onto a sensor. But before you do that, you have some tunable wavelength filter in place.

And so just to get you familiar with the language that I'm going to be using, data cube in this case really is this the standard thing. The volume is basically just however big your lens is and whatever its focal length is. Acquisition time-- the impact here is that you have to scan this filter. So you have to point the camera at a scene, and scan the filter and wait while that happens.

Photon efficiency here is probably one of the more relevant, or more interesting points. And this is what Ramesh was talking about earlier-- it's 1 over L . All right, when you're talking about visible light, 1 over 3 doesn't seem like such a big deal. 1 over $1,000$, that becomes a problem.

INSTRUCTOR: You want to explain how the tunable filter works? Or what it does?

MICHAEL STENNER: Yeah, so as Anka was talking about here, this might be a color wheel. So this might be a wheel that has different parts of it with different spectral filters. And each of the spectral filters in a device like this is going to be a notch filter, effectively. It's going to block everything but some narrow range, maybe 5 , 10 , 20 nanometers.

Or it could-- there are-- I don't even know how to describe them-- that cavity based tunable filters that are electrically tunable. So they got no real moving parts, exactly. But you can dial in the filter--

INSTRUCTOR: It can change its color response by usually changing the distance between the plates. And if the soap bubbles are floating, you will realize that depending on the thickness of the bubble, it has a different color, also. Or if you spill oil on top of water, then depending on the thickness of the oil water layer-- there you go.

MICHAEL STENNER: Yeah, so but by way of analogy, if any anytime you've got these two layers, you're creating a cavity, which acts like a filter. And you can think of that as-- is a loose analogy, don't take this too literally-- it is like an electronic domain, a single pull filter.

You can add additional layers, and you get to something like a Bragg grating, if that sounds familiar to you guys, or a multilayer dielectric stack, if that term works better-- where you can get even sharper cut-offs and narrower resolution, and larger stop bands on the sides. So you can end up getting very extremely selective filters in situations like this.

INSTRUCTOR: Unfortunately, the ones that are programmable are--

MICHAEL Typically not that way. Yeah, they're typically like just a single cavity [INAUDIBLE].

STENNER:

AUDIENCE: Yeah, so what's in acousto optic filter? How does it work?

MICHAEL Pardon me?

STENNER:

AUDIENCE: Acousto optic?

MICHAEL That's what-- acoustic optic, oh, boy-- oh, yeah, OK, acousto optic. So it's one of my favorite devices. Yeah, I'm going to try to keep myself from going too crazy here.

AUDIENCE: You can always come back next week.

MICHAEL Yeah, in fact, I will be back next week. But I had a plan for talking-- I don't want to steal Robbie's time.

STENNER:

AUDIENCE: That's OK.

MICHAEL I think he has a better idea.

STENNER:

AUDIENCE: No, it's missing [INAUDIBLE]

MICHAEL Oh, just working [INAUDIBLE]. All right, so acousto optics-- so you've got some physics, some piece of glass, or
STENNER: it's usually not exactly glass, but glass-like substance. And you generally have a piezoelectric transducer here. So it turns electrical into mechanical motion, or vice versa. In this case, we're going to do the former-- we're going to drive it with some typically RF signal. So 100 megahertz. If you work with one of these in the lab and try to listen to an FM radio, it gets really annoying.

And what this thing does, if you drive it with some frequency, is it will generate a traveling wave-- which this thing is usually built to dump it on the end-- it will generate a traveling wave of pressure wave, a sound wave, inside your device. And if you pass light through this, it will scatter off of that thing, just like any other diffraction grating.

I honestly do not know how you do this-- I'm familiar with these in the laser setting. How you get--

INSTRUCTOR: It's a standing wave, then you-- you get a standing wave pattern.

AUDIENCE: It's not a standing wave-- it's traveling.

MICHAEL It's generally traveling.

STENNER:

AUDIENCE: It has a Doppler shift.

