

MITOCW | mitmas_531f09_lec05_1.mp4

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

RAMESH So we should celebrate. The Nobel Prize is for imaging this time. Ah, of course, fiber optics, but who cares about fiber optics?

But it's also about light and imaging. And this is the point I was making just a couple of weeks ago about how in imaging, you keep on getting Nobel prizes. But somehow in other fields, just no Nobel prizes.

Of course, [INAUDIBLE] sent out an email. I don't know if I told you this. And he thinks prizes like Nobel Prize and all that have outlived their importance because they were always trying to feed this notion that science moves by a select group of people working in select places and coming up with amazing discoveries-- like, select individuals. And we know that's not how the world works now.

Even the development of fiber optics, it's not just one person or one team, but hundreds of people have made tremendous contributions and the same with CCD chip for which there was a Nobel Prize. And you know, somebody may have thought about it. Somebody may have implemented it.

But at the same time, there are a lot of people involved. So it's somewhat misleading to identify one or two people as the ones who made all the difference. By the way, the inventor of light fields, which we're going to talk about, Lichtman, also got a Nobel Prize. But he didn't get the Nobel Prize for light fields.

He got a Nobel Prize for inventing color photography, which is very closely related to light trace. It's surprising. We'll discuss about that today.

AUDIENCE: [INAUDIBLE]

RAMESH So the assignment actually has two options, of course, and the light field part has several parts, several sub-parts. But don't get intimidated by all the things you ought to do. It's mostly a matter of how much you want to do. And based on your background, so if I know you have a lot of background in light fields and so on-- so like Rohit over there, for example-- I would expect him to go quite a bit deep into the assignment. But those of you who don't have as much optics background or those of you who are not programming and using-- or don't have programming background and using some GUI-based tools to do the actual assignment, of course I'll expect you not to-- just finish the first part but not-- first sub-part but not all the parts.

So we will kind of normalize your performance based on your background and abilities. So you may want to send an email to me separately in case you run into problems and say, hey, I don't have as much background in this area. But the assignment is basically-- it's basically taking pictures and taking an average of those pictures. So it's just that when you-- you have to take a picture, shift it, and add it to some other picture.

So the concepts are really simple, and if you do the first sub-part, you should be able to do other sub-parts very quickly if you think about it some more. But again, given that you have limited time, you don't have to go all the way. All right, so what I'm going to do is actually start talking about the light fields, and then halfway through, we'll come back and talk about the assignment, OK?

So if you're in photography or you're new to cameras, people always throw the stumps-- at you-- the F stop and fast lens and slow lens and OK and depth of field. And Ankit covered it a little bit, but it's always very confusing. And photographers, just like any other cult group, always try to make it more exotic and more difficult to understand what they're saying.

So my general advice is as a researcher, as a student, as a scientist, just get rid of all the jargon and focus on really simple parameters. So F stop, for example, is really complicated, and it's not just it's complicated, but it's actually wrong. The numbers they use are actually wrong.

They're off by sometimes up to 10%. And increasing the F number or decreasing the F number is not very clear what it means because it's inverse of-- so an F 2.8 versus F 5.6, the 5.6 actually is smaller than 2.8. We're working with that. In microscopy and scientific imaging, they use numerical aperture, which is more meaningful because a numerical-- when an aperture becomes larger, the numerical aperture becomes larger, and it's part of order in a scientific way.

So let's see. So the whole concept of light field might seem very foreign, and you'll wonder why we need to study it. But once we have understood this relationship between 3D points to 2D pixels and rays and relationship between different rays and so on, all these other concepts such as depth of field and F number and all of that will become extremely clear, and you won't think of them as some abstract quantities but quantities with some real meaning to it.

All right, so this is how we started thinking about the ray space. And then we realized that rays are five dimensional because there's a point and direction. So three degrees of freedom for the point and only two degrees of freedom for the direction-- why is it not six degrees of freedom as we have elsewhere? Usually have three for translation, three for rotation. Here we have three for translation only, two for rotation.

AUDIENCE: Because there's not [INAUDIBLE].

RAMESH Sorry?

RASKAR:

AUDIENCE: There's no [INAUDIBLE].

RAMESH There is no role.

RASKAR:

AUDIENCE: Yeah.

RAMESH Ignore that. But let me realize that sometimes you don't have to use five dimensions. So you can get away with four dimensions. In which case does that happen?

AUDIENCE: Two plane approximation?

RAMESH Two plane approximation? That's good, but why does it allow you to go from-- why can't we go from 5 degree 2 degree? [INAUDIBLE] Raise your hand? OK, let's take the next one.

AUDIENCE: You're only caring about a single ray or a single point?

RAMESH Exactly. So if you don't have this uploader in the middle, the intensity on this side is going to be the intensity on the other side. So you can just use it as a four dimensional quantity. And now the four dimensional quantities, we can specify multiple ways.

We can either specified as a 2D position and 2D angle, or we can do it as a two plane parameterization where you have 2D point on one plane and 2D point on the other. And this is kind of in the 3D world. So UV coordinate on this plane and SD coordinate on that plane, and that defines the ray direction.

Now what's the disadvantage? I'll just go back to this. So you'll realize that when we're talking about all these optics and rays and geometric quantities, we'll always think about them in flatland. So the real world is 3D and rays are in 4D space. But for the sake of discussion, our world is just 2D, the flatland, and the rays are how many dimensions?

