MITOCW | mitmas_531f09_lec04_2.mp4

PROFESSOR: All right, let's get started. So I think what's interesting about light is that it's used in so many interesting ways, whether it's programmed, whether it's not programmed, how it interacts with the world, [INAUDIBLE] there's direct bounce, multiple bounces, different wavelength, modulation, time, and space, it's a lot of fun.

Like, for example, do you know how a TV remote control works?

AUDIENCE: IR pulses.

PROFESSOR: IR pulses. It's mostly optical, the output of RF-- but the LED of the remote is sending a code, optical code basically-- thanks-- over time to the photosensor on the TV. Now, why does it work in broad daylight?

AUDIENCE: [INAUDIBLE] spectrum.

PROFESSOR: It's a different spectrum. That's one benefit. That's not enough.

AUDIENCE: Just looking at the differences between light, like peaks--

[INTERPOSING VOICES]

AUDIENCE: --actual--

PROFESSOR: [LAUGHS] Actually, you can shine on the ceiling, and this will work.

AUDIENCE: Just the variation in time.

AUDIENCE: It has time. So it has the last--

AUDIENCE: Yeah, looking at peaks [? currently. ?]

[INTERPOSING VOICES]

PROFESSOR: Sorry, the last one?

AUDIENCE: It has less energy than [INAUDIBLE] difference of pulse or something.

PROFESSOR: Almost. What else [INAUDIBLE]?

AUDIENCE: The modulation.

AUDIENCE: Then the filter--

PROFESSOR: It's using modulation. So it's actually running at 40 kilohertz. So when it's 1, you're sending something 40 kilohertz, when it's 0, it's not sending anything. And so the AC component, which the carrier, is 40 kilohertz. And then the signal is 1 or 0. So it can record in presence of ambient light, because ambient light is mostly DC.

AUDIENCE: But it can't-- well, so yesterday, I was watching TV, and all of a sudden my remote-- or the cable box stopped working. And I thought it was [INAUDIBLE] one of those, and there was in the [? wrong mode ?] or-- nothing worked.

And then I noticed the TV light was blinking. And noticed-- I was sitting on the other-- at the TV remote, and it was just blasting IR light. But the cable box was just ignoring it, like it couldn't figure out--

PROFESSOR: Because you were pressing too many keys.

AUDIENCE: I was holding one button on the other remote without knowing it. And--

PROFESSOR: So two different remotes were conflicting with each other.

AUDIENCE: --remotes, yes.

PROFESSOR: Exactly. So it's just a simple principle that we always associate with the TV remote. But can that also be used for photography or imaging? So the signal peaks-- the photodetector on the TV is decoding the signal. But that's basically a single pixel.

Imagine if every pixel in the camera was made out of that photodetector that's decoding the 40-kilohertz signal. OK.

AUDIENCE: What about like a-- like in your ear, you have hairs that vibrate at different frequencies, do you have--

PROFESSOR: What's the analogy there?

AUDIENCE: Oh. Like so, if you had a pixels that were listening for light--

PROFESSOR: Uh-huh, at a particular wavelength, right?

AUDIENCE: Like listening for [INAUDIBLE] light.

PROFESSOR: So imagine you're trying to build a camera. So right now, I have exactly one pixel, which gets a signal that comes in at particular hertz, and then 0 [? at ?] 40 kilohertz, and then 0, and so on.

So thinking in terms of communication, we have a carrier, and we have a signal from that. The 0 [? with ?] 40 kilohertz, [INAUDIBLE] amplitude modulation and [INAUDIBLE] signal around that. That's how you are thinking in communication. And in case of a remote control, you send 40 kilohertz, nothing, 40 kilohertz. That's as simple as that-- a very simple signal.

Now on your-- so you have your [INAUDIBLE]. If you use a [INAUDIBLE] image [INAUDIBLE] instead of one pixel, imagine every pixel in a camera is able to decode the signal. So instead of taking the 40-kilohertz signal as a reference carrier, and it all becomes just 1, 0, 1, 0. That's all [INAUDIBLE].

Now imagine if I could put a camera so that every pixel has that property. So I'm going to put a camera. The [INAUDIBLE] pixel here can decode 40 kilohertz and just pick up what is [? illuminated ?] at 40 kilohertz and ignore what's in the room. So in a typical room like this, [INAUDIBLE] you know that this is sunlight. So sunlight, you know, a huge DC, and then a TV remote is giving a little bit of signal.

