

MITOCW | mitmas_531f09_lec02_1.mp4

NARRATOR: 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 --this easy. And it's mostly to get you set up, how Stellar works, how to set up your own website, make sure you have a

RASKAR: camera that you can use, make sure you have a way to transfer your images to your program, and as I said, it's completely up to you what environment you want to use. You want use Java, Flash, C++, Photoshop, Basic, Japanese, whatever language you want to use for your programming is perfectly fine. But make sure you show me your own images, not images captured from somebody else or from the internet, but something that you have taken yourself, some photos you have taken yourself.

Any questions over the assignment? [INAUDIBLE] for the [INAUDIBLE], for 531, we have four assignments, so three more after this. And for 131, we have two more assignments after this. While this is getting set up, I'm just passing around a UV light demo. Who has it right now?

STUDENT: I do.

RAMESH OK, so why don't you demonstrate the concept to everybody first? So when it goes around, they know what to look for.

RASKAR:

STUDENT: OK.

RAMESH You want to show the magic trick. You want to show [INAUDIBLE].

RASKAR:

STUDENT: All right, so you we almost nothing, and then we leave it alone. There are actually a lot of things that do this too. My driver's licence, the Massachusetts drivers licence has the Massachusetts State seal. [INAUDIBLE].

RAMESH So this [INAUDIBLE] paper in 2007 on how to create ultraviolet inks, inks that respond to ultraviolet light, that color.

RASKAR:

STUDENT: Yeah, definitely the car technologies are older, but this is color. Yeah, you actually see if many passports and visas, in particular, especially of the smaller the country is, the more colorful the UV image of their visa is. [INAUDIBLE], but people love to see red, and green, and purplish color. But I've actually never seen, basically, a full color before. That's pretty neat.

RAMESH Yeah, and also, if you go to a club, they use UV light. So if you don't use the right detergent on your clothing, or if you

RASKAR: are a stain that you think it doesn't look visible in light, but you go to a place, where you want to be hip.

STUDENT: I actually have a UV light LED on my key.

RAMESH Oh, so before you go stand in the line, you check yourself.

RASKAR:

STUDENT: Yeah.

RAMESH Yeah, you are too cool for this place. So yeah, if you become expert at playing with ultraviolet light and you like to

RASKAR: have this in your hand all the time, the loss of jobs. You could be the guy who checks at the entry at the airport at the security line. You know, he's checking, or you could be a bouncer [INAUDIBLE].

STUDENT: Or a banker, right? A lot a lot of the dollar bills have a UV strip.

RAMESH Exactly, so the theme for today is going to be lighting and different kind of illumination. And hopefully, it will make you

RASKAR: think with today, the budget is, I think, less than \$5,000. And, like, last time, I think it was a little over 16k last year, and we'll show you interesting things you can do with illumination. [INAUDIBLE].

What is it getting set up? Again, let me show you some other toys. So my collaborator, if you can just hold the book and show it to everybody. This is the book that I and my co-author Jack Tumblin, we are coming out with called Computational Photography, and the way Jack likes to describe the scenario today for cameras and photography is it's a lion who has been shackled for a long time. And if you just put this lion out of the cage, it doesn't know what to do. It just stands there and just looks around.

Even an unshackled lion doesn't know how to exploit the freedom, and I think it's the same situation with photography today. We have gone from film to digital. But even today, when we think about cameras, they're trying to mimic how this camera will look like a traditional film camera, a digital camera.

Even cell phone makers who have complete freedom on how to create cameras and form factors, they're trying to create an experience that's somewhat similar to a film camera, like they have a shutter at the same location and expect you to hold it with two hands. So very few of them allow a single handed interaction. And I just want a camera that I can squeeze, or I can tap, or I can shake, and then after five seconds, let's take a photo.

It can do all these other things, but it still want us to do this two handed interaction. But everything is stable. The button is off center, so you cannot even hold it with one hand. You have to hold it with two hands, so you can press the shutter and so on.

So this is how the world works, you know? You always are dealing with the legacy of what came before, and a really a classic example, something that makes me really squirm is look at this beautiful camera that came out from Nikon in the mid '90s. And this was one of the first kind of digital camera that professionals thought it's OK to be seen with, because they had this whole debate about, oh, digital image will never catch up with the resolution of the film.

These are the same people who still hold onto the LP saying, see, these are digital, and they don't sound the same as [INAUDIBLE]. But anyway, this camera, which is supposed to be digital, if you cut it open, what you realize is there is still place. This is a digital Camera There's still place for a film cartridge.

All they did was they removed the part, where the film cartridge is and where the film clips slap on a digital sensor. Extra electronics at the bottom added some electronic connectors, but why bother? Because people still like the same form factor and all that. There's still a place for film in that.

It's just mind blowing, and this is how it works. So a lot of times, when people say, wow, a billion people have cameras, and what is going to change about it? Hopefully, through this class, you'll realize it's just a wide, wide, open world.

There's certain space that peek. There's a lot of innovation and research, and then they mature. And there's not much more exciting about those fields, but there are certain fields that just continue to grow. And imagining is just one of them, but just look for the special cartridge here.

