

PROFESSOR: So let me hand it over to Ankit now. And we're going to talk about color.

ANKIT MOHAN: So I'm going to start with something that I discussed a few weeks back when we were talking about epsilon photography, and if you remember the slide I had back then. And this is one of the oldest ways of creating color images. You just capture the images over time, one with a green filter in front of the camera lens, one with a red one, with a blue filter. And from these three images, if you project them back on a white surface, again, through three different projectors-- one with a green, one with a red, and one with a blue filter in front of it, you will get what looks like a color image to the human eye.

And this is stuff that was done more than a century ago. And you can take these images now, these individual grayscale images, and color tint them and then again, add them back together, similar to what you did for your assignment when you were adding the lighting. And you can create these digitally. And now there's a whole library of-- Library of Congress has this huge database of images that they're, by registering them, they're recreating these color images from black and white images.

So this is one way of creating color in photography. Another related similar thing to this is the color wheel. I just wanted to-- this is something that, again, I talked about in the previous class. And it's something that's used more for projectors. You have this color wheel rotating in front of this-- probably has one if it's a DLP projector.

And at any given time, you're just projecting one component of the light. You're either green, red, or blue. And the human eye actually integrates it over time. Because it happens in rapid succession one after the other. So it's a little different from the previous case, in which you were projecting all three at the same time and light was adding in space. Here, light gets integrated over time. You get the effect of color.

So a third thing that I discussed previously in the previous task force, this concept of three CCD sensors, where you use these dichroic prisms and mirrors and you split light into the three wavelength regions. And then you have three separate sensors that capture the three wavelengths. And you get, using three monochrome CCDs, you can capture a color image.

This is probably what's the most popular in most still cameras, and also many of the digital cameras. And it's basically what's called the Bayer mosaic. That's because the person who invented it from Kodak, his name was Bayer. And he had this patented in, I think, mid or late '70s, which describes this technique. The basic idea is that instead of having monochrome pixels, you put these little filters, tiny filters, on top of the pixels.

And the filters, you have two green filters, and one red, and one blue filter in a 2 by 2 region. And then this 2 by 2 thing is tiled all over the sensor. So any one pixel is only sensing one wavelength, either red, green, or blue. When I say wavelength, it's actually a whole band of wavelengths. But let's just stick with color for now.

And so it's sensing just one of these three colors. And then they use clever demosaicing algorithms, which essentially interpolate between the image that sense. So for example, at a pixel with sensors only red, you would interpolate between the neighboring green pixels in order to estimate what the green color is over there, while still making use of the color, that center of the red color.

So what you see in most images is almost kind of you're hallucinating the higher resolution because you are going up by a factor of 3 in the resolution in the image when you do this kind of interpolation.

PROFESSOR: And so a 4 megapixel camera, and here it looks like we have for every four pixels we have two green, one blue, and one red. So if we take a 4 megapixel picture, how many pixels are actually taking each of those colors?

This four megapixel camera actually giving you 12 megapixels, because when image comes out, it's four megapixel in red, four megapixel in blue, and four megapixel in green. So what is going on?

STUDENT: Well, I mean, you get the resolution for each of those colors, since they're separated in the space by the other colors being in space, as well.

PROFESSOR: So in this 4 megapixel, there are only 4 million pixels total.

STUDENT: Yeah, total. [INAUDIBLE]

PROFESSOR: And then two megapixels are green-- one megapixel is red, one megapixel is blue. And so what's the benefit of, if you can pair this with a previous design, where you had three separate CCDs, one for each color. What's the benefit, and what's the disadvantage?

STUDENT: Money. Cheaper. It's cheaper. Yeah.

PROFESSOR: Why is it cheaper.

STUDENT: One CCD.

STUDENT: And it's less optics.

PROFESSOR: Less optics climate. Alignment is easy. What's the disadvantage?

STUDENT: Plus you're blocking light.

PROFESSOR: You're getting-- effectively getting one third light because every time you sense one color, the light for other two colors is being thrown out. And when Mike talks about a little bit later about some other multispectral cameras, this will become the biggest issue of the notion of [INAUDIBLE] light.

STUDENT: And also can you perfectly reconstruct the image.

PROFESSOR: So this one was an issue. So, most of the reconstruction, the interpolation that we just talked about is based on some assumption about the natural scene. But if your scene is not natural, maybe it's a black and white text or you have very fine stripes on your shirt, in that case, the interpolation would not give you the right results. And you see it often. You can see this strange [INAUDIBLE].

ANKIT MOHAN: Yeah. So, one of the things they do in order to avoid this kind of thing is that they place a low pass filter immediately on top of the sensor so that you get rid of any such high frequency details in a single color. But still, you get-- I mean, in many of these images, you do see color [INAUDIBLE] so you have some weird artifacts in one color channel and not in the others or rainbow effect. And that shows up, especially along near the edges, and so on. Someone-- did someone have a question?

So, one of the things as Ramesh mentioned, is that you are throwing away 1/3 or 2/3 of the light in doing the previous one. So, there is this recent pattern that Nikon recently introduced where what they're doing is combining these two notions that we just discussed. They have these red, green, and blue sensors at the pixel level.

But they-- instead of using three separate filters, they use dichroic mirrors in order to separate the light that's falling on the sensor into the R, G and B component. So, it's similar to the first three CCD case that they're using dichroic mirrors. They're not using any light, but it's all happening at the sensor level. So, conceivably they can build this thing in the semiconductor itself. And so it's much cheaper than using prisms outside of one.

PROFESSOR: Anybody familiar with dichroic materials? It's basically a type of [INAUDIBLE] in simplest words. It's a type of a glass where at the right angle, if you shine light at the right angle particular wavelength will pass through and all of the elements will be reflected.

ANKIT MOHAN: It's essentially, I think it's total internal reflection is what it makes use of. And the threshold is different for different wavelengths when the light is going to get total internal reflected or just keep going straight into the next material. I'm using a series of such mirrors they can separate the three weapons.

PROFESSOR: And in the break, we'll have the soap bubbles and then we'll demonstrate the same concept.

ANKIT MOHAN: So, one other popular-- semi-popular sensor design is what's called the Foveon X3 sensor. And instead of using three separate pixels, each pixel having a different filter on top of that, the design of this sensor is very similar to that, which is used in film. So there-- in film, they actually have three separate emulsions or three or four separate layers of emulsions each sensitive to a different wavelength of light.

And similarly over here, they have as you keep going down deeper into the pixel, different depths of the pixel actually are sensitive to different wavelengths of light. So, the top region is sensitive, more sensitive to blue and the next part, the green, and the bottom most is to red. And so a single pixel can actually sense all three colors, all three wavelengths, as they're falling down on it. And the advantage is that you don't need to put this low pass filter on top of it because it's traveling the spatial-- multiplexing-- you're doing this multiplexing in depth. So, that's a big advantage.

In terms of the resolution, yes, you do get-- you don't lose the resolution like do in the case of a Bayer filter. But so far, when they caught the number of pixels that they simply multiplied by 3. So, when sigma is the manufacturer, that manufactures these Foveon sensors now they [INAUDIBLE] over Foveon. So, when they say they have a 12 megapixel sensor, it's really a four megapixel sensor.

