

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.

RAMESH RASKAR: All right, everyone, let's get started. I am Ramesh Raskar. And this is--

AUDIENCE: I don't know.

RAMESH RASKAR: --MAS 131 and 531, computational camera and photography. And we have, I think, about \$50 to \$60,000 worth of equipment right here. So please hold on to the door if something goes on I really need your help to protect all these things.

It's going to be a lot of fun. We're going to see crazy kinds of cameras, crazy kind of photography, medical imaging, and applications in different domains. I just wanted to show this really beautiful picture that some of you may have seen on blogs. This is a real photo taken by an iPhone camera of a propeller blades on an aircraft. Anybody knows what's going on?

AUDIENCE: The sensor takes line by line images. And by the time it goes to the next line, next vertical line, those propellers have moved. So it faces the [INAUDIBLE] directions are not exact location--

RAMESH RASKAR: Excellent.

AUDIENCE: [INAUDIBLE] separated [INAUDIBLE].

RAMESH RASKAR: Exactly, so when we think about an image, we think about this as some kind of a snapshot, some kind of a progressive photo. But there are two motions going on here. One is the motion of the blade, which is circular, and with radial blades. And then there is motion effectively of the sampler, which is moving from up to down. It's a rolling shutter.

There's a very nice animation of that. So this is the rolling shutter the camera is exposing approximately just one line at a time. And as the blades rotate around, you can see that it's tracing this course. So this example should show you that we can't take anything for granted when it comes to this modern photography.

And this is mostly an artifact of a really cheap sensor that cell phone cameras use. Because of bandwidth constraints it's much easier to roll and read one line at a time than reading the whole buffer of the camera at the same time. So you get just beautiful artifacts.

And one open question is if camera makers start supporting or start cheaping out on rather than just one line at a time-- if they come up with even cheaper mechanisms where they're explaining and some random sequence or-- any artifacts they create, hover photographers exploit that to create this stunning and beautiful imagery.

More on the scientific side, let's see the lighting is [INAUDIBLE] here. Let's see if you can play with that. Yeah, I'll do that. See if there's another switch that doesn't like the back-- No. We're going to learn a lot about computational illumination in this [INAUDIBLE]. That maybe OK for now. And we'll switch it back.

So imagine you have somebody out here behind a shower curtain and you want to take a photo. And you may be able to come up with a photo where you just see the shower curtain behind the person behind without the person behind the shower curtain. But maybe you can also create a photo where you realize it was behind the shower curtain. And this is a trick that's achieved using a special type of a flash. And as you'll see, we will generalize the flash to a very highly programmable illumination. And with that, photography of this kind is also possible.

What else? Here's a great example I'd like to show. Imagine you have a scene and instead of a camera you just have a single photodetector. OK, It could be a light sensor of your SLR camera. Those of you like to think about these photodiodes, just one photodiode.

And instead of taking a photo, what I'm going to do is turn on one pixel of this projector at a time and record this. If the projector has one million pixels, I'm going to take one million measurements by turning on one at a time. What picture will I get here?

Will I get a picture from those 1 million measurements? Anybody? No, you get a photo that looks like the camera was placed here instead of over here. And it's exactly how a barcode scanner at checkout counter works as well.

The barcode scanner doesn't have a camera. It just has a sweeping laser. In the same apparatus there's also a single photodetector. So if the barcode scanner hits the 1010 patterns of the printed code, when it hits the black spot, it doesn't reflect much light and it hits the white spot. It reflects much light. And in this way, a barcode scanner without having to worry about focus and dynamic range and all those issues can figure out what the barcode is.

AUDIENCE: Isn't it important, though, that the barcode scanner is a laser, which keeps the pixels spatially localized?

RAMESH Very focused.

RASKAR:

AUDIENCE: Right. Whereas here if you have an incoherent light source, and you're illuminating a real scene, you'll illuminate the whole scene kind of diffusely.

RAMESH That's a very good point. So that's why we have to use a projector. So we are turning on only one pixel of the projector. So as in very simple words, it's turning on only one ray in the scene.

RASKAR:

The projector is not flat field. Not all the pixels are on it at a time. That's why we had to take one million readings while turning on one pixel at a time. So that's straightforward. This is very well known.

What are some other things you can do with this particular duality? If you replace the photodetector with a camera, you can do something similar. You can turn on one pixel at a time and take a full photo. So again, I'm going to turn on one pixel of the projector at a time and in this case, take 1 million photos.

Now, once you do that, what I will be able to do is create a relationship between what happens when I turn on exactly one pixel of this projector. One pixel is turned on. A photo is taken. I can measure what happens to this particular pixel of the camera. And in this way, you will create a four-dimensional relationship, 2D for the camera and 2D for the-- so 2D for the projector and 2D for the camera. So a million photos here and a million pixels here. So you're going to have a trillion measurements, 10^6 to the 12^{th} .

Now, we can invert that and ask for a single pixel of the camera, which projector pixels are contributing. So the question to ask is I'm going to turn on one pixel at a time. If I turn on just this array, for example, it will only contribute to, say, this pixel. But if there was some intersection, it will also contribute to some other pixel and so on.

So you're going to have this global transport of light. And from that you can ask this question, for a given camera pixel, what are the other projector pixels that are contributing to it? And so you can do this inversion, which is relatively straightforward to do.

And then you can do experiments like this, how can you read your opponent's cards from across the table? You can turn on the projector one pixel at a time, look at the reflection, and it's in the camera, then on the next pixel, and so on.

And after you allow your opponent to take a million photos of him, you'll be able to read the card although it [? will ?] not directly visible from your point of view. So you can kind of look around an uploader and what's behind a corner. What's the-- what's the flaw in this argument that you can look in a corner? Yes.

AUDIENCE: [INAUDIBLE].

RAMESH
RASKAR: Exactly. There's some device that you have to place that's actually looking at that card. So this is a beautiful project from Stanford. And one of their funding agencies, of course, the army-- and they would like to know what's around the corner. But of course you have to go and place a projector in the enemy line to be able to figure out who's out there.

So this is the message of this class. Pure digital cameras are extremely boring. And if you look at these two cameras, one of them is digital. One of them is phone. Can you really figure out which one is which? Left is digital. Right is digital. OK, and the rest are confused.

So hardly any difference between them, zoom, focus, aperture exposure, all the same old boring stuff. There's not much extra going on. And if somebody claims that digital camera have spawned a new art form, probably not. I mean, you could just take film cameras, scan them, and play with that. They're just faster, better, and cheaper.

So when people think about computational photography and computational camera, they start thinking about, wow, OK I have this camera that doesn't do a very good job of dynamic range and field of view and so on. So what I'm going to do is somehow improve the performance of this camera.