AUDIENCE: 2D.

RAMESH Two dimensions, OK? So something strange going on already, right? Because we went-- in a 3D world, the rays
RASKAR: are 4D. But in a 2D world, the rays are--

AUDIENCE: 2D.

RAMESH --2D, right? There's already a mismatch going on, and as we go forward, you'll realize where that comes from.
RASKAR: OK, so most of our discussion will stay in flatland.

So our world is 2D. Our rays are 2D. There is one degree of freedom for the position in this.

So in this case, one degree of freedom for the position and one degree of freedom for the angle, OK? So that's the two dimensions for the ray. We can either express it as position and angle or as position and position. Now that seems pretty good. There are some rays here, which cannot be represented very well with this representation, which, it's OK to represent them here.

But they cannot be represented here. Which ray is it here? Yes?

AUDIENCE: Ones that are parallel to the plane?

RAMESH Exactly. So there's something-- if a ray is vertical here, it will never intersect both planes. So it's difficult to
RASKAR: express that in this particular formulation. And if you are thinking about implementation, then any ray that's-- even if it's not vertical but close to vertical will have sampling issues.

You won't be able to represent that where you are. And here, you don't have that problem because you're going to get-- you'll use a different 3D position, and theta phi is always defined. So that makes it very simple. All right, so now if you can build this, a matrix, the machine that can see everything everywhere in every wavelength, how many dimensions is that?

AUDIENCE: 7D.

RAMESH Seven dimensions, right? Five-- so we go back to this one. 5 dimensions here plus wavelength and that one.
RASKAR:

AUDIENCE: OK.

RAMESH
RASKAR:

OK. Now we're going to get closer and closer to that matrix that sees everywhere. And we're going to start with a camera array. And we're going to say that an array of cameras can capture light field.

How does that happen? Here we have-- let's see. Imagine on this plane, I'm going to put a set of cameras.

And for every camera, its frame buffer I'm going to indicate as S and T. So every camera position here is given by UV coordinate, and the frame buffer of each camera is specified by ST. Like so is ST, OK?

So this is how we are. We have the UV plane. I'm going to put a camera here. And then I have ST plane. And what I'm going to say is for this camera, I'm going to show all these [INAUDIBLE].

So as you can see, the image starts becoming very messy in the real world. So we go back to find it. And they go right here. So this is my U plane, and this is my S plane.

I'm going to put a camera here. Camera here, camera here, and camera here. And when I say place a camera there, I really mean the center of projection, the pinhole of the camera, is placed here. So I put a pinhole camera here and the sensor back here. And what that is doing is for every pixel through this pinhole, they're shooting the ray over here.

Again, we're going to think like the Greeks. The camera should be placed out in the world [INAUDIBLE]. And the coordinate of these pixels now is even less.

But the position of each camera [INAUDIBLE]. So maybe this u equals 1, u equals 2, u equals 3, u equals 4, and so on. And similarly on x-coordinate, whatever, it goes [INAUDIBLE] coordinate.

And we do say, like, this camera can capture all the rays that pass through. And then this camera [INAUDIBLE] plus all the rays that [INAUDIBLE]. So very similar deal with [INAUDIBLE]. I'm going to the camera, shift, shift, shift, shift. And it's going to capture all of those rays.

Certain rays are missing here. They're not here to capture, usually [INAUDIBLE]. So the capturing-- if I take a photo from here, I will still capture 1,000 rays corresponding to each X coordinate.

If I go here, I captured another 1,000 rays. So in this picture, I have captured over 4,000 rays. And I'm using the rays without specifying what it means. But I think the whole notion of rays-- some type of bite doing that. So what's missing?

AUDIENCE: The point distribution [INAUDIBLE].

RAMESH
RASKAR:

Exactly. So I'd say that when q equals 0.5, those rays are not being captured. So that's what life is. The world is continuous. And somehow, it will discretize and [INAUDIBLE].

OK, but that's fine. It's close enough for us to start representing the world in a 4D resolution. So now if somebody asks you what's the dimensionality of the world, people of course say, yes, it's 3D. But what is this thing? The world is actually four.

And it's almost a difficult concept to get across. But some other ways of thinking about the same situation is imagine if you build a hologram. You have a hologram, and you're standing in front of it. And the hologram has a coordinate system ST.

And as you move your head, you see a new image coming out of this hologram, right? So if I move in the U direction, I see a different set of images. But if I move my head up and down, then I [INAUDIBLE] see new images in the v direction.

So the hologram also is recording four dimensional rays on particular displays, perpendicular [INAUDIBLE]. He's also recording a 4D version. So to record a view through a window, for example, I need to adjust to a four dimensional dataset in a two dimensional world.

And that's [INAUDIBLE] problem. And this one is actually inverse of that. We're going to capture the world. And when you capture the world through the window, again, we record a four dimensional [INAUDIBLE].

So try having that argument with your friends where the world is 3D or 4D. Maybe the [INAUDIBLE] have gone one more dimension.

All right, so now we have RF cameras, and we're going to think about how that can be achieved not within RF cameras but just with a single camera. So the RF camera is very much similar to this, just a set of cameras with a bunch of sensors behind it. And as far as photography is concerned, we want to think about how we can achieve that with just a single camera.