They're just call it a 12 megapixel sensor because there are three elements for each pixel. So, it's not clear whether that's really-- again, in terms of the resolution. But a definite gain is that you don't need a low pass filter over it so you don't get these-- you get-- you're able to capture much higher frequency information.

PROFESSOR: So, you can't win either way. When they say it's four megapixel traditional camera makers, you only get a two megapixel green image.

ANKIT MOHAN: Right.

PROFESSOR: And when they say it's 12 megapixel, you still get a four megapixel green image.

ANKIT MOHAN: Right. But one other disadvantage of this is that unlike the previous case where you can almost arbitrarily choose what filter you want over each pixel in this case, the separation between the red, green, and blue channels is not as great. You don't-- they apparently don't have as much control.

So, they need to do a lot of software processing on the image that's captured in order to separate the red, green, and blue ones. Possibly, they're doing some sort of patchworks based on some of image prior, so you would have some image artifacts over there which might not be these moire artifacts, but you'll have some color artifacts.

PROFESSOR: And the actual profile is very misleading because when you see this diagram on the left, the blue is basically getting everything. So, it's something like-- this is your blue, green, and red. Then the blue pixel is getting almost everything and the green pixel is getting a little bit less and the red pixel is getting whatever remains.

So, it's not-- it's not like blue, green, and red. So, the picture on the right is completely misleading. It's getting some free values that are highly overlapped. And from that, they're going to do some inversion and figure it out actually.

So in the very simplest words we can think of, the first one here is everything. The second one here is just this and the last one gets this. And from that, they can figure out what red and blue is.

ANKIT MOHAN: So, I should also say that there's nothing really holy about this RGB design. It's just that this was the one that was proposed first. And this is what's being used. It's the most common.

But there are a whole number of other so-called Bayer patterns that have been proposed which don't use this 2 by 2 tile. They have even bigger tiles. And the simplest one is red, green, and blue and clear. So, you get one pixel which gets all the wavelengths and then red, green, and blue interpolations.

And I'm sure there are studies that compare various ones because if you don't get RGB, you need to do some of inversion. And you would get noise in that inversion. But if you use [INAUDIBLE] magenta, yellow, you would be able to get more light onto the sensor. So, I'm not very sure of-- I mean, I couldn't find any real studies that compare the various sensors or various--

PROFESSOR: Kodak has done a few studies. They have it on their website.

ANKIT MOHAN: They do? I mean, it was always very-- I'm not sure. So, yeah. So, there are various trade offs between the different kinds of sensors but this is the only one that you find in practice right now.

PROFESSOR: This will be a great class project by the way of figuring out which spectrum to choose depending on the scene.

ANKIT MOHAN: OK. So, taking a step back from-- this was more a rehashing of what we already discussed before in terms of color sensing on cameras. It's important to look at what we are sensing when we sense light. And it's really a part of what's the electromagnetic spectrum that is actually much, much wider and has many other types of rays than what we are usually looking at in photography which is a visible spectrum. And the visible spectrum of light goes from 400 to 700 nanometers.

Usually, when you talk about light, you talk in terms of wavelength. And in different-- in radio and microwave, you would probably talk in terms of frequency or in terahertz radiation and so on. But it's-- for visible light, you always say the wavelength is 400 to 700 nanometers going from blue-- from blue to red.

And it's interesting that this is really the only wavelength region that is almost completely-- the atmosphere is almost completely transparent to it. So the sunlight that comes through is-- it's-- other than maybe radio waves, almost all other wavelengths are actually occluded by the atmosphere. And this is really the most of the natural illumination that you have is in this wavelength region which is probably why humans and most animals actually develop their-- are tuned for this region and not any other.

And another sort of interesting thing is as you keep going away from this or shorter wavelength, the more dangerous out of the way it becomes. Already starting from UV, you start getting cancerous rays and so on. And then X-rays and gamma rays are even more so. But if you go for a larger wavelengths, they're usually harmless.

PROFESSOR: And the way to remember that is wavelength is inversely proportional to the frequency. And the product of the two, what's the number, and they multiply the wavelength of spectrum to the frequency of the spectrum. That's the speed of light. What's the speed of light?

[INTERPOSING VOICES]

PROFESSOR: But in this class?

[INTERPOSING VOICES]

And so the frequency is increasing as you go to the right, which means it has more energy, [$h\nu$] which means it can penetrate deeper and damage more things. So, that's one easy way to remember what's going on the right. Unfortunately, the chart is flipped because we should think of limit increasing from left to right, but this is really showing the frequency increasing from left to right.

ANKIT MOHAN: So yeah, I guess the point I want to make here is that we are-- it's only this region that we are looking at, which is really, really small compared to the whole spectrum, EM spectrum. and there is lots of interesting stuff going on, especially in thermal IR and thermal and even beyond that in using it for imaging. and we are using Wi-Fi for imaging also and things like that, which is support for gigahertz or something like that.

So, you don't have to be constrained for even creating a photorealistic or a visual image of understanding what's around us. We shouldn't be limited to which will serve the visual spectrum. It's OK to think outside it. And I'll show a few examples of how thinking outside the visual spectrum actually enabled you to do a lot of things which you otherwise would not be able to do.

So before again, I just wanted to talk about what a spectroscope is. And we have one of those here. So, I guess you could pass it around and so you-- so, what spectroscope really is nothing but a prism essentially. And a prism takes a signal.

So, a prism is basically this optical element that bends any incoming ray of light. And but the interesting thing is because this is refraction, the refractive index of glass is actually slightly different based. It's a function of the wavelength of light. So, if you build a prism of the right material, right kind of glass, you can have it have a huge disparity in the refractive index between the 400 and 700 nanometers.

And so when the light bend coming out of a prism. So, it says grating here but can be a prism or creating any of those. When the light went from here, the red and blue-- red, blue, and green, the different wavelengths, actually bend in different directions. And if you have a detector placed in front of it, you can sense the intensity of the blue ray, of the green ray, and of the red ray separately. And you can sort of decompose an incoming source of light into its constituent wavelengths.

PROFESSOR: You have a grating?

ANKIT MOHAN: Yeah.

PROFESSOR: So, I'll pass this around and all you have to do is look through this hole here. And all it has is a slit in front. so the best path forward was to look at one of the bright lights up here. And very conveniently, there is a scale on this side that goes from 400 nanometers to 700 nanometers. So, the idea is that you can look at any point in the world and you can see its spectrum right away. So, unfortunately, we don't have fluorescent lights.

ANKIT MOHAN: Actually, we do.

STUDENT: [INAUDIBLE]

PROFESSOR: Right. And you'll realize the [INAUDIBLE] is actually very spiky. It has like sharp blue and then a little bit of green and very, very annoying spectrum. It's not very nice. It's more like the sunlight. And just pass it around. I was going to pass around the flashlight but I'm not sure how many people have good hand-eye coordination.

ANKIT MOHAN: There's also this, which is diffraction grating and it-- for the purpose of this dispersion, it's very similar to a prism. So, if you look through this, you will be able to see a whole rainbow of colors around any bright light.

PROFESSOR: So, when you're looking at this, look at it through a particular angle, a particular orientation and then rotate it. And as you rotate, you'll realize that it will fall down.