One common is by boosting the dynamic range by taking multiple photo by exposure bracketing or having a larger depth of field by taking multiple photo by taking focus bracketing and so on. So to increase the field of view, you're going to do a panorama and then stitch it together. If you want to increase the frame rate, you're going to play with exposure time and so on.

So this is what a lot of people think about as computation camera and photography. I just want to emphasize that this is not what we're going to talk about in this class. These are concepts literally easy to understand and I have a lot of lecture, notes, and videos online. You can look it up. In a couple of hours. You will get a pretty good overview of all the things that can be done in this space.

And if you go to Flickr, there are groups that are just exploring high dynamic range imaging, et cetera. So we are not really going to talk about that because all these techniques are just trying to improve the performance of a camera. It's not trying to change the game of photography.

So generalizing it even further, what we're going to look at is cameras that are not just 2D sensors mimicking film. But they could be zero-dimensional sensors that are, as I said in barcode scanner, and we'll look at time of flight, motion detection. One these sensors that you see in flatbed scanners or fax machines or line scan cameras that are used in photo finish at sports events, 2 and 1/2 new sensors. 3D sensors are a different kind and very quickly into 4D and 6D devices that are exploiting deconvolution and tomography in medical imaging or scientific imaging and also displays that are four-dimensional and six-dimensional.

So earlier we saw that you can look around the corner by placing a device in the line of sight, which is the projector. Well, that was done three years ago or four years ago, 2005. Do we have some new machinery now that will allow us to look around the corner? You're out here. The meetings are called. It's a good day.

You're out here, and you want to see what's inside with a door that's partially open. Is this possible? Or imagine your cameras that are right now in mobile devices. You have this thing that's shrinking every day, getting cheaper but also shrinking and actually degrading its signal to noise ratio performance.

But imagine if your whole LCD becomes photosensitive, so every pixel that's emitting on your LCD also becomes photosensitive-- and companies like Sharp and Planar are already doing that. So they have this LCD array where if you put a finger directly on top of the LCD, you can even look at the ridges on this finger.

So there's some very beautiful opportunities. And I just saw Kimo and Ted here. They have also a very exciting project where they can use an ordinary camera to look at right down to the resolution of your ridges.

Now the question is, if we redefine a camera not to be this perspective device that behaves like a pinhole camera-- but the whole screen is like sensing. There's a fusion of sensor and optics. Think about in the display domain a CRT. There is an emitter, the electron beam, and then the receiver, which is the phosphorescent screen.

There's always a separation between what's emitting it and what's receiving it. And over time, of course, we have a LCDs where everything is fused together. And if you think about cameras, it's the same. You have a lens that's collecting, and then you have a sensor that's receiving. And we haven't thought that there's always has to be some distance sufficient distance. If you have a 35 millimeter camera, you ought to have roughly 35 millimeters between your lens and your sensor.

And we can throw away all those constraints and all those assumptions, and we might be able to come up with devices where the lens and the optics and the sensor is all just one thing. And that would be kind of the LCD equivalent of a camera. So this is something we can think about. And these devices are coming mainly because they want to support touch sensing.

But imagine the future if you have your mobile device, and you can just wave it and take a picture-- and because your sensor is so large, you'll be able to collect a lot of light. And right now this it's not even megapixel yet. But this will be a megapixel, 10 megapixel. And you'll be able to multiplex those pixels for color and infrared and all kinds of beautiful things, different speeds, and so on.

So that's how we're going to think about in this class about cameras and photography. And the same with medical imaging. Right now medical imaging is thought as-- medical imaging today is very similar to photography in the 1930s.

The guy who took the photo was more important than the people whose photo was being taken sometimes, because those people have to go to the sky and stand for some time. And he had this one specialized device called a camera.

Everybody has to stand still for a few seconds. You take a picture, and then you go home. And after a couple of days, you get your photo back. And that's exactly what medical imaging is all about today.

We have to go to a special location to get our CAT scan or ultrasound scan done. And our time is completely-- our time is considered not so important. The guy who's running the machine, he's very important because we have to make an appointment with this guy.

So can we bring-- will medical imaging evolve to a stage where it's like photography today? Clearly, when we take a photo if somebody spills a coffee, we take a photo. It's a very casual way of thinking about photography. And then medical imaging evolved to a stage where we can do that. And it turns out we can. There are certain directions, certain computational methods that we will be able to develop that will get rid of older mindsets of how medical devices are being developed.

So a very brief introduction of where I come from and our group here Camera Culture. So in the past I have worked a lot on [? conventional ?] illumination, different types of projectors, creating multi projector displays, creating virtual reality setups, augmented reality, a lot of work on pocket projectors and augmented reality with pocket projectors, interaction paradigms, work with RFIDs and so on.

And in the last eight years or so, a lot of work in cameras playing with shutter aperture, light fields, illumination wavelength, and sensors and so on. And some of the questions we discuss in our group are this. And what will the camera look like in 10 or 20 years from now? I already gave you one kind of possible direction where it might be just a flat screen, or maybe it's just your credit card. But it could even be a retina display retinal sensor and digital sensors.

How will the next billion cameras change our social fabric or social culture? If you think about the-- if you think about the internet, it has been transformed by the ability to search. That includes indexing and segmenting and sorting and all those beautiful problems. But image search even today remains an extremely challenging problem.

So maybe we can change the game and modify our devices, modify our cameras, modify our displays, modify our storage devices and so on. So that image search can become simpler. OK. If they magically have a camera that's 100 megapixel, when I take your photo, I can Zoom in all the way down to your iris. And if I have an iris detector, then I have very easy identification. And all the photos I take of Dina, for example, they're all indexed with her iris code.

So that's kind of a very [INAUDIBLE] level way of thinking about simplifying image search. But we'll look at a very interesting example that exploit thermal IR imaging, electric cameras and so on. What will happen when for this recording is being used not just by the big brother, but it's actually being used for some beneficial purposes in commercial settings?

So imagine a scenario something like Google Earth live where right now you can go to Google Earth and fire up the browser, and you can go to any part of the world and see how it looks about six months ago or a year ago. But imagine if you can just fire up your browser and go to this location live as it's happening now. And of course, you can move your slider and go back and forth and back.

So if you have a camera on every street pole, every bus, every taxi, every person and all those data all that data is being streamed on the network, so you can go to any part of the world and see it, how will that world look like? I'm sure many of you are definitely scared-- and given all the privacy and security notions of it.

But this could become very similar to the safety we have a little bit in our financial system our financial transactions. When I use my card in a restaurant, a bunch of people actually see my credit card number and all the information. The waitress looks at the card. And the owner knows it. The bank knows it. The credit card company knows it. And the government, of course, knows it and so on.

