

ROARKE I'm a new student with the Camera Culture Group. But for a little over a year I worked at a company called Miter
HORSTMAYER: which, one of them's up the road. The other one's in DC.

And we on and off worked on this project, myself with a guy named Gary [INAUDIBLE] and then Mark Levoy who is from Stanford. And together, we eventually put together a nice project. But it took a bit.

So 2009, 2008, is kind of a general time frame we did this. But this is going along the lines of capturing multi-domain or multi-dimensional data spectral data or polarimetric information or high dynamic range with your camera but doing it in a single snapshot. So we'll discuss that in a little bit of detail.

But first, just some background. Just giving you an idea of what I mean when I say multi-domain. So we know there's a ton of pixels. Ramesh showed some slides of the increase in number of pixels per sensor and whatnot. And the thrust of this research was finding a new way to use those pixels besides just spatial resolution.

We think there's plenty of spatial resolution out there. What about color or some other interesting properties of light to capture it? So some of those include polarization.

So here a group used a polar emitter to distinguish a man-made object-- this is metal-- from some plants. The way it works is metal reflects more polarized light than a plant does, which actually reflects a little bit of polarized light. And they pick out objects like that. I'll show some examples some of that what we did later, I think.

Anyways, so you can see cool effects like [INAUDIBLE] in glasses or transparent media. And then, obviously, what's known for photographers is it helps-- I forget the word- but clarify haziness in images of water, because water reflects polarized light and so does the sky. So you get a clear blue sky with a polarizing filter for conventional photography.

You can also do multispectral imaging, which is just capturing a very specific spectral information instead of just red, green, and blue, which are huge bands in the frequency dimension. These are very narrow-band images.

These were both captured over time. So this one shows a aerial view. They're often used in airplanes to capture information about foliage or stuff. USGS is using them a lot.

And then this is an example of detecting blood oxygenation. So this is an image of a thumb next to another thumb. One has a rubber band wrapped around it. The other one doesn't. So you can obviously tell the difference.

And then high dynamic range imaging, which I think we talked about. So we're going to capture all that in a single snapshot. And basically what happens when you want to do that is you're presented with this huge dimensionality mismatch, which is true of all camera capture techniques.

You just have this two dimensional sensor. And it's small and it's flat. Soon, it might not be flat. But right now it is.

And you also have the time dimension to work over. But you have these seven or more, arguably, seven dimensions of information out there. You have a four-dimensional light field, which I guess you'll learn about more. Some of you might know about. He also has a temporal, the changing of whatever scene you're looking at. And you have spectral polarization information.

So there's a lot of different ways that people have come up with to capture this. And I'll talk about some. One of the most direct is just doing it over time. So this is a picture of one of the first color motion picture processes. It's called Kinema color.

What they did is they just put a rotating green and red filter in front of a regular old movie film camera. And it captured sequential images. Then when they projected it-- it didn't have blue for some reason. I don't know why.

But projected it, they just projected it through the same rotating color filter. And your eyes couldn't distinguish the fast frame rate. So it looked like a color code.

And then some of the things have been done with hyperspectral imaging. So the way I push through hyperspectral imager works is it's a two-dimensional sensor. One dimension is used to capture spatial information. And the other is used to capture spectral information.

And you put it up on a plane, and it flies. And it sweeps out the second spatial dimension. So you're capturing this three-dimensional data cube and over time, two dimensions. But you're capturing spectrum [INAUDIBLE]. It's a little bit different.

And then this is very similar to what RK was talking about. But there's tons of different ways to do that over spatial encoding. So you can work over time. You can also work over space.

The Bayer filter image, I think this is the exact image that you guys talked about. People have also done it with polarization filters. And then I could also talk about assorted pixels, which was from [INAUDIBLE] group. And we mix and match neutral density and polarization and spectral filtering, so-- all over a sensor.

The problem with that is it's extremely hard to fabricate. And once you make it, you're stuck with it. So you can't very easily take off your lens and peel off your filters and then put another one on it. Very sensitive to alignment and stuff.