STUDENT: Speaking of hand-eye coordination.

PROFESSOR: Exactly. I have an excuse. I was sick this week. You'll see that the image actually shifts around. So, if I'm looking at Daniel, I see him. Then I see a red copy of him and a blue copy of him shifted. I knew there was more than one of you. And if I rotate, it rotates.

And just remember this principle when we come back and talk, when Michael starts talking about different mechanisms for exploiting wavelength. Again, a great class project idea.

ANKIT MOHAN: Yeah. So, this here just shows one of our typical spectrogram. Kind of looks like it's-- you have the wavelength over here and the intensity over here on the y-axis. So, you can see those spikes over here at 550 nanometer and then again, 600 something. So this is probably a fluorescent light is kind of similar I guess.

PROFESSOR: Maybe we can pass around--

[INTERPOSING VOICES]

LED light.

[INTERPOSING VOICES]

ANKIT MOHAN: Will be very spiky.

PROFESSOR: Yeah. That'll be really good to see. Take a look at this, [INAUDIBLE].

ANKIT MOHAN: OK. So, while one of you is looking-- so, this is stuff that Michael is going to get into much more detail later. But I just wanted to mention-- are you going to talk about [INAUDIBLE]?

MICHAEL: It won't hurt the do it twice so go for it.

ANKIT MOHAN: OK. So, I mean, it's just something which is very high level of what he's actually going to talk about is the spectroscopy is really imaging the multiple spectrums of a single point in space like a light source over here. If you have a full scene, you are not going to get much out of a spectroscopy. And what you need to do then is what's called multispectral imaging or hyperspectral imaging depending on how many spectrums you end up getting. And so there are these standard ways of doing that.

And I just wanted to briefly mention what those are. So, usually this kind of multispectral imaging is very popular and remote sensing kind of applications where you have a plane or a satellite that's flying over the region and you want to get the multispectral or hyperspectral image or data set of what's on the ground underneath. So, what we have here is a lens kind of thing, which is on the plane image sensor behind it and the object space, which is on the ground.

And so the most straightforward way of doing this is what's called-- so, even this simple imaging where you're not imaging the complete spectrum, you can do it in various ways. The first way is where you have a lens which completely images the object space onto an image or 2D sensor. So, this is what just a traditional camera looks like. You capture the whole scene in one go.

Another way of doing this is what's called a push broom. And a push broom, the way I think of it it's you're pushing forward as the platform or the plane is moving in that direction. So first, when the plane is at a particular location, you are imaging this line on the scene. And then in the next instance, you will image this line and then you will image this line and so on. So as you're moving forward, you are imaging one line that's [INAUDIBLE] you.

The other one is called the whisk broom in which you are-- while you're moving forward, you're going from left to right and from right to left. And you're doing this whisk broom kind of thing, as opposed to a push broom. And what you do with this whisk broom is that when you are imaging this one element, you-- instead of just imaging onto one sensor, you pass the light through a prism, and you get a whole spectrum similar to a spectroscopy. And this gives you the complete spectral characteristics of this one element or one point on scene.

And then in the next instance, you are going to sense the next element right next to it as your whisk broom is going from left to right and so on. So, you're going to sense the whole scene. The way the push broom hyperspectral sensor works is that you-- so in this case, you need just 1D sensor in order to capture the complete data set of the hyperspectral image of the scene. Another way of doing it is using the push broom where instead of cap-- so, you put a prism in front of each of these elements, the image of each of these elements, and then you have a 2D sensor.

And the 2D sensor would sense along the x-axis is the point in space or the point in scene. And along the y-axis is the various spectral bands that you have for your hyperspectral imager. And then again, in the next instance, you will sense the next [INAUDIBLE] row in the scene.

A third way of doing it is basically something that's similar to the traditional camera is to put a filter in front of your lens and change the filter that's in front of your lens. Either have that color wheel or have a tunable filter which is-- whose response is changing. But at any given instant, you have-- you're either sensing green red or blue. And then in the next instance, you're sensing the next one.

It's similar to the first image we saw when I started the talk of capturing multiple images with different wavelength filters in front of the camera. It's similar to that. So, you capture the whole scene in one instant but only for one wavelength.

So this summarizes what I was talking about that you have this what's called the data cube or the object cube, which is the x and y are the same coordinates, x y . And λ is the wavelength where it goes from 400 to 700 or slightly beyond those extremes. You might also have near IR and so on. But this is what you want to get.

And unfortunately-- so this is a 3D object. And what you-- unfortunately, our sensor can be at most two dimensional surface. So, you want to somehow get the three dimensional data set onto a two dimensional sensor. And there are various ways in which you can do that. And you can-- the most traditional, most obvious way of doing it is to have a third dimension be time.

So, at any one given instant, you are either getting a slice like this or a slice like this or a slice like this either along x or along y or along the wavelength. And that in the next instant, you get the next slice and so on. And then you combine all this information together. So over a period of time, you build a complete object cube or data set.

So, this one is where you're using a filter. And this is where you're using a whisk broom or a push broom depending on how you name them, x and y . So, this is sort of the more traditional way of doing multi-spectral scanning. And Michael is going to talk about even more interesting fancier ways of doing it by just taking projections. But I let him get into that.

So, I think Ramesh briefly mentioned thermal imaging. And can we get the lights off or? So, this is what-- for those of you who've never seen a thermal camera, we have one upstairs. I think [INAUDIBLE] welcome and play with.

But this is what images through our typical thermal imager or thermal camera look like. This is sensing and usually in the wavelengths in the range of 1 micron to about 10 microns or 15 microns in that range. That's the range in which humans show up as warm bodies usually. So, 6 to 8 micron kind of range.

But you can also have-- actually, hot bodies, warm bodies show up like heat seeking missiles and so on. This is something that's explored a lot in the defense industry. This use of high resolution very fast thermal cameras, but it's something that's very rapidly coming and into the other applications also.

There was recently an article in *Time*. I think that they were looking at the various-- the thermal profile of a house from outside to find out where it's leaking and where the heat is escaping. And then they're using these kind of thermal cameras to actually test for whether pipes are leaking and your HVAC is working properly or not and things like that. So, it's finding lots of applications in areas other than just traditional defense. And so on. And that's one of the reasons why slowly the price of these cameras which used to be about \$20,000 each is hopefully coming down.

This is another image of just thermal images which I thought was interesting. The first one is-- what's showing is that when you think of thermal light, it's actually quite different from when you're looking at visible light. And one example, I think Ramesh already mentioned is that of glass, that glass appears opaque in thermal IR. So, you can't look through glass. And there are other objects which may appear completely-- excuse me, transparent in thermal IR but are opaque in visible light.

This is one example, which is sort of interesting, you have this fridge, which is brushed metal. And so in reflection, you see this very diffuse reflection. You can't really see what's on the other side. But if you just use thermal IR, you can very clearly see there is this rice oven on the other side that's really, really hot.

And the reflection is nice and sharp because the wavelength that you're using is now not 700 nanometers but it's something much larger. And the surface is actually very smooth when you look at that wavelength. But when you're looking at it in the visible spectrum, it appears very diffuse and you can hardly see what the reflection is.