And if you look sideways, this is where the film [LAUGHS] wraps around. So you can take pictures of this and accurate notes. I just want to spend the first few minutes going through some of the fast forward preview items we didn't get to cover last time and then come back and talk about this topic.

We were here, somewhere, last time, where you can create these cameras where, instead of having an expansion-- instead of having an out-of-focus blur that extends like this, you can create lenses where, when it goes out of focus the point spread function actually rotates. This is a relational point spread function. And again, this is just a teaser. So we will study this into detail. And this is very powerful, whether it's for photography, or for scientific imaging, or real-time [INAUDIBLE] applications and so on.

And then we discussed about the duality of particles and waves. And we'll be discussing that a little bit. And then we'll talk about this, new types of cameras that are being designed, where it can create a 35-millimeter lens, including the sensor-- the whole package can be built with a thickness of just about 5 millimeters.

And so they're building these lenses where light comes in from the annulus, around the edges of the lens. And light actually reflects around, and image is formed in the center. And this is how many of the Sony-- what's the Sony [INAUDIBLE] cameras?

AUDIENCE: The Sony [INAUDIBLE].

RAMESH [INAUDIBLE], yes, works as well, where the sensors is actually in the-- for the Sony cameras, the lens is over here, but
RASKAR: the sensor is actually all the way back in the bottom. So light reflects around, and the image is captured at the bottom.

That's why the cameras are so thin, so that the dream of a flat camera is something we'll be looking for. And again, by making flat-- I mean, the new iPod Nano is, what, 0.25 inches thick? And it still has a camera. And that's because they're using crappier and crappier lenses now.

So in short, just creating a straightforward design where you have a lens and, behind that, you have a sensor. If you have a lens where-- with addition of this concentric reflectors and images on in the middle, then, effectively, you have a 35-millimeter lens with optical parts that's folded in. So we'll be looking at a lot of this interesting designs that I'm sure you will see in products coming up.

And then we'll look at, again, some really interesting lenses. A traditional lens is some form of a convex or concave lens, but the new lenses are just flat. Their front and back profile is just flat. So it's very easy to stack them, easy to put them in modern devices.

And the way they work is not by changing the thickness of glass, but by changing the refractive index. So the refractive index is high in the center and very low in the edges. So by doing that, that-- effectively creating a lens. And that's actually how a lot of flatbed scanners are built, as well, so we'll be studying that.

Very briefly, photonic crystals and how they will change imaging. [? Then in ?] photography, we saw the picture last time, the study how [INAUDIBLE] photography can be built, polarization, underwater. And there are really interesting cameras that are coming up. I can see if the lights can play around.

[CLICKING]

--can see the [INAUDIBLE]. Yes, slightly better. So these are polarization cameras where, instead of adding a polarizing filter in the front of the lens, the polarization filters are actually on the Bayer mosaic. So a typical camera, as you will study, will have very tiny RGB filters in front of each pixel. So every [? other ?] pixel has red, green, and blue filter.

And here, actually, they have different orientations of polarization, vertical orientation, horizontal orientation, +45- and -45-degree orientation for polarization. And from that, it turns out, under certain lighting conditions, such as sunlight, light that's reflected from vehicles or from faces is partially polarized. And they can look at how the light's polarized in all these four different directions and, from that, estimate the orientation of the surface or the surface normal. And from that, they can create a 3D model.

So it's pretty exciting. They're getting lots of money from the government for detecting vehicles in complex backgrounds. So this might come in a consumer camera. It might come for scientific imaging. And there's a really interesting sensors. I think we talked about this very briefly, last time, [INAUDIBLE] sensing.

We'll spend a lot of time discussing [INAUDIBLE] sensing. There's a lot of hype about it. We'll try to understand where it works, where it fails, and what's the power of [INAUDIBLE] sensing, and where it can be exploited.

Some other bizarre photos you may have seen are near the photo finish in sports photography. So this photo looks like an ordinary picture. Here's the winner of this, I believe, 100-meter dash.

But if you look at the people behind, they have really strange [INAUDIBLE]. Now, look down here. Anybody has a guess, what's going on? Here's another one with even more distortions. Look at the leg of this. And that's the finish line. This is the photo finish picture in sports photography compared to [INAUDIBLE]. Yes?

Because the sensor is scanning, line by line.

Mm-hmm. Exactly. The whole photo is not taken at a single instant. The photo is taken, actually, one line at a time. As you can see, all the-- at least when the [? time ?] is finished because the shoulder-- the moment you cross the shoulders is when it matters.

So as you can see, when they finish, they all have their body leaned forward, right? So, clearly, not all of the time we're finishing at the same instant. But, so, this particular line, for example, captured first, and then the next one, and the next one, and the next one. So this whole picture was captured over a whole second, but it's actually running at 2,000 frames per second. So in one second, you're capturing 2000 lines, and that simply plays together to construct this picture.

And this is very useful because for the judges, for the referees, they can just look at this picture and figure out exactly when every player had crossed the finish line. And it's a nice summary of the finish of the race. But not finished at any given instant, but finish of every instant.