And then all the photodetector does is it just plants it-- that's just the frequency selection [INAUDIBLE] here. And this is the signal, and it ignores all the DC. Now, can I create a camera where every pixel becomes the same as that, and now I can shine the room with my remote so that whole scene is being [? flooded ?] by 40 kilohertz, and in bright daylight this scene will appear as if it was lit only by this flashing LED and nothing else?

Is that clear?

AUDIENCE: So there's a lens in here, right?

PROFESSOR: Sorry?

AUDIENCE: There's a lens in this case?

PROFESSOR: Yeah, there's a lens and all that. It's just a typical camera with a sensor and so on. This point is being focused here and so on. It's the same thing. It shows that the light I'm having here is [INAUDIBLE] 40 kilohertz and [INAUDIBLE] 40 kilohertz and so on.

AUDIENCE: But can you build cameras which operate at 40 hertz, like that [INAUDIBLE] per second?

PROFESSOR: [INAUDIBLE] you could. So we're not going to get cameras that look like this. It will happen as the silicon improves and so on. Of course, there's always people that [INAUDIBLE] 40 kilohertz.

Now imagine, somebody gives me a flashlight that actually runs at [? 2 ?] kilohertz. And this one runs at 40 kilohertz. And this particular pixel actually captures the signal across 14 and 15, and in software, it can decide what's the [INAUDIBLE] at 14 and what's the [INAUDIBLE] at 15.

What did you just [? look ?] here compared to assignment number 1? Two flashlights on at the same time. I want to know how the signals [INAUDIBLE]. This one and this one. So this is A and this is B. And the image I'm getting is A plus B.

But in software, I can decompose and say which part of the imaging-- which intensity came because of A and which intensity came because of B. So software, I can tune between this light source and that light source, just like I can-- on your car radio, you can tune between 99 megahertz station and the 80 megahertz station. So we will tune that on the camera.

And once we have that, imagine cinematography. You can put all kinds of lights through the movie, and then go in Photoshop and change any light, any color, any intensity. Again, beautiful light [INAUDIBLE].

AUDIENCE: But again, lots of data [INAUDIBLE].

PROFESSOR: Yeah, but it doesn't [INAUDIBLE].

[LAUGHTER]

[INTERPOSING VOICES]

And they'll be happy to clear [INAUDIBLE].

[LAUGHTER]

So there's a lot more to come. So every time you think about how light interacts with the world, say, how can I use that for imaging?

AUDIENCE: Is it crazy-- is that kind of how sonar works, or are they just like [INAUDIBLE]?

PROFESSOR: What's weird is I can take sound and create images like--

AUDIENCE: You meant sonar.

PROFESSOR: Sonar.

AUDIENCE: Boop. Yeah, yeah, yeah.

[INTERPOSING VOICES]

AUDIENCE: Or LiDAR.

PROFESSOR: LiDAR-- yeah, all those methods are basically using the principles for [INAUDIBLE]. It starts, it bounces, it has certain properties in terms of presence or [? absence, ?] position, color, space modulation, time modulation, and all these things.

AUDIENCE: Just a little [? silly ?] kind of question-- who was the first person to do computational photography?

PROFESSOR: Steve Mann right here [INAUDIBLE]. Steve Mann and Ros Picard were the first one to use the term "computational photography," although they used it in a very specific context for high dynamic [? range ?] imaging. And then later on, very limited-- very important people, pioneers in the field such as Shree Nayar, and Marc Levoy, and so on. They were [INAUDIBLE] even before the term was around.

Actually, when you look at all these papers and presentations, I would say over half of them are just because of those two guys.

AUDIENCE: Because when you talk about this, it's also when NASA explores planets, you also think of [INAUDIBLE], what does it mean, [INAUDIBLE].

PROFESSOR: Exactly. I mean, what we're talking about is really communication concepts. It's 100 years old. So a lot of concepts kind of get borrowed. And [INAUDIBLE], you couldn't think of decoding a particular hertz signal, right? And we have moved to digital only what, 15, 20 years ago?

And astronomy, all that math and all the techniques that we use in communication become possible in our world. So this is kind of a-- because at the same time, when you are in the communication world, the signals don't have very high dimension. It's usually a two-dimensional signal, and a number of stations, and a frequency range-- basically a two-dimensional signal. Every trans-- every radio station is transmitting from audio. And audio is one dimensional.