PROFESSOR: Behaves Like a mirror.

ANKIT MOHAN: It behaves like a mirror in the thermal IR range because of the difference in wavelength. So, this is sort of an interesting thing to think about of applications where you can-- imagine you had one of these thermal imaging cameras on each cell phone right next to your normal camera. What could you do with it?

Or let's say you had it right next to your webcam on your laptop, can you do use-- do something interesting with it? Another one is this paper from-- Colorado-- no, University--

PROFESSOR: Houston.

ANKIT MOHAN: Houston, right. Where they use thermal cameras to detect for lie detection. And they have this paper in *Nature* where they analyze the region around the eyes and how that changes when someone is lying in the thermal IR range. And they claim that they're getting as good performance as a traditional lie detector.

PROFESSOR: It's not easy to reproduce, unfortunately, the results. But it looks pretty interesting. And the whole concept is that the blood veins pump more blood as your emotions change. So, as long as you can detect subcutaneous changes in blood flow, you can detect the correlated emotions.

ANKIT MOHAN: I think the difference is really very, very subtle, and it's greatly magnified in this image that they show. So, when we tried to do this, we couldn't spot any difference [INAUDIBLE].

PROFESSOR: Of course, nobody lies in our group either.

ANKIT MOHAN: So, one other thing-- when we were doing this last year, someone talked about near-infrared photography. And this is more from a photography point of view that actually, it turns out most of the CCD sensors are-- they're actually sensitive to near IR or infrared that's just after 700 nanometer, from 700 to about 1 micron. And they are sensitive to that. But in fact, most manufacturers put an IR block filter on top of the sensor that blocks anything that's greater than 700 nanometers.

So, what you can do is remove that IR block and you can then capture nice 3D images like this. This of course, is all-- they're all fake images because you don't really get any sense of color once you go beyond 700 nanometers. So people usually just fill in fake colors based on an original visual image or something like that. But another place where you can get these colors is if you use IR film. Kodak has this color IR film which reacts differently to different wavelengths beyond the 700 nanometers. And that gives you these interesting colors.

PROFESSOR: So these are not nighttime photos. These are just daytime photos with some pseudo color superimposed.

ANKIT MOHAN: So, the one interesting thing is that I think sky becomes really, really black and opaque. So, that's because if you remember, sky does not allow IR to come through as much. But also, any vegetation becomes very bright and white. And that's why you have the snow like effect on trees and things like that. But your barks of the tree and the food actually does not-- it actually absorbs more light. It does not reflect so much back, so it's just interesting things.

And I mentioned this briefly earlier, one of the biggest applications of this thing has traditionally been in remote sensing also. And by capturing multi-spectral or hyperspectral images of scene of especially vegetation, you can actually classify what crop is what and what kind of-- what's the role and where you have plantations and what kind of plantations over here they're able to distinguish between all these different kinds of crops. And this is something you can do if you have enough resolution in the multiple spectrums that you're getting just because each vegetation or tree actually has a very different reflectance profile when you look at it.

In other words, they both appear green to the human eye. If you look at the actual spectrum response, it's quite different. Then you can distinguish between different materials based on that.

STUDENT: This is not based on IR though? This--

ANKIT MOHAN: This is--

STUDENT: Visible spectrum?

ANKIT MOHAN: This is-- I think, it usually goes into near-IR at least, most of the remote sensing stuff. Because I think most of the interesting stuff of this kind of distinguishing thing actually happens in IR. But they do include visible light.

PROFESSOR: They might have a band of 5 nanometers or 10 nanometers and capture 20, 30 channels and [INAUDIBLE] high dimensional signal [INAUDIBLE].

ANKIT MOHAN: I'm not exactly sure. Do you know Michael, what trains they usually use for this kind of [INAUDIBLE]?

MICHAEL: Usually-- I mean, I don't know usually, but they don't-- it is definitely common to go into the IR and sometimes [INAUDIBLE].

ANKIT MOHAN: Right. So, talking about UV, you can also do interesting photography in UV range. And this is an amazing website. If you get a chance, you should definitely visit it. They have-- this guy has like pictures of all kinds of flowers and both invisible spectrum and then in UV. And it turns out that the flowers look just amazingly different when you look at them in UV.

And you have these-- almost these landing strips that invite the bees to come and sit-- give directions. Don't sit here, sit over here kind of almost. Whereas if you look at it in the visible light, it's all yellow or it's all red and there's hardly any difference between-- you don't see that--

[INTERPOSING VOICES]

PROFESSOR: One is especially striking.

ANKIT MOHAN: Yeah. Yeah, there's [INAUDIBLE] portions which are-- I think, well again, this is all fake color. There really is no color in this UV.

But one thing I want to point out here is that you cannot do this type of photography with most traditional cameras because glass is actually-- it absorbs UV. So, they use these special rare Earth quartz lenses which are super, super expensive in order to do this kind of photography. OK, so--

PROFESSOR: What you could use, for example, what Professor [INAUDIBLE] was talking about instead of using a lens, you could use a [INAUDIBLE] zone plate which is like a pinhole camera except glorified. It has more interesting pattern. And then on the sensor, you could put a layer of fluorescent material, so that you will stimulate the fluorescent and then make an image for that. So, you can kind of go around some of the limitations in a traditional camera to do UV photography. That's what they do in a way for a lot of medical imaging [INAUDIBLE].

ANKIT MOHAN: Also a lot of these images that you find, they all use film photography. They rarely use digital cameras for these kind of things. I mean, even for near-IR, I think--

PROFESSOR: This [INAUDIBLE] is just one of the major craze that has been around for several years. So, remember the UV in the UV spectrum, these images are going to look just monochrome. It's black and white.

But artists like to start assigning some colors to different intensity levels to seem more interesting. And for NASA and for astronomy, [INAUDIBLE] sometime over 20 years ago, [INAUDIBLE] put in those pictures of the nebulae and all that is sort of just putting them in [INAUDIBLE] really boring monochrome colors, let's start coloring them. Beautiful reddish hues and greenish lavas and [INAUDIBLE]. The real thing doesn't look like that. But it allows ordinary people to start appreciating astronomy.

ANKIT MOHAN: Yeah, it's like the wallpaper on-- that comes with OS X has that. OK. So, now switching gears a little bit. This is more getting into the human perception of color. And I'm not going to get into too many details, just a very high level overview.

So, many of you looked at this figure before, this chromaticity diagram. So, what this really-- this is the thing that people use in order to see how humans are actually sensing color. And what you have along over here is each of the spectrally pure colors, so went from 400 to 700 nanometers and everything inside. So, you have this blue over here going to green and then red over here.

And anything that's inside is basically a combination of multiple of these colors that are on the outside. So, it turns out that most devices actually cannot-- so, this actually represents the space of all colors that the human eye can detect. It turns out that most devices actually occupy a much smaller space within this whole chromaticity diagram, this horseshoe shaped thing. And so sRGB which is the default color of most CRT monitors and it's the default color that's used on the internet is actually a much smaller region. You can only represent colors that are in between this triangle in sRGB.