MICHAEL I don't know--

STENNER:

INSTRUCTOR: [INAUDIBLE]

MICHAEL So the way you can select wavelength with these things is that it's basically a grading, right? So if you imagine
STENNER: one wavelength coming in will go this way, a different wavelength will go that way, and so on. And it's just like any other grading. And then you can simply put an aperture out here.

And then by changing the frequency that you're modulating with, you can change the grading period, and align different wavelengths with your pull and your aperture. So it's basically your standard grading, except you can tune the period.

INSTRUCTOR: In object-based media, my group, Jack Smiley, he's actually building these to make holograms. So if you guys want to see, it's like a little thing like that full of ways to get a change of frequencies.

AUDIENCE: For steering the lasers.

MICHAEL Yeah, one of my favorite devices that are really awesome-- I'm going to say one more thing, just because I have
STENNER: to because it's so cool. As a physicist like myself, you can think of exactly how these things work in terms of how much angle you get, and it actually shifts the wavelength of the light as it does this. So what comes out is not the same as what goes in.

And you can figure out all of that if you simply conserve energy and momentum, and treat the incoming light as photons with known energy momentum, and treat the sound wave as phonons with known energy and momentum. You do that, then it's all just a standard physics 101 billiard ball type stuff. It's very cool.

OK, back on the path here. Yeah, so that's pretty much our baseline scanning filter thing. Any other questions on this guy? If you haven't--

INSTRUCTOR: This was simple as possible.

MICHAEL Yeah, if you're lost now, you're in trouble. So speak up.

STENNER:

OK, all right. Next one. So our second baseline-- Anka also talked about this one. This is the standard push broom. This is basically very similar to just a standard spectrometer. In fact, the thing that you are passing around here in the little box is pretty much exactly this. I mean, there's slight differences in how it's made.

INSTRUCTOR: It's in the box [INAUDIBLE].

MICHAEL Yeah, so in that box, as you were looking through it--

STENNER:

AUDIENCE: It's in here.

MICHAEL Yeah, great. As you're looking through this hole, you see that there is this slit here, going in the wrong place, a
STENNER: slit there. You can see it right there. And the slit, basically the slit is-- the light is in this architecture, light is imaged onto the slit.

That light is then made parallel again. As it hits the grading, it's spread in different directions now. So the direction of the light is related uniquely to its wavelength. And then, that is refocused on the sensor. So now, if we have some flat, uniform scene, what we're going to see on the sensor is vertical stripes-- exactly like you saw through this thing.

And where each stripe corresponds to the wavelength. So if you measure how much intensity you get in each of those stripes, you know how much intensity, or how much power you have in each wavelength.

Now, if the scene has some structure vertically, then the structure along the vertical line is related to the object's physical structure, vertically. If the object is uniformly dark at the top, and light at the bottom, then that's what you're going to see on the sensor. And the horizontal structure of what you get on the sensor is directly related to the wavelength. That's completely it.

So you get-- what this gives you, if you have this 3D data cube, a single frame of the sensor will give you one solid plane through that data cube, except instead of x and y, like a normal camera, it gives you x and lambda. OK?

INSTRUCTOR: So y is lost.

MICHAEL y is lost.

STENNER:

INSTRUCTOR: But averaged.

AUDIENCE: What did you write--

MICHAEL No, it's not averaged, it's--

STENNER:

INSTRUCTOR: Lost, only one point.

MICHAEL It's one point.

STENNER:

INSTRUCTOR: That's it.

MICHAEL Exactly, and so normally what you do in this case is you scan it. So you either move the camera and take
STENNER: multiple frames at different locations, or Anka was completely correct, a lot of these standard applications are airborne, either for military purposes or for agricultural surveys, or whatever. So you're flying on an airplane, and basically you get the scanning for free. You just use the airplane to do the scanning.

OK so the volume is a little bit bigger, you get more optics. There's no reconstruction here, you're building the data cube up directly. And the other interesting note is this is also not particularly photon-efficient, because you are getting light from other locations, other x values. They're simply being dumped by this slit-- that is if they don't hit the slit, but hit other side of the slit, we're throwing the light away.