So other people have tried capturing multispectral information or other types of information with just many cameras. So you guys saw the profusion camera, we had on the first day. This is a very simple, just similar design called periodic or [INAUDIBLE]. We also learned about where you just put different color filters over each little camera and you can get a single snapshot from all the cameras.

The problem with this is there's issues of registration and alignment with the cameras and it's expensive. That camera costs \$10,000. So it's a good method, but our method is a little bit different, I guess.

The third way, I think the coolest way, of doing this is with this method called co-division multiplexing. There's only a couple of examples of this. But basically, you computationally capture a three-dimensional data keyed on two dimensions and then computationally try and guess the third dimension.

So you're trading off in computational power and also error associated with your estimations. But some examples of this is this CASSI which was developed at Duke. And it's a single snapshot spectral imager similar to what we do. What they do is they put a coded mask in an intermediate image plane and it shifts the data cube in three dimensions. And they can estimate color, many color channels from a single image.

And a similar very famous example of that it's called CTIS. It stands for, I think, computed tomography infrared spectrometer. But it was invented in Arizona. And what they do is they use a very, very special, cool holographic plate. And with that you can take an image, and the holographic plate makes this very interesting spectral description, which they know exactly how it's being dispersed so they can reverse the spectral dispersion to get a full dimensional image but also have all this color information.

Problem with it is that the hologram they use is very sensitive to directionality of light. So it has a very, very narrow field of view. So basically they have a really cool movie of some guy lighting a lighter and you see all this spectral stuff. But that says why does the field of view is it's like one little lighter. But it's a really cool.

So just the basic idea of our approach-- I won't get into too many details. I'll go over this kind of quickly. But you have a camera, and you're capturing two spots on an object, on one red one. And at the center, the rays-- let's pick the red spot-- are both integrated at the sensor so you can't tell which ray came from where, any other part of the lens.

If you misfocus your sensor, you can tell, once again. So you know this ray is not integrated with that one. So you know it's coming from this spot in the lines, likewise for that one.

But when you misfocus something, everything just blurs together. There's no separation, spatially, of the rays. It's hard to distinguish them unless you're doing astronomical imaging where the stars are surrounded by black and you have a chance.

What we did is we put a pinhole array on top of the sensor. And that has a couple of interesting properties. It's very similar to a light field. I won't really get into it. But basically what it allows you to do is it allows you to take your Bayer filter that you originally had at the sensor [INAUDIBLE], and you just move it up into the lens.

So now instead of having a repeated Bayer filter pattern over your focal point, you just have one filter and you can stick it in your lens. And what each of the pinholes are doing is they're effectively imaging this aperture. So you can think of each of these little pinhole camera.

And each of them, all you can see is this filter. So you're creating a style of filter. But it's also still capturing the information about the object.

So for example, if you put a red filter in this part of the ray and a blue filter there, the red filter will attenuate the blue ray. So you'll figure out it's red there. But this is a grayscale sensor. Only one of the sensor areas will light up. The same with the [INAUDIBLE] area on the other pinhole.

So that's the general idea. But the trade-up with this is now you have a reduced spatial resolution. And your spatial resolution of your output image is going to be given by the number of pinholes you have in your pinhole there. And your filter resolution or whatever-- it's not really color resolution anymore, because we're going to put different types of filters in here-- it's going to mean even by the number of filters you stick in your filter.

And there's an interesting trade-off there, which I might have time to talk about. I won't really talk about the specifics. Basically, just since we have a digital image now on a film set up, this wouldn't work at all, because you can't really cut out the little pieces of different spectral areas and stitch them together. It'd be really difficult. You could reverse project it, which people have done with them.

But anyway, we just take different areas. So let's say-- and this is just a piece of an image, which I'll show you later. And each of these is not projecting well.

Each of these white boxes is the area where the filters are being projected. So there's a 3 by 3 filter at the center one, I think. It's a very non-dense filtering.

So let's say this is a green filter from this area. You can take all these green filters, stitch them together, and make the green image. So let's say this is the red filter here. You can do that in the red image.

And you just do that digitally. It's just like a lookup table. It's very straightforward.

And this is what our setup looked like. We just had a regular lens. It's a Nikon lens. It was \$100. And they're really easy to take apart.