So, if you look at the background, you'll see that there are no vehicles. There are no big signs. Same with the [? crack ?] here. And that's because we're capturing the same exact line in the camera and just capturing the same one again. So imagine taking a photo, throwing away all the pixels, except the center line, taking another picture after one-- by 1/2000 of a second, again, just keeping the center line, and just putting it together, so very interesting photography.

And I believe I saw a product on-- maybe it was an iPhone app or one of the mobile phone cameras, where they're creating this new, fun photography where, instead of taking the photo in a single instant, you release the shutter and it actually does a very slow rolling shutter. So it's actually a video camera that exposes the first line, then the second line, then the third line, and so on, for every frame. So I can release the picture, and I can simply turn in front of it. And what you will see is, the top of the picture, my forehead is being seen. In the middle is side of my face, and at the bottom, it's the back of my head. So it's this very beautiful pictures, that can be created with this.

But in this class, maybe you can convert a video camera into some other interesting projection of this x y [INAUDIBLE] right, because a video camera is x, y pixels, and we are capturing it over time. So it's a three-dimensional data set. And a final photo is actually some 2D projection of that. And this is one type of project-- one way of projecting it, but we'll study how there are some other ways in there for scientific imaging or for artistic photography.

AUDIENCE:[INAUDIBLE].

RAMESH For sure.

RASKAR:

AUDIENCE:Do you know what type of lens they use?

RAMESH They usually use standard lenses because they still haven't create-- if you're looking at a very narrow line of view

RASKAR: direction, you still have to form an extremely high quality image. So you cannot use a [INAUDIBLE] [vertical ?] lens or anything like that. So they still use a traditional lens for that.

AUDIENCE:Is this the main difference between a video camera and standard cameras?

RAMESH In terms of its-- I think there are subtle differences in-- when it comes to electronics, of course. So, for example, a

RASKAR: camera like this and on a cheap device is actually a video camera. And when you release the shutter-- or any cheap digital camera is actually-- still camera is actually a video camera. And when you release the shutter, it takes the next frame and captures that as a real photo and gives it to you. But it's constantly running as a video camera, and that's why, on the viewfinder, you can see the video of what's been captured.

So even if they're selling it as a digital still camera and even if they don't have a video mode, it's actually a video camera. And then there are more purists who build SLRs and so on. And they say, no, no, it is a digital still camera, and you must see through an optical viewfinder. And, but, just like the people who make the camera we saw, those people are a minority. And All this snobbishness about, when you have an SLR camera, you must have a mirror is also gone, fortunately.

So that's the fundamental difference between a still camera and video camera. But there's still an issue of bandwidth, and that comes to storage and processing. How quickly can-- even if you're exposing the whole sensor-- so, in particular, you might have a digital camera that's six megapixel, but the video is only less than a megapixel. And that's mainly because of bandwidth, and storage, and so on. So, but it's not really-- it's just an electronics issue, so there's no reason to-- in the future, it will be all fused into-- once the bandwidth and storage issues become not so critical, they'll just fuse into one single device.

There are already cameras, I believe, that, once you release the shutter-- you can be in the video mode, and then you release the shutter and just store some of those frames. So that's straightforward. So, in terms of optics and image formation, there's no difference.

There are some other differences, of course, in terms of noise. So if you-- you know that, for a digital still camera, when you release the shutter, usually, a mechanical shutter opens and closes to integrate the light for a finite duration. But, clearly, for a video camera, there's no mechanical shutter opening and closing for each frame.

So why don't they just do that for digital still cameras? And why not have a-- why not get rid of the mechanical shutter, altogether, and just do an electronic shutter? And, again, purists want that sound, ka-chung, ka-chung, but there's absolutely no reason to do that. And, nowadays, it's true that a lot of cameras have focal-plane shutters that sometimes are mechanical, and sometimes they're just electronic.

AUDIENCE: Is there absolutely no reason to have a mechanical shutter?

RAMESH So, exactly. So sometimes you have issues like just thermal noise. So if you're taking a picture that's 15 seconds long,

RASKAR: then you don't know how the camera's going to behave, how much noise is going to collect over 15 seconds. So a camera maker has a very good model of what the noise will be if you take a very short exposure for it, like under a second. But if it's a 15-second long exposure, depending on what conditions you are in, how warm your hand is-- all these things are going to change the noise properties on the sensor.

So, typically, they're two pictures. They'll take a 15-second long picture with the shutter open. And then they will take a 15-second long exposure with the shutter closed and then subtract the two images to get rid of the noise because there could be noise that varies over the frame of the sensor.

So, every once in a while, you have to-- it's good to have completely dark conditions created mechanically. But, again, that will change over time. I'm sure they'll come up with solutions where you don't have to do this extra measurement to figure out the noise in the system, of doing the subtraction.

I mean, communication world, right? I mean, you have to deal with noise all the time, and you don't always transform [INAUDIBLE]. So I'm sure we can come up with some intelligent encoding. Maybe every fifth pixel is being used for measuring noise, and other pixels are being used for capturing photo. So there will be some interesting encoding mechanisms there, I would imagine.