And a lot of people know about your financial transactions. But somehow you are completely comfortable sharing your financial data. And similarly can we create cameras and photography and imaging infrastructure so that if you walk down the street, and all these cameras are looking at you, anybody in the world could be looking at your live?

You feel completely secure that it's being used only by the right people for the right reasons. And if you don't want them to see it, you have some switch on you that says, I'm completely invisible. And you should be able to walk down the street.

Can we create such an imaging infrastructure? Think about high speed cameras and high resolution cameras. Maybe we'll have microscopes and nano scopes with us that will again change the way we think about medical imaging. And what about movie making and news reporting? As you know it has dramatically changed over the last three or four years because, again, a billion people have cameras out there.

Whether there is an in Tibet or there's a satellite image of what's going on in Burma or a plane that's about to crash, you get amazing videos. And they're not captured by CNN. They're captured by other people. So how are we going to, again, change this imaging infrastructure, the whole pipeline, to think about the future of moviemaking and news reporting?

So overall in this course, we're not going to think purely about software. We're going to think about how we can change the camera, not just use the camera. So that involves optics and illumination and sensor and motion of the sensor, motion of the optics, the different wavelengths, 3D cameras, polarization, probes and actuators and also priors and online collections, the network.

And one kind of theme you'll see is after years of research in computer vision, one could argue that we have exhausted the bits that are available in pixels. There's still a lot to be done. There's a lot to be done. Still challenging even today with sophisticated computing algorithms, it's challenging. So maybe we can build feature revealing cameras that will go hand in hand with existing or modern [INAUDIBLE] algorithms so that we can process photons and create this meta structure for our imaging pipeline.

So what we're going to do today is I'll briefly describe what this course is about, do some introductions. And the second half, we will come around and do a fast forward preview of the [? core ?] groups, OK. So in the next few slides, I'm just going to give you a quick kind of rundown of what this course is about, the layout.

Here's a nice overview of the biological successful vision in the animal eyes. And there's a nice paper in Science. And you can take all the successful biological vision and place that in kind of eight categories. You have eyes based on shadows, based on refraction, based on reflection, single chambered eyes, or compound eyes sometimes with apposition, sometimes with superposition, and so on.

We'll come back and discuss this more. If you look at the eyes of a scallop, it's based on a mirror, not on a lens. So you have a mirror down here, and the sensor is up here. And light actually reflects from this concave mirror. And the image is found on the detector.

And the future camera in your mobile device on a flat sensor could be on this architecture. It doesn't have to be on this single chamber high refraction. So the human vision is in this corner here. It's a single chamber design with lenses. And all the cameras, at least all the standard cameras that we know of are in this particular part.

But this is a lot to explore. And that doesn't even explore computation. This is pretty much raw imaging that these animals are using. So film light, traditional photography, light comes in through the lens. It falls on the detector of the film, and that's the end of the story. That's your photon. And you just transfer it digitally or by chemical processing, and you should see it.

Computational camera is a little bit different. We're going to have some crazy optics that's going to think about how rays and wavelengths are manipulated. We're going to have a sensor that's not just mimicking film. It's not going to be a [? flag ?] sensor. It'll have its own geometry and spectrum and so on. And what you see finally will not be just a raw image but will have some reconstruction and computation on top.

So that's computational camera. That would be what's in the conventional camera. But there's also an element that's outside the camera, which is like-- and it's sort of just a flashlight. You're going to have a flashlight with sophisticated modulators that are changing the intensity and phase and polarization of light in different directions, different kinds of additional optics, and so on.

So once you have a programmable light and a computational camera, we have a framework to really exploit and understand the light transport in the scene. So something interesting is happening in the camera work. We are here 2008, 2009.

About a billion cameras are being sold every year, which is fine. Most people know that. But if you just go back about six years, 2001, 2002, 0 cameras were sold embedded in a mobile phone. So we have gone from 0 to a billion in just six years.

So it's just an amazing time to think about imaging. And this is very much like all the fun in networking and communication back in the '90s when millions of people were coming on board and thousands of websites were going on. And that was computing. And this is the time for visual computing.

And because this game is changing so fast in terms of cost, in terms of performance, applications where imaging was considered not the perfect solution is changing. We are seeing cameras being used in some really casual and very strange ways. And here's an indication of where that's going.

So where are these cameras? Remember, cameras are not just 2D sensors for photography, but they're used in various ways. So we are here, 2008, 2009. The pink here is because of mobile phones.

About a billion sensors are being sold. Any guesses about what these other slivers are, the blue one, the green one, and so on? There are a couple of million here. About 100 million here are being sold. I guess [INAUDIBLE]. Sorry.

AUDIENCE: Optical mouse.

RAMESH Optical mouse, very good, very good.

RASKAR:

AUDIENCE: Gaming.

RAMESH Gaming. So if you think about-- we, I think, sold like 30 million remotes. 30 million remotes were sold. And this
RASKAR: chart is actually made before that. So that's not even here, but that's going to be big, yes.

And of course, there is traditional photography. But I'm glad you got optical mouse because that's one of the largest markets. It's basically a very low resolution, usually 20x20 or 32x32 pixel camera that's running 1,000 or 2,000 Hz. It's a high speed camera that's doing optical flow to figure out where your mouse is even if you put it on a very clear surface.

And the mobile phones and digital cameras-- and all this worry about big brother, think about security. It's a very, very tiny sliver. And if you think about the first three categories here, optical mouse mobile phones, and digital and video cameras, these are all personal devices. And this is going to scale with the number of people in the world. So this will easily become 6 billion or 6.5 billion, whatever we have.

But these other slivers may or may not grow. Gaming, for example, could still grow very rapidly because almost every person individually might own a game console. So very interesting here. If we look around the internet, cameras are being-- sometimes it's just silly.

You may know about this do it yourself green screen effects company, U-Start. And they want to be the Guitar Hero in the visual domain. So instead of playing music that's synchronized with some prerecorded data, here they have pre-recorded video segments. And then you can star in this movie. It's pretty big.

If you go to their website, a lot of people are uploading their videos with their software and their experience screen. All they give you is a camera, a simple camera, not even a studio pair, just a simple camera and a big plastic green screen. And they have some software for doing screen mapping and so on. You may be familiar with this video. Which company is that? I forgot. It's starts with M. It's a studio webcam.

AUDIENCE: [INAUDIBLE].

RAMESH It's a Japanese company, OK. Some other interesting ones, look at this one here. Panasonic Beauty Appliance,
RASKAR: what does it do? Any guesses? It's a pretty big camera.