Now, what's interesting with this color space is that once you have these three points, any point that's inside this triangle can actually be represented by a sum of just these three color primaries. So, that's what I mean by fixed color primary is that most cameras that have a single green, single red, and single blue filter on the sensor actually have-- the green is somewhere over here, blue somewhere here, red is somewhere here. And that allows you to represent any color that lies within this triangle.

Now, that's what the color response of traditional film, and that's what it looks like for most cameras. So, it's very similar. There is very little difference between the color response of the two. And in both cases, the response is made to close the [INAUDIBLE] out of the human eye itself except for in firmware, like this firm [INAUDIBLE] I think is something that's very good for nature photography. So, they stress more on the red than what the digital camera does just because-- it's a little unnatural, the colors.

PROFESSOR: [INAUDIBLE] for ocean photography.

ANKIT MOHAN: It's any landscape kind of thing, sunrises and sunsets, nature. It doesn't work very well for skin tone, for example. So, this works fine as long as your color is within this triangle. But once you want to represent a color that's outside, it becomes-- it's not possible to do that because this RGB, the values, can only be between 0 and 1. So, there are a number of algorithms that you can use in order to estimate what the color should be.

But each of them is a [INAUDIBLE] lose information because you may project it to the nearest point or you may project it to the perceptually nearest point and all those kinds of things. But eventually, you end up losing this information. So, it turns out like all this region over here, it's very hard to represent colors over here just because you are triangle is actually-- it's outside this triangle. So, one alternative you can do is that instead of using-- putting the color primary over here, you can put a color primary out there.

So, now you have this really big color gamut and you can represent all colors that lie inside this color space. Unfortunately, what that means is that your green is now very close to, let's say 520 nanometers here. So, it's a very spectrally pure-- it's just one very narrow range of wavelengths like that over there. And that means you have to use a very sharp wavelength profile, LED, or a laser or something of that illuminating it or a filter that's really dark. So, it turns out that these optimal color gamuts that you have there are very good compromise between having a wide gamut and having filters that allow a large amount of power to come into the system.

So, if you have these primaries very close to the edges, you end up throwing away too much of the light. You get very little light coming in. So, what you'd ideally want to do is have these adaptive color primaries. So, for a scene like this, you don't have too many reds in this for example.

You want color primaries to be like this that's just some arbitrary shape. And for a scene like this, you want them to be more like that. Rather than have the same set of colored primaries for every image that you capture, you want-- you want to be a little more adaptive about it.

STUDENT: [INAUDIBLE]

ANKIT MOHAN: So, it's kind of like the equalizer, you mean? So, you want to be able to tweak the various wavelengths-- what wavelengths should be sensed more and what should be sensed less.

PROFESSOR: And if you're playing pop music versus classical music, you may want to change your synthesizer, more bass, less treble, and so on. And because those settings on your equalizer are also boosting certain frequencies and attenuating other frequencies. In case of cameras, the RGB has already been fixed [INAUDIBLE] three points. And those three points have been fixed. But maybe you ought to change the frequency.

ANKIT MOHAN: So, that's what I'm going to briefly talk about how one way of doing that. And-- is everyone aware? Have you talked about [INAUDIBLE]?

STUDENT: I have a question. So, here in this graph, the chromaticity diagram, it seems like you can represent it in a 2D plane so which would mean that two primaries are enough to define any color, because it's just a 2D plane.
[INTERPOSING VOICES]

ANKIT MOHAN: There's an intensity, which is-- this is a projection of the whole--

STUDENT: Oh, so it's actually a 3D thing.
[INTERPOSING VOICES]

PROFESSOR: But in terms of the color, you're right. You could represent with two numbers and the third one is the intensity. And that's why you have LUV or LAB or what is luminance and there's A and B which is chromaticity A and Chromaticity B or [INAUDIBLE] is the old one.

ANKIT MOHAN: One of the things that these-- so this-- when I said this is the RGB color space, this is actually incomplete. Just defining RGB is not enough. You need to define the white point inside, where is the white.

And that's also something that's defined over there. And that-- if you look at the third dimension, that's what represent the various levels of gray. So, in order to represent the intensity, you need that third axis.

MICHAEL: And let me emphasize something that it's easier to lose I think in this. We tend to use the words color and wavelength interchangeably in some situations. And wavelength is a truly physical phenomenon. It's a property universe of the light itself.

Color is specific to human vision, maybe other animals too. But it's different for them than it is for us. So, when we talk about turning wavelength into color, what we're talking about is the process of human perception. And so this plot, for example, is specific to human vision. If it's for some other creature or some physical device that we built, it would have a completely different plot.

PROFESSOR: Exactly. So here's-- before we get onto the color sensing, here's an interesting puzzle. If you take a red laser pointer, which-- they're using a red color for right now. And as Mike just reminded us, don't think of color, think of the wavelength. So, let's say it has certain wavelength-- I don't know-- 680 nanometers, something like that?

MICHAEL: 630 probably.

PROFESSOR: 630. And if I just take the laser pointer and shine it in a piece of water-- a piece of glass or a tank of water, what's going to happen to it? [INAUDIBLE]. [INAUDIBLE] shine laser in it. What happens to it?

[INTERPOSING VOICES]

It bends. And why does it bend? Does anything change about its physical properties? Of course, it does.

Either the wavelength or speed or something has to change. The speed is actually decreasing. So c in air versus c in water is related by a factor of what?

STUDENT: [INAUDIBLE]

PROFESSOR: The refractive index. So, this is actually reduced by a factor of $2/3$. Now, we know that here, we have a wavelength of light times the frequency of light.

I'm just talking about a for air. This is for a laser. And similarly here, we have wavelength in water and frequency in water. So, this has gone down by $2/3$. Something here also has to go down by a factor of $2/3$.

Yes.

STUDENT: [INAUDIBLE] frequencies fixed.

ANKIT MOHAN: Yeah, the frequency [INAUDIBLE].

PROFESSOR: So, the wavelength has changed. So, that means from 630 you say? From 630, we have gone down to $2/3$ of that. So, that would be what 420.

So now, this red laser is actually-- has a wavelength of 420 nanometers. And so we started from somewhere here and the light wavelength is actually now even further down. So does it mean that when you shine a laser in water, it's actually blue because it's wavelength now is 420 nanometers? Is it blue?

STUDENT: Well, we can't ever see that because we are [INAUDIBLE].

PROFESSOR: So, when it comes out it's back to 630.

[INTERPOSING VOICES]

But when it's inside, it's 420.

MICHAEL: You could immerse the CCD in water, for example.

ANKIT MOHAN: I guess we sense frequency more than wavelength.

PROFESSOR: Very good. [INAUDIBLE]. So, remember that message. Don't think about color, think about the physical properties. And the way to think about that is you may have different wavelengths but they mean different things in different media. In air, if you really want to think about colors, in air, 630 is red. In water, 420 is red. So, it gets too confusing. So, you can use colors when they're convenient. But when we start talking about physical interaction, it's better to talk about wavelengths or really something that remains constant, [INAUDIBLE].