And since it's a two-dimensional signal that's in our world, and we capture that on our antenna as a one-dimensional signal over time, and we decode that and record back the [INAUDIBLE]. So usually it's not very high dimensional. And even if it had high dimension, they're multi-scale. So just if I'm sending 500 channels on a fiber, that's just 500 separate signals. They are not intermixed like we have here.

So although [INAUDIBLE] similar, the problems in imaging are more complex. The high dimensional have [INAUDIBLE] problems [INAUDIBLE]. But research is all about fusion of the similar. So if you can learn ideas from communication, and optics, and quantum computing, and signal processing, you put all that together and mix, and you all this-- you can create magic.

And almost every project we see has some element of magic. And that's what makes it very exciting.

AUDIENCE: So 40 kilohertz seems quite fast. But it seems like maybe we don't actually need to be quite that fast.

PROFESSOR: Yeah, yeah.

AUDIENCE: Yeah. So I'm just trying to think, what would the-- how slow could you go and still, basically, eliminate the DC component? I mean, it's part of the fluorescent lights that are 60 hertz or--

PROFESSOR: Yeah. I mean, [INAUDIBLE] like, what, 25 kilohertz to remove the flicker? But yeah, you could use-- I mean, if the camera's 60 hertz, you could just use a 60-hertz strobe and turn it on in one frame and off-- on every odd frame and even-- off in every even frame. And that alone will allow you to do this subtraction.

So the only problem is that if you do a pure subtraction, you're going to subtract really two large quantities, two large numbers. So in the first image, you have sun [INAUDIBLE] [? mesh. ?] And the second image is just the sun. And this is very, very small compared to this. You're subtracting two large numbers and expecting to recover the contribution because of the [? mesh. ?] Is that [INAUDIBLE]?

AUDIENCE: Sounds like an error accumulation problem.

PROFESSOR: Yeah. But that's [? exactly ?] the problem, communication [INAUDIBLE]. The carriers and the signal is so tiny that's riding in free space over large distances that they use really clever coding mechanisms so that your imaging [INAUDIBLE] increase [INAUDIBLE].

So I just want you to think very broad. I know many of you here have very interesting backgrounds in communication, and chemistry, and interaction, and so on. So try to make the best of that. So temporal modulation, actually, is not used that effectively right now in imaging.

So [? certain ?] projects, I'm not going to go into detail, but they're on that wiki that I sent you. So please add more information there. Add your own experiences, some of the things you are mentioning, some of the projects you're mentioning. Please go and add all those things to those wikis.

All right. So sometimes, you can't control the illumination, but you can just exploit natural illumination. OK. So here's a project from Washington University, St. Louis. And what they did was they took webcam images all day long at a given time of the day.

So on the x-axis, you have time of day. So it's dark in the night, then daytime, and again dark. And on the y-axis, you have day of the year. So those are how many-- yeah, I guess just day of the year. I don't know after many days each was calculated. If the top is 1st of January and bottom is 31st of December, what can you say from this data set?

AUDIENCE: Winter has shorter days.

PROFESSOR: Winter has shorter days, which means where is this camera?

AUDIENCE: In the northern hemisphere.

PROFESSOR: It's in the northern hemisphere, right? And you can probably say more about if we just had the ratio of the smallest day to the largest-- longest day, that will tell you the latitude, because when you're on the equator, the longest and shortest days have equal length. But as you go away from that-- there are already a lot of data embedded in this natural illumination.

So this project is really beautiful. They did all kinds of interesting things. So they have hundreds of static-- thousands of static cameras, variation over the year, over a day. They put all that together. They can do really interesting things.

So it turns out in a traditional lighting, in a typical scenario, light is linear. What does it mean? It means that if I have a scene, I light it with particular brightness, a particular intensity of light, I get certain brightness. If I make my light twice as bright, everything will become twice as bright. As simple as that.

This is not true at all the intensities of light. When you go really, really bright light, it's not true. The world starts behaving in a nonlinear fashion. If you have your speaker on your synthesizer, if you pump twice the power to your speaker, do you always get twice the loudness?

AUDIENCE: It only increases 1 decibel.

AUDIENCE: It begins to saturate.