So there are a few concepts we are going to go in to look at more carefully. One is that an image that's created for a 3D point is actually a sum of multiple images created by different parts of the lens. So if we have a whole lens and you cover all these parts and then just do this more often, it will create one image.

Then if I look at one where we see this one open, you'll get another image and so on. And in this case, I will get, I guess, five different images. Now what's the difference between those five images. You can always kind of see the hint here.

It's as if I placed a camera here, then I placed the camera here, then I placed a camera here, here, and here. So it's just like opening a part of the aperture or sub-aperture is just like moving a camera but within the aperture of your oculus. If your aperture is only about 25 millimeters, the whole shift is only about 25 millimeters.

But there's a difference between this situation and this situation, OK? So because all of these are pinholes and there's no optics in it and this one has pimples-- but there's some special shape. If you just look at any one section, what does it look like?

Don't look at the whole lens. Just look at one section. Sorry?

AUDIENCE: The trapezoid?

RAMESH
RASKAR: Trapezoid, which is a truncation of a prism, right? So you can think of a lens as being made up of sections of a prism. And of course, this prism is truncated here. And this prism is truncated here. And this one is just a slab of glass, just a [INAUDIBLE]. So it's just a flat piece of glass. It doesn't really do anything.

And so in high school, when they say-- when the light goes to the center of the lens, it doesn't change light. It's a slab of glass. Of course light goes through it.

And here, the whole prism. So light goes in, and it deflects. So we just have a set of prisms.

Here we just have a set of pinholes. So let's assign some of those coordinates one more time. We're going to say the camera position, right?

The camera position here is given by U and the pixel position is given by S . You can ignore the second variable here. And we're going to try and understand this four dimensional space that the camera has captured.

I'll give you a punchline here, which is if a camera can capture this four dimensional representation of light, it has captured everything that's coming through that lens, and it includes phase and all that information. So all the geometric information about the world is completely captured by this four dimensional representation. So once you have that, of course, you can go back and do anything you want.

You can do digital refocusing. You can estimate 3D shape. You can insert new objects in it, whatever you want. That's all you can get.

So just by going two dimensions higher, from 2D to 4D, it's a complete representation. And that's why it's so important to capture that for the lectures because if you want to think about any sophisticated photography applications, that's the maximum information you can get. We still have wavelength and time and polarization and so on, but those are not geometric dimensions. The geometry dimensions are just space and angle or two space and so on.

All right, so, as I said, we're going to think of the lengths as being made up of a sequence of prisms. And when you capture a light film, you end up getting an image that looks like this. So the picture from here versus here versus here versus here will have a small parallax as you move around.

It will look almost the same, but there'll be a very tiny parallax. So if you look at this set of images, they look almost the same, and they're captured by a light field camera. But if you see very carefully, then the distance between the eye and the ear of the body is changing from top to bottom.

Here it's almost touching the eye. And over here, there is a significant gap between the-- so there are these small gaps, and I'm sure when you do this experiment where you look at something with left eye versus right eye, it seems to shift.

Your finger seems to shift. And that's the parallax that you're introducing. And that's the same effect that's happening here because parallax between your eyes is only about six centimeters. But that's sufficient to create shift and relative to certain positions.

AUDIENCE: What's a method for capturing the phase?

RAMESH
RASKAR: So we'll come and discuss that. It's a very important question. All right, And Marc Levoy and his group at Stanford are world leaders in thinking about many types of light field cameras and also their applications in different areas. So I've taken a lot of slides from his presentations, and he has also applied it for microscopy and so on.

All right, so very briefly-- I'm not sure I'm correct. OK, very briefly, there are other ways you can represent this. So right now, we're thinking about a UV plane and an ST plane and just the intersection of that. But depending on your application, you may have some other ones.

Maybe you have a sphere. And you cut the hemisphere, the two hemispheres. And any point on this hemisphere and any point on that hemisphere together will create a ray as well.

Or in some cases, you want to create some convenient 2D manifold, and the position on the manifold and the theta field respect to dive can also be electric. So just those of you who think about it mathematically, there's a continuum between flexible light fields to two plane parameterization [INAUDIBLE]. All right, so let's come back to talking about light field inside the camera.

So why does a traditional camera capture this 4D light field? So you have a point in sharp focus. The reason of that [INAUDIBLE] they bend at the lens, and they converge to a photodetector. And all the radiance coming from different directions is integrated together to get to-- and that's recorded as the intensity of the [INAUDIBLE]. So what you get out there is just two-- so fly sensor, you get 2D image. The question is, what do we get extra when we capture a 4D light field in terms of [INAUDIBLE]?

AUDIENCE: So right now, it's compressing all the angles into [INAUDIBLE] extract the angles.

RAMESH
RASKAR: Exactly. So this is the most important question. Remember, I'm repeating it multiple times. But now we want to figure out not just the total of all the space but also the value of each of these rules. OK, and where the rays comes from individually?

They're coming from different parts of the lens. So there's a very interesting relationship between the lens and the sensor. And if realize that geometric relationship, you may be able to recover this 4D lecture.

So [INAUDIBLE] and his student came up with an idea in 1992 to try and build a compact device. The original idea was again presented by Lippmann in 1908, more than 100 years ago now. So the concept of light field has been around about 100 years.