ANKIT MOHAN: I think the important thing is that the energy of the ray remains the same. It's E is equal to $h\nu$ -- it's the frequency. So, that's what the [INAUDIBLE] it's the energy of the electrons that get displaced. But when-- so, that's why I said for whatever reasons when you're talking about visible light, they always like to talk about it with wavelength. And it's not entirely correct because as Ramesh pointed out, this figure is only true for air. And it won't be-- it would be completely very different if you were underwater. But it's just a convention that people like to follow. And that's what we are sticking too.

PROFESSOR: [INAUDIBLE] to ask this question to your photographer friends because they like to talk about colors.

ANKIT MOHAN: So again, going back over here. As I was saying, in different-- that's why they have both the frequency and the wavelength over here. In different fields and in different-- for different purposes, you would use either frequency or wavelength. It's just that for photography and for especially for this kind of imaging and [INAUDIBLE] the visible region, it's something that people usually use wavelength and not the frequency or the energy.

So, we wanted to try to come up with a way of having adaptive color primaries. And so I'm going to go through a analysis of an optical system that we developed. And it's going-- everything that I'm going to discuss is going to have be in what's called a flashland case. So, it's just in two dimensions but it scales up to a real 3D or 4D case also.

So, we start with a simple 1D signal and this 1D signal is arbitrary intensity. And along-- this is the x position of the signal and this is the intensity. And it's a white signal, which means it has all the visible frequencies between 400 and 700 nanometers. And the intensity of each one of those wavelengths is actually the same.

So, that's what the wavelength profile or the color profile, spectral profile looks like at any point between a and c . So, we take this signal and we put it in front of a pinhole. So, here's your signal and here's a pinhole. So, the pinhole essentially creates an inverted image of the signal.

So, you have this a, b, c here, you get a' , b' , c' over here. It's just a pinhole camera. So now, instead of positioning a film or sensor over here, we put a lens in this plane. And this lens essentially [? complements ?] all these rays, so you have a' , b' , c' , just like an orthographic set up so-- except it's inverted which is not so important, but you have a ray coming in for each point in the scene.

So, next we place a prism in front of this. And now as I talked about-- mentioned earlier that a prism actually bends the incoming ray where the bending angle actually depends on the wavelength of light. So, what I'm going to show for simplicity over here is that when you have a white light coming in, the green-- the red corresponding to the green wavelength are 550 nanometers or something. Actually goes through straight and bend. And red gets bent upwards and blue gets bent downwards.

Now, once again, I'm going against what Michael was saying. I'm calling it red, green, and blue, but it really is the wavelength, 400, 500, and 700 nanometers. I'm just going to call red, green, blue because it's easier to talk with it that way. And the other thing to notice here is that I'm only drawing these three rays, but really, it's a whole fan of rays because it's a continuum and it's a white object--

PROFESSOR: A whole rainbow basically.

ANKIT MOHAN: You have the whole rainbow here. And I'm just drawing three of those rays. So, you would actually have a ray going down here and up there and anywhere between this blue and the red ray that we have. So, now looking at this prism more closely, along this axis, we have the spatial points of the scene. So, you have a prime, b prime, and c prime.

And coming out of each of those points, you have this wavelength angle or λ . Now, this figure should remind you of something. And anyone wants to guess? Does this figure look familiar to--

[INTERPOSING VOICES]

It's exactly like a light field, except the only difference is that instead of having λ -- θ over here or the angle, we replace it with λ . So, this is something we call the spectral light field. It's-- the angular light field or the spatial light field-- the spectral light field. So, any point over here in x or λ will represent the intensity of a ray in the space coming out of some point along a and c and going in a particular direction.

And since we started with a white light source, it's going to have the same intensity along each of these along each wavelength. So, now-- since we've reduced this to nothing more than just a light field, it turns out we can use the various properties of a light field in analyzing the system. And the one-- again, this is something that you should be familiar with if you know what a light field is is that if you place a screen somewhere in front of a light field, you get a projection of the light field on that screen in a direction perpendicular to the screen in terms of the light field.

So, what that means in the flashlamp cases that now if I place a light-- place a screen over here, this yellow thing, what that does to the light field is basically get a vertical projection, which is a direction perpendicular to the direction of the screen, which is along the x -axis. So, what we're getting for every point over here is an integration over the various all the wavelengths for that particular x position. And you essentially get the shape of the signal in this projection.

So, if you were to place a screen over here, not surprisingly, you would get an image of the signal itself on the screen because it's so close to the prism it's-- these rays haven't dispersed as much yet. It's almost the same thing as if there was no prism over there. But now as you move the screen away from the prism, the angle of projection changes and it becomes more-- you get this sheer thing.

And because of the sheer, you get the various wavelengths, the signal corresponding to the different wavelengths is now dispersed and it's overlapping and shifted. And once again, I've shown just the blue, green, and red wavelengths here. But clearly, it's a continuum and you're getting this rainbow like smear over here.

And when you were looking through the diffraction grating, you could almost see this. You get a rainbow coming out of every point when you look through it. And that's similar to what you're getting over here. Now, if you move this plane away to infinity, when it's at infinity, you're going to get this horizontal projection through the spectral light.

And now, for each point on the screen, we are integrating over all the spatial positions, all the points on the signal for each wavelength. So, what you get on the screen is nothing but a rainbow because we started with a white light source. And you have something going from blue to red. Any questions?

So, what we've done is that we have this one plane where if we put a screen on that plane, we get the spectral characteristics of the signal that we started with. So, the problem with that, of course, is that this plane is infinitely far away. But we want to move it closer so we place another lens in front of the prism.

And this lens does two things. The first is that it creates a copy of the prism itself on this plane at some distance from the lens where basically this is imaged, c prime is imaged at c double prime, b prime, and b double prime and so on. So, similar to the x that you had there, you get an x over here.

But you also get this plane in the middle where each ray for each scene point of a particular wavelength actually converges and meets at a point on this plane. So, all the red rays are meeting here, green rays are meeting here, blue rays are meeting here and so on.

So, once again, looking at this, when you place a screen at this plane at the end, which is an image of the prism itself, you're getting this vertical projection. And this is an image of the scene itself. This is where we place our sensor. And we call this the sensor plane of the optical system.

And if you look at this plane, which had this nice cubed property of all the reds coming together, greens coming together, and so on, we actually get the horizontal projection. And this is the plane, which was infinitely before.

It's moved much closer than that. And you get this nice looking rainbow at this plane. And we call this the rainbow plane.

So, the nice cube property of the rainbow plane is that all the rays of a given wavelength coming from all the points in the scene actually converge to a single unique point in this plane. So, just to give you a couple of examples, if instead of the signal being wide if it had been completely red, let's say between 650 and 700 nanometers or something like that. When you took this horizontal projection at TR or the rainbow plane, you would only get this part of the rainbow. You won't get the other remaining rainbow.

Similarly, if your signal was half blue and half red, you would get something like this at the rainbow plane. So, now if you place a mask in the rainbow plane, let's say you block out all these red rays that are going through, you actually put an occluder in that plane. When you-- what that essentially does is that it multiplies the incoming light field at this plane, which is this, with a light field which is all zeros corresponding to the wavelengths that you block and all ones corresponding to the all-- to all the other wavelengths.

And so what you get is basically this-- what looks kind of like this. And now when you put the sensor at the screen at the sensor plane and you take the vertical projections, you don't get any of the red components. You only get the green and the blue component.