PROFESSOR: It tends to saturate. And eventually, you'll run into nonlinear behavior. And the same thing is true for light as well.

But as far as sunlight is concerned and the type of world we are involved, everything is linear, so we don't have to worry about it. And because everything is linear, mathematically it can all be expressed as just linear transforms, and linear algebra, and so on. That's why background in linear algebra is very useful when you're doing any imaging work.

So they did some very simple things. Like they took all these images, just did a PCA, Principal Component Analysis. And that image allows them to figure out the haze, and cloud, and the orientation of the surfaces. So this essentially lists-- and I believe they can figure out that this building is facing one way versus this building, and so on, just without even analyzing and doing any sophisticated computer vision, just from the sequence of images.

And then they can segment the scene. This is something close. And with distance very far away, they can encode that. And they can even figure out where a webcam is, its latitude and longitude.

And Robert [INAUDIBLE] told me that they can do-- just based on the sunrise and sunset data set that we saw earlier, they can localize with 50-mile accuracy. And if you have some speed cameras where you know the locations, then you can interpolate and go down to about 25 miles.

And in addition, if you have satellite imagery, so you know how the intensity is changing, then you can do around 15 miles. And then the people at CMU such as Srinivasa Narasimhan and Alyosha Efros, they recently did a paper where they can just look at a patch in the sky. And if you look at a patch in the sky in broad daylight, it always has a gradient. And depending on where the sun is, the gradient has a particular orientation in x or y, the intensity ramp. And that actually localizes the direction of the sun.

So now they can look at webcam images and click on the part that shows the sky, and they can localize the cameras down to, again, a few tens of miles. I forget exactly what the numbers are, but pretty fascinating. And they're not even using polarization. If you use polarization, it can get even better, because the sky is highly polarized.

AUDIENCE: I have a question about it, actually. Do normal digital camera sensors or film or anything, do they have any polarization dependence at all?

PROFESSOR: An ordinary sensor doesn't. But you can always put a polarization--

AUDIENCE: You can always put a polarizing filter, but there's really no correlation [INAUDIBLE]?

PROFESSOR: As far as I know. Yeah. Even the human eye does not have very strong sensitivity to the polarization. But there are some results-- and if you talk to Matt Hirsch, he claims he knows-- he can see polarization. He even has experiments where if you see one way, you see one color, and you the other way you see a different color.

He's shown it to me dozens of times, but I don't see the difference. But he's been able to recruit a lot of people to say yes, they see it. And there are very few actually animal eyes can sense polarization. There are some underwater creatures that can do a pretty good job, for instance.

So again, they can do the encoding of that, how far other things are, orientation of surfaces. So here, you can see that orientation-- this is different from orientation. That-- how would you figure that out, by the way?

AUDIENCE: Shadows.

PROFESSOR: Shadows and sunlight, because some faces will be lit [INAUDIBLE] than others depending on time of day. So you don't have to process it in an individual manner. You just throw it in a big matrix [INAUDIBLE] PCA, and [INAUDIBLE].

OK. So let me-- we saw this example last time. So I'll skip that. Let me switch to light fields and talk about our assignment. All right. So light fields. It's one of the most important concepts we're going to learn in this class. And again, realizing that the appearance of the world is higher dimensional, not two dimensional. You have a 3D world. You project it on a 2D image, clearly a lot of information is lost.

Now, if you build a so-called plenoptic function, which is-- what is set of all things we can ever see? It was a name, actually, given by Ted Adelson, a professor here in the early '90s. Then it turns out it's a very high-dimensional world. If I stay in one place and think about the bubble around me, I have the azimuth and elevation of every direction, just the bubble-- on Google Street Map, you have a bubble for every location. That's [? three ?] times [? too. ?]

And that's over time and over wavelength, different colors and over time. So that's four dimensional. Now, I can put these bubbles in different places. And every bubble can be placed in x, y, z. So there are three additional degrees of freedom.

And if you can capture all that information, then you can recreate a movie from any viewpoint at any time at any wavelength. But it's extremely high dimensional. This is seven dimensional. So the world is actually seven dimensional. And if somebody built this magical device, it will have-- if we make a [INAUDIBLE].