CityBlock Project, that's Guarav Garg. He was my intern. And this was a project that started with Augusto-- I forget his last name. Augusto Roman, I think.

And they started this project when they were grad students at Stanford, did an internship at Google, where they came up with an idea of mounting a camera in a truck and just driving in Palo Alto and, again, using a very similar idea of just taking the central pixel of every video frame. So if you can imagine, you take thousands of pictures, and take on the center column of each image, and just put it together. Then you have basically created this orthographic camera. It's a panorama that just goes on and on for the facade of a street.

And they applied some other interesting algorithm. Because if you have streets, or people, or cars that are moving or in front of the part that you care about, then, because of the parallax between multiple views, you can eliminate them, as well. So they put all those things together.

That was their CityBlock Project and, in it, had very interesting distortion artifacts and so on. But it was beautiful, and what Augusto tells me is that this is what pitched for Google Street View, as the representation. Because, right now, it's all discrete, right? You jump between one bubble to the next bubble.

And what it tells me is that the reason why they didn't choose what he did in his thesis and they went for this traditional bubble model is, apparently, in user studies, they realized that people get really confused when you show them pictures like this. People are more comfortable with jumping to the next bubble, and then looking around, and then jumping to the next bubble, and looking around. It's odd to me because I would have been more comfortable looking at something like this.

AUDIENCE: Yes.

RAMESH But this is what they have right now, of course, the new Google Street View system actually has a lot of additional

RASKAR: sensors. They have GPS. They have compass. They have a LiDAR that's actually taking 3D images of all the cities, so if a Google truck is going by, cover your eyes because they're shooting lasers.

[LAUGHTER]

I'm just kidding. It's eye safe. But they have 3D models, and the reason why the 3D models are not available to us, on the browser, is because they haven't figured out how to use this data, tera-- petabytes of data and how to transfer that and make it available in a streaming fashion. But, again, when that problem is solved, we will see 3D models that are exploited for all kinds of interesting purposes.

And this year, I was on the SIGGRAPH committee. And there were just a ton of papers on what you can do with street-level imagery because this is becoming a very hot topic right now, which is good and bad. If it's already hot, it's probably not worth pursuing it because there are too many people doing it. But, for a course project, it's perfectly fine to do, and you might be able to come up with a great idea for a SIGGRAPH paper.

We'll be studying a lot of work-- motion deblurring. This is actually a picture we took in Kendall Square. And we have cameras now that can take highly blurred images, motion-blurred images, and recover a sharp photo in post-capturing. So, last time, we saw focus deblurring. This is motion deblurring.

And the way it works-- and here's another example where you might have an aerial imagery scenario, an aircraft that's flying sufficiently low, and with a sufficiently long exposure, you'll get a blurred photo. But, again, using this technique I'm going to describe in the next slide, you can do motion deblurring for cars. And the basic idea was, instead of taking a photo where, when you release the shutter, you open the shutter and keep it open for, say, 100 milliseconds, and close the shutter, and get 1 photo-- instead of that, you release the shutter, and then you open and close it multiple times. And at the end, you still get one photo, but it has stopped and started the integration of light multiple times in between.

OK? And, of course, if you do that with a mechanical shutter, you'll probably void the warranty, and the camera will be unusable very quickly. So, instead of a mechanical opening and closing of the shutter, we used an LCD, actually, a ferroelectric LCD that becomes transparent or opaque. So when the LCD is, of course, opaque, you block the transmission of light, and then it's transparent. The light goes through. And that [? incurs ?] the motion of the object.

So if you have a point light, with a traditional camera, if I move the light very fast, you will see a streak in the photo. But with this camera, you'll see a ham code. You'll see a dash, dot dot, dash, dot in the screen. And that just preserves high-frequency information in the scene, and, as we'll study later, it allows you to reconstruct the original image and recover the sharp features.

And we'll study this in the frequency domain. We'll study this using linear algebra. And we'll also study this using pure [INAUDIBLE]. So those of you with different backgrounds, hopefully, will be very comfortable with many of these concepts.

Now, this is something-- this is a very worrying trend for sensors. So, 1994, pixels were about nine microns wide. OK? A human hair is about 50 microns, so it's 1/5 of the width of a human hair.

And the way you sense the color is you put a filter, green, blue, green, red. And that's the Bayer mosaic. And we'll study this data again. But over time, the size of this pixel is shrinking. Because if you want to buy a phone that has five megapixel, but it has really tiny focal length and-- that's fine.

So the sensor is shrinking. And the camera makers want to produce a lot of sensors from a given wafer. So you have a wafer of a certain size. You want to slice and dice it and then put it into each of the cameras. The more you slice it, the more pieces you will get, of course. And so, the logic is, if you keep shrinking these pixels, they can keep shrinking the size of the overall image sensor.

So, right now, we are down to 2 micron, 1.5 micron. I just saw a paper from Sony, yesterday, where they're claiming 0.9 microns. And what's the wavelength of visible light?

AUDIENCE:[INAUDIBLE].

AUDIENCE:0.4 to 0.7 microns.