So if any of you ever wants for whatever, your project for this class, to put stuff in the aperture of a lens, come to me. I can teach you how to do it in five minutes. And it'll take you 30 seconds each time to undo the lens and put it back together.

But we put color filters in there. You can't really see them that well. But we just put these little color filters discs. And then I'll show you some results for something I did this summer. I put a spectral filter in the aperture, which changes our design to a snapshot spectral imager.

And then here we have our [INAUDIBLE]. This is the CCD. You just sandwich a pinhole array right up against the CCD. And the covered glass over the CCD, which is like a millimeter thick, dictates the distance between the two.

So here are some example images. This is just a image, a raw data image. And here's a zoomed up version. You can see each little pinhole region close up.

And then you could do some [INAUDIBLE] division, which is just normalizing the image. So basically we wanted to use a lenslet array, which is just almost the same thing as a pinhole array. I could show you some slides from Stanford where they have the lenslet array.

That's just really expensive. And once you get it, it's fixed. Pinhole arrays you can print off, you can vary the distances between pinholes and the sizes of pinholes.

So it's a much easier way to check when you're doing an experiment. But they have imperfections, obviously, especially from printing. So that's why we did that version of it, sort of get rid of those imperfections.

And so this was done on a grayscale sensor. So it's just a gray sensor, no color information. But we put six filters in the aperture for this picture, a red, green, and blue filter and then three polarizers.

So we took an image. And the images of-- there's a big TV back here. And there's a color chart here and a book and then a polarizer sheet over here.

But you can see we've got a color image, which approximates the color of the scene pretty well. Yes.

AUDIENCE: When you say "printed", why you say "printed"?

ROARKE So there's this company called Pageworks. It's actually in Cambridge. And, what's the process called? I'm
HORSTMEYER: blanking.

AUDIENCE: Because it's just really just [INAUDIBLE]

ROARKE So it's on our transparency. So it's just like printing a transparency.
HORSTMEYER:

AUDIENCE: OK.

ROARKE Yeah.
HORSTMEYER:

PROFESSOR: OK. It's just easy to make. And the contrast isn't great, but it's because there is, [INAUDIBLE]

ROARKE Yeah, yeah.
HORSTMEYER:

AUDIENCE: [INAUDIBLE]

ROARKE No, it wouldn't have. So the pinhole sizes are about 50 microns. So I don't think the laser--
HORSTMEYER:

AUDIENCE: [INAUDIBLE] shrink.

AUDIENCE: So that's what we were trying to do last year.

AUDIENCE: And then, what?

AUDIENCE: Well, it does work.

AUDIENCE: OK.

AUDIENCE: For what we wanted to do, we needed still much more.

AUDIENCE: So it's sort of-- OK.

AUDIENCE: But, yeah, it worked. You print on a [INAUDIBLE] and then--

AUDIENCE: And it sort of-- it shrinks and it--

AUDIENCE: It's not uniform but--

AUDIENCE: OK.

ROARKE Yeah. So yeah, you just print out on transparency. But you have to use, you can't just do it on a normal printer.
HORSTMEYER: You need a high resolution printer. Well, we did. Because we needed 25 micron resolution. 720, I think, dpi.

But anyway, so in one image, we get a color picture. But we also get a degree of color polarization picture. So you can see the TV's polarized, because TVs have polarizers in them. There, you can see down there, there's a polarized down there over a resolution chart. So that was a simple little tip.

Then we did one of 16 filters. So basically I bought every filter I could from Edmonds that was cheap and put them in a filter array. So we have red, green, blue, yellow, magenta, cyan. I got an IR filter.

I got a

bunch of different polarizers going to different angles to get a precise degree of polarization measurement. And then I put three neutral density filters. And so these are the 16 images you get from one single image. You just separate them with a lookup table.

So you can see up here is the color information. It's kind of hard to see from the projector. But the spots are the color are changing, the checkerboard pattern's changing for the color.

The near-IR is really dark. And I've got a lamp here, which is saturated on some. But you can see the resolution chart with the near-IR.

The polarization filter, you can see the TV screen screen's changing and how bright it is and stuff. And then neutral density is just darker.

AUDIENCE: [INAUDIBLE]

ROARKE What's that?

