MITOCW | mitmas_531f09_lec04_1.mp4

RAMESH RASKAR: So we'll finish up the computational illumination that I couldn't finish last time. And then, we'll talk about light fields. We'll talk about assignment number two, which is on optics. We have two options. Some other announcements. Make sure you're checking the wiki and the reading material. There's a lot of material we will not be covering in detail, for example, how a camera works, and depth of field, and apertures, and so on.

There's plenty of information online. This particular tutor on YouTube seems to be pretty reasonable. The guy looks scary and has very shady lighting. But other than that, he seems to do a good job of explaining for free. And this particular one is about different SLR concepts. I don't know why he has a pink curtain behind him. Very strange.

[LAUGHTER]

But he does a very good job, basically, understanding all those things. This URL is up there. Good. So there are lots of ideas for final projects. I would like you to start thinking about them. Just come and join us on this side as close to this corner as possible.

[LAUGHTER]

Yeah, lots of ideas. Just come and talk to me. And we have several mentors listed for the class. Professor Mukaila is there. He's one of the world's leaders in computer vision. So you can talk to him about your ideas. Or Ankit Mohan, who presented last week, is also a mentor. So there are a lot of people who you can talk to. We also have Mathias here. And you can brainstorm some ideas with him about sensors and so on.

And I just wanted to give you an update. So these are some of the class projects from last year. So as I said, actually, [INAUDIBLE] photography was the best project award. The bidirectional screen got the student research prize at SIGGRAPH. And we just heard last week that the looking around the corner project won the Mark Prize, which is the top prize in computer vision. So it's extremely prestigious.

And Kermani, who was working on this project, received the number two prize. They gave out two prizes, the best paper award and the second best paper award. And he won that. So hopefully, your project will also be with that level of fame and fortune. All right.

[LAUGHTER]

So homework, there are three things you need to do. And I don't think those are very clear to everybody in the first assignment. You need to create your own website where all the information is kept. You need to submit the link onto Stellar, and then to submit the input and output photos on the Flickr group for the class.

Right now, I think I have only about six people on the Flickr group. That means not everybody was able to post their pictures on the Flickr group. And the reason why we had it on Flickr group is that we can comment on each other's results and so on. Of course, the rest of the world can as well. But we can.

And hopefully, as we go along in your assignments, you'll be able to create pictures that are not just visually interesting, because the guy has spent a lot of time with a lot of patience, with very expensive care, but because they're really beautiful computational techniques. And you'll be creating magical photos.

Last time, we were talking about people are-- there was a fascination with high-dynamic range. Right now, there's fascination with the tilt-shift. But hopefully, the next fascination that you'll see on Flickr will originate from some technique that you invent in this class. Maybe a light-filled picture could be the next one, which is the assignment for us this week.

So there are two assignments, two choices, sorry, for the second assignment. The first one is extending Andrew Adams, who's a PhD student at Stanford. You'll hear his name all the time. He's just done wonderful, wonderful work.

And we'll be extending his so-called virtual optical bench. So the way it works is you go to his virtual optical bench, which is built in Flash. But you can use anything you want. You can use Java, C++, MATLAB, anything you like.

And it allows you to do operations on light as if you're on a vertical bench, so optical bench. So you can put things here. So you can put lenses. And then, you can put mirrors. And you can put blockers. You can put, I believe, diffusers and so on. So you can put any of these elements and, of course, rotate them and so on, to create a really, really useful set of tools to understand how light propagates.

So I don't know, if you're building an XCI system, rather than having to draw everything in Illustrator or on a piece of paper, you can just fire up this application. And you can design your whole thing and also understand how light behaves, insert the focal length, and so on. And you can very nicely build your optical setup.

So what Andrew has done, and he's graciously offered his source code with us, is he's provided some basic functionality of inserting lenses, and mirrors, and blockers, and diffusers. But it's not complete. So for example, if I change the focal length-- so let's see our optics 101 here.

So as you can see, here's our plastic optics. If I put the lens here, then all the rays come into sharp focus. As I move the lens to the left, the focus points also start moving. At some point, the rays start diverging. And at what point will they be compatible if I keep moving the lens to the left?

AUDIENCE: [INAUDIBLE] focal length?

RAMESH RASKAR: At the focal length.

AUDIENCE: Yeah.

RAMESH RASKAR: So if I move the lens all the way to the focal length of the same length, then all the rays coming from that will become parallel. And even if I add a second one here, then all those rays will become parallel. And as I'm being even closer, they'll start [? damaging. ?]

And you can use this to demonstrate how to build a telescope or a microscope by adding a second lens and so on. So as you can see, you can play with multiple lenses, and things, and so on. So it will be a lot of fun.

So what you have to do for the assignment is add more elements such as prisms, and gratings, and so on. And right now, it's just creating-- just shooting rays. But it's not forming an image. If you go to optical bench, you'll be able to actually form an image. So you have to write a very simple routine to integrate the light from multiple rays and actually form an image, which is relatively straightforward because if I put a sensor here, all I have to do is go to this pixel and sum up all the rays that reach that point. And that's the intensity of the picture.

And as I move this around, the same pixel has no rays arriving there. So the intensity is zero and so on. And if I add some other point here, then this should look-- this should add up to become a blue pixel, this should add up to become a vertical pixel, and so on. So it's a really simple operation. And maybe you have a window here that shows the form image.

And also, it doesn't have an ability to actually enter the specific numbers. So if I want to say, the focal length should be 30 millimeters, it doesn't have a way to do that. It's all based on GUI. So a few minor additions you have to make.

And it's an open-ended assignment. The more you add, the better. I'll give you some suggestions on how you can improve it. So that's option one. And this requires programming, [? section ?] programming.

Option two is synthetic aperture photography. So we saw this in the very first class, where you have an array of cameras. And you can see behind an occluder. So with an [INAUDIBLE] camera, if you're seeing through these trees, by doing some software operations outside of images, you can actually see around-- see through the foliage, and the leaves, and so on. So this is what you'll be doing for your assignment.

Now, of course, you don't have this million dollar camera array. So you'll have to come up with a shortcut. And the way we will do that is we will just translate a camera and take multiple pictures. So instead of having the camera array, you can just take your XLR or webcam.

Whatever camera you have, even your cell phone camera, just transfer it and take maybe about 30 pictures or 40 pictures. And then, you're going to do the same operations. And we'll learn about how exactly it's done, how you can see through some of those. So you'll be creating these see-through effects by eliminating the foreground pixels. So it's a lot of fun to do this assignment.

And there's a third option which you may want to take, which is same-- so for this one, you actually have to make sure that you're moving the camera in a reasonable way. If you have experience with building LEGOs, you can just build a LEGO robot that moves the camera and takes pictures.

Or just manually, you can just mark it on a ruler and just take the camera here. Take a picture. Take a [INAUDIBLE] camera here and so on. But all of this involves a little bit of physical work. If you want to just stay in your office, and look at the screen all night long, and do the programming, and never have to go and play with real things, you can do the whole thing just in software.

But again, using this virtual optical bench. So you can synthesize the images maybe in OpenGL or using the virtual optical bench with different viewpoints, and add them together, and so on. Not as much fun, but maybe you can create even more interesting effects. So it's just an option. I would suggest going with the real thing as opposed to doing it all in software.

All right. So let's go back and finish the few topics we left. We couldn't cover develop-- well, it's coming up. The reason why I wasn't here for the last class was I was at an event called Gadgetoff in New York.

It's a really great event if you get an opportunity. It's the gadgets equivalent of [INAUDIBLE]. And people are blowing up things. Crazy displays. Crazy robots. Crazy cameras. It was a lot of fun. If you go to YouTube and just type Gadgetoff, you will see many of the [INAUDIBLE].

While that's coming up, let me show you something else. So if you remember, we were talking about Google Earth live. And imagine if you can fire up Google Earth and you can go to any part of the world and see it live because, very soon, we'll have cameras on every street light.

We'll have cameras on every bus, every taxi. And even people carrying their own cell phone cameras will subscribe to some service. And they'll be broadcasting. They'll be-- we don't know for what. But they will.

And when that happens, when you can really fire up, and go to any part of the world, and see it live, it will be a very different notion than what we think about as mounted cameras. Nowadays, we think of them as mostly for surveillance. But over time, they will become-- maybe we'll be able to use them for beneficial commercial solutions.

So the question is when it's going to happen. And some of the things are already happening, like this particular project. I'm sure it scares a lot of people, that you'll be able to go to any part of the world and see it live. But you can also imagine it will be used for some good reasons, definitely for commercial reasons.

You can have an index of, what's people's health in this area? If you just put two points and see how long it takes for people to walk across those two markers, that tells you how healthy people are in this town. Or at least, how much in a rush they are all the time. If you are a real estate agent, you can figure out, what's the property value of this area based on what foot traffic you get at a particular restaurant? If you want to go to a restaurant, you can find out, is there a long wait there? And so on.

So this was a project at Georgia Tech, [INAUDIBLE] group. And it's an interesting concept. So remember, we had a poll about, when do we think a digital camera as a standalone device will disappear?

And most people were between five to 10 years. There were a couple of people who said even two years. I think the mean was somewhere between five to 10. So here's another question. When do you think we'll have Google Earth live when you can fire up, and go to any part of the world, and see it live? You won't be able to do it in the middle of Iowa.

AUDIENCE: What percentage of the world are you saying?

RAMESH RASKAR: Let's say at least one city, let's say.

AUDIENCE: Wait. At least one city?

RAMESH RASKAR: Where you can fire up and go to any one particular city, let's say Manhattan, and see it live.

AUDIENCE: So one particular city or any--

AUDIENCE: One city, or any city?

AUDIENCE: Not one [INAUDIBLE].

RAMESH RASKAR: Any one particular city. Any one city.

AUDIENCE: One city.

RAMESH RASKAR: Yeah.

AUDIENCE: It's going to have to be--

RAMESH RASKAR: It could be Tokyo. It could be New York. It could be Sioux City.

[INTERPOSING VOICES]

AUDIENCE: You probably need one [INAUDIBLE].

AUDIENCE: Oh, yeah.

RAMESH RASKAR: Yeah.

[LAUGHTER]

RAMESH RASKAR: But remember, it's not serverless.

[INTERPOSING VOICES]

RAMESH RASKAR: It's not serverless.

AUDIENCE: It's any [INAUDIBLE].

RAMESH RASKAR: It's A commercial service, the same way people are thinking about creating free Wi-Fi networks.

AUDIENCE: Can't you do that-- they have a [INAUDIBLE] TV show. Didn't they have something with web cams, publicly available web cams, all along the border that [? essentially ?] watched [INAUDIBLE]. He watched the--

[INTERPOSING VOICES]

RAMESH RASKAR: *Minuteman. Minuteman* [INAUDIBLE] or something.