RAMESH 0.4 to 0.7 microns, so let's write this down because we'll be talking about this all the time. So this is 50 micron. That's

RASKAR: [INAUDIBLE] your hair. This is 10 micron, 1 micron. And it's [INAUDIBLE] here, 500 nanometers plus 1 times 10 minus 6 [? meters, ?] 500 nanometers.

And blue is here. The green is here. The [? amber ?] is here [INAUDIBLE]. So the best [INAUDIBLE] all the way around here.

Now, the pixel, itself, is 0.9 micron. OK? And the [? ribbon ?] of light is invisible or it's really, really close to it. And this is a challenge, as we'll study later, because of the limitations of diffraction and other laws of physics. It's getting very challenging to do this [INAUDIBLE] with these type of sensors.

And, again, we'll study how people are thinking about this in a very traditional way. And some other teams are actually trying to exploit these sensors in completely unique ways. Now, there's a new design where people are claiming that, by using a pixel that's less than 1 micron-- by the way, compared to this pixel-- let's say this is 10 microns, and this is 1 micron. The amount of light one pixel will capture is reduced by what factor?

AUDIENCE:100.

AUDIENCE:100 [INAUDIBLE].

RAMESH By a factor of 100, right? So, with my phone, in 1994, if everything was the same, if I could take a picture in 100

RASKAR: milliseconds, now, how long we'll be my exposure time? It will be 10 seconds, right?

But, clearly, over the last 10 years, the exposure time of a typical photo hasn't changed. Otherwise, you'll have a camera shake because of the jitter in your device. So, clearly, the technology has improved, where the pixels have 100 times less area to capture light, and, still, we're able to get photos with roughly the same exposure.

So, in a way, the light gathering capability has improved by a factor of 100. Not really, but it's about that order of magnitude. And there are a lot of software tricks that are being played to compensate for the noise in such tiny sensors. So we'll study this issue quite a bit.

One common theme toward imaging is the two things you need. You need a lot of light. You'll always want a lot of photons, to take a good picture.

And number two is would like to have negative light, and this may sound strange. But people who work in radio and other fields, that are blessed with negative energy, and people who work in optics always have to worry about creating negative light. And as we'll go through the class, you'll realize that if somebody invents negative energy for light, it's like the invention of CO, to just represent nothing.

AUDIENCE:Negative what [INAUDIBLE]?

RAMESH Negative light.

RASKAR:

AUDIENCE:Negative light.

RAMESH Yes, a photon that has a negative energy.

RASKAR:

[LAUGHTER]

You have an idea? Think about it very hard. Yes, you'll get multiple Nobel Prizes.

[LAUGHTER]

All right. And then there's some really interesting biological creatures that we'll be studying, animal eyes. So a dragonfly or a krill has these compound eyes that-- this is the simulation of what this guy is looking at. He's proliferating thousands of images, probably not of that good quality, and they're used for very interesting applications. So we'll study that, as well.

So there's the project called Tombo, in Japan, which, I believe, stands for-- I mean, it stands for thin observation module by bound optics. But the word tombo means shrimp in Japanese. What does the word tombo mean in--

AUDIENCE:Dragonfly.

RAMESH Sorry?

RASKAR:

AUDIENCE:Dragonfly.

RAMESH Dragonfly. Sorry, dragonfly. So it's a nice play on the dragonfly. I should remember this. It's right on the slide before.

RASKAR:

And so, again, you have a single sensor. And you have multiple tiny lenses, and this is placed really, really close to the sensor. And the idea is that, if you have an object and you have plenty of tiny sensors right next to the-- plenty of tiny lenses next to the sensor, then, again, it'll form thousands of images. And from that, you may be able to do something interesting.

[COUGH]

We'll also look at a time-of-flight cameras. We saw, last time, one of the 3D [? V ?] cameras that uses time of flight to compute depth. So, how fast does light travel? 3 times 10 to the 8th meters per second. But what does it last-- what did it--

AUDIENCE:One [INAUDIBLE] nanosecond.

AUDIENCE:One foot--

RAMESH One foot per nanosecond. And sound?

RASKAR:

AUDIENCE:One millisecond.

RAMESH One foot per millisecond.

RASKAR:

AUDIENCE:OK.

RAMESH Just good numbers to remember. And so these cameras will study how they work. And it's quite possible that these

RASKAR: type of cameras will be available in really cheap devices. I'm talking about devices that are less than \$100.

So Microsoft Natal is likely to have a 3D-- a depth-sensing camera very soon. And Sony EyeToy is also likely to have one. In fact, Sony EyeToy has been-- Richard Marks, who is a kind of spiritual leader of EyeToy, has told [? me ?] multiple times that they will come out with a 3D camera anytime now. And they were testing it. They just want to make sure the cost is low enough to make it happen.

So this is a very exciting time for imaging. And those of you who think, oh, we already have a billion cameras-- how much is it going to change? You'll be extremely surprised, what you will see, just in the next two to three years.