And then Brenning at Stanford in 2005 created this really beautiful device that can very compactly capture a light field as well. The basic idea is very straightforward. You move the sensor a little bit further back, just a few micrometers, and then replace this plane with a set of microlens array. And if you just look at these two here, it's very similar to the lenticular display or microlens-based display that you see on cereal boxes and rulers and so on.

But now we're going to use that for imaging rather than for display. And this is how they built the device. They placed these microlens arrays.

The pitch of each microlens is 125 microns, and the pixel itself is about 9 microns. So under each microlens array here, you have about 9 by 9 pixels-- sorry, 14 by 14 pixels. I'm sorry. And then, so what they're going to do is they're going to record the incoming light in 14 different directions.

That basically means that they're going to slice the lens into 14 segments, and they're going to capture the light that would have appeared if I just opened the first part of it and blocked the remaining 13, and then I unblocked only the second one and blocked everything else. So the 14 pictures that I would have taken by exposing only one sub-aperture at a time, you can capture that in one shot. Yes.

AUDIENCE: So I have a question about the [INAUDIBLE] array. Do any of the rays ever go outside of that 14 by 14 area and hit a different sort of thing? So--

RAMESH
RASKAR: So it's not coming from here. It's something from here, but come and--

AUDIENCE: Yeah.

RAMESH --and spoil image?

RASKAR:

AUDIENCE: Yeah.

RAMESH That's a great question. That's a great question, and we'll come to that. And it basically relates to the F number

RASKAR: or the relative opening angles of the main lens and the microlens.

AUDIENCE: But the goal is for that not to happen, right?

RAMESH Exactly, yeah. Not just it should not happen in terms of overlap, but you also don't want any black lesions here.

RASKAR: So you want to put the sub-aperture images as tight as possible. And these are some examples of that. Yes.

AUDIENCE: So why can't I just think this as like [INAUDIBLE]? I mean, it's quite a nice array you're working on. Maybe I can think of it, the system, as my negative equations [INAUDIBLE], and if I know the parameters on the unknowns, then I can just solve. I don't really need to-- an extra--

RAMESH Excellent. That's a very good question. So those of you who think of this as some kind of a projection of a 4D space onto a 2D sensor will say, OK, there must be some way I can recover the 4D image, 4D representation, from this 2D image, right? And fact of the matter is depending on the scene, you may be able to do it.

But in general, this inversion is very unstable because-- I'll explain. Because if you have-- so what does it mean to actually capture this image? If I have something that's really sharp in focus, then all the rays, have converged on that. So there's almost no information recorded about the relation between this and this and this. And let's see if I go out of focus.

Then I will get a blur here. And we'll look at that more. And all the blurs will go on top of each other. So basically, we are reducing the dimensions of the data set and hoping to kind of look up and recover a higher emission data set.

AUDIENCE: But even if something in terms of-- OK, so what you are saying is that you're probably not solving this important because you get linear [INAUDIBLE] equations.

RAMESH Yes.

RASKAR:

AUDIENCE: [INAUDIBLE]. I'm trying to understand why [INAUDIBLE] just because in focusing on the summations, the coefficients of all the time of all the variables are going to be very small on every single sensor?

RAMESH It's a bigger problem than that because remember, we're going from four dimensions to two dimensions. What you're saying is valid when we're going from two dimensions to two dimensions. But if I'm going from four dimensions to two dimensions, there is a significant-- it's a lossy representation of the signal. If I have a very unique scene which itself doesn't have a 4D light field going out, then I can recall it [INAUDIBLE].

AUDIENCE: [INAUDIBLE] long [INAUDIBLE] sensing [INAUDIBLE] did not only have full information--

RAMESH Exactly.

RASKAR:

AUDIENCE: But does that mean that I can't actually solve?

RAMESH In most of the cases, you can't.

RASKAR:

AUDIENCE: [INAUDIBLE] four pixels as-- or four [INAUDIBLE] original scene as one pixel.

RAMESH Yes.

RASKAR:

AUDIENCE: Can I solve for that?

RAMESH So to rephrase your question, you're saying, how can I queue up some resolution and recover a four dimensional representation, right? And that's exactly what's happening here with a hell of some physics. So instead of thinking of this as just bare bones lens and sensor, if you take some help from physics and optics, then it turns out that that inversion is possible.

AUDIENCE: But why can't I [INAUDIBLE] like, just theoretical stability? Because [INAUDIBLE] small [INAUDIBLE] physical process [INAUDIBLE] equation can have independent?

RAMESH It's just numerical stability. Yeah, I mean, first of all, you're going from 4D to 2D. Sometimes even going 2D to 2D, it may not be invertible. And here, because you're going from 4D to 2D, it's even more challenging.

AUDIENCE: [INAUDIBLE] did not have complete, full [INAUDIBLE].

RAMESH Exactly.

RASKAR:

AUDIENCE: [INAUDIBLE]

RAMESH Exactly.

RASKAR:

AUDIENCE: It should be possible to solve that.

RAMESH So one good example of that is in astronomy, and we'll see it later. We're looking at a star. And clearly, if you take a picture of the star, within the aperture of the lens, the star doesn't change that much.