HORSTMEYER:

AUDIENCE: [INAUDIBLE]

ROARKE Yeah, yeah.

HORSTMEYER:

AUDIENCE: But then you stack them from each other.

ROARKE No, they're not stacked on top of each other.

HORSTMEYER:

AUDIENCE: Oh.

ROARKE So here are the filters. It's hard to see on a projector. But here you can see it's just this little 3 by 3 filter array.

HORSTMEYER: Here you can see.

So basically it's just literally a flat 2D array of filters. And I put it right on the aperture stop. You have about this much room on the aperture stop, a little bit bigger than a quarter or so. So you can put as many filters in that space as you can.

AUDIENCE: So this also decreases resolution.

ROARKE Yeah, definitely.

HORSTMEYER:

[INTERPOSING VOICES]

Definitely. So your spatial resolution is going to be decreased at least proportional. If you did it perfectly and you put one pinhole image of each filter to 1 pixel it would be decreased by the number of filters. So I have 16 filters. My spatial resolution is going to be at least 1/16 the original spatial resolution. That's a critical point.

So these are very low spatial resolution images.

AUDIENCE: The effect is very similar to what they are--

ROARKE Exactly. So it's the same thing as a Bayer filter. You have a red, green, and blue filter. It's a third.
HORSTMAYER:

AUDIENCE: But you get this on the lens, right? So it's not on the sensor. So how do you know which actually pixels are affected by those guys?

ROARKE Yeah, so each image--
HORSTMAYER:

AUDIENCE: Because if you go back to the slide when you show, OK, this part of the [INAUDIBLE]

ROARKE Yeah.
HORSTMAYER:

AUDIENCE: So you can actually know exactly each pixel was effected by--

ROARKE Yeah, so I'm sticking a pinhole right over the sensor. So the pinhole array is going to create multiple images. Each
HORSTMAYER: one of these dots is a pinhole array image essentially.

And we zoom into that. I'm highlighting one pinhole array this white box. And so I know just from a priori knowledge, when I put the filter array in there that this center filter is a red filter. I know the one on the upper left is a green filter.

The one in the upper right's a blue, for example. Because I put the filter array in there. So now under each pinhole, I have a, ideally I would have nine pixels. And I would know which each pixel corresponds to just from the orientation of how I put the filter array in there.

AUDIENCE: And the [INAUDIBLE] has to be fixed for this to work, right?

ROARKE Yes. So I glanced over that. I skipped that slide.
HORSTMAYER:

But yes, you basically have to be imaged at infinity. You can't change the focus. Once you change the focus, you're basically crossing those rays in front of the sensor plane, and that's destroying your detection diversity.

AUDIENCE: So to scene has to be planned out.

ROARKE Yeah.
HORSTMAYER:

AUDIENCE: And it has to be [INAUDIBLE]

ROARKE The scene has to be planned. And lambertian. Well, not technically lambertian. It has to emit polarization or
HORSTMAYER: spectral information.

AUDIENCE: But they should be the same.

ROARKE And all the same rays.

HORSTMAYER:

AUDIENCE: Did you try removing the effect of reducing the spatial resolution by trying to demodulate? Because see, the pinhole array is basically a [INAUDIBLE] which has a modulation function on it. So if we demodulate it on computers once the image has been captured, you can actually regain the spatial resolution reduced by modulation.

ROARKE I don't really know what you mean modulation. Basically, since you have filters in your aperture, those are
HORSTMAYER: attenuating light.

AUDIENCE: Operating a function over the image.

ROARKE Yeah, but you're also mixing in all this diversity. So now you have all this spectral and polarization diversity
HORSTMAYER: mixed in from the different filters, which are going to change depending on your spatial location. So if I have a polarized object up here, it'll be modulated by 1 polarizer and not by another.

PROFESSOR: It's just like there [INAUDIBLE]. You can take off some sophisticated [INAUDIBLE] information. But it wouldn't just be hard [INAUDIBLE]

ROARKE Yeah. OK. I'll just try to finish up. So just an example of using those 16 images. People ask, well, why do you want
HORSTMAYER: 16 images of all these different things? So I try to come up with a clever example.