So, you essentially-- by including the red channel over there or the red color or the red wavelengths over there, you remove them from the light field that gets projected on the sensor. And you only get what looks like cyan. It's just green plus blue.

If you were to do that to the green channel, you could just occlude that you have a big 0 over here and you only get this plus this and this plus this and so it looks magenta and so on. So, essentially what you can do is place any arbitrary mask you want in this rainbow plane and that will influence what colors get sensed by the sensor and what colors are not sensed by the sensor.

So, it-- by placing a mask in the rainbow plane, it allows you to control effectively the spectral sensitivity of the whole imaging system on the camera by modulating what rays go through and what rays get occluded over there. So, that's what the whole optical system looks like. We have the lens, the pinhole, the prism, another lens that images the rainbow plane and we place a mask over there and then the image sensor itself.

PROFESSOR: So what's the benefit of this and what's the disadvantage compared to the other schemes we have seen before? Let's look at the disadvantage [INAUDIBLE].

STUDENT: Well, starting with the pinhole.

PROFESSOR: That's a bad idea.

[INTERPOSING VOICES]

It's always a bad idea. [INAUDIBLE] pinhole when you're blocking light, it's always a bad idea. So, what would be-- does everybody see that? There's a-- what would happen if the pinhole was made larger?

STUDENT: When you have to deal with focus problems [INAUDIBLE].

PROFESSOR: So, you'll get blur. But what kind of blur [INAUDIBLE]?

STUDENT: You're getting a blur between your spectrum pieces exactly.

PROFESSOR: The blur will be the wavelength as opposed to its space. Because even if you make this pinhole larger, conceptually, I guess you want to follow this trend.

STUDENT: Well, that's--

PROFESSOR: Well, just to get your thinking, if you increase that pinhole, you will still get an image at the end that looks sharp just like [INAUDIBLE] image. Because we created two images.

The prism plane is basically some kind of a virtual image. We formed an image on that. And the image of that is being formed on the sensor. So, there's no problem in terms of the focus blur. But the ability to control specific wavelengths has now decreased. So, there is a blur in terms of wavelengths, spectral copies are still missing.

STUDENT: So what if we put some kind of like pattern to light in more light but then we can [INAUDIBLE] out later.

PROFESSOR: Some kind of [INAUDIBLE] aperture or something like that?

STUDENT: Or mask or anything.

PROFESSOR: Very good. Very good Thank you. We haven't covered aperture yet so--

ANKIT MOHAN: I briefly talked about it in the end. But yeah, you're right. You could do something of that to deepen all the effects of the blur arising from this. But it turns out that this is a different kind of trade off that you from what we saw earlier like where we were doing the scanning in time. When you're scanning in time, it's almost like have a pinhole in time. And then if you have scene motion, then you would have blur due to that.

In this case, you have a pinhole in space. So, you're getting all of this in one shot but you're turning off a light over there because you need this-- you can't have an infinitely-- very large aperture. And you can have some sort of-- reasonably not very tiny aperture, you can get something like six or 10 different wavelengths over here, which is not great, but it's something that gives you a control of what you can control the wavelength. The aperture size and control the fidelity of the wavelengths that you're going to get here.

So, that's sort of the setup that we build. You have the image sensor over here, which is nothing but standard camera. This is the [INAUDIBLE] the mask that controls what wavelengths go through and what don't and then a bunch of lenses over here. And that's the diffraction grating that bends everything.

So, one thing you would notice is that there is a bend in the optical axis, something that I didn't show in the optical diagram so far. But it's-- again, something you have to take care of when-- if you're actually building the system. So, you want me to go through all the examples or I mean?

So, this is a simple test setup that we built. And the idea was we had a spectral rainbow generator. So, we have this rainbow which is going from red to blue. It's kind of like that.

And then we are imaging this with our agile spectrum camera. Now, notice that the color that you see here is because of the color of the Bayer sensor on the camera. And it's really-- if you were to use a monochrome sensor, it would be all gray and just everything with equal intensity. So, first when we block off a certain wavelength, so let's say about 606 nanometers over here.

You see a corresponding gap in the image that's captured by the camera. And it gives you an indication that it is blocking off in the right range and the Bayer filter actually helps you see that here. And if you put some other arbitrary mask that's actually blocking off in two regions, you get similar-- a corresponding image on the sensor. It looks very different over here. It's like more blue over here and so on.

So, one of the applications you can do with this thing that we worked on was trying to reduce the layer coming out of this LED. So, it's almost impossible to see here but you actually have text written in the background, which is some EG over here in the background. And then you have this right LED in the foreground. And you want to be able to capture both the background and the foreground at the same time.

And the background is much, much darker than the foreground. So, if you were to do traditional high dynamic range imaging, you would-- if you increase the exposure much more, this halo that's coming out of this LED starts to occupy even more parts of the scene. And you can't really see the background.

If you decrease the exposure, then this halo or this artifact reduces. But then the background becomes even darker. So, what you really want to do is just occlude this LED, the effect of the LED, and be able to see the background. And you can do that by-- you still don't see it.

So, you can do that by blocking out these wavelengths corresponding to the LED. And now the LED is much dimmer. And actually, if you look at it over here, you can clearly read the background. So--

PROFESSOR: This is strange. Now, we can see [INAUDIBLE].

[INTERPOSING VOICES]

ANKIT MOHAN: So, you can see the background a little bit over there but this way. So that's-- it's doing high dynamic range but by modulating the spectral profile of the scene rather than by cutting off all wavelengths equally or cutting off just certain wavelengths that you know [INAUDIBLE] the scene that are causing this disturbance in the scene. So, you can-- it turns out that you can build a similar projector also. [INAUDIBLE] was all for a camera but you can also use this as a projection system.

So, now we have a traditional projector diffraction grating, this lens, and then you have this R plane, you have the mask and then the screen is up there. And the advantage of doing it with this projector system is that you already have a pinhole it turns out inside the projector. And most projectors actually have a very long optical path between the light source and the projection lens in order to increase the depth of field that they get then once they project the image.

And it turns out we didn't even have to stop the lens or anything of that sort. And it just worked for a projector, a standard projector, without modifying it much. So, this is an acute example where the thing I wanted to mention over here is this concept of metamers. I don't if that's been discussed. But metamer is anything-- two colors which appear the same to the human eye or a camera when viewed under a certain type of illumination.

So, this is the wide illumination of the scene and you probably can't see this, like this orange cloth or blue cloth and so on. And many of the colors over here, you can probably see there. This-- it's very hard to distinguish between the start and the [INAUDIBLE]--

[INTERPOSING VOICES]

PROFESSOR: Is that 3M stickies, right?

ANKIT MOHAN: Yeah. These are all stickies. So, these are actually fluorescent, these envelopes I think. And if you look on the-- so, if you look over here, you'll see as we project different wavelengths, how these colors they appear almost black and very, very similar to one another under different illumination. And now, one of them is going to become brighter than the other. And you can actually see the difference.

So, it's very interesting to see the progression of how-- it gives you an intuition into that there's much more to the wavelength profile than what the human eye simply sees in white light.

STUDENT: So, can you just explain again what happened I mean, during this animation?