And the 3D cameras are also being used in TV studios, where you may want to insert some virtual objects in the scene. So, a traditional [INAUDIBLE], you just replace what's the background. Maybe you have a blue screen, and you replace the background. But, here, you can put something in front and behind the person, with the right occlusion order. So this really is nice.

We'll be spending a lot of time about cameras for human-computer interaction. This is the topic of a lot of interest. And we look at different types of cameras, camera looking at people, camera looking at fingers, such as the frustrated total internal reflection FTIR, how optical mouse works, and [INAUDIBLE], and so on, different types of motion capture, [? V, ?] and so on. And we'll also look at what are the type of camera [INAUDIBLE] which are interesting, and some of them are just very 20th century. OK?

And this is very interesting because if you go to places like SIGGRAPH emerging technologies, where people are really combining the latest generation of algorithms and hardware, you'll see, over time, which projects are interesting and which projects are just stupid and boring. OK? And there's some really, really common ones, and those of you who are doing this for projects, it's perfectly fine.

But if you're thinking about using it for research, then think about it. This is something you should avoid because it's been done to death. The most common one, I will say, is where something moves and the music changes.

[LAUGHTER]

Just get out of the business. All right? Another one is you're going to write some Excel application where it's going to depend on detecting some skin color or something, something. It's not going to work.

OK? You might be able to show it as a demo. But as research, it's not going to work when the lighting changes, or the orientation changes, and so on.

Another very common excitement is I'm going to take a light, and put a studio camera, and [? crack ?] it, and create a 3D part of it. Done to death. There are things you can buy for \$100 that'll do it for you. And then there's a product that's not worth doing research, around the same problem. Again, segment a finger or face by putting some kind of a glow or something for segmentation.

Some artistic interactive displays, I can be behind and change something. Don't worry about it. A lot of people know how to do it.

And the problem with a lot of these demos, as you'll see, is that you can build this demo and you'll be able to get some people excited about it. And that will give you this positive reinforcement that, oh, this is the kind of stuff I should be doing. Because you can impress some people all the time.

[LAUGHTER]

--and you can impress all the people for some time, [LAUGHS] but not all the people all the time. And, of course, the world was not impressed in the original [? code. ?] And just remember that. So that's the problem. Let's see.

[INAUDIBLE]. So what we're going to do in this class is--

AUDIENCE:What's the solution?

[LAUGHTER]

RAMESH Yeah, exactly. Exactly. Well, [? you'll have to ?] attend the whole semester to see.

RASKAR:

AUDIENCE:[INAUDIBLE] the [INAUDIBLE] camera lens.

RAMESH Right.

RASKAR:

AUDIENCE:What's the current status of the gaze estimation [INAUDIBLE]?

RAMESH Right.

RASKAR:

AUDIENCE:[INAUDIBLE].

RAMESH Yes, we'll talk about that, as well, in the class. And there are some really good solutions. But I think it's still-- gaze

RASKAR: tracking is still pretty challenging problem. So I think that's worth still exploring.

And, again, this is not an offense to anybody. If you're doing it as a class project, it's perfectly fine. But, everyone, let's just be honest. When you are an acquaintance, we praise each other's work, and when you're friends, we criticize. And we're just going to be very, very honest about it and see what we can learn and what we can do next, right?

So the solution is to [INAUDIBLE] change the game. You want to try to build things that are robust, that do the kind of things that nobody else can do, and they allow you to do things that are not possible to these cameras. All right?

So we you just want these smarter sensors, smarter processing. And what you create should be just magic, not in terms of its application, but in terms of its basic building blocks. And, hopefully, through this class, you'll realize there are hundreds and hundreds of solutions that you could be using, instead of using some cheap camera that's available. And just because you have the SDK for it, just because the device is available, you're using it.

Let's get away from that, and let's try to build something that's unique and new. OK? Any questions on this one? All right.

AUDIENCE:- Question?

RAMESH Yes?

RASKAR:

AUDIENCE:So if you want to build stuff like the [INAUDIBLE] camera or [INAUDIBLE] like changing the motion of the shutter and stuff like this, so what other components are available?

RAMESH That's a great question. So let me just repeat [? Vena's ?] last question. If you really want to do all these things, what

RASKAR: are the components that are available? Can I just slap together a sensor, and a light source, and all this electronics?

The answer is it's not always that easy, unfortunately. It's not-- I wish Canon would just come out with a LEGO for cameras and just sell that. And I'm sure they'll sell millions of those. Because then you can create your own things, but, unfortunately, that's not available.

At the same time, what I've seen, especially at Media Lab, are people are extremely innovative. And they go and pick pieces from different places. I think [? Dan ?] has been looking into that. The other--

And SparkFun has been very supportive. And they sell camera modules for \$3 to \$9 now. You can buy a full-fledged camera for \$3. Unfortunately, there's not much you can change about it. But then you can go to companies like Point Grey and buy a little bit more expensive solution, maybe \$500 to \$1,000, and they'll give you more access to it.

So, I mean, all the projects that we do here, we are not building our own sensors. We are not building our own chip, processing chip. We're putting pieces together from different places.