But basically you can create a color image and you have this huge saturated area from the sun or from a lamp or whatever. So you can just make a quick HDR image using the three HDR filters on top of the color information. So now you can see there's like a resolution chart back there.

And then what you can do is take a degree of polarization measurement using the five polarization filters. And that'll find things that are reflecting polarized light or emitting polarized light. So for example, the man-made object example I showed you before. You can pick out which objects are man-made, which are these or whatever.

And you can also pick out things emitting IR information in the IR filter. So people or well, I guess plants emit near-IR more than people. But if it was extended into the IR range, you could identify people. But you would need a better sense than a regular CCD.

So you combine all those and I did it with 12 different filters out of the 16. And I created this sort of phony region of interest extended, dynamic range color image. And you can see there's some errors, obviously.

And that's because what I was trying to find polarization, for example, if it's a low intensity area, it's not going to emit any real information about polarization. It's just going to be a dark area. Things like that causing error.

And also from angular or the different planes perspective, you can see that here that there's a bar sticking out. And that's because as things get closer, error associated with essentially having different perspectives on an object becomes apparent. So it only works for plane or objects towards infinity, or it works best.

So the next example, this is I just did this over the summer for a quick paper. Instead of putting a filter array, we put a spectral filter in the aperture down here. And what this does is essentially like a grating, just splits up the incoming light into all different wavelengths.

It went from 400 to 700 nanometers, so across the color spectrum. And it had a resolution of about 10 nanometers. That was pretty good.

And they're cheap. They're like 200 bucks. And they're small enough to fit in an aperture. So they're really useful to do experiments with cameras.

And so I just took a picture of some crayons. And picking out different-- So this is a reduced-resolution image of some crayons. By picking out certain pixels, you can get a roughly 25-pixel resolution of the 400 to 700 nanometer range. So you can see this is magenta. That's this red and blue parts, but not the green, teal, orange, blue.

So yeah, I have one more example of that. So this is same crayons, different pictures. So I noticed when I was doing the experiment that if I had the lights on or if I had another light on, I would just get really different results. And I was scratching my head. Why does that happen?

Well, it's obviously because fluorescent lights emit so much different spectral information. Or I just can't really tell, but these lights versus these fluorescent lights have a very different spectral distribution.

So I didn't really label them. But these are definitely just fluorescent lights. Or maybe one's fluorescent lights and one's another desktop fluorescent light.

So you can see the sharp peak through this. Most pixels just have a sharp peak. Our eyes just integrate over all of that. So we don't really notice it. Just looks like light.

But if I just put on a regular desk lamp, like an old fashioned bulb, you get this distribution. So these are all from the same exact pixel. But I was just changing the lights in the room. So it was pretty interesting.

PROFESSOR: [INAUDIBLE] the light coming in as a modification of the incident--

ROARKE Incident lighting and the color of the-- exactly. Exactly. So, yeah, that's it. Any questions?

HORSTMAYER:

AUDIENCE: How many [INAUDIBLE] you have?

ROARKE From this? Roughly 25. It just depends on your pinhole array pitch.

HORSTMAYER:

Another thing-- so I really want to do this and make a better filter. Because the filter I use is just a one-dimensional filter. It's generally used for laser experiments. So you hit the laser on it and you can tilt in different ways and select a specific wavelength if the laser is more broadband than 10 nanometers.

So it was only one-dimensional. So it would be cool, much better and more efficient if I had a two-dimensional filter so that I could use all the area. But yeah, roughly 25. 25. Any other questions?

AUDIENCE: So those graphs were made out of 25 [INAUDIBLE].

ROARKE Yeah, yeah, yeah.

HORSTMEYER:

AUDIENCE: [INAUDIBLE] in your experiment, so each pinhole was making an image of the--

ROARKE Of the filter, yeah.

HORSTMEYER:

AUDIENCE: --of the filter? How much in your experiment, how many of the pixels were just wasted on--

ROARKE On the border? Yeah. A lot. A lot.

HORSTMEYER:

So I started with a, I think it was 10-megapixel sensor. And my spatial resolution was around 300 by 300 for each of the images. So that's for something like 16 filters or nine filters.