It's a personal beauty appliance where it provides humidification in a very local spot. So you can sit there and sit all day around and it I guess maintains the right level of moisture and humidity around. So your skin won't get very dry.

What about this one? The camera is not anywhere near the eyesight. But the camera is actually on the ear. Can I look inside this? I was reading up on this. It wasn't very clear what the applications are. But you can think of it. I mean, if you're walking down the street and you want to be safe, something's coming at you. Maybe it tells you. Or if there's somebody you don't like-- they're walking towards you-- don't turn around.

Yeah. And that's going to happen more and more, more crazy places. I like this GigaPan Epic 100 which is for creating panoramic imagery. But it's actually basically a tiny, tiny robotic platform about \$300 for an [INAUDIBLE]. And it has a physical lever that presses on the shutter as it rotates around. It's pretty amazing. And of course Fuji is doing some fantastic work just starting with the studio pair. And they're trying to provide a whole pipeline all the way down to printing on a lenticular screen so you can enjoy this studio box.

OK, so after you look at all this camera, they say, OK, what else is there to do because you don't have to do much research to do any one of these projects. You don't have to take a class to build these applications. So let's think about something trying to improve the camera just a little bit.

All I want is depth per pixel, something humans seem to be doing very well. We can use our two eyes and figure out how far things are. At least we think we know how things are. There's a lot of prior information.

But for a camera just to get depth is extremely challenging. After 34 years of research in computer vision, still you can't go out there and buy a camera that's a reasonable cost, which runs at a reasonable speed and gives you that perfect [INAUDIBLE]. It's amazing.

So we'll show you what the state of the art is using this multi-thousand dollar camera. This was donated to us by [INAUDIBLE]. And [INAUDIBLE]. So this is camera that uses time of flight to [INAUDIBLE]. Oh, we don't use [INAUDIBLE] cameras as much. OK, so this is time of flight.

All those alleles are being-- those of you who have a cell phone camera, just look at this lights through your camera. You'll see all of them are lit up in the air. And as you can see here-- what is it looking at? Is it [INAUDIBLE]?

OK, so the people in the front are marked in green and somebody in the back over there. Yeah, Martin. So it's giving you some estimate of depth by doing a time of light calculation. Now that is the camera. And it works OK.

But as you can see the quality even in this relatively easy configuration, most of the objects are diffused. There is no resistance. The movement doesn't have some light to overwhelm this active light. And even then the quality is OK. It's not that great. And I forget the exact cost of this one, but they cost about \$10,000, these cameras. Go ahead.

AUDIENCE: What's is it called?

RAMESH This one is Canesta. Yeah, but during the later classes we'll learn about all the different 3D cameras, whether they're based on active light or stereo or structured light, polarization, and so on. So this one is going to start.

RASKAR: The other company is 3DV, which was recently bought by Microsoft.

And then there's a couple of companies in Germany who are building 3D cameras. And apparently this market is now being driven by game consoles because PlayStation and Xbox are all interested in using 3D cameras in the house for gaming, for interacting with gestures.

So they want to distinguish this from this, for example, because right now all the games are in the gulag. The hands are away from your body to play the games. I-toy and Xbox and so on. But the new ones will allow you to do gestures that are more intricate.

So there are too many challenges just to get 3D because we have to use some kind of active light to compensate for ambient scenes. You have to stitch geometric capture from multiple views. If you have something like marble or skin, it has subsurface scattering. So that's difficult to deal with. I cannot do triangulation. And objects that are diffused have diffuse reflection component. If you go into something that has glass or are dark, it's just out of caution.

So computational camera, it should do something more than just capturing a 2D image. And those of you who are here building real time and C applications for robotics, anything like that. There's no camera right now that can deal with objects like this. This is amazing. So we have a billion cameras out there. But none of those cameras can solve this problem. So that's what I want to look at in this class.

AUDIENCE: Can I ask a quick question about that?

RAMESH Yeah.

RASKAR:

AUDIENCE: So I'm just thinking back to the traditional film type camera that has autofocus sensors that are contrast based or a rangefinder type thing. So that in a way gives you depth by moving optical elements. But this doesn't have any moving parts, I guess. Does that count as a way to try to make a camera that perceives depth by having something that scans physical optics--

RAMESH Exactly.

RASKAR:

AUDIENCE: --around.

RAMESH I think you're asking a very important question. Let's see if I have a slide. So if you think about a different way of scanning in 3D. And this is by the way-- it's one of our visiting students Doug Lanman and Gabriel Taubin. They have a beautiful course at SIGGRAPH on all different ways of doing 3D scanning.

And they're contact based and then non-contact based. And right now we're talking about mostly active by using time of flight and so on. But they're of course passive as well studio and motion and so on based on focus and defocus. That may be a good segue way to actually look at this camera here. And Ms. Emily, you want to run it?

AUDIENCE: Yeah, [INAUDIBLE]

RAMESH Yeah. So this camera is completely passive. It's like a full camera, which has-- You want to talk about it.

RASKAR:

AUDIENCE: Yeah, basically it's 25 separate images, each one of these is a separate camera.

RAMESH Which is on a single chassis.

RASKAR:

AUDIENCE: Yeah, and they're just 25 to 30 minutes. But each one's from a slightly different perspective. So you can do things like add all the images together but slightly shift them. So I can focus down at infinity towards the end of the table. Or I can change the focus by just simply by shifting or moving the images to focus right there.

RAMESH So remember he's doing all this operation in software. You can take those 25 images. And then in software you can refocus it anywhere you want. And then, again, based on maximum contrast or one of those operators. You could potentially figure out what's in front or what's behind.

RASKAR:

But of course, if you put in some transparent object, it's going to be quite challenging to figure out where this is. If I just put something that's really flat, then it wouldn't know whether it's in focus or not in focus in the middle of the paper. It might be able to do an OK job on the boundaries of the paper. Or in the middle of the paper it always looks like it's low frequency.

So these kind of cameras are coming. You can buy this camera for about \$20,000 from U-Plus. So far we are up to what, 30, up to 30K. We did 10, 120. The next one will be-- we'll get to 60s very quickly, 60K. And if you, again, think about your cell phone camera, the current mindset in camera makers is that they want to shrink it, something that's smaller and smaller and smaller.

But if you want to create enough baseline, create some kind of focus based depth extraction, for example, as we're looking at here, we must have some baseline between them. And that's what this camera does. It's still pretty compact. I forget. I think 6 centimeters total, something like that.

I don't know. It could be on your phone. It does have that much space. So if they would just turn it around and make all the pixels as sensing pixels, you have enough baseline from the left edge of the camera, left edge of this cell phone camera to the right edge of the cell phone camera to create those effects. But that mindset has to change.