[LAUGHTER]

AUDIENCE: People running across--

RAMESH RASKAR:	Yeah, I think that's what it was called.
AUDIENCE:	So I guess that's a film.
RAMESH RASKAR:	Right. But again, think of it from surveillance. All the efforts we have seen so far are surveillance-based. And the video that you just saw is not about surveillance. It's about watching sports games, getting a sense of what the traffic is on the street, so making-- right now, all these street level maps and aerial maps are completely lifeless. They're just static snapshots. And you want to add some dynamic element to it.

It may not be-- it's not realistic. Maybe you have-- there are technologies that shows a flow of people. But you cannot recognize who's who. There's all these other overlays on top of that. So just like when Google Street Maps came around, you could see the people. And now, the faces of the people are blurred or the license plate number. But in the beginning, everything will be free for all.

[LAUGHTER]

Pretty scary. But-- |
| AUDIENCE: | Yeah, it brings up a lot of queer issues of voyeurism.

[LAUGHTER] |
| RAMESH RASKAR: | Certainly. Certainly. The same issue that Google Street Maps. Yeah? |
| AUDIENCE: | I think [? Google Labs ?] were developed [INAUDIBLE] on satellite-- |
| AUDIENCE: | Which is when?

[LAUGHTER] |
| RAMESH RASKAR: | But satellite, that's in the geosynchronous satellite or somewhere that's like a blimp that's floating around?

[INTERPOSING VOICES] |
AUDIENCE:	[? The media have ?] their own satellite. And they can take their own picture of the [INAUDIBLE].
RAMESH RASKAR:	That's too far away.
AUDIENCE:	Really?
RAMESH RASKAR:	And the satellite is too far away.
AUDIENCE:	You have a small satellite [INAUDIBLE].

RAMESH RASKAR:	But you also want to exploit infrastructure set up by somebody else. So you might send your own satellite to look at your ex-girlfriend.
	[LAUGHTER]
	But you want to also exploit existing networks. So maybe you won't have to because there'll be enough cameras. Today, Google might come and say, I'll pay you $5 a month. Just aim the camera outside your office, or outside your dorm room, or outside your shop. And we'll pay you $5. Just give us the stream. And for me, it doesn't matter to me. I can just aim it outside. But by using that network, they can provide this service.
	There's a bus-tracking proposal [INAUDIBLE] cameras outside your house.
RAMESH RASKAR:	See? Yeah.
AUDIENCE:	They were on route.
RAMESH RASKAR:	Yeah.
AUDIENCE:	You can use that to figure out where buses were in the city, because there wasn't enough money in the city to pay for traffic buses.
RAMESH RASKAR:	See?
AUDIENCE:	But also, I imagine that-- imagine getting a camera on every cab or something and [INAUDIBLE].
RAMESH RASKAR:	Yeah, exactly.
AUDIENCE:	And then, it's like you're having a fleet of Google [INAUDIBLE] drive your car.
RAMESH RASKAR:	And the light because they can figure out the condition of the roads. They can tell about traffic. They can tell about potholes. They can tell about rain or no rain.
AUDIENCE:	Yeah.
RAMESH RASKAR:	In fact, there was a very nice project in Tokyo where they placed just on-off detectors-- not on-off detectors, but detectors on the wipers of taxis. And from that, they figured out-- they got a map of the whole city and how much it is raining in which part of the city.
	And that was, of course, much more accurate than the weather reports that they were getting because they just measure it in some bucket, how much it's raining. And here, just, they're not measuring rain. They're just measuring how fast the wipers are moving all over the city. So that's an indirect way of capturing visual data. So one city, how many years it would take?
AUDIENCE:	You're going to do a poll?

RAMESH RASKAR: Yeah, we'll do a quick poll. And you're going to put your vote on the line.

AUDIENCE: I'd say two.

RAMESH RASKAR: Two years?

AUDIENCE: Yeah.

[LAUGHTER]

AUDIENCE: Yes.

AUDIENCE: Either be a European city--

[INTERPOSING VOICES]

[LAUGHTER]

RAMESH RASKAR: Were you two years last time as well?

AUDIENCE: Yeah.

RAMESH RASKAR: Yeah.

[LAUGHTER]

All right. So you're a two-year guy. All right. 10 years?

AUDIENCE: Yeah, one city?

RAMESH RASKAR: Just one city, yeah.

AUDIENCE: One city in 10 years?

RAMESH RASKAR: Yeah.

AUDIENCE: So which city, Las Vegas?

[LAUGHTER]

[INTERPOSING VOICES]

AUDIENCE: It would be probably some city that wants to--

No, [INAUDIBLE] doesn't want to be in that [INAUDIBLE].

RAMESH RASKAR: It has to be a city where something interesting is happening.

[INTERPOSING VOICES]

RAMESH RASKAR: So I'm guessing it's not in the US, first of all. We've got too many privacy issues.

AUDIENCE: Yeah.

RAMESH RASKAR: It's going to be--

[INTERPOSING VOICES]

AUDIENCE: It would be like a Hollywood, but maybe in another country, the equivalent of--

[INTERPOSING VOICES]

RAMESH RASKAR: Yeah, I can imagine a city like Hong Kong or something.

AUDIENCE: Yeah.

AUDIENCE: Singapore would do it.

AUDIENCE: Yeah, Singapore.

[INTERPOSING VOICES]

[LAUGHTER]

RAMESH RASKAR: Yeah, all right. Yeah, any city in the world.

[INTERPOSING VOICES]

RAMESH RASKAR: So 10 years?

AUDIENCE: Sure.

RAMESH RASKAR: Yeah, and we should go the other way. Five years? Sorry.

[LAUGHTER]

Wow. 10 years? And never?

[LAUGHTER]

Wow, Jamie. Only satellites?

AUDIENCE: Yeah.

RAMESH RASKAR: [INAUDIBLE]

[LAUGHTER]

Don't cross your path with Jamie because he'll still-- he'll put his satellite on your trail.

AUDIENCE: $8,000?

AUDIENCE: What?

AUDIENCE: It cost $8,000 to put in.

AUDIENCE: It's much improved if you have a small satellite. And you guys, you can implement [INAUDIBLE].

RAMESH RASKAR: But the imaging is not that easy.

AUDIENCE: The optics aren't easy.

AUDIENCE: It's not easy.

RAMESH RASKAR: Horrible image from 3,600 kilometers.

AUDIENCE: Yeah, that's the problem.

[LAUGHTER]

RAMESH RASKAR: Yeah.

AUDIENCE: [INAUDIBLE].

AUDIENCE: [INAUDIBLE] software reconstruction.

[INTERPOSING VOICES]

RAMESH RASKAR: Yeah, just zoom, zoom, zoom, yeah, in software.

[INTERPOSING VOICES]

AUDIENCE: In software.

[LAUGHTER]

[INTERPOSING VOICES]

RAMESH RASKAR:	Is that it?
AUDIENCE:	It is also a lot of data storage. You can store the whole world's history--
RAMESH RASKAR:	Right.
AUDIENCE:	--based on this Google Earth. If you have the records and archives of this, then you can actually track back all the crime scenes to [INAUDIBLE]. And you need super huge data storage for that because storing in complete words, live history is difficult. So it depends what's important and what you really want to store.
AUDIENCE:	Store it on [INAUDIBLE].
RAMESH RASKAR:	Certainly.

[LAUGHTER]

Storage on the [INAUDIBLE] would be happy if you come up with ideas where you will store more and more. I remember back in the days when Mosaic was around in '92. I remember if I opened the Mosaic homepage, they listed every website in the world. Do you remember this, Mathias?

[LAUGHTER]

Right? And over time, of course, there's no way you can make a list with everything so dynamic. You still have places like archive.org, I believe, where you can time shift and look at how a page looked over the last 10 years or so. I don't know how far back they go. But you can look at-- you can go to my homepage, for example. And you can see how it looked every few days over the last 10 years.

AUDIENCE:	There's a [INAUDIBLE], though. You have to be a certain popularity, don't you, to make it there?
RAMESH RASKAR:	I'm sure, yeah. But still, the fact that they are keeping a snapshot of the whole world wide web over time and over time will make a record of the whole world visually, not just on the websites. So it's going to happen.
AUDIENCE:	[INAUDIBLE] with the book? Because based on it-- because all the URLs [INAUDIBLE].
RAMESH RASKAR:	Yeah, everything printed or anything.
AUDIENCE:	Yeah, that's something that would be like, [INAUDIBLE] telephone. You couldn't pick up.
RAMESH RASKAR:	Right.
AUDIENCE:	[INAUDIBLE]

[LAUGHTER]

AUDIENCE: Yeah, [INAUDIBLE] like the telephone book. I imagine if there's equivalent of the internet, most every year, you've got this bigger and bigger *Yellow Book.* It's awful.

RAMESH RASKAR: So the complexity of data, and bandwidth, and computing, and memory just grows exponentially. I remember when I was in grad school, the hot topic was how to assign IP addresses from mobile devices like mobile laptops. That was a huge research problem. And now, we are flipping through it all the time completely seamlessly. So we don't even think about that as a problem. So it's a similar situation here in the visual domain. Remember, this is going to be the decade for visual computing.

But we're not just talking about social implications or business opportunities in this class. We really want to think about, what, kind of, cameras make sense? If you're going to put a camera on every street light or on every taxicab, what, kind of, cameras should be developed so that it's compatible?

If you put a camera on a taxicab and you keep taking photos, it's all going to be blurred. It's going to be totally useless. If you put a camera on a streetlight, it'll be a wider field of view. But you will not be able to see anything specific. There are all these issues that come up.

Maybe it should exploit different wavelengths. It should use different optics, different processing. Maybe they should talk to each other, and do some coordination, and so on. So these are some really interesting challenges. So it could be one of your final projects.

All right. We can go finally up here. So last time, we were talking about computational illumination. And we saw several examples like creating cartoons, [INAUDIBLE] for CI, or [INAUDIBLE] flash matting by taking pictures in the foreground and background. And we'll cover just a few more projects before we switch over to optics and light fields.

So here is a [? sensitive ?] paper where if you wanted to create an image that's more comprehensible, like on a leaf, where if you take any one of the photos, it may not show the structure very well. You can just put an object, keep the camera fixed, put it in the video mode, and just move the flashlight around just randomly. And you collect a set of photos, maybe three, four, five, 10.

And then, you do some computation on that to create an image, where you can enhance the shape or you can-- see here, you're just enhancing the shape. Here, enhancing the detail and so on. So there are lots of techniques that they showed, for example, there are five and four images. Yes?

AUDIENCE: What do you mean by the shape and [INAUDIBLE]? I don't see the difference [INAUDIBLE].

RAMESH RASKAR: So in case of the leaf, you want to show the leaf, how the folds--

AUDIENCE: But literal shape versus for the texture shape.

RAMESH RASKAR: The terrain.

AUDIENCE: All right.

RAMESH RASKAR:	That's what I mean by shape, not just the outline. But for example, here, it's the shadows. It's illuminated as if the light is at a very crazy angle. So the relief of the terrain is more clearly visible here. So you can see shadows here and so on. So we know that this particular structure is much higher than this particular area. And so the height field of that is enhanced.