So a ton is getting thrown away somewhere. And each filter is not being imaged to one single pixel. It's being imaged to a 3 by 3 or 4 by 4 array of pixels. So I mean it's incredibly wasteful but it's just an idea. So you can really fine tune it, probably make it more efficient.

PROFESSOR: I think the big advantage of this, you can do all of those things while putting a filter in the lens and not have to put it on the sensor. So someone was asking, can you change those kinds of things. Yes, you can just take it out and put a new one, a different spectral response if you need it. That's what's interesting.

ROARKE Yeah.

HORSTMEYER:

PROFESSOR: Can you say something about the diffraction issues? Do you have any?

ROARKE Yeah, a little bit. So some pinholes are really pretty inefficient, pretty bad at resolving light, which you guys

HORSTMEYER: might learn about. So lenses, everything has an associate. It's called a point spread function. I don't if you've heard that term, but basically a blur spot size.

Pinhole blur spot sizes are more of a geometric projection of the size of the pinhole. So you want to let a lot of light through with a pinhole. But depending on the size of the pinhole, you can just imagine light just passing straight through the pinhole and not changing its direction. So that light will have the blur spot size roughly the size of the actual pinhole.

So that blur spot size, in my case, I was using 50 micron pinholes. It's going to be 50 microns, which is relatively large. Because a pixel is about 10 microns. So that's why I was imaging to about five pixels.

If I use the lenslet array, it would have been much, much better diffraction effects and whatnot would have been much less.

PROFESSOR: I think we're going to be talking about using the next array in other projects.

ROARKE

And another interesting thing is that, so the point spread function in a camera or the blur spot in a camera, is

HORSTMAYER:

given by the shape of the aperture. So if you have you're taking a picture and you go out of focus, the things out of focus have a circle, like a little blur spot. That's because the aperture is a circle. With a square aperture, the things would have a square and so on.

But when you're in focus, even the little blur spot is given by the shape of the aperture. So I was putting all these crazy filters in the aperture. And each filter is going to change the point spread function depending on the color of the scene.

So basically across my image, I was getting many different types of point spread function. Some were strange shapes because of the different properties of the objects in the scene. So it's very different problems to try and fix or analyze because of all the filters in the aperture.

PROFESSOR:

So I have just a couple of more points I want to make before we end. We've seen right at the start, beginning of the class, I had these four or five things that we wanted to improve about cameras. And one of them was improving the spatial resolution.

So in the case of film it's sort of ambiguous. Because there is more, it's hard to define what resolution is for film. But for sensors or for digital, it's usually just a number of pixels. And it would be hard to imagine how you can increase the information you're getting from a fixed number of pixels.

But there are techniques for doing that. And the most popular ones are what's called super resolution. I'm not going to go too much into the details of what super resolution is and how you do it. But just to give you an intuition, the idea is that you take an image with your camera fixed at one position and then you move the image sensor by a fraction of the pixel size. And then you take another image.

And since you've moved it by a fraction of a pixel size, the information that you've captured is actually different from the first one. And you can combine these two images together to get a higher resolution image. There are, obviously, issues with that, one of which is that you have to move the sensor by less than a pixel size, which is usually one or two microns. And it has to be very precisely controlled.

So you find this in really high end camera that you sign out [INAUDIBLE] cameras which do this medium format, big medium format cameras. And there are also fundamental limits to how much you can do using this. You can usually not go beyond a 3x or 4x or even just a 2x increase in resolution. Anything beyond that is what they call hallucination. It's you may have create an image which is 100x resolution, but it doesn't have any information in those higher frequencies.

So that's super resolution. So panoramas over time is just taking an image and just scanning it like this and stitching them all together, create a large panorama. You have a big high resolution image and you do so. And in doing so you also get a wider field of view.

So recently, this technique was adopted by people at Microsoft Research and the University of Konstanz where they built a device that essentially scanned the whole scene and took a whole bunch of images. So this one was created out of 800 images that were captured over the period of a few hours. And then they found correspondence points and essentially just stitched them together to create this one large 1.5 megapixel image, I think. Or 3.5 megapixel image.