Right now, as we will see in the class later on the section on sensors, there is tremendous innovation in how photons-- sorry, image sensors are being built that are wafer level cameras back start elimination, 3D VLSI, and so on. I mean, it's all great, and it's going to help us.

But unfortunately all of them have a single tract mind. They just want to make a higher signal to noise ratio that must collect as many photons as you can. And they want to shrink the sensor as small as you can so that they can sell it for a cheap cost.

But if you think about-- if somebody had told you 30 years ago that the TVs are really expensive, so what we're going to do is start building TVs that are smaller and smaller because they'll be cheaper and cheaper, all right-- but that's not how it works.

People want larger and larger TVs. And they're willing to pay for 40 inch and 50 inch TVs. And there will come a time when people say, it's OK if I have to pay a little bit more, but make my sensor and make my sensor array larger and larger so that it's almost the size of the whole device.

Even for a camera like an SLR camera, not the whole camera is sensor. Only a small part of it is sensor. But once we get around this notion that silicon is really expensive. And we have to shrink it, and we have to create this wafer. We can slice and dice it into millions of tiny sensors for cameras. Hopefully things will change around and we won't be watching those tiny TVs as they have predicted some time ago.

So that's kind of where some of these things are. So let me go back-- so getting 3D very challenging, and that I would say is kind of version 0.1 of a computation camera. I mean, it has to sense the three-dimensional world to do anything interesting. And it is being used in other scenarios, I mean, [INAUDIBLE] challenge.

This is an animation for movies.-- where you can build vehicles that can navigate through any kind of terrain, including urban neighborhoods. So the first version of the [INAUDIBLE] challenge actually had a lot of cameras. But the sad aspect from a camera point of view is that the second version that was-- Stanford was the winner in that one.

In the traditional definition of a camera, there were no cameras on the whole vehicle. There were zero cameras on a car that's built to navigate in a city. It's pretty sad. I mean, if you think about it, a human driver driving through a city-- I would say almost all our input, all our actions are based on visual information. But cameras are so primitive to build this self navigating vehicle that no traditional cameras are used.

Now, of course, I'm taking it to the extreme because the kind of devices they did use were similar to range scanning devices. They had laser scanners LIDAR, reducing time of flight and so on. And in a way they're capturing information that's sensitive to different directions. So in a way, it's a camera, but it's not a traditional visible range camera that they're using. It's a detector, a single pixel detector that's measuring light coming from different directions.

And on the other extreme, that's one extreme where cameras are not good enough to self navigate. On the other extreme, you have lots of people online-- again, this is a slide from Doug Lanman building their own 3D scanners. So of course, you can take a Logitech webcam and put some [INAUDIBLE]. You can calibrate that and move it around to create 3D models and a whole bunch of--

We'll look at this during the-- people are using just a wineglass to create laser stripes. So I can take a laser pointer and shine it at a wineglass to create a laser stripe. I don't even have to buy expensive optics. And they just scan it and develop a cheap DV Cam to create 3D scanners.

This one is probably most interesting. They want to create, again, a very accurate 3D model from a LEGO rig. And this is how it works. You can probably guess it from the picture. There is a milk pot. They want to take a character and scan it. They'll put this character in a milk pot. And take a photo from above doing a very simple segmentation.

And then over time, they're going to put a little bit more milk so that the level of the fluid will rise in this square bucket and this rectangular bucket. And they'll continue to take pictures and they'll create this 3D model section by section. Pretty amazing.

AUDIENCE: They did it in New York with humans.

RAMESH Sorry?

RASKAR:

AUDIENCE: They did it in New York with humans. They go in the milk and scan you. [INAUDIBLE].

RAMESH Excellent, I would imagine going the other way is easier. Just drink. Just straight up [INAUDIBLE] and drain it. It

RASKAR: will go out almost concentrated.

AUDIENCE: True.

RAMESH But, yeah, this could be a great class project. So here's a slightly different direction that some people are taking.

RASKAR: And I really like this work from [INAUDIBLE] Washington. Some of you are-- OK, let me state the question first and see if anyone of you has an answer to that.

Let's say you go to Rome, and you are in front of the Trevi Fountain. You take a photo, and you don't know anything about Trevi Fountain and all the hype about it. You want to take a photo and figure out which part of this photo is interesting according to everybody else.

So why am I here? Should I be looking at-- if I go to the Old Town Square in Prague, I'm not going to look at a fountain here. I'm going to look at another part of the castle.

But somehow when I'm at Trevi Fountain, I'm not looking at the buildings. I'm looking at the Fountain. So how would I determine without looking in my guidebook which part is most interesting? Any clues, any answers?

AUDIENCE: I don't know if I should answer.

RAMESH Go ahead. You're [INAUDIBLE] yeah.

RASKAR:

AUDIENCE: One thing you could perhaps do is take a consensus for all the photos that have been taken of this object and basically see where the feature tracks overlap. An alert [INAUDIBLE] sort of.

RAMESH Excellent. So this is actually an offshoot of the photo tourism Microsoft Photosynth project where you have
RASKAR: millions of photos of the same tourist location. And once you have registered them in 3D, then you can just shoot the rays back to see which rays will intersect the pixels. If most of the rays seem to be shooting you just kind of just a histogram, you'll realize that most of the photos are looking at the fountain. And in this case, what the photos are looking up at the top part of the [INAUDIBLE].

AUDIENCE: So maybe it's what is popular here. Everything might be a different [INAUDIBLE].

RAMESH You're right, you're right. So the next question is, how can we-- it's like popular versus interesting on Flickr.

RASKAR:

AUDIENCE: Yes, that's true.

RAMESH So here's our question, how can we ask-- how can we answer this question?

RASKAR:

AUDIENCE: There may be certain features that the human visual system finds particularly interesting, things with high levels of non-repetitive detail or certain shapes or ratios.

RAMESH
RASKAR: Yeah, certainly. So I think there's a lot to be done. So maybe if we want to build a camera that directly detects the building pattern, it doesn't care about finding it, capturing a real photo. But it does a very good job of finding out if there are patterns we find interesting whether it's symmetry or repetition or the right scales the right aspect ratio and so on.

And the same project. So this is the Pantheon again in Rome. These are all the places from which people take photos. And if I just tell you this is the view of the Pantheon, and you usually are-- up here is a fountain. People start from here and then go in and roam around and take a picture of this big hole in the Pantheon.

So the question is if I take a picture here at the entrance and from inside the pantheon looking out at the door and a picture from outside, what's the part that would connect these two photos? It's not a straight line path. I don't have an image. I don't have it. But if you actually have this, again, voting scheme, you've realized the best way to go from this view to this view is follow this particular path.