So that light is all getting wasted. And so this is also not particularly photon efficient.

All right, so now we get into the first wacky version of this. It is also perhaps the most complicated. So sorry, the graceful introduction is over.

So this, architecturally, is completely identical to the thing that you just saw, except we have now replaced the slit with some sort of code. So now, all of the light that gets imaged onto this thing, it will be modulated by the code. But then, everything else is the same.

So the light gets modulated by the code. You can think of, if you prefer, by the way, if this helps you, you can think of the slit as a particularly boring code. All right, it's a perfectly legitimate code. It's just a really boring one.

So this is a more interesting code. In general, they're going to be half filled and half blocked. But then, all of the usual stuff happens-- we read columnate, go through the grading, and the bits get separated and fall on our sensor.

INSTRUCTOR: OK, so we have the notion of multiplexing, where we talked about if you want to measure nine bags, we can measure them one at a time. Or we can take random linear combinations and then invert. So it's the same concept, again, done for light.

MICHAEL STENNER: Right, so this device was originally built to be a spectrometer just like this. So we will conceptually step back for a moment and think about it in that context. Don't worry about imaging, don't worry about x and y . In fact, you can assume that all of the light is maybe hitting a diffuser in front of this thing, for example. So there's no structure here.

So imagine for a moment that this were just some interesting spectrum-- maybe one of these fluorescent lights, or whatever else. But spatially uniform. The benefit of that is that now, instead of just the single slit through the middle of this thing, now we're collecting a lot more light-- we've gone from a factor of 1 over n , or-- sorry, 1 over n_y to 1 over 2 , effectively.

All right, so we modulate this guy. And then now we have this problem. This is the problem in general with the slit-- I'm going to back up one more time.

Why don't we just make the slit bigger if we want to collect more light? Any takers on that? Go for it.

AUDIENCE: Because then you're probably not sampling just one light source that you're interested in, you're probably sampling a larger portion of the scene, which won't give you the specific [INAUDIBLE] you're looking for.

MICHAEL STENNER: Yes, so if we look at our sensor-- if we look at our sensor plane, and we look at one column of pixels along here, if we have an infinitely narrow slit, we know that this physical location on the sensor is associated with one wavelength. OK? If we now have a wider slit, then this is associated with one wavelength from the left half of the split, but a different wavelength from the right half of the slit.

So now, we have this mixing of spatial and spectral information that is problematic. Now, the assumption that the whole scene is spatially uniform helps, but that's not generally a realistic assumption. So what this does is help us get around that. So now, what we have-- because we've coded this-- is we have a way of disambiguating this otherwise ambiguous spatial and spectral mixing.

So now, let's go back. We have this single line on our detector. We really do have a different spatial locations on our code, on our aperture, mixing-- they're all combining on this column of pixels. Except that this one over here has one wavelength contribution, because that bends say less. And this one over here is contributing with a different wavelength, because that one bends a lot.

Right, so what do we do? We take our measured values here, and we take the dot product with the appropriate code-- and these codes, if you've talked about Hadamard codes, for example-- are designed so that the dot product of any two mismatched codes will just be zero. And the dot product with the single lonely correct code is some large value.

So what we do is for this, we take the dot product with each of these, and that tells us how much light came through this part of the aperture and landed in this place. And you do that for all of those, and you get to figure out what each of those contributions is, and you can reassemble them.

OK, so that's how this thing works as a spectrometer. Now, turning it into a spectral imager is just a little bit different. We've only got one-- we're only trying to make this a push broom spectrometer, so we've only got one spatial degree of freedom that we're trying to recover.

And in this case, we're going to actually make that the vertical direction. Is that right? Did I get that right? Scanning wide-- no, make that the horizontal direction.