RASKAR: So clearly, it's a redundant data. And then light fields are used there to actually figure out the aberration in air. So that's an example we look at. But that's a very important question because we'll always get into this. We always want to do more with less, right?

So one motivation is how a 2D sensor somehow will capture more information. And we're going to either use some optics and some physical techniques, or we're going to use some computational techniques. And the best case scenario or the best strategy is actually to combine the two. Combine a physical and computational approach to recover the small one. And the things we saw in the last class, for example, being able to tell whether an apple is real or fake or being able to read a playing card and all that, is a similar problem, right, because the world is high dimensional, and we only have a 2D sensor.

AUDIENCE: OK, so one question about this. Wouldn't it be similar if I take the photo sensor and the something and just move it forward so that all the rays don't converge? I mean, I'm waiting to separate them. When would I get them before they converge?

RAMESH
RASKAR: You're thinking the right direction. There are multiple ways you can recover all the 4D rays. You can-- the simplest one would be to simply block part of it and take multiple pictures. So that's time multiplexing.

This one is space multiplexing because I'm giving up some resolution here to be able to recover the ray direction. So I'm trading off basically spatial resolution for angular resolution, right? And then the technique you're describing says, why not just move the sensor back and forth and take multiple pictures?

And that's actually what's used in astronomy to recover the aberration in the sky and so on. So if I had just come up with that a few tens of years ago. Actually, 1967 is when phase diversity was, I believe, invented by a professor at Tufts.

All right, so just a piece of that thing we saw just earlier that we can take the 16 megapixel image now, 4,000 by 4,000. But then because each microlens now will be treated as just one pixel-- one megapixel, rather. But under one megapixel, you have 14 by 14 pixels. So in terms of megapixels, you only have 292 by 293 image. But because each megapixel has 14 by 14 underneath, we can do a digital refocus from a single shot.

So this is spatial multiplexing because we are trading space for angle. So here's an example where you said, OK, I want to give up four pixels to get angles for only one pixel. So now we have given up 14 by 14 for one megapixel. But still, this involves using additional optics, and your dream is can I do that without changing anything, right? Is that correct?

AUDIENCE: Can you put the left part of the image in focus up close and the left part of the image [INAUDIBLE] way in the back?

RAMESH
RASKAR: Yeah, you can do that, and that's part of the assignment. Having a slanted focus is now possible. As opposed to just planar change in focus, you can have focus that's not at a constant depth but at some slant. And the tilt shift photography, which we saw in the very first class, where your angle between the lens-- so typically, when you have a lens at the center, what's in focus over here.

And using [INAUDIBLE] law, if I shift the sensor from here to here, then the planar focus shifts from there to there. So different things come into focus. But actually, if you tilt these guys, then the actual planar focus was on the [INAUDIBLE]. And the [INAUDIBLE] principle says that this and this case will all meet at one point.

But with the light field camera, you can do this directly. And you can decide this effect afterwards, after the problem statement. The plane of focus doesn't really shift [INAUDIBLE] and so on. So I recommend you to do that part of the assignment. You can skip the first one and just go for that part.

Because the way the refocusing is achieved, as we discussed last time, you're just going to take all these 14 images in case of flatline and simply shift and add them. And how you shift and add them decides that planar focus. If I don't shift and add, I focus at infinity.

If I shift a little bit and add, I start focusing closer. If I shift more and add, I focus really close. But then I don't shift all of them the same way.

I can shift some of them more and others less. And then I can focus on [INAUDIBLE]. Go ahead.

AUDIENCE: [INAUDIBLE] to assume that the pixel blur will be constant or [INAUDIBLE] definitely for a given step.

RAMESH What do you mean by that? Can you repeat?

RASKAR:

AUDIENCE: So as you dive for a given [INAUDIBLE] the focus.

RAMESH Right.

RASKAR:

AUDIENCE: The pixel [INAUDIBLE] to see what all [INAUDIBLE].

RAMESH Mm-hm.

RASKAR:

AUDIENCE: Is it a reasonable assumption?

RAMESH It's a reasonable assumption. It's a reasonable assumption that every point at a given depth will have the same amount of blur. When you start thinking about the length separation and so on, you will actually realize that this is true only in the middle part of the image. But as you go to the periphery of the image, you start getting some length separation effects. But the first order approximation, you can assume that the whole chain has the same overflow.

And that's what distinguishes a cheap camera from an expensive camera-- because for a cheap camera, the center is in good quality, but the periphery is not. And an expensive camera will try to use multiple lenses and all kinds of tricks to make sure that particular statement is true. All right, so now we come to the part that kind of suggests the concept between scientific computing and light fields.

So when you want to-- there's this whole concept of adaptive optics which is used in astronomy and telescopes and so on. And what this simply says is that if I'm looking at a star which is very far away, then the wavefronts that are coming from the star should be mostly planar. So imagine that for a stone in water, and it starts creating the wavelengths where if you go sufficiently far away, they become mostly planar [INAUDIBLE] vibrations.

So everything is clear between the star and where you are. All the wavefronts should look planar. But if there is some disturbance because of hot air and so on, this will be distorted. So let's say this one is-- there's hot air here. There's nothing over here, which is conceptually the same as putting a piece of glass here and nothing over it.

So here, the light continues as it is. But over here, the light slows down and takes slightly longer to come up. And so the wavefront will look something like this.