So this is data that's been captured from visual media, like photographs, but inherently it's more about geometry. And I would call it-- it's almost non-visual. So that was computational camera, and let's think about computational photography.

And computational photography and my collaborator Jack Tumblin at Northwestern. This is how we like to define the two parts. We want to capture the scene, and we want to synthesize this thing. And when we capture, we want to capture it in an extremely rich fashion so that it's machine readable and the machine can understand what's out there and re-synthesize. We want to synthesize in a hyper realistic manner so that it represents the essence of our visual experience.

And within that there are three major teams. One is epsilon photography, which is basically generalized bracketing. So whether I want HDR or panorama, I'm going to take my camera, Google's exposure bracketing or focus bracketing or view bracketing and so on and just create a very nice picture and mostly to overcome the limitations of a camera.

That's just epsilon photography. I'm going to change the parameters of the camera within an epsilon neighborhood. And then people think of that as the ultimate camera. But again we're not going to focus much on this part in this course.

The next part-- that's quite interesting-- is so-called coated photography. So the comment we had earlier I really care about some mid-level feature. I want to know if they're symmetric. I want to know if there's a repeating pattern. I want to know where the edges are. I want to know how the regions can be segmented and so on.

And it's worth taking multiple photos such as in bracketing. I just want to take one or maybe two photos which are reversible and of course the information of what the world in my image. And the image may not be a single 2D image. It could be a light field camera where you have 25 images or a single snapshot or you could have a time of flight camera where you really can't call it a photo because it's actually measuring the amount of time it takes for light to travel back and forth in a given direction and so on.

And it would be very useful for scene analysis. So a lot of our course would be actually discussing this part. And the last part is [? instance ?] photography which is really to see if we can go beyond this low level of features, pixels and mid-level features such as motion and foreground and background and symmetry and so on to a higher level understanding. So that it's not just mimicking human eye, but it's doing possibly things like this, telling me what's popular for example.

And I claim that only when we have computational cameras and computational photography supporting this, we'll be able to create new visual art forms. So within that, this is a chart that we'll be talking about throughout this class. I won't go into too much detail right now.

But we have certain goals, the mid-level features, low level features, high level features, and so on. And we have certain tools that can capture the raw image. We can capture the incident angle and spectrum such as UV and thermal IR. We can capture high dimensional reflectance field. We have non-visual data such as GPS and identity, metadata, image priors and so on. And with that we're going to explore this whole space of camera and photography.

Of course high dynamic range and so on is right here. We're not going to spend too much time on that. But we're going to think about how can we insert a virtual object, how can you take a photo and relight, change the lighting in the scene and so on. Material editing for example from a single photon, Ted Olson is here, and he's done a lot of work in this space.

And if you look at this chart, you realize that even the human vision is not at the top right of this diagram. With human vision I cannot look around the corner. I cannot see what's inside the body. I cannot tell you what's behind the curtain. I cannot tell you when I'm in Rome what's interesting and so on.

So you really want to create this augmented human experience by using these tools and using the mechanisms that we have available. So we'll be coming back to this. Just keep this in mind.

And a lot of this is actually available in the book that I and Jack Tumblin have published. The PDF of this book will be available throughout the course. And the real book should be out any time.

All right. So just a couple of more slides, and we'll take a break. Let me skip over this one for now, some of my favorite examples. All right. So we also got to spend a lot of time thinking about cameras that, again, show what's not seen. And I realized that Mathias is here. His company is the founder of Redshift, a thermal imaging company. And you want to say a couple of words?

**MATHIAS
OMOTOLA:**

Sure. So this was my second startup, and the idea was to build an ultra low cost thermal imaging camera and do it using a standard cell phone camera as, actually, the image sensor and using a little thermally tunable film as a translator between infrared and visible.

**RAMESH
RASKAR:**

Exactly. And the goal is to dramatically bring down the cost of a thermal camera. So let's bring in our next toy. Will you need some time to set up?

**GUEST
SPEAKER:**

Oh, yes I need--

RAMESH

RASKAR:

OK. Yeah. I we'll set up this other time, which will very quickly take us towards the 60K. And you can do some amazing things. So this is not from Mathias's company. Is that a maker or just a-- they actually make cameras or they're just integrators?

And they can use it in sports analysis? This is in the game of cricket where you want to figure out if the ball was hitting the pad or the pad and it's a very difficult call for the referee but if they have a thermal imaging camera, then the ball here if it hits the bag, it will leave a hot spot on the bat. But if it hits the bat, it will leave a hotspot on the pad.

So just by looking at an image a couple of seconds after the ball has left, the player, you can figure out if the ball hit the-- the ball hit the bat or the bat. And this is major. People get really angry and all the sports analysts will be writing or the referee next day in the press. This really saves their life.

But the goal for Mathias and other companies is to make it available possibly at the same cost as your traditional camera. I mean, if you think about a cost of a digital SLR, it was \$30,000 dollars. And now it's worth about \$500.

It's just silicon. It's not that expensive. If Canon or Nokia decides to have a thermal camera in every cell phone, they will be available for very low cost, very extreme low cost. What is your guess? What would be the cost in five years?

AUDIENCE:

Five years? It's still \$500.

RAMESH

RASKAR:

\$500. That's not bad, right? And as you'll see later in the class and you'll see in this demo-- we can do this in the break as well. So take your time. It can do some amazing things.

All right. So that was computational camera and computational photography. And in this class we're going to look at both aspects. This is how we're going to do it. We have two numbers 131 for undergraduate and 531 for graduate.

The main difference is we have four assignments for the graduate version and three assignments for the undergraduate version. And we have a midterm exam where there will be fewer questions for the undergraduate version other than that, the rest of the course is very similar. So in the assignments we're playing with hands on with optics and illumination and sensors and other elements.

And we'll have all these toys available for you. We'll have projectors, different light sources, lasers, and things like that to play with. I will also have a best project award. At the end last year there was exactly one undergraduate student, and he won the best project award, pretty impressive, pretty impressive. He did an amazing job because last year we did not have an undergraduate version. It was only a graduate version.

Midterm exam will be early November. A big component of the class and your work would be a final project, which should be novel as well as cool unlike a lot of other fields. It is possible to come up with an idea that nobody else has even thought about in this field. So you could come up with ideas that are not just incremental but that could be game changing.

And I'll tell you later that three projects, final projects from last year led to SIGGRAPH or ICCP papers. And two of the projects are becoming multi hundred thousand dollar projects. So it is possible to come up in this class with ideas that are game changing.