On the other hand, this one shows all the texture in great detail. But it looks like a very flat leaf. So you can have knobs in your software which says, show me more of the detail or show me more of the shape. And you can get it off the internet.

And they came up with multiple methods. This particular method, again, put this 3D object, move it around, and create this particular-- so this is my older method, the one I showed for doing the day and night images. So they compared with that.

And their claim is that if they use my method, the shadows are preserved, which is true. And in their methods, they can create more beautiful combinations by using multiscale decomposition of bilateral filter. And we'll talk about that at some other point. So this is very unique. As a photographer, you will never think of setting up the lights in such a way so that in post-capture, you can decide whether you want to highlight the shape or the detail.

So that was, so far, light position we were changing. But there are lots of other parameters we can change for light. So let's look at this project called dual photography, one of my favorite projects. And we also saw it as a teaser, where the experiment they want to do is read this card from this camera. Although, it's facing away from the camera.

So what they're going to do is place a projector in the line of sight and some reflective surface. If you shine light from this projector on just one spot, it's going to bounce around on the book and eventually arrive at the camera. And by shining one light at a time, this is what the camera will see directly.

So again, this card is facing away from the camera. You're going to shine different pixels on this card. And you do this a million times because you have a million pixels in the projector. So if you're taking a million photos, this is what you'll be able to see from the camera. So this is a dual photo.

AUDIENCE:	What does the [? plan it look ?] like?

RAMESH RASKAR:	[INAUDIBLE] looks like this.

AUDIENCE:	Oh, OK. So that's it.

RAMESH RASKAR:	Yeah, any one of them, all the photos look almost like this. Except when you shine the red part of the card here, this book will look a little bit reddish because it'll be cast. And when you shine the yellowish part, the projector shines the yellowish part on the card, there will be a yellowish glow. That's it. And from that, you can figure out how this works.

AUDIENCE:	It still requires line of sight.

RAMESH RASKAR:	Sorry?

AUDIENCE: It still requires line of sight of the object, right?

RAMESH RASKAR: The projector has to be line of sight.

[INTERPOSING VOICES]

AUDIENCE: So you had to put something in the line of sight.

RAMESH RASKAR: Yeah, but it's cute. I agree. It may or may not be practical. But it's very interesting. There was a very-- a project that received a lot of press about five years ago, where some guy came up with a technique to figure out what you're reading on your monitor by just looking at the window.

So imagine it's a nighttime scene. And you are in the room working on your monitor. And if you stand outside-- somebody stands outside the window, they can see the glow from your monitor in the room. Now the question is, can they figure out what's on the monitor by just looking at the glow in the room?

AUDIENCE: Yes.

RAMESH RASKAR: How?

AUDIENCE: Because they can synchronize with the scan of the monitor.

RAMESH RASKAR: Exactly. So it only works for LCD. But if it's a CRT monitor, only one pixel is being eliminated at a time. So if you have a photodetector that's aimed at the window, and it's running at the speed of the CRT, and you synchronize it, you'll be able to actually read what's on the CRT monitor. So for TVs, at least old TVs, it's running only at about 30 Hertz or 60 Hertz. So it's really easy to figure out what's-- if the image is sufficiently high-contrast, you can figure out what was on the TV or the monitor.

AUDIENCE: This is how they try [INAUDIBLE].

RAMESH RASKAR: No.

AUDIENCE: So if that light from the electromagnetic radiation [INAUDIBLE] component on your screen is tracked by an optical, then you can deduce what [INAUDIBLE] is only for two devices emits any radiation. So that [INAUDIBLE]. license. So it [INAUDIBLE].

RAMESH RASKAR: Yeah, but they're looking at RF. They're not looking at visual spectrum.

AUDIENCE: Oh.

RAMESH RASKAR: Right?

AUDIENCE: [? One is ?] spectrum RF.

RAMESH RASKAR:	Yeah, RF spectrum.
AUDIENCE:	Yeah, and actually, a while ago, there was Tempest Proof Fonts that were supposedly able to mask your online activities from Tempest spying.
AUDIENCE:	Yes.
AUDIENCE:	Yeah, I don't know how good those work, though. There's new methods also to directly read memory off of people's computers. So you don't even need to think about what they might be running. Just go err. I was like [INAUDIBLE].
RAMESH RASKAR:	Even disk memory or RAM?
AUDIENCE:	No, RAM. Wireless.
RAMESH RASKAR:	Based on the--
AUDIENCE:	EM.
RAMESH RASKAR:	Based on electromagnetic.
AUDIENCE:	Yeah.
RAMESH RASKAR:	So there's a lot more coming.
AUDIENCE:	Yeah.
RAMESH RASKAR:	Does-- at Google Earth live, you're worried about only what's in the line of sight. There's much more.

[LAUGHTER]

All right. So how does this work? So again, you have a camera. You have a [? projector. ?] You're going to take-- this is how the camera looks at it, and this is how the projector looks at it. But eventually, you'll be able to compute the image that looks like the one on the right. And all the [INAUDIBLE] taken shadows, and refraction, and reflections, they're all captured that as well.

All right. So I'm going to go through very quickly because we covered it. But here's the point, that you're going to turn on one point and some light will reach the eye. And that's just a regular photo in the primary domain. Now it should replace the light and pi. And this is the Helmholtz reciprocity for in the optical domain.

If you replace the two and shine the same spot that I was looking at, you will receive the same exact intensity. And it takes into consideration the bidirectional reflection distribution function, the R-squared falloff, and all those things. So the duality is extremely powerful. And it can be exploited in many ways.

So this is how a barcode scanner works. When you go to the aisle, you have a scanner that's just scanning a laser stripe across the barcode. And there's just a single photodetector.

You hit a white spot on the barcode, light gets dispersed. And you can detect it. If it hits a black spot, then you don't see it. Now instead of using a photo sensor-- sorry, let me go back. So this way, by scanning light and using a photodetector, what we have created is basically a situation where you have a camera and an omnidirectional light source like a bulb. So it gets a barcode scanner.

It's as if the barcode was lit by just a flashlight and you took a picture with the camera, except the barcode scanner can check out eyes. It does not have any camera inside. It just has a laser scanner and a photodetector.

So that's the simplest version of dual photography, that you can record light by the so-called flying spot principle. You eliminate the spot. And you see how much light was reflected from it in aggregate. So that's the basic reciprocity.

Now, so let's look at the math for this very briefly. And it's very straightforward. So just follow me. One step at a time, So let's say I turn on the first pixel of the projector and record the intensity of light.

There is-- can't get my [INAUDIBLE]. I record the intensity for first pixel. Then I turned on the second pixel. I record the intensity per pixel and so on.

Now, if you replace the photo sensor with above and the projector with the camera, the claim is that you will see the same exact intensities. So with the light, I'm going to just floodlit the scene directly. And the claim is that the first pixel will receive the same light that the photodetector would have received over here and so on. Is that clear?

AUDIENCE:	But you said, for example, the first pixel is getting reflections from all over the rabbit? Yeah.
RAMESH RASKAR:	That's the point, that it will reflect all over the place. But if you look at that point directly, because of the duality of light, you can replace your eye and the light source. And then, you see the same exact intensity in that particular direction.
	In case of the analogy I gave you of being able to see what's on your monitor by looking out the window, you're just-- the light from your monitor lights up the whole room. And some of the light leaks through the window. But it doesn't really matter. There's some proportion of light that leaks through the window. And when the next pixel turns on, the same proportion of that light leaks through the window.
AUDIENCE:	Yeah, I get that, because every pixel is being turned on and off in order.
RAMESH RASKAR:	Right.
AUDIENCE:	But here, all the pixels are being [INAUDIBLE] at the same time.
RAMESH RASKAR:	No, in the experiment, only one pixel is being turned on at a time.
AUDIENCE:	But the light bulb is showing-- it's also flickering the same way the projector is--

[INTERPOSING VOICES]

RAMESH RASKAR: No, it's not. It's not. That's a good point. So here, you have to understand that only one of them is turned on at a time. And we are recording it on the photodetector. Now when you have the light source, the point you're making is that light will reflect-- light will be coming from here. It will bounce from here. And it will reach here. But light could also bounce from somewhere else and go in the same direction.

[INTERPOSING VOICES]

RAMESH RASKAR: It will not because when you shoot away from this projector, it will hit only one point. And that particular point is lit by this light in only one direction. There is still just [INAUDIBLE] mapping between the two.

AUDIENCE: Pictures of light, it bounces at some spot. If it's a mirror, it bounces on that spot again.

RAMESH RASKAR: That's a very good question. So when you have inter-reflections, we'll see that in the second half, how we deal with it. So this is a very simple demonstration, where we have only one-to-one correspondence between the direction of the light and the direction of the pixel.

AUDIENCE: So you could do a camera just with the projector, and ADR, and the sensor?

RAMESH RASKAR: Yes, basically.

AUDIENCE: Because you have prior knowledge of the scene layer, imaging may dramatically reduce the number of--

RAMESH RASKAR: Measurements.

AUDIENCE: --measurements they need. For that card example, you didn't actually need to do a million to figure out what color it is and how they do it.

RAMESH RASKAR: How would you do it?

AUDIENCE: Figure out where the symbol is first.

AUDIENCE: Adjust the corner.

AUDIENCE: Find the corner by--

[INTERPOSING VOICES]

RAMESH RASKAR: So you're saying that--

AUDIENCE: --figure out

[INTERPOSING VOICES]

RAMESH RASKAR: You're saying the white part of the card, I don't need to illuminate that at all because I know it's only white. Is that what--

AUDIENCE:	Sorry?
RAMESH RASKAR:	The white part of the card, playing card, you're saying, I don't need to shine those parts because there is no information there. Is that what you're saying?
AUDIENCE:	You don't have to shoot the whole card. Hell, you can just shoot the little tiny corner where the K or the Q or the seven is. And then, that's all you need. You can--
RAMESH RASKAR:	But then, you capture only the K. You'll not capture the rest of the card.
AUDIENCE:	Well, that's all you need to--
	[INTERPOSING VOICES]
AUDIENCE:	Yeah, but to get--
	[INTERPOSING VOICES]
RAMESH RASKAR:	All right. OK.
	[INTERPOSING VOICES]
RAMESH RASKAR:	Now you're thinking about actually playing bridge or something.
AUDIENCE:	Yeah, right.
	[LAUGHTER]
	Irrelevant information is--
RAMESH RASKAR:	Yeah, it could be. But you don't know where it is.
AUDIENCE:	Well, if you use the projector to first scan one horizontal and one vertical, you get enough information to figure out where the--
RAMESH RASKAR:	Where there's some--
AUDIENCE:	--information is.
RAMESH RASKAR:	Certainly. You can come up with some optimizations. If you really want to--
	[INTERPOSING VOICES]
RAMESH RASKAR:	If you really want to read up on those cards, if that's all you want to do, I'm sure you can come up with interesting things.