And as you go further, they all start to connect. Now you can imagine there's another pocket of hot air somewhere here. And then this guy will slow down. So you'll start looking something like that, like so. So the wavefronts is bent.

Now the way it's thought about in science and imaging is that if I want to take a picture of this star which is far away and the wavefronts are bent like this, the image that I will get will not be sharp or actually will be blurry. And this happens to you, right, if you are just looking at hot air or a foggy scene-- OK, not foggy scene, but just, if you look across the river to Boston, even on a clear day because the temperature variations on the water versus land, the lights will not be sharp. They will be blurred.

And actually, over time, they seem to shift a little bit [INAUDIBLE]. And that's happening because at one instant, it's bent this way. Some other instant, it's bent the other way because the pocket is changing. The pocket of air is changing up and down.

So this is a big problem for telescopes because they would like to see this very clearly. And this is an example where we know that we should be looking at something that's very far away and really doesn't [INAUDIBLE] because the [INAUDIBLE] is all black here, and there is some galaxy or some [INAUDIBLE] over here very, very far.

So what I would like to do is take all these guys. And if I somehow can by a different method figure out what this disturbance is, I want to create a mechanism so that when it comes out, it's again back to nice and clean. So before that, I should say that in a normal situation, in a friendly situation, from a same point, light goes up.

It is parallel. And then remember the certain piece of glass which is thickest, right? So it slows down quite a bit. And here, the glass is very thin. It goes pretty fast.

All right, so what comes in as a planar wave actually just starts becoming a concave because remember, this one slowed down the most and this one slowed down the least. And this will converge down to one. So this is kind of the real propagation way of thinking about how the image is formed.

And the ray representation will be our lens, our point. They go parallel. It's a prism. And each prism, the middle piece of this glass, it goes straight through.

This one is a prism with this light film. So the ray is bent a little bit. This one is basically prism [INAUDIBLE]. And eventually, it forms [INAUDIBLE]. But then [INAUDIBLE].

And by showing these two examples, what we realize is that we don't think about how different parts of the lens impacts the incoming image. So going back to this example where you want to deal with looking at something that's very far away, then I can-- so we mentioned this distorted wavefront here. I can first reflect it off of some mirror.

And what I'm going to do is because the shape of this mirror is going to be exactly opposite of this particular thing. So I will do something like that. And when that was true, even though it was spent this way here, when it goes out, it will [INAUDIBLE]. And then I can put a lens and capture the image.

So that's how it's done in astronomy. And this deformable mirror is deforming at thousands of Hertz. And how do they calculate what the perturbation are? Anybody knows?

They actually shoot a laser which ionizes at a particular height. So they basically create a virtual star. Sometimes they call it pilot star.

And they take a picture of that. And they know that star is supposed to look like a point. So any change in the appearance of that star is an indication of how the air between the telescope and the stratosphere is impacting the incoming light. And so you can use that mechanism to correct for-- so that pilot star basically acts as a calibration, and they will feed that signal to the deformable mirror.

And it will correct that incoming wavefront. So if we have the same thing, you can go out in hot air where you have those mirages on a street, on a highway, or over water, and you can correct for it in real time. But this is really, really expensive.

AUDIENCE: To make a window, [INAUDIBLE].

RAMESH So you want to correct for it, or you want to create an illusion that it's hot air outside?

RASKAR:

AUDIENCE: Illusion, yeah.

RAMESH Yeah, I think you can. Yeah, that's pretty easy. It'll be a very interesting effect. In fact, what you can do is you

RASKAR: can take a piece of glass. And I mean, the one key property of mirage is that it's not just that you have an inversion. Everybody's familiar with mirage here where you have-- it's Walter [INAUDIBLE].

AUDIENCE: I think it's John.

RAMESH John, sorry. You like also with this?

RASKAR:

AUDIENCE: Yeah.

RAMESH OK. Just tell me, and I'll switch over to the other side. In case of mirage, you have the highway. You're driving

RASKAR: over here. I'm sitting here.

And so the ray is going straight. The [INAUDIBLE]. OK, and sometimes, you basically create the lens [INAUDIBLE].

And what you see in the picture is that you have the road that's going towards you. Then you see the blue sky, right? Let's look back here. And then you see the car that's coming from the other side.

So you have this inversion. And that's very similar with this inversion because the air near the road is very hot. It's [INAUDIBLE].

AUDIENCE: [INAUDIBLE]

RAMESH Sorry?

RASKAR:

AUDIENCE: Is that because of the Bragg scattering [INAUDIBLE]?

RAMESH The Bragg scattering? Yeah, that reflects as well. The [INAUDIBLE] reflection and so on because again, changing

RASKAR: the [INAUDIBLE].

Very minor change spread over kilometers. So anyway, the reason why I'm bringing this up is the way that's computed-- now you shoot a pilot star, and you want to figure out how the air is filtered. How do they do that?

AUDIENCE: Ramesh.

RAMESH Yes.

RASKAR:

AUDIENCE: How are they going to image those pilot stars? So say you shoot them with a laser, right?

RAMESH Yes.

RASKAR:

AUDIENCE: It should [INAUDIBLE] so that they could see it.

RAMESH No, it ionizes the air, and you just see this bright spot. [INAUDIBLE].