If you are taking the class for credit, you'll also take class notes for the lectures. We'll have plenty of guest talks in line and online discussion. This class, most of the materials is on slides. But starting next week, about half the material will be on slides, and the other half will be just on whiteboards and demos and so on.

So if you're looking at a slide, you'll get an idea of the type of material we're covering, but what I realized last year was most of the discussion was actually not captured in the slides. And if you're a listener-- and I a lot of you are going to be listeners in this class, which is perfectly fine. But what I would like to do is participate in discussion. Bring in a new viewpoint. I want to learn from you.

I'm glad Mathias is here, for example. He's the world expert world's expert on thermal imaging. So I want all of you to who are taking it, who are just listening in, to contribute to the discussion and give a new viewpoint. I will not be offended at all if you jump in and give a new perspective or a new reference about what we are discussing.

And if you are a graduate student or a postdoc, then I would also like to spend some time in sharing with us one short presentation, maybe a cool idea or some new work you're doing. But that's what I would expect if you are not taking this class for credit.

In terms of the credit breakdown, we have about four or three assignments depending on which version you are taking 40%, final project, 30%, midterm 40% and class participation. That includes discussions often online as well as taking the notes.

Pre-requisites. Let's see if I have-- I'll come back to that, another slide. The emphasis for this class is really on fundamental techniques in imaging. Yeah, it's coming up. It's going to be fun.

And in class as well as in homework, emphasis is on techniques. And they include all these key words, signal processing, applied optics graphics vision, online photo collection, statistical techniques, electronics, visual arts, and so on. And so in that sense it's not a discussion class. We're going to learn about techniques.

All right, there is entertainment here. It's very addictive to play with cameras. OK. So if you can just focus here for a couple of minutes. And there are three areas we want to focus on. And just keep this in mind and know this. We're going to focus on photography. We're going to focus on active techniques, real time Techniques. And we're going to focus on scientific imaging.

So within photography, we're going to think in higher dimensions. We're not going to think about just HDR or focus stack and so on. We're going to think about high dimensional imaging light fields, thermal imaging range cameras, and so on. And active computer vision, we're going to think about HCI applications, robotics, tracking and segmentation and how we can change the game by using feature revealing cameras. And a lot of concepts from scientific imaging such as composition sensing, wavefront coding, deconvolution, tomography, point square function, and so on--

And at first glance, these three areas might look very distinct and have very different techniques. But fortunately they all use very similar principles. And what we realized over the course is that this fusion of this dissimilar ideas-- I mean, tomography and wavefront coding seem very far away from traditional photography or HCI. But you'll realize that many of the problems you may be encountering and many of the visual arts that you might be interested in could be impacted from some of these new techniques.

So in terms of the prerequisites. The last time I taught this course, there was a request to be supportive of students with different backgrounds. So this is what we're going to try. We're going to try to tracks, a software only track and a software, hardware track.

And software intensive track is for those of you who just because of your interests or want to do things in a particular way-- and you may be able to use some key based software. And the software hardware track you will be using a lot of programming, OpenCV, MATLAB, C++ , Flash, whatever you like, Java. It doesn't really matter to me. But you'll be doing a lot of programming.

And those of you who might be thinking about the software only track-- and you might be trying to use GUI. Actually it will be a lot of work also. And those of you who use Photoshop know very well that if you want to do simple edits to a photo, you take an hour, two hour, sometimes six hours. And you'll realize that sometimes writing a small program you can do that task much faster, but it's your cost.

What is helpful but not absolutely necessary is some knowledge in linear algebra, signal processing, image processing. But what's critical is that you should be able to think in 3D. And this is a skill that's absolutely necessary. If you're going to think about in higher dimensions. 3D is just a beginning, but [? we're ?] thinking about 40, 60, 80 and so on.

Now we're trying to keep the math to basic essentials. I like to use a lot of diagrams and visual analogies in the visual analogies to explain the concepts. But when you're doing your own assignments, you will have to go back to the questions. I might just flash the question, but I'll try to explain that again by drawing on the board.

But it's possible for you to go through the whole class without actually writing down a lot of equations. At the same time, many of these concepts are complex, and they will arrive at a very fast pace. And we'll be discussing a lot of concepts.

Now, if you're the kind of person who can just sit back and watch a presentation and grasp a concept, where you ask questions. It's interactive. This is an ideal class. On the other hand, if you're the kind of person who likes to really look at the math and see what the relationship between the variables is and so on, that's fine as well. So you need to have one of those two skills to be able to do well in this class.

Now assignments for this class is also going to be a little bit different. As I said during the class, you'll be listening to a lot of advanced and complex concepts. But the assignments are structured in such a way that they have increasing level of sophistication.

So you can do pretty well without much background for up to the first 60% or 70% of the assignment. But to do the last 10% or 20% you will need a good background in some of these areas. So that's another way of thinking about how you might be able to take this class, even if you don't have very strong mathematical, algebra, or signal processing background.

So you can do pretty well, up to the 60% or 70% of. It and again in the superior of supporting students with varying background, what I will do is I will normalize your performance in those homeworks based on what background you have. So if I know that you don't have a linear algebra background, and you're taking this class, I will think that if you reach 70% of that assignment you have done a pretty good job. So I'll normalize it again by how much you know in that particular-- the knowledge required for that particular assignment.

So you can really pick your level of how you want to do it. And we did this last time, and it worked out pretty well. So we'll see if we can do the same thing. And those of you who are taking the 131 class for undergraduate, please come and talk to me. And we can similarly figure out based on the classes you have already taken, how we can structure those assignments.

For all those assignments, the four or three assignments, there's always an option. There are two assignments, and you can pick any one of them. So that gives you, again, an ability to pick an assignment that's appropriate for your level of understanding.

All right, so let's see. Any programming environment is fine, OnePlus notes. Send me an email, raskar@media.mit.edu, to put you on the mailing list. And we also have a sign up sheet that I'll pass around in the break.

And, remember, our class runs into the happy hour on Friday. So after the class, we'll all go over to the marriage house and continue our discussion over beer, if you are over 18. No, 21. Is it 21? So that's at 4:30.

This is a rough outline of how we're going to proceed. We finish on December 4th. It's funny because the classes end on December 10th. And our last class would have been December 11th. So I think having a Friday class with no final exam means you finish really, really early. So the semester ends on December 4th, which means final projects.

Unfortunately, the week before that is thanksgiving. So many of you will be thanking me for the delay and the procrastination that led you to work through the Thanksgiving break. Yeah, there's just a list here.