Now we're going to simplify that. And let's say, OK, for all these bubbles shown in the blue, all the rays are emanating, and if I think of any point in the world-- and for now, we're going to ignore time and wavelength-- it becomes five dimensional, from seven to five, because we ignore time and wavelength. I can take a point in 3D. And from that point in 3D, I can think of a direction.

And the direction is only 2D, not 3D. Why is that? x, y, z for position--

AUDIENCE: Theta [INAUDIBLE]

PROFESSOR: But only theta phi for angle. Why is it not three dimensional?

AUDIENCE: What would you use the third dimension for?

PROFESSOR: Because the roll along the ray does not really matter. So you have yaw, pitch, and roll, but the roll can be ignored, because the intensity is the same even if you have roll. So it's only five dimensional.

But then if you have an occluder here, then the intensity of this ray is different from intensity of this ray. If you have no occluder, then the ray intensity remains the same. So now, actually you can go down to just four images rather than five. So the space of all lines in 3D is actually four dimensional if you want to express all the rays, then it's just four dimensional-- ax plus b y plus cz plus [? p. ?] There's four [INAUDIBLE].

Now, you can simplify that further for the camera world, where we're going to assign the plane of a sensor, and the plane of the lens, and so on. So that's what we'll see very briefly. So let's say there's a light field in this room. Rays are traveling from light sources, bouncing around everywhere.

If I just cut a plane in midair, I can parameterize that plane as x- and y-coordinates. And for every point on the plane, I have, again, the theta and phi. So this becomes four dimensional. And that's what we're showing here. The positions is s and the direction is theta.

So often, we will think about flat land. So we'll just think about the plane of the screen as opposed to the 3D world. So in the 3D world, we have x, y, and theta phi. But in flat land, we have just the position and angle. So it's just two dimensional.

So this is called a-- so that was one-plane parameterization, where you had position and angle. And another common way to think about that is-- another common way to parameterize the light field is a two-plane parameterization, where you have one plane that has position, and the second plane that, again, has position, and a ray that connects those two, again, represents the ray space. The coordinates for that represent the ray space. So this is so far two-plane parameterization. And this is very commonly used in computational camera and photography.

So let me jump ahead a little bit because of the time left and explain how we're going to do it for our assignment. So remember, we're going to create an effect where we'll put a whole bunch of cameras or take an array of cameras like this and be able to see 12 [? rulers. ?] And the effect is relatively straightforward.

And we go to the so-called synthetic aperture photography. We're going to create an artificial aperture to be able to see through [INAUDIBLE]. So if you have a point [INAUDIBLE] focus versus a point that's out of focus, the green point will create a very bright spot. The red point will create a blurred spot. That means its intensity will be correspondingly reduced per pixel.

And if you stopped on the picture, what will happen is that the green spot will become slightly dimmer, because less light is reaching the sensor. But the red spot will also focus-- also blur in a smaller region. If you go in the opposite direction and have a really, really large aperture, then the green spot will be very bright, because it's capturing-- a lot of light is being captured-- and that will be over here-- but the red spot will be highly blurred.

Now, building such a large aperture is very challenging. So what we're going to do is create that using an array of cameras, like this. And it's the same as synthetic aperture radar, where they use an array of antennas to create, effectively, a much larger antenna. Again, analogies with communication, and RF, and so on.

So again, the whole point that-- so again, we're going to subdivide this lens into multiple apertures as opposed to one large aperture. And that will be effectively created with a set of cameras like this. And then if you sum the images from each of those apertures, that's the same as creating an image with this very large lens. And for a different point in 3D, will correspond and create a different image. That's it. [INAUDIBLE].

OK. So how does this work? How are we going to create an effect where something that's out of focus effectively is going to be completely blurred? And we saw that if-- even with aperture of my eye, which is only about 6 millimeters or so, if I put an object really close, then I can see the world through this eye so that this is basically-- it doesn't impact. If I put a needle in front of me, it gets completely blurred.

And that's the same effect we're going to see. So you take an array of cameras or camera at different positions, take-- collect, say, 25 photos. If you simply take those 25 photos and sum them up, what will happen? So if I just take a camera-- and for simplicity, we're just going to [INAUDIBLE].

If I have, say, five cameras this way, and I will point at the [INAUDIBLE], this coordinate in each of these cameras is going to be added [INAUDIBLE], because it's [INAUDIBLE]. It's like when you're driving and you're looking at the moon, it always appears to be the same position.