Because we know how much-- we've now figured out exactly how much light of each wavelength came through each of these vertical columns, so we know how much light came from that location and what the wavelength of each contribution was. So we're almost there-- there's only one remaining problem, and that is, if the scene has structure, then that's a problem. The Hadamard code assumes that the underlying thing we're trying to reconstruct is smooth, is flat, is unstructured.

So the cleverness here comes from the following. We can slide the code vertically and wrap it around at the bottom, right, and take each of those versions where we slide it one, take a frame, slide it again, take a frame. And reassemble them at the end, so that you can, say, take this-- pick it right here-- this spot in this code. And look at what looked at what happened when it was here, and then look at what happened when it was here, and look at what happened when it was here, and so on.

Or equivalently, you can look at this spot on the image and say, what happened when this part of the code was there, then this part of the code was there, and this part of the code was there, and so on. And you can synthesize a full version, a full-frame version of this that is actually using only this row of data over and over and over again. That's horribly unclear, and I apologize, but that's just about the best I can do.

The point is that you can get past the limitations on spatial structure that this spectrometer generally has by scanning this thing in the y direction in this case. And that's OK, because we're planning on doing that, anyway in these push broom architectures.

So we do that, we get the vertical direction. We already had the horizontal direction. And so you get the two spatial degrees of freedom. And you have your spectrum.

The upshot of this guy is it only throws away half the photons, which for these things is actually pretty good. And that's pretty much it. There is a little bit of reconstruction-- every point on this thing is actually created by taking a dot product with the code, but that's actually not that bad. So this thing is actually a pretty cool system. This is, in some sense, kind of my favorite.

INSTRUCTOR: And I keep going back to the comment that Mike made, this can be thought of just as the simplest-- the previous one of a line, which is just a boring code.

MICHAEL That's right.

STENNER:

INSTRUCTOR: And the operation would be identical. You would assume, as if you didn't know what the code was-- I mean, you didn't know that the code was so simple. And you would still do all the same operations, except now, instead of measuring multiple quantities at the same time, we're actually measuring only one row, I guess. Only one row at a time.

MICHAEL Yeah, let me try one more time this way. If this thing were-- if every column had a uniform value here, we're all pretty confident that this would work, right? The Hadamard code does a good job of allowing us to distinguish how much each column contributed to each column on the sensor.

STENNER:

So that tells us, for a given column, what came from what place in the code, and what had what wavelength in the code. So that's pretty straightforward. The only real trick is what happens when it's not vertically uniform, and that we can fake. We can fake by sliding the--

In fact, let's do it this way. This is a much easier way to say this. We've got some structured thing here, and we're going to fake just a single frame that has uniform structure vertically. How we're going to do that is I'm going to say that we are only interested in this line. As we're scanning the object past this thing, when this line is here, I'm going to collect this row. And I'm going to ignore everything else.

A moment later, when that line is up here, I'm going to click this row and I'm going to ignore everything else. And so on, do that all the way up. And one step at a time, you construct exactly the frame that you would have if the whole thing were uniform with this line.

So that's how we fake it. Except we're not doing one line at a time, we're doing a whole thing.

INSTRUCTOR: You're doing Hadamard multiplex [INAUDIBLE].

MICHAEL We're doing Hadamard, exactly. So this guy does a pretty good job, and it's pretty--

STENNER:

INSTRUCTOR: As you'll see over and over again, this trick is constantly being used of taking linear combination of multiple quantities. Because we want to make the photon efficiency go as close to half as possible.

MICHAEL This is a big factor, here, the photon efficiency. Because the main problem with those two baselines is they both have lousy photon efficiency.

STENNER:

All right, yeah, so I think I've already talked about all of this stuff. So this is the reconstruction. The scanning options, by the way, we can either move the scene over the code, like with the airplane moving, or we can circularly scan the code in a circle snapshot kind of way. That is, just point the camera-- nothing moves but the code. Yeah, go ahead.

OK, so this is out of the same group. So by the way, my notation here, the person in parentheses is the professor in the group, and the name in front is the person who did all the work.

INSTRUCTOR: [INAUDIBLE] really get the credit.