All right. So one of the things we will not cover or cover not in detail are the art and aesthetics of photography, the 4343. Software image manipulation, there's a great course by Fredo Durand on digital computational photography which is completely focused on software. So these two classes, first class and my class, is actually a very good complement of each other because in this class, the emphasis is more on hardware and optics and sensors and of course, software. But emphasis is on the hands on elements.

And I believe Fredo's not going to teach the class in spring. But Will Freeman, who's also a great instructor, is going to teach the class in spring. That's what I hear. I haven't confirmed it. And there are excellent classes in computer vision. I think, Ted, you are teaching a class on scene perception.

TED: Shape.

RAMESH For shape perception that's on Mondays, just Mondays.

RASKAR:

TED: No, we don't know when.

RAMESH All right.

RASKAR:

TED: First class is Monday, [INAUDIBLE].

RAMESH

OK, but I already sent an email about Ted's class on the mailing list. So once I hear from you, I'll send it out again. There are a couple of great optics classes. This class is not about optics. You'll be learning a lot of concept. But the emphasis is not on optics. The emphasis is on imaging and photography.

RASKAR:

We will not learn about Photoshop. Actually, I don't even know Photoshop that very well. So if you're going to do your assignments in Photoshop, please come and tell me. I'll look forward to it. And we won't talk about anything that's included in the instruction manual of the camera. So you will not learn about how to set the exposure and how to change the aperture.

I'm happy to do a separate crash course, one of the evenings. And we can all sit together and do a course. On the many people here it is super expert with all kinds of cameras, and he'll be happy to do that as well. So there are a lot of resources available if you have those questions. But this class is not really a program.

So as I said a few classes I teach a class that's more discussion oriented in spring, the computational photography class and optics class. And Professor Han, I don't if he's teaching it in Spring. Yeah, I don't know. All right, so any questions about the structure of the class and what you expect? Rose, yes.

AUDIENCE:

Will you be posting the assignments in advance so we can kind of see what sort of work will we be doing and engaging with during class?

RAMESH

RASKAR:

Excellent question. The question was, will I be posting assignments in advance? What's going to happen is based on the feedback I receive. Some of the assignments will change. But if you look at the OCW page from last year for this course, which is only an indication by the way because of course it's changing significantly this year, you will see the-- let's see the projects.

You'll see the type of assignments that were given out. It has all the details as well. So four assignments, writing and so on. So you can already get a sense of the type of assignments that were given. And so, again, the options. There is 4A and 4B for example. So you can choose which one you would like to do.

Again, the OpenCourseWare page is for last year's course. And we are recording it for this year. But it's not up here till next fall. So there's a long delay before the model appears. And next year we'll be doing something entirely different again. Other questions?

All right, so let's go around very briefly. I realize there are quite a few people. So maybe 30 seconds per person no more than that just to get a sense of who's here and why you are here. So before we get started. Are you a photographer? All right, almost everybody.

How many of you are videographer? You create videos. It's funny. Why there's a distinction between photographers and videographers? Do you use cameras for computer vision. Pretty good. Do you use cameras for real time processing like HCI or robotics and so on? OK. Do you have background in optics or sensors? OK, all right. So we have pretty good distribution here. So so let's start and very quickly go to the just your name, department here, and why you're here, just 30 seconds.

I take this opportunity to say that Professor Oliveira is the top scientist in graphics and vision from Brazil, and he'll be in our group. And we are sending him and Professor Mukaigawa who is a famous researcher from Osaka University and both of them, as well as Ankit Mohan, who is a scientist, those 3 and some additional people will be the mentors for this class. So if you have any questions, you can come to me or if you want to brainstorm about projects or ideas or have questions, they will be available as well during the week.

Oh, this is amazing, just an amazing set of people, chemistry, arts, communication, vision, HCI. It's going to be fun. Night vision goggles. Night vision goggles. That would have taken us over 100k.

All right. So let me just quickly make a couple of announcements, and we'll break. So there's a great opportunity for those of you who are taking this class or credit. And there's a conference on computational photography which is the second edition of Fredo Durand who some of you may know from CSAIL. And Mark LeVoy and Eric Zaleski from Microsoft did the first one, which was this year in April.

And myself and [INAUDIBLE] from Toronto and Raphael Bastion from Colorado was one of the world's experts in wavefront coding. Three of us are doing the second edition of this conference. And it will be held here in the new building in the last week of March. And the papers are due in November.

So if you're doing well in this class, and one of the project ideas is interesting, you could even attempt to submit a paper for this particular conference or at least as a demo or a poster or something like that. So it's right here on campus. And most likely we'll have additional opportunities who are not in the peer review track to show off your work as well.

So this is really catching on. A lot of big companies like Nokia and Samsung and Canon and HP all have started big research groups in computational photography and computational cameras. So there's a lot of interest. Another thing we're going to do in this class is-- as I said, it is possible to come up with completely new ideas in this field. It's such a new field because of the intersection of lots of interesting domains.

So we're going to learn how to come up with new ideas. And we're going to learn how to write a good paper. We're not going to do it exactly it in the class. But toward the semester during discussions with me and your mentors, we will help you actually write a good quality paper.

And writing a good quality paper actually has really simple methodologies which is amazing. We never learn about it in a formal class. So we'll try to do that in this class because the final project is really important. So just deciding if an idea is worth pursuing is half the battle.

And we'll help you do that, you can actually just use the Heilmeyer's rules which the military uses to decide whether they should pursue a project. And raise simple questions. If we answer those questions, you can very quickly make a decision whether you should pursue that project.

And as I said last year was extremely, extremely beautiful. There were papers that became SIGGRAPH submissions. Matt Hirsch, who took this class last year, went to SIGGRAPH and also in the student research competition this year based on the last project he did. And there are two major research teams that are coming out one in mechanical engineering. Some students, they're starting this multi hundred thousand dollar project based on the class project here and so on.

So we really want you to focus on novel ideas that are cool and publishable. Of course, those of you who have a design background or art background, we also will try to think about how they can be given the right exposure. And when it comes to technical publication, there are some simple rules that you can follow to write a reasonable quality paper. So will help you do that.

So let me stop there with this image to make you go a little bit dizzy. And then we'll reconvene in about 10 minutes. And in the second half, we will have a fast forward preview of the whole class. So we'll spend about 2 to 3 minutes on each of the projects. We have 12 classes. We'll spend about 5 minutes on each class, and we will see all the pieces.

So in the break, we'll also have the IR. camera you can try that. And Dan [? Sarkis ?] suggested for all of us to build this-- what do you call it? A pinhole camera from just a piece of paper. So you can just take a piece of paper, cut it up, and build a pinhole camera. So if you want to do some projects in that space--

AUDIENCE:

Hi.

[SIDE CONVERSATION]