If you're looking at a [INAUDIBLE] which is very, very far away, its coordinate in the camera is going to be the same [INAUDIBLE]. So if I just take these five photos and add them up, sum them up, I basically get the same exact [INAUDIBLE], because I [INAUDIBLE].

Now let me [INAUDIBLE] something that's nearby. This coordinate in the top camera will be the top of the point of infinity. But in the bottom camera, it will be below the point of infinity. So now you have summed these five images up, this point will be completely blurred, because the coordinates in each of the five images are very different. On the other hand, what's at infinity will be the [INAUDIBLE].

And that's exactly what's happening when you focus. You're basically taking an image from every part of the lens and summing it up [INAUDIBLE]. Here we're doing it in software.

And of course, mathematically we're not going to sum it up as is. We can shift each image and sum it. So if I wanted to focus here, what I would do is I would take this picture, keep it as is, and I'd take this picture. And I know that from here to here, this one is shifting left about 5 pixels. I'll shift this whole image by 5 pixels and then add it.

I'll take the next one. And I know that's going to be about 10 pixels [INAUDIBLE]. I'll shift the whole image by 10 pixels and then add it. And the bottom one, I have to shift by 20 pixels and then add it. If I do that, then this point will be sharp focus. And the point at infinity will be completely blurred.

And using this very simple shift and add mechanism, we're able to focus very close [INAUDIBLE]. Is this clear?

AUDIENCE: I have a quick question about this.

PROFESSOR: Yeah.

AUDIENCE: Does this set some minimum focus distance?

PROFESSOR: It could. It could, yes

AUDIENCE: Or what [INAUDIBLE] minimum focus distance does it need?

PROFESSOR: The field of view is set-- for example, if I get really close, then these cameras can see the sky, [INAUDIBLE]. So that's [INAUDIBLE].

AUDIENCE: So is that the only thing, though? Is it just the field of view that [INAUDIBLE] focal point?

PROFESSOR: Mainly the field of view. The resolution [INAUDIBLE] as well. But mainly the field of view. Otherwise, this technique can focus on [INAUDIBLE]. It can even focus beyond infinity. So first [INAUDIBLE].

[LAUGHTER]

Because you can-- you're sort of adding them up. I can add them up in the reverse direction, minus 5 pixels, and focusing at infinity beyond. So this is all we're going to do. But that's-- as you'll realize in your assignment, there are a few things [? you ?] have learned. Here, I just choose some numbers, and shift by 5 pixels, and add.

What will end up happening in your case is you'll realize that if you put these cameras and take the pictures and then figure out what distance this should be, what the projection of these points are going to be, and if you don't use some meaningful numbers, you'll never get [INAUDIBLE], because either your parallax, which is the distance-- the change in coordinates as you switch from one [INAUDIBLE] to the other-- too large or too small.

And it's really easy to do by just kind of eyeballing it, but eyeballing it not with camera, with your own eyes. So you can just stand in one place and see if you move by 10 centimeters, do you eventually see the point behind? And in the case of the Stanford project, that was a really challenging example of a set of trees and people behind them.

You don't have to choose something that complicated. You can choose some set of objects in the front and then some painting in the background. If you want to set up the scene, and you want to put some-- the best would be really just put a pencil first in the [INAUDIBLE] to a [? fence, ?] and then there's a painting in the background. And then if from any single camera this painting is occluded, by taking multiple photos, you can see--

I can do it on a table. You can do it outdoors, for example. There are some trees here. You can see through it. Choose your situation, and you'll be able to answer this. And there will be more instructions on the website.

AUDIENCE: Are we allowed to do it all computationally?

PROFESSOR: Pure software?

AUDIENCE: Yeah.

PROFESSOR: Just OpenGL, you mean? Absolutely, yeah. Yeah. But I mean, where's the fun that? You are perfectly welcome to do that, yeah.

AUDIENCE: What's the tolerance for the parallelism?

PROFESSOR: So you want to be as close to parallel as possible. But-- we are not going to discuss it here, but as you know that once you have-- for example, if you misalign this camera and collect your [INAUDIBLE], then you know that this image to this image is just a pure one model, a pure, single [INAUDIBLE]. So you could just fix that mathematically if you wanted to.

But you should just avoid that for this assignment. Try to keep it parallel. Just put it on a ruler and slide it. And I'll give-- there's more information on the website.