And just to give you a sense for how much information is there, you can actually see the person sitting inside that crane, whatever that is. And there is a website. I think it's called GigaPan. They sell this device now. And Microsoft gives you the software. I think it's called HD View, which you can use to both view these images and also generate your own images.

You buy this thing. You put your camera on it. It sits there and takes images for a few hours. And then you just stitch them all together. There is also a group, there's a husband and wife pro couple who've been doing this for a while using film photography. They have this large format film camera that they go out and take these huge gigapixel images. And they have their own custom scanners. And they will scan them and get a digital image out of it. And a number of other people who have done this kind of thing.

AUDIENCE: But if this is a consumer device, how I imagine this actually having the camera move around a lot.

PROFESSOR: So the way they have is something like this. If you can see it, it just rotates and takes multiple images. It doesn't have to move at all. So it's this thing.

The one that they're selling is it has a very cute shutter release button also. So you can just use any camera. And as long as you can just move the arm and have it so that it can pretty much [INAUDIBLE] button. And it will go out. And you can, I think, program it to say how many images you want and so on.

So either you can just create a panorama or you can create a very, very high resolution panorama like that. So some of the challenges in this thing that I just want to point out is there's a huge variation and intensity. And this is something that we discussed the high dynamic range issues. When you're moving the camera and you're doing all of this, you could have parts of the image that are much, much, much darker than other parts or much brighter.

And then how you stitch them together, what kind of exposure issues do you have to take care of? Because you just do auto exposure then one part of the image may be too bright than the other. And it might not line up. You might not get the alignment right or you might have issues with finding correspondences. And all of that is taken care of the software from Microsoft called HD View.

Another way of doing high resolution, this is again going back to those three basic things. This one I just talked about this was epsilon in time. But you could do epsilon in sensors. You can have 500 cameras and you just take lots of images at the same time.

So this is similar to the camera array we saw earlier. But the difference in this is that the cameras are all looking out parallel to one another. In the previous case, they were all looking at something in between. So there was a lot of overlap. In this, the overlap is much less. It says for people-- and I think it's even less than that.

And what you can do with this is you can get all of these images and then, just again, find the correspondences and stitch them all together and get something like this. This is a very, very high resolution, and it has lots of information. And then again, you can again, find those correspondences, like them and just stitch them together.

But there's also issues of geometric and color calibration and also high dynamic range issues, which I briefly discussed earlier. So this is what you get once you fix all of those. For some reason, it's really dark here. And again, it shows you can really zoom in and get a high resolution image.

I don't think this is anywhere close to [INAUDIBLE] Because this is a much older project. And again similar to the assorted pixels, they actually have different exposures for different cameras. And so you can combine all of that information together to get not just a high resolution image but also a high resolution, high dynamic range image.

The last thing I want to talk about is increasing the temporal resolution. So far, we've seen how we can increase the spatial resolution. We've seen how we can increase the dynamic range and focus depth of [INAUDIBLE] and the other thing that remains, the temporal resolution.

And you can now buy cameras that have 1,000 frames per second or even more. And you can just hit the shutter and it takes a whole series of images. But this is from a few years ago where what they said was that instead of taking one camera with very high frame rate, what if you took multiple cameras with more reasonable frame rates and then you combine that information together.

And sure enough, they came up with their own camera array. This is also from Stanford. And you can combine, I think they triggered them in a way that they can then later combine the information.

But again you need to calibrate the images and also color correct them and so on. It's just a [INAUDIBLE]. You can see this off like artifacts as the thing moves and because the calibration is imperfect.

And each of these is a relatively low resolution camera, I think about six [INAUDIBLE]. So that's basically it. So that's epsilon photography. It's how we enhance film-based photography.

And the idea is to modify the exposure settings, spectrum of color, focus, camera scene, illumination, basically anything else that is a parameter of the camera and take multiple images or take images [INAUDIBLE] over time or sensor or pixels. And the end result is get a better camera.

And as we'll see in the next class and future classes that this is not the interesting part of computational cameras. This is just something that it's good to know about. And there's still research going on on many of these topics. It's just, it's furthering or making even better cameras the way cameras were for more than a century rather than coming up with cameras that do new things. I guess that's it.