MICHAEL STENNER: That's right-- that's right, so it's the real person, and then parenthetically, don't forget about this guy. He paid for it.

So this is a similar architecture in some ways. In this case, we do the same thing. Well, wait-- I thought it was on the wrong one.

So this is exactly the same piece of hardware, All right, so this is going to take much less time because everything I told you about before, it is completely true here, except for one thing. We're not going to move it one line at a time and take every measurement. We're going to take fewer measurements. How many fewer? Eh, it's up to you.

So in this case, we know that we don't have enough information to fully reconstruct all of the full data cube. So we're going to use some clever algorithmic tricks, and that's where this horrible scaling comes in, to try to reconstruct that full data cube.

All right, so you guys haven't talked about compressive sensing. I'm going to give you my one-minute version of compressive sensing. You have some linear algebra problem here, or you've got some big 1V vector, and you're going to operate on it with some non-square matrix.

INSTRUCTOR: Is that a little bit colored [INAUDIBLE]?

MICHAEL STENNER: That would be handy. And I can't erase very well either, so apologies for that. Let me try this.

This is about as complicated as it gets. Oh no, dear God.

AUDIENCE: I hope that meant yes.

MICHAEL STENNER: All right, so this is as complicated as my little diagram gets here. We have some number of parameters that describe a scene, and we have some matrix that is performing an operation on that. And as a result, we have some number of measurements.

If this guy is not square, and in particular, if it is wider than it is tall, then the number of measurements, we have is smaller than the number of parameters that fully describe the scene.

So this-- so what does this tell us?

INSTRUCTOR: In a traditional numbered multiplex thing, it's exact square matrix. With a number of unknowns, the number of knowns is equal.

MICHAEL Right.

STENNER:

INSTRUCTOR: And now, we have fewer measurements than this.

MICHAEL So even if you don't know anything about the structure of this guy, what do we know from this? And don't think

STENNER: too hard. This is inter-linear algebra.

AUDIENCE: Under constraint for all of it.

MICHAEL Yeah, it's under the constraint, or under determined.

STENNER:

AUDIENCE: Do you know the thing you're capturing on the right?

MICHAEL What's that?

STENNER:

AUDIENCE: You know the thing you're already capturing?

MICHAEL We do not know it. That's why we're doing this.

STENNER:

INSTRUCTOR: Right, it's the unknown.

MICHAEL Yeah, so this is the scene, this is the world. And this is the data that our camera, in this case, spits out. So our
STENNER: goal is to-- from these measurements, figure out what this thing was.

We do know this matrix-- I mean, it doesn't matter what the structure is at the moment, but in general, doing these problems, you do know what this matrix is. And you have your measurement, but you're trying to get back at this guy.

All right, so this is an under-determined problem. And there are many algorithmic approaches to estimating this, and anybody who's ever used a Moore Penrose pseudo inverse, you can try and invert this thing. And then, with that, this guy will give you an estimate of this.

In general, it won't be that good. It depends on the details. But there's this-- and that's as much detail as I'm really going to go into. There's a relatively new discipline called compressive sensing, which is devoted to finding better ways to do this.

In general, they assume you have control over the nature of this matrix. But the idea is that you're going to collect far fewer measurements than what you're trying to reconstruct. And the basic way that they do that is by using algorithms that assume something about the object. And the thing that they generally assume is sparsity in some domain.

So for example, if you're looking at stars at night, you can assume that compared to the number of pixels, we don't have that many stars. Most of the pixels are black, only a few of them are going to be white. So if you get lots of blurry stuff, you can generally assume that, in a very simple case in the center of each blur, you've got a single star.

In more general imaging applications, you can assume that the object is sparse, say, in the wavelet domain. We don't have white noise, but we've got nice, round faces, and eyes, and hair, and that sort of thing. And if wavelets didn't work so well, then we wouldn't use them for JPEG 2000 and other things like that.