AUDIENCE: But for other applications, where you know what you're looking at, but you're looking for some very specific thing.

RAMESH RASKAR: Exactly.

[INTERPOSING VOICES]

RAMESH RASKAR: So you could do some binary search.

AUDIENCE: Right.

RAMESH RASKAR: So for example, when JB asked, can I just use a projector and this photosensor to look at something, if all I want to do is locate some retroreflective dot, then, yes, you might be able to just shine a whole vertical line, sweep across, and then another line and sweep across. And that will quickly give you the x and y-coordinate of that spot.

So yeah, you could come up with quick-- some variations in the [? scan. ?] All right? But this is a standard technique that we have seen many times. We're going to go beyond that now. But is this part clear now, that we're going to measure the intensities and they'll do the same in both directions?

Now, we're going to go a little bit further. All right? So this is what they get. This is the images they get by using this mechanism. All right? Not great quality, but reasonable. There's a bunny there. All right. And the flying spot camera, that's exactly how it works. You have a spot that's moving around. And that's how scanning electron microscopes work. It's not a high-quality camera, but a high-quality light source that's very focused.

So you just have a sample. But you just eliminate the sample at only one location. And then, you have a detector that integrates over the whole sample. And so you can do a very good job of that. Now instead of-- for the sensor, we're going to make it a little bit more complicated and replace that with a camera.

So what you're going to see throughout this class today is how we can start thinking in higher dimensions and realizing that the appearance of the world is not just flat 2D. But actually, it's much higher dimensional. In this particular case, how many dimensions do we have? When we had a photodetector, the photodetector is zero-dimensional. And the projector is two-dimensional. So the measurement was two-dimensional. And we got an image out. How many dimensions do we have here?

AUDIENCE: Four.

AUDIENCE: Four.

RAMESH RASKAR: Four. We have two photoprojectors-- x, y, or the frame for the projector, and two for the camera, the UV.

[INTERPOSING VOICES]

RAMESH RASKAR: Sorry?

AUDIENCE: What's the time dimension?

RAMESH RASKAR: There's time dimension as well. But right now, everything is static. And things are not changing over time. You could say there's wavelength and so on. But we are ignoring that as well. We're just assuming it's just [INAUDIBLE]. So now, we're going to think in four dimensions. All right?

So now, what we're going to do is exactly the same thing. We're going to turn on-- we're going to send some coordinate system, PQ for predictor, NM for camera. And we will realize that just the same way we swapped the eye and the light source, we'll be able to swap the camera and the projector.

Coming back to [INAUDIBLE], same trick. We're going to turn on one particular pixel of the projector and [? re-word ?] it, one particular pixel of the camera. Except because it's the camera, I can read all the NM pixels simultaneously in one snapshot. I don't have to read them one at a time. But the projector, I can only turn on one PQ pixel at a time.

All right? So that's a 4D mapping. I have four functions, four parameters, PQ, NM. And for every PQ, NM, I have one value. It's a function of four values. All right? And then to capture that, it's relatively straightforward. I'm going to turn on one pixel at a time of the projector. And I'm going to record one image, m times n.

And this is a very standard method in linear algebra, where instead of representing it as a 2D, you represent it as a single vector. So let's say-- let's take some numbers here so it becomes easier to remember. Let's imagine the projector is 1,000 pixels by 1,000 pixels. And the camera is, again, so let's say it's 2,000 pixels by 2,000 pixels.

So this particular vector is now going to be 1,000 times 1,000 times one. So you'll have one million entries corresponding to each pixel of the projector. The image that you get is not 2,000, but 2,000. It's four million. So four million entries in this particular vector. Good so far?

Now we're going to come up with a relationship that maps these values to these values because, remember, when I turned on one pixel here, it made contributions to all these 4 million pixels out here. And we're going to try and represent that mathematically. And it's really, really straightforward.

We'll just go slow, one at a time. And that transformation we're going to specify as this matrix team, which is now, what's the dimensionality here? M times N is--

AUDIENCE: 4 million.

RAMESH RASKAR: 4 million. And P times Q is?

AUDIENCE: Million.

RAMESH RASKAR: 1 million. This is a huge matrix. 4 million by 1 million. So this is already in terabytes. Sorry, gigabytes [INAUDIBLE]. So 4 terabytes. All right.

Now this is how we're going to build this matrix. We're going to fill up the values in this matrix. Remember, every time I turn on a projector pixel, I get 4 million values. Next time I turn on projector pixel, I get 4 million values in the camera. So 1 million times 4 million is what we would need.

So very simple. I'm going to turn on the first pixel of the projector and see what image I get. And that value will go here. All right? So the 4 million values will simply go here. And that will be measured. That will be captured. So again, turn on the first pixel of the projector. See what image I get, which will be these 4 million values. And those values are simply the first column of this transform projector.

Then I turn on the next pixel of the projector, record the image in the camera. And that becomes the second column of this transformation matrix and so on. Now, it's very easy to see what's going on here. If I turn on both pixels, the first pixel as well as the second pixel, what image should I get?

AUDIENCE: I have two.

AUDIENCE: Sum of both the columns.

RAMESH RASKAR: The sum of the first two images I got. So again, if I turn on the first pixel, I took a photo. I turn on the second pixel, I got the second photo. If I turn on both pixels at the same time, I will simply get the sum of those two photos. And that's what is being represented here mathematically. So if I put a one here and a one here, then the matrix multiplication will say, these two values multiplied by one and one here, the sum of that goes in the first value.

And again, the same two values in second row [INAUDIBLE] is different two values in this row multiplied by, again, the one and one here for the second value here and so on. So we're going to set up this huge linear system by just probing the scene one pixel at a time. And so filling up this T-matrix is relatively straightforward.

I turn on the third pixel, take a photo. And that photo becomes my third column. Very simple. And you keep doing that. And you can build your T-matrix. Again, a lot of data. But that's research. All right.

[LAUGHTER]

Now let's say you have this wall or a simple screen. And you put a projector and a camera. What would happen? If I turn on one pixel of the projector, how many pixels are the camera will get a non-zero value?

AUDIENCE: One.

RAMESH RASKAR: Only one. If I turn on the next pixel of the projector, some other pixel of the camera will get the value. But only one pixel in the whole image will [INAUDIBLE] non-zero. So if you have a scene that's really, really simple, maybe just flat or a convex object, in T, every column will have only one value that's non-zero, or maybe a couple of values if the camera has a higher resolution. So the matrix T will actually look very sparse. Most of the values are zero. Only some values, almost along the diagonal, will have non-zero values.

On the other hand, if you have a scene with lots of inter-reflections, if you're looking at a corner of a room or there's a glass bottle and all this complexity, when I turn on one pixel of the projector, many pixels in the camera actually get the intensity. So first because of the bottle. The next because of the screen. The third because of inter-reflections in the screen. And so multiple pixels in the camera will be lit.

And so the matrix T will be very nice because, remember, every column there, more than one entry will be non-zero. So that gives us some-- if I just look at the matrix T, if somebody just shows you the matrix T, the gigabytes or terabytes of data in just a visual representation, I can tell you from what, kind of, scene it came up. So just to make sure it's clear.

So the projector and the camera is a transform matrix. But not the first pixel. The photo that appears because of the first pixel, I'm just going to put it here. If the scene was very simple, we had to have a projector, and a camera, and I turned on one pixel, and it matched only one camera, then only one of the values will become zero. The rest of the values will be all zeroes.

And in fact, on the next pixel projector, only one of the camera pixels will be lit and so on. So maybe you'll have a value here, a value here, a value here, value here, maybe a value here, and so on. Every column will have only one value non-zero. Everything else is just zero. And that's a sparse matrix.

On the other hand, you've got a scene that has a lot of complexities. There's a glass bottle here. There's an inter-reflection here and so on. But you turn on one pixel of the projector, it reflects on here. It reflects on the back of the bottle. And some other pixel, it hits this wall. And some other pixel, light reflects off [INAUDIBLE] pixel. So maybe 10 or 15 different pixels of the camera [INAUDIBLE].

So in this case, when I write the matrix, for the very first entry in the first column, multiple values will be non-zero. And when I turn on the second pixel, similarly, a lot of value will be non-zero and so on. So this matrix will actually look pretty nice. A lot of values are non-zero. So by just visually looking at the structure of this matrix, I know how complicated this is. This is very useful. And this will keep coming up as we go along. Is this clear?

So this is what we call our primal space, which is done on one pixel of the projector. Take a photo. Now we want to create the dual space. That was our dual photography problem. How will the picture look like if I put a camera where the projector is? And we're going to solve that by doing very simple operations on this matrix.

So again, we spent all this time taking a million photos by turning on one projector pixel at a time. Now is the time to see the money. This is the problem I want to solve. If I put a camera where the projector is, how does it look? So I need to come up with a transformation, T, a modified transformation that will map my camera pixels to projector pixels. Right now, we're mapping projector pixels to camera pixels. That's the question we have to answer.

And let's call it k double prime. And it's dimensionality is now-- the upper one was one million by-- sorry, four million by one million. The bottom one is one million by four million. That already gives you a clue of what you could do. So if I turn on pixel j of the projector, pixel r of the projector gets a certain value. And it corresponds to this column and this row. Now, let's call it, well, pixel [? back. ?]

And now, when I put the camera at the projector, I want to say, which pixel of the camera-- when I look at a pixel i of a camera, which projector pixels contribute to that? So I took these million pictures. Only when some of the project pixels were turned on, I got a non-zero value at this point [INAUDIBLE] pixel. And that's also a very similar structure.

So the hint for that is the matrix T double prime is simply a transpose of the matrix T. You just flip it along its-- you just split the horizontal and vertical coordinates of the matrix. And that gives you the T double prime. As simple as that. And from that, we can compute an image that appears as if the camera was placed at the projector.

So here are some examples. So here's a picture taken by when the projector was floodlit. This is how the camera looks at it. We're going to probe it multiple times. And we compute an image as if the camera was placed at the projector. In this case, it's still all diffused, only [INAUDIBLE] mapping. There's no inter-reflection. There's no class and so on. So it's really easy to compute.

Here are more challenging examples. So where is the projector in this case? It's a little bit to the left. You can see all the shadows going back very deep. And now, we're going to create a grid from this photo as if the camera was placed at the projector. Will we see shadows when we create that picture? What will we see?

AUDIENCE: Well, you won't [INAUDIBLE] shadow.

RAMESH RASKAR: You won't see the shadows because now your camera is placed where the projector was. But you'll see something else.

[INTERPOSING VOICES]

AUDIENCE: --shadow on the horse.

RAMESH RASKAR: Sorry?

AUDIENCE: You can see that shadow on the horse.

RAMESH RASKAR: Exactly.

AUDIENCE: [INAUDIBLE]

RAMESH RASKAR: So why do we have a shadow on the horse?