And we are definitely on the bleeding edge. So if you want to build something unique, there is no solution available, you can just buy a kit, right now, to do it. I know [? JB ?] has been interested in building a-- what do you call it? What kind of camera, open-source camera, or--

AUDIENCE:Yes. And I'm actually working with a company in New York City that [INAUDIBLE] camera [INAUDIBLE].

RAMESH Excellent.

RASKAR:

AUDIENCE:But it's very slow. And it's still [INAUDIBLE].

RAMESH Right.

RASKAR:

AUDIENCE:It's not this small module that you can program in environments and everything.

RAMESH Right.

RASKAR:

AUDIENCE:It's obviously a good start to go to, to think of [INAUDIBLE].

RAMESH Right. Yes, so there are a lot of efforts in that. So I'm glad-- good to hear. I'm glad to hear about your work.

RASKAR:

And the Stanford, Professor Marc Levoy and his group, has been also proposing an open-source camera architecture, and they also ran into the same issue of how can they get some of the chip makers, and lens makers, and all that excited to put this together. So they're just getting started, and I hope that movement will continue. We have a lot of our own plans here, which we'll be disclosing in coming weeks. And it should be possible to--

There was a time when it was very cool to hike around with your car, right, in the '70s and '80s, if you're a cool guy, if you fix your car and put new tints. Now people say, OK. And then there was a time when it was very cool to build electronics, build cool electronics, do something with robots. People say, that's OK. I can just buy a kit.

And what's going to happen now is we're going to get into the physics of these things, not just the chemical engines and the electronic hardware. But we're going to get into the physics of it, whether it's in a UV light, whether it's chemical elements, or whether it's the sensor. And the next generation of kids who want to be cool are going to build things that just create magic. So I'm really excited about this whole area.

In fact, there's a group at Columbia University, Professor [? Schneier, ?] and he's putting together a-- I forget the name of the project, but it's also kind of a LEGO for cameras. And that's going to be a lot of fun. They're trying to create a whole high school curriculum based on campus, so it's very exciting.

So we look at how Jeff Hahn's project was developed. And remember, his paper came out only in 2005. That's the very first time he disclosed it, and after that, as you know very well, it's everywhere. John King is using it on CNN, to see how Obama is doing versus McCain, and it's just everywhere. Beautiful piece of technology, very old idea that's used in lots of other environments, and to study that.

AUDIENCE:Can I ask you a quick question?

RAMESH Yes.

RASKAR:

AUDIENCE:Yes, so if this is a camera, are all touchscreen displays essentially cameras?

RAMESH That's a great question. That's a great question. So this one, in a way, is still using a traditional camera.

RASKAR:

When John King is playing on CNN with all the [INAUDIBLE], what you don't realize is that, behind him, he requires a lot of space to put a camera and a rear projection screen. Right? It's not just looking-- walking up to some [INAUDIBLE] and playing with it.

AUDIENCE:I mean, so if I have two pieces of glass that I'm pressing on and I'm detecting which wires are actually touching each other, I've never thought of it as a camera before. So I'm thinking about it right now. And maybe it's kind of like a [INAUDIBLE].

RAMESH Certainly. That's the way you should be thinking. A camera is not a two-dimensional sensor. It's zero-dimensional,

RASKAR: which is a point; one-dimensional, a line or a curve; two-dimensional; three dimensional. You'll see eight-dimensional sensors.

AUDIENCE:But we usually think of cameras as sensing light, not sensing pressure.

RAMESH Mm-hmm. But you can convert pressure into light. So [? Kimo, ?] I don't know if he's here. He'll come and give a talk

RASKAR: about how his device works.

And what they basically do is they have created this surface. It looks like Jell-O. And they have built some really, really beautiful demos. I think-- I believe it's going to come in a couple of weeks. And I'm sure it's--

And what this Jell-O-like device does is I can put my finger on it and it transfers that into a highly visible impression. And then, in very simple words, this [? predicts ?] a photo of that, with changing lighting direction. And they can create this very millimeter or micro-- multimicrometer-scale objects and then they take photos of that.

So, yes, you always need a transducer. A camera is converting photons into electrons. But you could have other transducers, pressure into electrons or pressure in the photons. Yes.

AUDIENCE:I think what he's talking about is [INAUDIBLE].

AUDIENCE:Well, I mean, but a camera is sort of some generalized-- we sometimes think of any kind of sensor that can give visual information as a camera.

RAMESH Right. Maybe the other word of phrase is that what is-- what sensors can give you visual information? And then, yes,

RASKAR: whether it's a resistive, or positive, or inductive, it still could give you some geometric information. So, but, although, we won't be studying them as such. We'll be studying more based on whatever happens to photons.

But, then, going beyond that, how can you create something out of a thin screen? And this is [? Miparg. ?] That was his class project, last year, in this class. And he started thinking about it in the class, and he built an initial prototype for his final project.

And this, of course, how it now has become a SIGGRAPH paper, part of this master's thesis. He won the student research competition this year, just out of this class, last year. And we'll study how that works.

And then we'll study things like this, an [? Autopen, ?] which has a grid of dots that are slightly misplaced with respect to the center of its position. And every block of, I believe, six by six is unique. And just with this grid of six by six, that can create sufficiently unique cores, so that if you create paper with this core and just laid it out, it'll cover half the land area of the US.