So if you assume the thing is sparse in some domain, then you can generally do quite a lot better in terms of reconstructing this guy. So that's the basic idea. You design this thing well, and then you reconstruct with these clever algorithms, assuming sparsity, to get the original object. Any questions on that really, really fast intro to compressive sensing?

So that's where this lousy scaling comes from. Unfortunately, the algorithms typically used for compressive sensing are not especially fast. They scale badly with a number of points. So what this guy does is, using the same hardware we were just talking about, we can take a compressive measurement where we simply don't scan as much as we might like to. The signals get mixed together in a way that we cannot uniquely unscramble them, but using these approaches, assuming sparsity and whatnot, we can do a pretty decent job.

Again, this is a good example where I'm not talking about image quality, because there are problems with it. There's no easy way to characterize it and compare.

OK, let's move on. So this is a similar one out of the same group. I'm not going to dwell on this too much. The architecture will feel familiar. In this case, we image-- let me make sure I get it straight-- so we go through a standard spectrometer, except put our-- now we have our code here. And then we go through a completely identical grading and lens formation to basically remove this wavelength smearing that we had before.

INSTRUCTOR: Optically.

MICHAEL STENNER: Optically. For example, if we just removed this coated aperture and just left it open, left nothing in that space, then what we would get at the end here would just be a normal picture. That is this grading separates the wavelengths, and this grading puts them back together again. This would do absolutely nothing but take a nice normal grayscale picture.

What changes things up is that we have this coated aperture in there, and it's there, it's in there in an interesting place. It's in there where the spatial and spectral information has been mixed. That is some wavelengths have been shifted more than others. And so things get all scrambled up. And then put back together.

So what happens is-- and again, maybe we can kill those fluorescents again, thank you-- so this is a little bit hard to see. But you have what looks like a relatively normal picture here, except that at different pixels, you have different codes that have been applied to the spectrum.

So yeah, so you can again decode on the back end to reconstruct the image. This is also a compressive sensing problem. It's half photon efficient, again, which is going to sound familiar if you hear more about this coded aperture stuff, because roughly half of these guys are closed. It's compressive, in that you take a single spatial shot, and that's got all of the information you're interested in.

So we can compare these guys a little bit in terms of how they work. If we look at the same scene for both of these-- and this is intended to be three colors, some red color, some green color, and some blue color, with a little bit of overlap in between-- as in the single disperser, this thing gets modulated by the mask before any shifting.

In the dual disperser, you shift first and then modulate. And then in the dual disperser, you put things back together. In the single disperser, you modulate and then shift.

So this is just giving you a feel for how things are mixed together.

Going to the next one-- some of the implications of that are if you have three white sources, in this case, the dual disperser, because you split them up first and then modulate, you might lose some of the color bands associated with the single point source. So the red and green for this point source are just gone.

But the odds are that when you put something back together, all three of those point sources will be represented.

For the single disperser, on the other hand, because the modulation happens first in the image domain, if a point source happens to fall on a closed point in the mask, it's just gone. Tough luck. The good news is, you do retain better spectral information about the locations you do see.

So it's really just kind of a trade-off in terms of what information is higher priority to you, and what works better for your application. And the last example, same basic thing-- you lose spatial structure here, but you get good spectral information. Here, less spectral information, but you still have a better idea of what the spatial structure looks like.

All right, what's next? OK, great. So here's where the tomography comes in. So this is another architecture that works in a completely different way. The way this guy works is, we're going to have some sort of lens come through our imaging system, columnate the light, and then go through a diffraction grating.

OK, now this is a specially created diffraction grating that scatterers in many different two-dimensional directions.

AUDIENCE: 16 [? bit? ?]

MICHAEL STENNER: Yeah, so yes and no. It's one-- it's hard to say, exactly.

INSTRUCTOR: 3, 4, 5--

MICHAEL STENNER: I think it's six directions, but then you get combinations. So this one here is a combination of this and that. And this is two orders of this thing. So this one-- these six around here are the first orders of diffraction.

INSTRUCTOR: Show those same effectiveness--

MICHAEL STENNER: OK, yeah.