[INTERPOSING VOICES]

RAMESH RASKAR: So now, although the horse was lit where the projector was, we swapped the projector with the camera. And from the camera's viewpoint, the horse was actually occluded. So when I placed-- when I saw the two, now I don't see the light on those [INAUDIBLE] see shadow.

AUDIENCE: So are those real? Or are they simulated?

RAMESH RASKAR: This is real. This is-- yeah, they spent days and days capture--

[INTERPOSING VOICES]

RAMESH RASKAR: Days and days capturing this data.

AUDIENCE: So this is in all the shadows.

RAMESH RASKAR: Yeah, you can think of it as an occlusion. What is occlusion?

AUDIENCE: The [INAUDIBLE].

RAMESH RASKAR: How do you know--

[INTERPOSING VOICES]

RAMESH RASKAR: How do you define occlusion?

AUDIENCE: You [INAUDIBLE].

RAMESH RASKAR: How would you define it more precisely?

AUDIENCE: Light can't reach it.

No light.

RAMESH RASKAR: Light cannot reach. So if I shoot a ray, we're going to talk about rays a lot in the second half, when you shoot a ray, a ray is occluded. And it stops the path of the ray. And that's occlusion. So when we think about light, it's very obvious because the ray starts from the light source, and it's occluded, and there's a shadow on the other side.

But even for cameras, there's occlusion. If I'm looking at this monitor, I cannot see what's on the other side. So this is the inverse ray because I'm shooting the rays from my eye. In fact, the Greeks used to think that the way we look at the world is we shoot the rays from the eye. And they were right.

[LAUGHTER]

Because the reality of light, you can just shoot the rays from your eye. And it's the same math after that.

AUDIENCE: So is it like in computer in the rendering graphics, [? raycasting. ?]

AUDIENCE: Raycasting, [? feature ?] install. And the question here is that, was there totally dark before they illuminate all. [INAUDIBLE]?

RAMESH RASKAR: No,

[INTERPOSING VOICES]

RAMESH RASKAR: There are certain.

AUDIENCE: --shown everything.

RAMESH RASKAR:	They're going to [? live with ?] it.
AUDIENCE:	So we're going use this in a room like this with standardizing unless we [INAUDIBLE].
RAMESH RASKAR:	No, no, no.
	[INTERPOSING VOICES]
RAMESH RASKAR:	Whatever additional term you have, it will remain. Or you can take two photos, one with the light source, one without the light source, and subtract it away. So you can do all those things.
	But you're right, that I cannot do this in broad daylight because the light-- it's like taking a flash photo in broad daylight. The flash doesn't make any difference because it can't override the sunlight. Or it doesn't make a significant contribution about the sunlight. It's the same thing here.
AUDIENCE:	So obviously, talking-- how how easy it is-- is it to get that [INAUDIBLE]?
RAMESH RASKAR:	If you have enough memory.
	[LAUGHTER]
	[INTERPOSING VOICES]
AUDIENCE:	[INAUDIBLE].
AUDIENCE:	You can do it in MATLAB. But It's [INAUDIBLE].
RAMESH RASKAR:	So it's not that easy to take the-- if you have a small matrix, taking a transpose, it's very simple. Just M squared. Just take the value and just change its x and y-coordinates. Make x y and y x. It's very simple. Even a million by-- 1,000 by 1,000 image, 1,000 by 1,000 matrix, no big deal. But eventually, you'll run out of memory. If you load the whole matrix and try to invert it, it's going to be a challenge.
	Just think about taking a sheet of paper. And all I want to do is do its transpose. If the room is limited size, I can take bigger and bigger sheets and do the transpose. At some point, if you give me a really large sheet, I won't be able to do the transpose. Although, the task is very easy.
AUDIENCE:	You don't even need to do that. You just omitted that issue?
RAMESH RASKAR:	Sorry?
AUDIENCE:	All you need to make the images you have are each of the columns. So can't you just--
RAMESH RASKAR:	So you can do it mathematically.
	[INTERPOSING VOICES]

RAMESH RASKAR:	But you may not be able to even load the whole thing in memory.
AUDIENCE:	[INAUDIBLE].
AUDIENCE:	But you don't need to. You just access--
RAMESH RASKAR:	Exactly.
	[INTERPOSING VOICES]
RAMESH RASKAR:	Exactly. So you just do this--
AUDIENCE:	--forget about it.
RAMESH RASKAR:	Exactly. So you won't be able to do this in real-time manner. But if you do some disk access, then you'll be able to do it. And in the paper, they describe all these challenges. And if there's time, I'll just show you very briefly what approximations they came up with. I'll ask them to load this quickly.
AUDIENCE:	Very briefly, does it exist, projector and camera in the same package? Then you can switch between shining light through and [INAUDIBLE]? Because It's always different in [INAUDIBLE] projector and the camera.
RAMESH RASKAR:	Right.
AUDIENCE:	But now that they are friends--
	[LAUGHTER]
	--they're-- and you could easily see through, probably seeing things, and--
RAMESH RASKAR:	Right.
AUDIENCE:	But does it exist in one package or [INAUDIBLE] one?
RAMESH RASKAR:	Just to clarify, in this particular case, it doesn't matter where the projector is and where the camera is.
AUDIENCE:	Sure.
RAMESH RASKAR:	You can just put them arbitrarily.
AUDIENCE:	Yes.
RAMESH RASKAR:	But they can't be too far away. But as long as one can see the impact of that. Now are you saying that they should be optically coaxial?

AUDIENCE:	[INAUDIBLE] it could be interesting for [INAUDIBLE].
RAMESH RASKAR:	Certainly. My thesis, by the way, my PhD thesis was projector as a dual of camera.
AUDIENCE:	Can you build one?
RAMESH RASKAR:	I built many versions. But actually, in Hershel's group, there's a concept called I/O bulb. And then--
AUDIENCE:	And this concept [INAUDIBLE] the work. And it was more asking of a physical device [INAUDIBLE].
RAMESH RASKAR:	But I remember John and others actually built projectors and cameras that are coaxial and for interaction and so on. Other people have built it as well. But they show a lot of interaction on top of that. So the concept of putting them together for some specific task is well-known. But what this group showed was that mathematically, you can create these magical pictures. [INTERPOSING VOICES]
AUDIENCE:	Did you get a resolution to any?
RAMESH RASKAR:	Of course. So that's a great point. So you might start with, I don't know, a 10-megapixel camera. But if your projector is only a megapixel, then your dual image is only going to be megapixel. And there's going to be a lot of aliasing issues and so on. So that's a very good point. And the projector's color properties are going to impact how you capture the image as well. It's not purely dependent on the camera quality now, camera sensor quality now. It's also dependent on the projector's illumination quality.
AUDIENCE:	Is it a real challenge [INAUDIBLE]? Did that [INAUDIBLE] of that?
RAMESH RASKAR:	I didn't hear the last part.
AUDIENCE:	Real technical challenge.
RAMESH RASKAR:	Right.
AUDIENCE:	Did the author of this paper ever say, [INAUDIBLE] light could be useful for this-- [INTERPOSING VOICES]
RAMESH RASKAR:	Yes, I'll show you their motivation. But as you can imagine, this can be used in many other ways. It's not always practical because you'll take a million images.
AUDIENCE:	Yes. [LAUGHTER]
RAMESH RASKAR:	But at the same time, I'll show you the reason where it does give a lot of benefit. Any other questions? All right.

AUDIENCE: You don't always have to take a million. Don't they share how to do the--

RAMESH RASKAR: Subdivision.

AUDIENCE: Yeah, so subdivision?

RAMESH RASKAR: Yeah, I'll show that very briefly. It's a second-order effect. And I don't want to go too much into detail. And now, you can do really complicated scenes such as some global illumination. So the [? crossticks ?] and so on is all natural. And then, we can start doing some special effects. So remember, the picture on the left was taken by the camera. The picture on the right was computed as if the camera was in the projector.

On top of that, now that we have swapped the projector and the camera, I can convert the projector into a slide projector mathematically and see how the image looks. So if you go back to this matrix, I can just put all ones here, that will be flat field of the projector, and see how the image looks.

And similarly, I can put all ones here to see what happens when I switch the camera and the projector. But I don't have to put all ones. I can put 10101010. That means as if I put a slide projector where every pixel was on and off. Every argument pixel was on and off. So you can create slide projectors to create those effects. So you can create interesting effects like this.

AUDIENCE: But you can't do that backwards on the image?

RAMESH RASKAR: You can do both. Yes, of course. In the primary domain, it's very easy because I can turn on one pixel of the projector, take a photo. Next pixel, I don't turn on. The third pixel, I turn on.

AUDIENCE: So you [INAUDIBLE].

RAMESH RASKAR: And I can just take addition of those 1/2 a million photos. And I'll automatically get this effect. So in the primary domain, it's very easy to do.

AUDIENCE: I think in the parallel domain, it's even easier. You can just literally project that slide.

RAMESH RASKAR: Exactly.

[LAUGHTER]

Exactly. But the point is that if you had collected this terabyte of data from that, you don't have to think in advance which slide you want to project.

AUDIENCE: [INAUDIBLE]

RAMESH RASKAR: You can change that slide in software.

AUDIENCE: That's good.

RAMESH RASKAR:	You understand?
AUDIENCE:	Yeah. I--
	[INTERPOSING VOICES]
RAMESH RASKAR:	There's the same effect as we saw earlier, where the woman was in a dorm.
AUDIENCE:	Yeah.
RAMESH RASKAR:	In the Milan versus LA. And post-capture, I can decide how she should look for a light. All right. So let me skip over the rest of the slides, it will be available for you, and go to the next part of-- the motivation for that was this. And again, thinking in even higher dimensions. So we just realized that the lighting and appearance-- lighting is two-dimensional. The photo is two-dimensional. So it's already four-dimensional.
	But we're going to go now in much higher dimensions. So instead of putting just one projector, imagine if we started putting multiple projectors. So you start turning on one pixel of one projector of the given projector. Now I have four million pixels. A million pixels in each of the projectors.
	So this adds an additional two degrees of freedom because if I have a projector, I can place it anywhere in 2D. So you can just think about the hemisphere. Any azimuth and any elevation, I can put this project. And every projector has a buffer with x, y. So the illumination itself now becomes four-dimensional, two degrees of freedom for the position and two degrees of freedom for the pixel coordinate. And the camera is still two-dimensional. So how many dimensions do we have now?
AUDIENCE:	10.
RAMESH RASKAR:	Sorry? No?
AUDIENCE:	[INAUDIBLE].
AUDIENCE:	Yes.
RAMESH RASKAR:	No? We have two for position of the projector, two for coordinate of the pixel of the projector, and two for the camera.
	[INTERPOSING VOICES]
RAMESH RASKAR:	It's six. So we're increasing our dimensionality here. And the problem is getting more and more crazy. So imagine if you want to do this. How will you do it? It will take four million steps. There are four million photos. But there is a shortcut based on what you just saw. How can you reduce the total number of pictures? So this is thinking in primary domain.
AUDIENCE:	Duality.