So it can print many, many, many papers with this [? cores ?] printer on them. And a pen has a camera that looks like this six-by-six core and figures out its unique location in a coordinate system that could span 1,600 kilometers by 1,600 kilometers. So it's really unique. And then, this way, we know where the pen was, on which page, and which x-y coordinate. And this is a way to basically record the strokes that you [INAUDIBLE]. Yes, [? Jim? ?]

AUDIENCE: So [INAUDIBLE] 200 people, and more than half are lawyers. There are 250 patents on this. So it's interesting to see the context of this. So it's very difficult, actually, to have it to be able to go beyond that [INAUDIBLE] or what [INAUDIBLE].

RAMESH Right.

RASKAR:

AUDIENCE: So [INAUDIBLE].

RAMESH But I think [INAUDIBLE] has license.

RASKAR:

AUDIENCE: Yes, yes. You have license, but you cannot go, and especially in terms of hardware, beyond what--

RAMESH Exactly.

RASKAR:

AUDIENCE: --[INAUDIBLE] maybe you would like to [INAUDIBLE] things.

RAMESH Right.

RASKAR:

AUDIENCE: So [INAUDIBLE].

RAMESH So maybe that should be a project, how to get around the patent.

RASKAR:

AUDIENCE: [INAUDIBLE].

[LAUGHTER]

RAMESH How to invent technology that can [INAUDIBLE] a patent. [INAUDIBLE]?

RASKAR:

AUDIENCE: This is it.

RAMESH Oh, you have one.

RASKAR:

AUDIENCE:Ah.

- Oh,
excellent.
Do you
mind
passing it
around?

AUDIENCE:No.

RAMESH Yes, just pass it around. All right. And then we'll talk about [INAUDIBLE], which is a project with [INAUDIBLE] and

RASKAR: Professor [? Herrera, ?] who was here last year.

And the idea is how can we create-- how can we exploit properties of cameras for objects that are very far away? How can we add intelligence to the world, so that the world is more compatible with billions of cameras people are carrying. And I won't go into the details, but the basic idea is to convert a point that's in sharp focus, looks like something that is a size of three millimeters by three millimeters, and take an out-of-focus photo [INAUDIBLE] and convert a circle of confusion into circle of information.

And we look at motion capture solutions for [? XCI, ?] some other solutions we have built here for inverse optical motion capture and some other [? XCI ?] devices. So we'll spend a lot of time on that. So, an announcement here. [INAUDIBLE], who's one of the leaders in using cameras for [? XCI, ?] is giving a talk, actually, on Monday-- this changed-- on Monday at 4:00 PM, I believe, in the [INAUDIBLE] Room. And he's the inventor of-- If you're familiar with Microsoft Surface, which is a tabletop surface with a projector and camera underneath, he built a version that's a variant of that, where they put a screen-- the screen is actually not diffused, but it's switchable. It switches between a diffused screen and a transparent screen, electronically.

So, in one frame, you're projecting the image on it, and you can see it on the tabletop. In the next frame, it switches to become completely transparent. And the camera underneath can see the world through this diffusion, through the screen, and it can do some gestures on top. And again, in the old frame, it goes back to being a diffuser. So he'll be talking about that.

He just received a TR35 Award from Technology Review. So he's here in town for that. And he'll give this presentation on Monday at 4:00 PM. I believe I sent out an announcement, but I'll send out one more.

And our own [INAUDIBLE] mystery, who brings this little beautiful sixth-sense display, also got a T35 Award. So those of you who are not familiar with his work, really great [? XCI ?] projects can lead to TR35 hours.

[LAUGHTER]

All right. Then we'll spend quite a bit of time talking about scientific imaging and [? conversion ?] imaging in sciences. And this is something that I've heard a lot of times, new instruments lead to new discoveries. And in the 20th century, the most important instrument was a computer, right?

And what we might see in the future, in my biased opinion, is the most important device we'll have is a really important imaging mechanism. We don't know what it will be, but it could be some permutation of a combination of what we are studying here. So computational imaging, we mentioned, has led to-- it has just transformed our world, unlike a lot of fields in-- we know, which are really important.

I mean, if you think about just Nobel Prizes, my background is in computer vision and graphics and there have been no Nobel Prizes in computer vision graphics and not even Turing Awards in computer vision or graphics. Pretty sad. But if you think about imaging, there have been tons of Nobel Prizes in just purely imaging mechanisms, and we'll be studying them, phase contrast microscopy, a lot of CT scanning, and MRI, and so on. I don't know why important fields like graphics and vision are not getting as much attention because we are solving very important problems, as well. But maybe it's not being pitched right, or there's something more there.

Anyway, so we'll study computational imaging, in terms of medical imaging, astronomy, applied physics, and biology. And a lot of the ideas are [INAUDIBLE] applicable across different fields, whether it's photography, [? XCI, ?] computer vision, and so on, so tomography, confocal microscopy, and so on. All right. So let me switch over to this topic.