INSTRUCTOR: So think of one of these as this guy here. So if you shine light on this, you'll get a sphere. And as you rotate this guy, these other guys will show. So [INAUDIBLE] orientation this one. But imagine if there are final gradings that will mirror even more.

MICHAEL
STENNER: Actually, what I think this is physically, is-- yeah, I'll pass this around in just a moment-- if you took three of these guys and rotated them each 120 degrees from each other, this is basically what you would get. So this is the scattering off of one of them.

Go for it. Oh, hey, good. I can do better than that-- look here.

INSTRUCTOR: Now we just have monochromatic.

MICHAEL
STENNER: Yeah, that's true. But you can get the idea of multiple orders here. So one of these, that one, is my zero order beam through this diffraction grading.

INSTRUCTOR: The little one here?

MICHAEL
STENNER: And each of them, going further--

INSTRUCTOR: 1, 2, 3, minus 1, minus 2, minus 3.

MICHAEL
STENNER: Right. So what we're looking at here, if we just had one of these guys, we'd have 0, 1, 2, and the 1 and 2 are just merging into each other. It looks like one long one, but it's actually two.

So 0, 1, 2, negative 1, negative 2. So this is one, two from a grading at this kind of angle. This is almost certainly one from this, and one from grading that's at this angle. So you get these multiple gradings all on the same material. And you get this scattering in different angles.

We can pass that around if people care to look at it again.

INSTRUCTOR: And again, if you just look at the world through this, and you have zero [? tech, ?] you'll see there are the copies of [INAUDIBLE].

MICHAEL
STENNER: That's right. So all right, now, here's where the magic happens. What good does this do us? Right?

Think about what it means to smear these wavelengths, and then let that light fall on a camera, or on a sensor. We don't have just a single unidirectional beam here, but we have an image. We have a scene that's being propagated through this thing.

And if we imagine the scene has multiple colors, or multiple wavelengths, as it must if we're interested in it, then what's going to happen is say, a blue wavelength part of the scene, and a red wavelength part of the scene-- which started out on top of each other-- are going to get shifted and added together on the sensor.

So if we imagine the red-- let me do this-- OK, we can maybe turn the lights back on for a moment. Imagine this is a red layer, and now I'm going to take the blue layer. It's going to be shifted and added together. These two things are going to be shifted with respect to each other and added together.

What's another way to think about that? I'm going to draw my data cube here, which I've been talking about. And I love to draw-- here we go, OK.

If this is x , y , and λ , what does a normal monochrome camera do? It gives us a value. So I'm representing the scene here-- a normal monochrome camera gives us the integral vertically through λ of that scene. It just takes all of the different values at every different λ , adds them all together, just gives us the total amount of power for all wavelengths in the scene.

What our shift now does is it gives us the integral along some line, at some non-vertical angle.

INSTRUCTOR: Which is the same concept as the light field for assignment two. You've seen that [INAUDIBLE].

MICHAEL
STENNER: So for a single-- so if we measure this out here, for example, what we're getting is a full image made by many line integrals at some angle through this data cube. Right? And now, if we measure this one way out here on the edge, where things have been shifted a lot in the other direction, that we get some very steep line integrals in the other direction.

OK, now if this is starting to feel like tomography, it should, because that's effectively exactly what we're doing here. We are taking many line integrals through this data cube, at many different angles, and we can use standard tomographic techniques to reconstruct them.

INSTRUCTOR: Let's just make sure everybody's with that one. You see the analogy between X-ray tomography and what we have done here?

MICHAEL
STENNER: The wacky thing is that this is no longer a head, or some physical, three-dimensional object. It's no longer x , y , and z , that are three dimensions in this case.

INSTRUCTOR: It's x , y , and λ .

MICHAEL
STENNER: It's x , y , and λ . And our path integrals are no longer density, but they're just the amount of energy represented by that x , y , and λ voxel. All right, mathematically, completely identical. We're just taking path integrals through at a bunch of different angles. We get enough of them, we can reconstruct that entire data cube.