RAMESH RASKAR:	Duality. How will you exploit that?
AUDIENCE:	If you switch the-- instead of the camera and the projector, switch [INAUDIBLE].
RAMESH RASKAR:	Exactly. And the benefit--
	[INTERPOSING VOICES]
AUDIENCE:	What was that? What was that?
RAMESH RASKAR:	So now, if you switch the cam-- I think someone [INAUDIBLE]. So that's 4D. That's 2D. So the total is 6D. And other people have done things like this, where they have a camera, and they have a projector, and they move the object and capture the 6D.
	So people have built these rays all the way back in 2003 in [INAUDIBLE]. Now, again, it will take 4 million pictures if you're to do that. [INAUDIBLE] and so on. All right? But now, if you switch them, what's the benefit?
AUDIENCE:	[INAUDIBLE]
	[INTERPOSING VOICES]
RAMESH RASKAR:	Because all the cameras can work in parallel. All the projectors could not work in parallel. Only one projector could be on at a time. But all the cameras can take the picture at the same time. So now, all we have to do is turn on one pixel of the projector. And all the cameras take the snapshot at the same time.
	So instead of taking four million pictures, I can just take one million. And I can put as many cameras as I want. And by switching the role of the projector and the camera, I reduce the total number of samples I have to capture, total number of time slots [INAUDIBLE]. Total number of samples remain the same. Kevin is not happy.
AUDIENCE:	Well, I have a question. But I think it's going to be answered on the next slide if I remember correctly.
	[LAUGHTER]
RAMESH RASKAR:	All right. So we'll do that. So it's just 6D. So this is how they did it. Instead of creating 16 cameras, they just took one camera, 16 megapixel, and they just aimed it at these mirrors.
	So that becomes a virtual camera. And then, one projector. And from this, they're going to create an illusion that they have 16 projectors and one camera. So results, hopefully? So card experiment, which-- so do they need a camera for the card experiment?
AUDIENCE:	No.
RAMESH RASKAR:	What do they need?
AUDIENCE:	Light or a color-sensitive photodetector.
	[INTERPOSING VOICES]

RAMESH RASKAR: They just needed one photodetector.

[INTERPOSING VOICES]

RAMESH RASKAR: They could have done it. But for some, that is the [INAUDIBLE]. [BANGING]

The vending machines. All right. So there's some other experiments that show how they use that multiplanar mirror [INAUDIBLE]. Any questions on that one? Yeah?

AUDIENCE: Maybe I just forgot this. But so if you have a 4D projector set up, they can actually project where the camera might not be able to see? And maybe that would reflect onto the ground. And then, the camera could see it. But if you flip them, then the projector can't illuminate occluded stuff.

RAMESH RASKAR: That's the same example in the playing card.

AUDIENCE: Yeah, it is.

RAMESH RASKAR: So it still works. So do you understand this question?

AUDIENCE: Yeah.

RAMESH RASKAR: When we're writing this whole thing down, we had this implicit assumption that when the projector and the slide-- the camera can see it. And if it doesn't see it, we're going to ignore it. But that's not the case. Even if the camera cannot see it, the example we have there is-- we have a card and-- let's see. How does this [INAUDIBLE]? All right. Let's go back to this.

We have a card here. And you have a projector. It shines light. And then, you have a camera. And none of the cameras can be reduced to these pixels because it's on the other side of the playing card.

But it's going to see the reflection that comes from those areas here. So you're still going to see that yellow or red glow on the book on the scan. So remember, you don't have to see-- the camera doesn't have to see that pixel directly, that [? flash ?] directly, to be able to compute how to look when you switch the camera in [INAUDIBLE].

AUDIENCE: But just as in the example where the horse was occluded by the emblem, there can be cases where none of the rays actually reach the [INAUDIBLE].

RAMESH RASKAR: It's possible, yeah. It's possible, yeah.

AUDIENCE: And is it possible to send the ray from the projector in an incident, something, something that would reflect the ray [INAUDIBLE]?

RAMESH RASKAR: That's fine. If you put a mirror here--

AUDIENCE: The mirror is OK. But not the defusing surface because--

RAMESH RASKAR: No, [INAUDIBLE] question.

AUDIENCE: Because there is-- I understand it in one way. But there's a way to be complex on it. So if you shine light not directly on the card, but on the book, would it still work?

RAMESH RASKAR: Right. Here, that-- let's take that example. So I have this. I have the card.

AUDIENCE: Yeah.

RAMESH RASKAR: I have the projector. I just shine the book.

AUDIENCE: Yeah.

RAMESH RASKAR: And let's say that thing is still not visible from the camera.

AUDIENCE: Exactly.

RAMESH RASKAR: What will I-- I will read the book. I will not read the card.

AUDIENCE: Yeah, oh, OK.

[INTERPOSING VOICES]

RAMESH RASKAR: Yeah, they say the book has a bluish card shape, all filled with blue. And I'll say, oh, yeah, it must have been [INAUDIBLE].

AUDIENCE: In theory, you should-- in theory, you might be able to still--

RAMESH RASKAR: Leave the card?

AUDIENCE: Yes.

[INTERPOSING VOICES]

AUDIENCE: And a blue light [INAUDIBLE].

It's an--

[INTERPOSING VOICES]

AUDIENCE: --then you have a very diffused light source lighting up the cards from different spots on the book, effectively, and reflecting back--

AUDIENCE: Yeah.

AUDIENCE: All right, somewhere in the matrix?

AUDIENCE: [INAUDIBLE]

AUDIENCE: Oh, I was going to ask if-- is there any clever way to compute the answer? Instead of having to take all, a million, like take some pattern?

RAMESH RASKAR: Exactly. So you should know--

RAMESH RASKAR: --that certain parts of the scene--

AUDIENCE: Or even without any knowledge of the scene.

RAMESH RASKAR: Yeah, you can do some probing. I can shine a light at one part and see, do I get light from here? I don't. So maybe that [INAUDIBLE] the future, I will shine both my lights there. So you can definitely do those probings and quickly figure out how to.

AUDIENCE: Yeah, in the paper, they even have the algorithm for it. It's only this long.

RAMESH RASKAR: Yeah, by the way, so we-- our next assignment, assignment number three, dual photography that's one of the options. So you can either take a million photos or you can use the shot algorithm to reduce the number of samples. But it still takes a long time.

[LAUGHTER]

You've got to run it from-- you'll have to leave it running overnight and come back in the morning. All right. So to me, a lot of projects in conversion photography are just sheer magic, being able to see a playing card that you can't really see from a camera. The next one, which I was involved in, is also a lot of fun called visual chatter. So the concept where the two batches are talking to each other, we call it just visual chatter.

If they can't see each other, there's no chatter. They can see each other, there's chatter. And this concept is actually very unique. By the way, when I meet photographers or professional photographers who are technically inclined and they say, what do you mean by computational photography? And I give them all the buzzwords and big definitions. And they don't get it.

I try to explain to them dual photography. And they immediately get it. You should. They say, wow, OK, with all the techniques I have in my bag, all the polarization filters, and all the flashlights, and all the umbrellas, and all the fancy lenses, I can't think of a way to create those, kind of, photos. And the next one you'll see is similar. So giving these concrete examples, at least for me, it has been helpful to communicate to photographers, what's different here with computational cameras and photography?

So now, we're going to look at a little bit more about how light bounces around in the world. And we're going to think about direct illumination and global illumination, which is unfortunately wrong terminology because what we're really talking about is direct reflection and global reflection. But because of some reasons I won't get into, people call it direct and global illumination.

So direct reflection is straightforward. The global reflection is when light bounces around and reaches the rotation. So the path a is just direct. Path b is reflect, reflect, reflect. So if I shine the spotlight here, that's direct. If I look at this shadow region, that's indirect.

The third one that's interesting is subsurface. Light actually enters the surface, it bounces around, and then it comes out at some point. And marble or skim is a good example. This is subsurface scattering. There are some other ones such as participating mediums. If you have fog or water, then light is going to scatter around and then come and you can look at that.

And then, finally, you have transmissive or translucent. Light is going to transmit to something and then come back to you. In this particular case, light actually is coming through this curtain through [INAUDIBLE] without diffusion and [INAUDIBLE] scattering. Now what we're going to do is distinguish between direct and everything else.

And I showed this example in the beginning of the class, where you have somebody behind a shower curtain. You can find out what the direct bounce is. And unfortunately, it's too dark. But you can only see the face here. And that's the indirect because of transmission. This is one of the options in the third assignment.

How does it work? Very simple idea. Very, very simple idea. You turn on all the pixels of the projector, take a photo. That's easy to capture the direct path. Now let's introduce a little bit of terminology. Keep it really simple. Before I got there, let me give you some intuition. And then, we'll go here. So imagine if I--

AUDIENCE: [INAUDIBLE]

**RAMESH
RASKAR:**   You'll just come on this side because you're more inclined to go this way.

[LAUGHTER]

So imagine if I have a [INAUDIBLE] and I take a laser pointer and shine a spot here, what's going to happen because of that is light will, again, bounce here. And that light can bounce here.

And when I look at the corner of this room, I will see a very bright red spot, but also see a red cast everywhere else. And so the red spot is because of direct. And everything else is because of indirect or global transport.

Now let's say I go from this patch i to the next patch j. So I sh-- I change the direction of the laser just a little bit. So now, I'm going to shine here. Here in the bright spot, I'll-- [INAUDIBLE] noticed that the bright spot has moved from i to j. What will happen over here? Will I see a change?

AUDIENCE: [INAUDIBLE]

AUDIENCE: Very [INAUDIBLE].

RAMESH RASKAR:	Very little change. Same here. There'll be hardly any change. What does it tell you?
AUDIENCE:	It tells you separate direct and global.
	[INTERPOSING VOICES]
RAMESH RASKAR:	We are going to separate direct and global.
AUDIENCE:	It tells you--
RAMESH RASKAR:	But this very simple experiment tells you that when the light is bounced indirectly, it's a very localized result. But when the light is bouncing or bouncing globally, it's a very low-frequency effect. And it doesn't really matter whether I shine here or shine there. So now, let's say I wanted to figure out from these two photos, photo when the light is shining at i and light is shining at j, which one is direct and which one is global? How will you do it?
AUDIENCE:	You would go what doesn't change.
RAMESH RASKAR:	You go for what doesn't change. That's it. The part that did not change is the [INAUDIBLE] part, the global part. And the part that changed is the direct part. That's it. From those two photos, I can tell you which is the direct component and global component. So let's look at it a little bit more precisely. So photo one and photo two.
	So photo one is direct plus some global. And the second one is some direct plus some global. Now we know that when I should change the light a little bit, the global did not change a lot. So I can just call this i1 and i2. If I subtract i1 minus i2, what will I get?
	These two i, I'll just get [INAUDIBLE]. And simply from that, I can figure out a photo has a distinction between d1 and d2. So I have two equations. And right now, I have three unknowns, d1, d2, and this [INAUDIBLE].
	Now I can take multiple photos. I can take a third photo. I can shine the laser slightly next position. I'll have d3 plus [INAUDIBLE] and so on.
	So I'll have n equations and my n plus 1 on this. For every equation, I introduce a new direct. And the global is the same. Of course, once I start shining somewhere else, if I go sufficiently far away from this, the global will change eventually. If not, you'll see the same. But it's a very simple observation.
	So going beyond that, let's do some real simple experiment. Instead of-- what we're interested in is not shining a laser at a time. But I'm going to put a bulb here. I'm going to turn on the flashlight. And I want to know when I take a photo, which part is direct? And which part is global? So I don't have this choice of doing one pixel or one direction at a time. How would you do that? That's the [? secret ?] [INAUDIBLE].
AUDIENCE:	What do you mean you can't answer it?
RAMESH RASKAR:	You can't answer it because you're at Columbia.
AUDIENCE:	Yeah.