RASKAR:

AUDIENCE: And then compensating only for the actual spherical variations.

RAMESH Exactly.

RASKAR:

AUDIENCE: Because there are other deviations due to mass and [INAUDIBLE].

RAMESH It's changing very rapidly. Every second, the air is changing. So only if you compensate for it once, that's not good enough because it's changing every.

RASKAR:

OK, so now we have the pilot star. How do you figure out how to deform those mirrors? They capture a light field.

But this is not how they explain it. They explain it using the notion of [INAUDIBLE] wavefront sensor. That's the model they use. They are really expensive because they're very high quality. All they're doing is this wavefront is coming, and they have a lens very similar to the one we saw from [INAUDIBLE] at Stanford.

And the image that's found here actually tells you how the light is bending. So let's go back to our picture here. We know that if you have a lens and the wavefront is traveling in parallel, the wave image form is--

AUDIENCE: Focal.

RAMESH Focal. Exactly with focal.

RASKAR:

If the wavefront shifts, where they will form? [INAUDIBLE] formation, right? So a particular-- This is A. This is B. The image of A will be formed here. B will be formed.

So just again in wave optics, if the light is coming from here, the image will come here. If the light coming from here, the image will come here. [INAUDIBLE].

Now what's happening over here is we don't have this one lens, but we have a set of lenses. So we will have lens, lens, lens, lens, and so on. If you were looking at a very clear sky up on the pilot star, the wavefronts will compound.

And all the images will form exactly at the same time. But you mentioned because of perturbation, you have, I think, a very simple example. You have something that's going straight and something the same. That is one example.

Now what's going to happen is for this path, when it hits these lenses, it [INAUDIBLE] made at the same location. But for this particular one, which is tilted, the image will be formed slightly off center. So I can look up this image, right?

Originally, the point was in the middle, middle, middle, middle. And now after the aberration, the actual point will be shifted. So these two guys are fine, but these two guys, they've some shifted in the middle.

And that tells you how the waves are bending. You just feed this differentiable signal. You feel the mirror, and the mirror will [INAUDIBLE].

It's as simple as that. But this is also [INAUDIBLE] because you can think of this as a perturbation of this. So in this case, the waves are traveling in a straight line. So the optical axis.

And in this case, the waves were traveling [INAUDIBLE] like this. And so that's why these guys get here. They go slightly off.

And this guy goes [INAUDIBLE]. So the notion of light field is also very compatible with the notion of wavefronts. There are a few more details, but again, keep it simple.

It's also used in ophthalmology to look at any aberrations in the lens of the eye. Same exact idea. You have-- let's see.

You had the retina here. You could somehow assign the beam and create a bright, small spot. If everything is well, the light will go through it, and it will create this set of dots in the middle. If there is an aberration in the lens, this wavefront will bend, and the points would not be at the center, but they will be offset [INAUDIBLE]. So in a single snapshot, I can tell you what the variation of the string is.

There's one critical element here that I'm not talking about, which is when you get this curve out, you don't really get the whole thickness or the depth of this curvature, but some piece of information is missing. What is that? Perfect.

AUDIENCE: [INAUDIBLE]

RAMESH The phase or-- you get the phase, but you only get--
RASKAR:

AUDIENCE: It's 2π , right?

RAMESH Exactly. It's modular 2π . It's modular 2π , which means that if you-- let's say I put a piece of glass here, and the
RASKAR: weapon goes here like that.

And if I put-- so let's say there is a piece of glass floating down exactly by one wavelength. I can put a piece of glass with twice the thickness. And light be delayed by two wavelengths.

That's why this is concern. The phase delay is computed only by more [INAUDIBLE]. So there's a phase wrapping that's going on. So if you have a shape like that, if you have a shape, let's talk about it.

So this picture here is actually showing that it's phase wrapped. So it might look like a lot of discontinuities. But basically, it starts from the edge, and there's zero difference. And there's-- this curvature here is bent.

And then it seems to be jumping from, say, here to here. But actually, it's a continuous surface. It's just phase wrapped.

And this concept is used in many other fields. If you're familiar with synthetic aperture radar or arranging. You get some values that are [INAUDIBLE] specific things. So here we have parallels that you see on top of transparencies.

If you have a lens like that, disrupting a very thick lens, you can create a very thin lens which basically helps the slope. And you take this slope right here. You take this slope right here. You take that slope, put that here. The middle one is just-- start with us and so on.

And this one and this one-- again, in central cases, [INAUDIBLE]. So the flat magnifying glass is [INAUDIBLE] or the large [INAUDIBLE] on top of the [INAUDIBLE] projector. They all look [INAUDIBLE]. It's always at the edges like [INAUDIBLE].

So here, there's a small surface. They're not actually-- this is just conceptual models. But if display [INAUDIBLE], it doesn't really matter. It's still that density. But here, the ratings of the boundaries [INAUDIBLE].

There are some issues there. But a lot of those issues based on this [INAUDIBLE]. And here, what you're [INAUDIBLE] are saying I'm just going to chop this glass into multiple slabs, and I'm going to keep one in the last slab. So I'm going to keep this one, this one, and this one, this one, this one, and then I'll go to the last slab, this slab. That works the traditional lens.