INSTRUCTOR: And if you have, say, I forgot how many was it? 25 here, roughly? Let's say you have 25 such arrows here, that corresponds to what in the next [INAUDIBLE]?

AUDIENCE: Angles?

INSTRUCTOR: Corresponds to 25 angles. As if you put X-ray position 1, 2, 3, 4, 5, 6, and 1, 2, 3, 4-- 1, 2, 3, 4, 5, 1, 2, 3, 4, 5, 1, 2, 3, 4, 5, and you basically put those 25 projections. That's exactly what it is. Which also means in terms of resolution, what's going on here? In terms of your sense of resolution?

You're losing resolution by at least by a factor of 25. Usually much more, because you cannot pack these guys very tightly.

MICHAEL
STENNER: Right, yeah, so that's one of the problems with this architecture is this is the sensor-- this square box is the sensor. And you can get some 12, 14 megapixel sensor, throw a lot of pixels at this thing. But you're wasting a lot of them. No light's ever falling on them.

And that's because you got to make sure that these don't overlap with each other. That's a real problem. So as a result, you end up with dead space in between them.

INSTRUCTOR: But it's a very clever scheme.

MICHAEL It's a very clever scheme.

STENNER:

INSTRUCTOR: When you think about how tomography [INAUDIBLE] is being used for hyperspectral imaging.

MICHAEL All right, go ahead.

STENNER:

INSTRUCTOR: Let's stop here, because we are running over time. And let's see if there are any questions on that. And then, you're here next week, right?

MICHAEL I am, actually. Just do me a favor. Just flip through to remind me what's left and I'll see if I have any comments.

STENNER: I'm not going to talk too much about yours. It's just summary table.

[INTERPOSING VOICES]

INSTRUCTOR: Are there any questions on the [INAUDIBLE]? The last one?

[INTERPOSING VOICES]

AUDIENCE: So the only one we didn't get to was Isis?

MICHAEL So yeah, this is very similar to see, so that's not so interesting. Isis and then the edge of spectrum, which Anka

STENNER: talked about a little bit-- well, a lot, actually. More than I will. I'm going to talk about it.

So these-- I'll just say, these two aren't traditional hyperspectral or spectral imagers.

INSTRUCTOR: So the goal there is not measuring the cube, traditional cube.

MICHAEL That's right.

STENNER:

AUDIENCE: So all of these are the images that your company developed?

MICHAEL Oh, no, this is a survey of academic literature. We did-- Roark did work on something that could be a light field

STENNER: architecture, that could be used as a spectral imager. That didn't even get in here.

AUDIENCE: I presented it a couple of weeks ago.

MICHAEL So you've seen it otherwise.

STENNER:

AUDIENCE: You probably don't remember.

INSTRUCTOR: So your assignment number four is multispectral imaging. And we won't be building hardware, because that takes too long. But we'll give you a data set that has 32 channels, 10 by 12 each, for flowers, and people, and beer, and so on. And then we'll do something, like what Anka was saying, where the stuff using standard RGB color response, you will be allowed to mix and match those spectrums and create interesting images.

So this should be a relatively simple assignment, because you are simply adding up the start of the images and creating the three-channel image at the end. But hopefully, it'll get you intrigued about how different things look in different spectrum.

So the assignment number four is actually open-ended. This is just a suggestion. If you want to do something very simple, something that takes a little more than six hours to do, propose that to me, as well. We can support that as your fourth assignment. Or if you don't want to take too much, then you can just do this one and focus your creative energy towards the final project.

So next week, we'll be talking more about scientific imaging-- microscopy, tomography-- we did it in a slightly opposite sequence. We wanted to do the tomography first, but hopefully now you already have some idea, deconvolution and so on. That's what we'll talk about next week.

And we also have a guest speaker. And then, we'll do very brief overview for the exam on the 13th, which is open book, open laptop, open everything. So you don't have to study for it. You should really be focusing on your final project.