[LAUGHTER]

RAMESH RASKAR: And I'll give you a hint. Instead of taking thousands of pictures, we're going to take exactly two pictures. We're going to replace this bulb with a projector.

I'm going to project one pattern, take a picture. I'm going to project a second pattern, take a picture. i1 and i2. The global still remain the same. And the direct will also be related. And from those two pictures, I'll be able to figure out what's direct and what's global.

AUDIENCE: Can you make the two directs inverse grids?

RAMESH RASKAR: Exactly. So in fact, somebody asked the same question. Remember I said in two patterns? There has to be some symmetry--

[INTERPOSING VOICES]

[LAUGHTER]

So I'm glad you're catching on that. So the answer is very simple. Did everybody get it? You didn't? There's more than one solution. But they all have the same principle, that they are inverse of each other.

So the simplest one would be to show a checkerboard pattern. And I'll show the inverse of the checkerboard pattern. So let's take that a little bit further. And if I shine now, instead of one laser, I'm going to shine multiple lasers because I'm going to turn on every alternate pixel of the projector.

So I'm going to shine this. I'm going to shine this, and shine this, and so on. So let's say I turned on all the even pixels of the projector. And take a photo. Certain pixels will be bright. The next pixel, some patches will be bright. The next patch will be dark, bright, dark, bright, dark. And that's not what I want. Now if I project the exact inverse of that, the pixel that goes on is off. And the one that is off is on. What will happen to global?

AUDIENCE: The same.

RAMESH RASKAR: It will remain the same based on the appearance from the [? Sun. ?] So in the first picture, I have some global. In the second picture, I have some global. And in case of direct, I put a direct here.

And here, I had basically one minus direct because I just need more direct. And you can make it all even simpler. And therefore, the equations [INAUDIBLE] so we go through it very quickly.

You project the pattern. Look at the pixel. In this case, it's lit. It's receiving the direct component and 1/2 of the global component because, remember, if I turn on all the pixels of the projector, I'll get certain intensity in the dark patch. If I turn on only half of them, then the global intensity also would [INAUDIBLE]. So the alpha here is just 1/2. So what I got in the first one is direct plus 1/2.

If I switch the checkerboard, I'm going to go-- again is deep because it's just inverse of that. One minus these are [INAUDIBLE] because it's just 1/2. And again, 1/2 of [INAUDIBLE]. That's it. So we have two equations and two unknowns, a direct component [INAUDIBLE]. If I take the subtraction, I will get--

AUDIENCE: [INAUDIBLE]

[LAUGHTER]

RAMESH RASKAR: What will I get?

[INTERPOSING VOICES]

AUDIENCE: Are they exactly the same [INAUDIBLE]?

RAMESH RASKAR: Yes, so [INAUDIBLE] all of them are related. So all I'm going to do is take those two pictures, find the max of the two. And that's going in my direct. And to find the mean of the two, that's going [INAUDIBLE] global.

So if I look at this particular patch when it was not lit, it's the same. So think about this patch, particular patch. So i of x. Let's call it patch x.

This is the same, direct plus 1/2 global. But unfortunately, this particular one did not get a direct path in the first picture. In the second picture, again, you [INAUDIBLE] plus 1/2 g.

In the second picture, I did get the direct. So my first equation is 1/2 g. Second equation is [INAUDIBLE]. If I just apply the two, I get the direct. And then, if I do additional [INAUDIBLE] or if I just take the minimum of these two, I will get the global. And if I subtract the two, I will get it. That's it.

So from these two pictures, I can tell you what bounced straight from the wall and what bounced here. This is almost like magic because physicists build really expensive equipment, laser equipment, to solve this problem of what's going to bounce and, what's a [? global ?] bounce? And now, we have come up with a very simple technique, where we just change the illumination, which we call computational illumination, to figure out what's direct and what's global. Yes?

AUDIENCE: So by figuring this out, you eliminate the global illumination. And basically, you have a [INAUDIBLE] in a very high-coordination device, like a laser. So you remove all the [INAUDIBLE], all the indirect and global.

RAMESH RASKAR: Right.

AUDIENCE: So maybe we solved it. And so we can use a light [INAUDIBLE] like a [INAUDIBLE] high-collimation.

RAMESH RASKAR: I don't know what you mean by collimation. You--

AUDIENCE: That you [? inside the ?] laser.

RAMESH RASKAR: Just the [INAUDIBLE] you mean?

AUDIENCE: Yeah, [INAUDIBLE] of the spot is--

[INTERPOSING VOICES]

RAMESH RASKAR: Is very high, yes. But as you can see, even if you have a very narrow spot, it still has global scattering.

AUDIENCE: Sure.

[INTERPOSING VOICES]

RAMESH RASKAR: So collimation alone does not give you--

AUDIENCE: Sure. Sure.

RAMESH RASKAR: So the right way to think about this is imagine if somebody was sitting in this room and there's a wall. Now think about two different pictures. Imagine this one is white. And I turn on the light switch.

What will happen is-- and this is what's the difference between professional photographers and consumers. If you go to a professional photographer, they'll put a nice white drop. I think, Marcel, you have that box in your office? And the reason the box is white is because you want to see the object nicely lit from multiple directions because of scattering. You will get 1,000 pictures.

Now imagine if this box around the object is completely black. So that light hits this thing. It goes to the reflector. But nothing comes back because it's completely black. The picture that you would get when the rest of the wall is completely black would be the direct illumination. There's no scattering going on. As opposed to putting it on a white wall, well, you would get this one.

AUDIENCE: So it's a bit like being in an echoing chamber for some.

RAMESH RASKAR: You can think of that way. So if you go to, yeah, an [INAUDIBLE].

AUDIENCE: [INAUDIBLE].

RAMESH RASKAR: Yeah, you feel that it's strained. This wall, I can hear my echoes. And that helps me change my tone because I know with that, I'm talking too loud and so on.

AUDIENCE: Yeah.

RAMESH RASKAR: But sometimes, when you're at a performance, you go on the stage and you can't hear yourself. That's a similar feeling. So if you eliminate all the other bounces, that's the direct component.

And so as you can imagine, in photography, the global component is very critical. It's very, very critical. When you're shooting a movie, and even if it's a nighttime scene, the director will always put some small light sources here and there so something is still visible on screen and you can figure out what the characters are saying.

A real nighttime scene doesn't look like that. But by creating this global illumination, they just created enough brightness so they can do so. So it's extremely powerful [INAUDIBLE] across the globe. And let's see how we can use it. So before we see that, let's look at some real-world objects and how, for them, direct and global components really matter.

So we have marble. We have this very diffuse candle wax, water with some-- a glass with milky water, and inter-reflections between the walls, and so on. And this one, if you do this trick, you'll see something really interesting. What will you see, let's say, at the corner here? Brother, do you know why the corner of any room looks really, really bright?

[INTERPOSING VOICES]

RAMESH RASKAR: It's all inter-reflections. There's more light reflecting back and forth there than anywhere else. That's [INAUDIBLE]. What about this wax here? Most of the light is actually scattering inside. So it's almost all global.

AUDIENCE: Yeah.

RAMESH RASKAR: Same with marble and so on. So we'll see how it works [INAUDIBLE]. But this is how the world looks. This is all the light that's coming straight back.

If you look at wax or milky water, none of the light actually comes straight back. It all bounces around before it comes back. On the other hand, over here, you can see that light is actually reflecting. And also in the corner, you can see that everything is because of inter-reflection.

There's a bug here. He's in the corner. And you see it's very bright. It shouldn't be because the corner here should not be any brighter than the rest of the wall because it's just a direct reflection. But there's a limitation to this algorithm in terms of its resolution. That's why it doesn't--

[INTERPOSING VOICES]

AUDIENCE: Yeah, so you get misategorized because your checkerboard size isn't small enough.

RAMESH RASKAR: Exactly. Your pixels are not small enough to capture that card.

AUDIENCE: Yeah, another question is, so in the previous example with the shower curtain, the direct component lets you see through the scattering curve?

RAMESH RASKAR: No, the global component

[INTERPOSING VOICES]

AUDIENCE: Oh, it's global. [INAUDIBLE] here. OK, never mind.

RAMESH RASKAR: Yes?

AUDIENCE: Well, it's a good technique, also, to take a picture of an aquarium.

RAMESH RASKAR: Aquarium?

AUDIENCE: Yes.

RAMESH RASKAR:	Yeah, certainly.
	[LAUGHTER]
AUDIENCE:	[INAUDIBLE]
RAMESH RASKAR:	Yes?
	[LAUGHTER]
AUDIENCE:	Is there a [INAUDIBLE] to adjust how much flow [INAUDIBLE] happen or directly happen [INAUDIBLE]?
RAMESH RASKAR:	Beautiful. So you can play those tricks in your Photoshop when you aren't just changing the brightness, but you're changing the global versus direct. So the wax will start glowing. But the regular [INAUDIBLE] surfaces will not. You can play all those tricks. Also, the scan. It turns out, we'll see later that if you lived in a world where direct and global could be separated, there'll be no racism.
	[LAUGHTER]
AUDIENCE:	Sorry. Can you go back? How about the highlight on the glass on the direct side?
RAMESH RASKAR:	This one here?
AUDIENCE:	No.
	[INTERPOSING VOICES]
RAMESH RASKAR:	Oh, this one?
AUDIENCE:	Yeah.
RAMESH RASKAR:	Yeah, so you will see later that this technique doesn't work for highlights. And I want you to think about why, why that is. All right? So let's look at some fun examples. So this is a weed that grew in a wall, which we already saw. It should look flat. And only the corner, you'll see inter-reflections. Here's a failure case.
	[LAUGHTER]
	If you have a shiny mirror, then if I shine a laser at the mirror, I'll get-- so let me try it properly. I have a mirror. And I have a bulb. If I shine a laser here, it reflects here. On the other hand, if I had a bulb, then if I shine a laser here, it reflects everything. So it only reflected one spot. If I change the laser direction to slightly here, it goes in a different spot. And the global effect of here versus here is quite different as opposed to here.
AUDIENCE:	Yeah.
RAMESH RASKAR:	If I'm shining this part versus this part, this will have the same ratio, same ratio. But that's not true for everyone's. [INAUDIBLE].