[INAUDIBLE]. And similar concept is also applicable phase for [INAUDIBLE]. So then your doctor can see that in a perfect eye, all the points are in the center. But if there are aberrations, the points will shift on the center, and all you have to do mathematically is this direction, so keeping with the gradient, 2D gradient.

So you have to do a 2D integration to recover the surface. You have to solve the Poisson equation. And for those of you who are familiar on how to go from 2D gradients to 2D height fields, it's a straightforward process.

AUDIENCE: When they test your eye, which test is this? Is the one where they scan the light on the back of your eye?

RAMESH I don't have glasses. And I don't have contacts. Anyone else?

RASKAR:

AUDIENCE: [INAUDIBLE]

RAMESH Right.

RASKAR:

AUDIENCE: [INAUDIBLE]

RAMESH Exactly, yeah. That's a very good point. That's a really good point. So all these discussions we are having, where
RASKAR: we have on both so far, is for one specific wavelength.

So let's say it's a red color or black. It's linked to this particular type of distortion. But it had a green light, blue light, [INAUDIBLE], then the distortions are a function of the refracted areas. The refractive index is a function of the wavelength.s

So the distortions are also a function of the-- and so the aberrations that you will see or focusing mechanism you'll see will be different at different wavelengths. And we'll go into that a little bit later when we talk about [INAUDIBLE].

So that was a bit of a small detour of how waves and rays are connected and how the concept of light field is used in many areas. OK, so let's start having some fun with this sequence of images.

So we saw last time that you can see through occluders. And this is what you are going to do for your assignment. And think a little bit about how these different rays are allowing us to achieve these effects, OK?

Now some of the slides will have some math and some geometry. So those of you who are interested, stay awake. Those of you who are not interested, you can ignore these slides mostly.

But for those of you who really want to think about light rays, this is an extremely powerful concept. And when you visualize light field, you're talking about position and angle. Some people call it phase space. In Fourier objects, it's space and spatial frequency or spectrogram.

And all it's saying is that again, if you back to the ray representation flat line, there's position, and there's angle. So I can take this green ray and plot it in a position and angle space. Position is black, like so. There's x here and p there.

I'm going to take this green ray, find its x-position, which is x_1 . And it's at a particular angle, θ_i . And I'm going to place that. There may be another way that starts from the same point but at a different angle-- so the same position for a different angle.

What about this one here? This one has different position for the same angle as the green ray. So x_2 .

It's a different position and an angle, the same as the red. So it's the same as the red ray. Is this diagram very clear/ Because we keep coming back to this, this particular diagram.

OK, now something-- we can start thinking about how life will propagate. And if we're doing the first assignment, the virtual optical bench, this is what you'll be doing. And once you see this couple of slides, it will be crystal clear how this is done. Now let's see this ray as propagating from this plane to this plane.

If it travels by a distance of z , I can write down the new coordinate as simply the original coordinate plus the angle times the distance. Typically, you would have I guess a [INAUDIBLE]. But a small angle [INAUDIBLE]. And then the new x-coordinate is just old x-coordinate [INAUDIBLE].

So how can you represent that over here? So if you look at the original points, OK, I'm going to take the green gray and change its x-position. But its theta position is not changing because it's still going in the same direction.

So its position is x_1 prime. But the angle is still theta. So I will take this green guy and just shift it to the right to represent this critical path. I just want to shift it from here to here.

And these other guys will shift as well. They'll shift to the right if they're above this plane, and they'll shift to the left if they're below this plane. So you create a shear of the x theta representation.

And this is unnecessary. Why are we thinking about this dual space when we can just think about the primal space of the real world? But again, it becomes extremely easy to analyze it in this dual space of x theta. Is this clear?

AUDIENCE: I didn't get why the others were affected if you were only moving the green ray [INAUDIBLE].

RAMESH
RASKAR: So if I take now this particular ray, think about what will x_2 prime for that will be. So we started with x_2 . Started with x_2 and theta g.

And now we want to find the new x_2 prime. It's going to be x_2 plus similar equation-- z times theta g. So it also has no sheer.

And depending on how far up you are here, the sheer moving on. So this means that if the ray was actually parallel to the optimal axis and perpendicular to this plane, x will not change and theta will not change. So if you are lucky enough to be on this equator here on the optical axis, nothing will change, really. Does that make sense to everyone?

All right, so let's do a couple of quick exercises to make sure we are all on this. If there's a point and a light source-- and here's an example where the like is very simple. I have a bulb, LED, and it's emitting light in all directions. How can I represent that index theta space?

So there's a particular x. I'm going to put it somewhere here. And it's going to span all the thetas.

AUDIENCE: The vertical line?

RAMESH So it'll be a vertical line, very good. Now if I take, if I let it propagate-- or let's say I have a new light for here.

RASKAR: Where the light is going very far away and all the rays are coming at 10 degree angles with respect to the optical axis, what will this say?

AUDIENCE: [INAUDIBLE]

RAMESH
RASKAR: A horizontal line because they're spanning all the x-coordinate but only with the theta of 10 degrees. That's what we're going to do. All right, so they're two very simple examples-- point that's very close, vertical line-- point that's very far away.

Think about a star-- horizontal line. This is why you think there's two. We'll be able to build a lot of machinery to understand how we can create some truly amazing effects for using light fields. OK.