AUDIENCE:	But would you even call that global illumination?
RAMESH RASKAR:	So global domination is anything that has more than one bounce.
AUDIENCE:	Oh, OK.
RAMESH RASKAR:	So by the definition of global reflection, anything that has more than one bounce of a photon is global. So yeah, it fits. Depending on your application, you can say [INAUDIBLE] is not part of my world. And I can [INAUDIBLE].
	It has its issues. So now, we have some interesting examples. I really like this part. So make a reflection, x. What should you see?
AUDIENCE:	[INAUDIBLE].
RAMESH RASKAR:	Because between those x's, it's got a lot of light bouncing around here. So [INAUDIBLE]. It's bottom right. You can also do a trick, where you can do it in sunlight. You can just do the inverse effect. I can just take a stick and move it. And remember, if I take lots of images, the minimum of all the images is the global. And the maximum of all the images is direct plus global. That's it. Those three equations are all done.
	It's very easy. So if I move a stick and shoot a video, at any given pixel, I look at all the frames. And the minimum of that is the global component. So you can do this outdoors [INAUDIBLE]. Sorry.
	Then the shower curtain, our favorite example, [INAUDIBLE] reflect. And here, light that bounces straight from the curtain is the direct component. Any light that actually bounced around and came back to the camera is the global component.
	So anything that's behind the curtain will actually play a bigger role in the global. Now the problem is that even the texture of the curtain actually has some subsurface character. It has some thickness. And light is bouncing around and then coming up. So that's kept here as well. But you can capture those. Yes?
AUDIENCE:	What was the mesh diffuser that you were using?
RAMESH RASKAR:	So yeah, that's a good point. So there are multiple ways you can do it. I can either use the projector. Or I can just take a mesh. At least buy it at Home Depot. Buy a high-frequency mesh for the light source. And just move the mesh in front of it and shoot a video. So to cast-- to create this slide projector that projects a checkerboard and it's inverse, I don't have to use the projector. I can use the light source and a mesh or a stick, whatever it is.
AUDIENCE:	For you-- so for Photoshop, do you-- need a video. For the stick, you need video? Or do you need to [INAUDIBLE] just two pictures befor [INAUDIBLE].
RAMESH RASKAR:	You're going to take more pictures.
AUDIENCE:	OK. OK. Good.
RAMESH RASKAR:	But the minimum number of pictures you need is two because you have two unknowns, so direct and global.

AUDIENCE:	And global.
RAMESH RASKAR:	And then, this is your fish tank, where you will see very clearly in murky water. That's direct. Actually, it turns out most of the light is bouncing around and coming [INAUDIBLE].
AUDIENCE:	Yeah.
RAMESH RASKAR:	And the pixel on the left-- maybe I should lay on the lights here.
AUDIENCE:	There's some pixel [INAUDIBLE].
RAMESH RASKAR:	Yes.
AUDIENCE:	[INAUDIBLE].
RAMESH RASKAR:	Is that better?
AUDIENCE:	Yeah.
AUDIENCE:	Yeah.
AUDIENCE:	No.
AUDIENCE:	No.
	[LAUGHTER]
AUDIENCE:	[INAUDIBLE].
RAMESH RASKAR:	I don't want to disturb those [INAUDIBLE].
AUDIENCE:	Yeah. I moved it.
AUDIENCE:	Yeah, exactly.
RAMESH RASKAR:	And this looks like some computer graphics rendering. Why is that?
AUDIENCE:	Because in the old days, they didn't have global.
RAMESH RASKAR:	Exactly. Because chip [INAUDIBLE] graphics rendering is basically just a direct bounce of light. It doesn't think about RenderMan, Maya, doing all this bouncing of light around and so on. So when Pixar decided to make animation movies, which movie did they make first?
AUDIENCE:	"Toy Story" [INAUDIBLE].

RAMESH RASKAR:	"Toy Story." Why?
AUDIENCE:	Because [INAUDIBLE].
AUDIENCE:	It's all plastic.
	[LAUGHTER]
	It's very easy to render.
	[INTERPOSING VOICES]
RAMESH RASKAR:	They cannot do ants. And they cannot do fur. And they cannot do rain. And they cannot do fog. It's just toys. Toys are the easiest thing to render because they have very simple surfaces. And they all look like this on the left. So you can--
AUDIENCE:	So what explains the color shift then that's in there?
RAMESH RASKAR:	Because if you know how light reflects inside some scattering medium, then it has a preference or certain wavelengths. If you're underwater, then red is suppressed and blue, green is not. So similar thing here. So it always has a blue, green tinge because what you're scattering suppresses one wavelength but not the other.
AUDIENCE:	Well, so the scattering goes, well, maybe to the border. So the blue should be scattered more
	[INTERPOSING VOICES]
RAMESH RASKAR:	So yeah, in the global component, you see more bluish. And the tinge, the bluish tinge on the red cup, there is gone on the left.
AUDIENCE:	Oh, OK. I see.
RAMESH RASKAR:	Because this is just a direct bounce. We're ignoring the scattering. So imagine the projector pixel comes. It scatters through the water. And it hits the red cup. And I'm just trying to measure that. I don't want to worry about how everything else is contributing back to that red cup.
AUDIENCE:	In that example, you show you either use a light bulb or a projector. Could you use the light of the sun?
RAMESH RASKAR:	So the example of the stick was with the sun.
AUDIENCE:	But can you motivate it? Could you via--
RAMESH RASKAR:	Yeah, I can take a grid or a mesh, and just move it, and shoot the video. It's the same thing.
AUDIENCE:	But then, just out of curiosity, would this whole scattering colors that I've been able to get, what's the actual real color of the green card, blue and green.
RAMESH RASKAR:	It's that. If I take it out and put it in free space--

[LAUGHTER]

No, it's that color.

AUDIENCE: Green.

RAMESH RASKAR: Green.

AUDIENCE: So there's the one that [INAUDIBLE].

AUDIENCE: Is it green? I don't know.

AUDIENCE: New age group. [INAUDIBLE].

RAMESH RASKAR: And then, you can play with adding more tinge and so on. As you were saying, can I have more global component or less global component and so on? Now here are some really fun examples. You can even figure out what is real and what is fake because one fruit is real and one fruit is fake. Let's see. Let's start with the banana. Which one is real? Which one is fake? Top is for real?

AUDIENCE: Yeah.

AUDIENCE: Yeah.

RAMESH RASKAR: How many think top is real? The bottom is real? Wow. Everything bottom is real. What about the apple? Left is real. Again, left is-- I said, right is real. So for bottom and right. And the lemons. Left is real? Right is real?

AUDIENCE: Yeah.

RAMESH RASKAR: All right. Remember your choices.

[LAUGHTER]

Pretty good. Pretty good. But how would a camera do it? If I just take this picture and you would do a camera, like a computer vision algorithm, it's almost impossible to figure out which one is real, which one is fake. And the reason for that is when you make a plastic or a fake fruit, you only try to match its outer appearance.

You're trying to match its inner appearance. So in case of the apple, you realize that light bounces around inside in the real apple. And it becomes reddish. That's what gives the reddish color. But in a fake apple, no light bounces around. So you know which one is apple and which one is Microsoft.

[LAUGHTER]

Over here, it's a little bit more complicated because although it's fake, I know, it has-- first of all, the direct bounce from a real banana is actually a very different color. And this is all because of the ripening that's happening inside. And when it comes to a lemon, this one is very complicated for me because in the real lemon, if you look at the direct component-- and this lemon, maybe it's good quality, a fake, because even the internal reflections inside that fake lemon.

And I think I've seen those fake lemons. They are a little bit soft and fuzzy so light can bounce around inside as opposed to apples, which they're just hard. The way to figure that out is if you look at the direct component of a real lemon, first of all, you realize that it's green. It's not blue-- it's not yellow. And the real lemon actually has a texture that you never see directly. So the real-- direct bounce, you'll see direct reflection.

So a lot of interesting things. And I wish I had it in my cell phone camera so I can figure out-- so this is, by the way, a great trick to figure out if something is ripe or not. If it's more ripe, light will bounce around more inside because its permeability is different. So I can build a cell phone camera that can tell you if your food is ripe or not.

[LAUGHTER]

AUDIENCE:	And so in photospectrometry, you can see the [INAUDIBLE].
RAMESH RASKAR:	Exactly.
AUDIENCE:	[INAUDIBLE]
RAMESH RASKAR:	Yeah, you can do it with the direct as well. But this is another tool you have--
AUDIENCE:	Yeah,
	[INTERPOSING VOICES]
RAMESH RASKAR:	--which can be implemented easy.
AUDIENCE:	Even with 3D image, the face of somebody, you should actually put everything inside because if this has worked, then it's [INAUDIBLE].
RAMESH RASKAR:	I'm sorry. Repeat the question.
AUDIENCE:	You showed that [INAUDIBLE] maybe something like that, something that is not made from the real
	[INTERPOSING VOICES]
RAMESH RASKAR:	Right. So whether there's a makeup or not.
AUDIENCE:	[INAUDIBLE] or not.
RAMESH RASKAR:	Yes, so unfortunately, yeah, this trick will also tell you if somebody is wearing makeup or not. Let me just talk about ethnicity. So as you know, that hands of different ethnicities look very different by skin. But if you look at the direct and global, the direct is almost always gray. And all the pigment is in the global. So if you take people of different ethnicities, I hope I have the example, the direct component is almost the same.
AUDIENCE:	Are these normalized at all?

RAMESH RASKAR:	They may be normalized, yeah.
AUDIENCE:	Because I--
RAMESH RASKAR:	But the global component is what actually gives the real color.
AUDIENCE:	OK, So when you add [INAUDIBLE]--
RAMESH RASKAR:	So there is no color. The direct reflection has no color. So if an alien comes in and can see only indirect illumination, that's [INAUDIBLE].
AUDIENCE:	You should also look on Flickr for infrared photographs.
RAMESH RASKAR:	Yeah.
AUDIENCE:	They show people are really wrinkly.
RAMESH RASKAR:	Exactly, yeah. And we'll see some thermal [INAUDIBLE] in there. That's also [INAUDIBLE]. Anyway, I'll go through it quickly.
AUDIENCE:	Yeah, what happens if you do this with IR?
RAMESH RASKAR:	It's the same effect.
	[INTERPOSING VOICES]
RAMESH RASKAR:	Sometimes, light-- at some wavelengths, there's more in inter-reflections than others.
AUDIENCE:	Yeah.
RAMESH RASKAR:	So if you can imagine, some surfaces are actually darker in a particular wavelength. So the situation I explained, where you have a person in a white room versus a black room, that appearance looks very different because there's direct plus global in the white room but not in the dark room. So at a particular wavelength, the room may be black. And you'll see just direct. All right. So let's take a break. And then, we'll come back and talk about optics and light rays.