ANKIT MOHAN: Well, this-- I'm just projecting monochromatic very narrow wavelength of light.

PROFESSOR: [INAUDIBLE] projecting red light.

ANKIT MOHAN: Yes, Yes. So red, green, and blue. So, about 12 or 15 different wavelengths.

PROFESSOR: So, it starts with very reddish, slightly yellowish, then greenish, then magenta, and then finally, bluish-- not magenta, cyan and then bluish.

ANKIT MOHAN: It just shows you some of the-- I mean, I think if you look at the background, does this looks so different from this envelope whereas over here, it's actually very similar color in white light.

PROFESSOR: And these papers are very fluorescent intentionally, so they look very beautiful when they're on your tabletop. But they're also responding to very narrow wavelengths of light. So, that's why some of these envelopes become completely dark when you eliminate all your particular wavelength.

STUDENT: I have a question. Could it have been possible for you to just take an image and then use different channels to gender those 15 images that you generated? Did you have to definitely take them at different illuminations or--

ANKIT MOHAN: Or if you have a camera that captures all 15--

[INTERPOSING VOICES]

MICHAEL: Broadband illumination. Those are the two things you need, broadband illumination and a camera that captures--

[INTERPOSING VOICES]

PROFESSOR: 15 spectrums. If you take only three spectrums, then you only have three numbers of pixels. So, there is no way you can recover the 15. And if you have a camera like this, then [INAUDIBLE] other applications like distinguishing between a fake vegetable and real vegetable and freshness of skin and looking through fog-- all those-- not fog maybe, but some of the problems become manageable.

And film and even digital photography is just trying to mimic human eye, three colors, in fact, three fixed colors as Ankit was saying. But what you really want is something like [INAUDIBLE] equalizer. You can tune any frequency, any wavelength so you can see the world in interesting ways.

ANKIT MOHAN: So, this is another example of where adaptive color parameters would help. And as I mentioned earlier, there is very little-- cyan is one of the hardest colors to predict for most projectors because your color gamut is usually over there. And there's very little representation from there.

And it turns out there are two definitions of cyan. One is the traditional printing cyan and the other is what's called the electronic or electric cyan. And this color is-- so what I'm projecting here is just a ramp between blue and green. So, this is blue here and green and colors in between. And this is what the computer thinks of cyan at the top and bottom.

And you can see there's clearly a leak over here. And that's because when you-- along this line is what I'm projecting over here, this line. And so this cyan is nothing but a point that lies somewhere on this line.

So, you cannot project a color which is out here using a standard projector or a display. But if you use this [INAUDIBLE] spectrum, you can tweak it so that your color over here is actually something that's outside your color spectrum that the projector could have displayed. That's because you're using-- you're not using the filters that the projector was using but actually displaying something outside it.

PROFESSOR: So, some companies are also trying to sell you four or six color projectors. So, they'll have this point here, maybe another point here, another point here, here, here, and here so they can cover more colors than a standard three color projector. Because there is a three color projector, you must pick some three points.

And you cannot represent this shape with three points. If you start going to close to over here, if you take this part too close here, then as Ankit said earlier, it has to be very pure green. And that's difficult to generate unless you have a laser projection.

STUDENT: So, do the laser projectors have a wider color gamut?

PROFESSOR: Exactly. So, the laser projections can go off-- they can stay almost on the pure colors. So they can go all the way here and here and here. So, they can cover a larger gamut. But there are other issues with that.

STUDENT: So, they usually end up with higher contrast but they are less brightness or-- how do they usually?

PROFESSOR: They have to have good color rendition.

STUDENT: Yeah.

PROFESSOR: [INAUDIBLE] a problem.

[INTERPOSING VOICES]

Yeah, so in general, a laser projector is going to create a very nice rendition. It's just that it's not compatible directly with the human vision. So, sensitivity about the human eye for pure color is not so great, unfortunately.

ANKIT MOHAN: And there Are some constraints in the wavelength that you can use. You can't just choose any wavelength. They have to have a very specific problem. Then they use frequency doublers or something to get the other colors.

It's just a little more cumbersome to do that. And then there's the issue of power that you need if you want the laser projector to be bright enough. You need to have very bright lasers.

MICHAEL: Yeah. The amount of optical power that a projector throws out compared to a typical laser is quite large. So these things actually put out a lot of power compared to your laser pointer for example. So, putting a laser bright enough to generate this much light-- yeah, that is challenging.

ANKIT MOHAN: Yeah. Yeah, so-- I mean, this looks bright because it's illuminating just one point, the laser pointer. But if you were to actually distribute this over the whole region, you would barely see it

STUDENT: Most of them are scanning a single point, aren't they? Aren't most laser projections--

PROFESSOR: Yes, exactly.

ANKIT MOHAN: Yeah. So, it's as though you're actually distributing this over the whole space.

PROFESSOR: So if you have a million pixels, all one frame, every pixel is illuminated only one millionth of a frame.

MICHAEL: That's the trade, right? Is you still need a lot of power because you're not spending much time in any one place.

PROFESSOR: [INAUDIBLE] problem? I'm very, very hopeful that solid-state lasers will--

MICHAEL: Actually, that's [INAUDIBLE]. If we assume for the moment that laser pointer is as bright as the red chunk of that slide you were showing us, like on the previous one, if it's roughly is bright, same ballpark, but that's roughly a megapixel image--

PROFESSOR: A million times more--

MICHAEL: About a million times more power.

ANKIT MOHAN: Yeah.

PROFESSOR: Which is true. This is about one milliwatt. And the project is 250 watt. It's 250 million times more the [INAUDIBLE] compared to the [INAUDIBLE].

ANKIT MOHAN: I think laser projectors work great if you are in this room where it's not too bright and you had the projector here [INAUDIBLE]. It works reasonably well, but not if you have the lights on you probably won't see anything. So, another thing you can do with this is don't-- since you're not stuck with those three color primaries anymore. So, this is just a scene again. It doesn't show up there as well.

You have blue over here and you have yellow over here. And so traditionally, you would have a RGB filter where for yellow, you're turning it on in the red and the green color filter is in front of your projection. And for the blue, you are turning it on only when the blue part of the color is in front.

But if you know that your scene has just as yellow and blue, you can actually just project a yellow. And you can use a yellow and blue projector. And you don't have to use the same traditional fixed color primaries. And this will give you a brighter scene and one that has more saturation also.

So, one other cube example is that of colorblindness simulation [INAUDIBLE] Professor Emanuel has done a lot of work also in this area, I think. And the idea over here is so one of the most popular types of colorblindness is this different-- you can't tell the difference between red and green. And I think it's called [INAUDIBLE].

And so when you have white light illumination, you can-- I think, on the right and it's-- so, the lower ones are due to [INAUDIBLE] so you cannot tell the difference between the rose and the leaves and they both appear at the same color. Again, it's much easier to see over here. But if you actually project a certain wavelength of light on the object, it becomes much easier to tell the difference between these two. And it-- you can again, clearly see on the--

[INTERPOSING VOICES]

PROFESSOR: [INAUDIBLE]

ANKIT MOHAN: Yeah. You can actually see the difference between those two when you have the certain projection. So conceivably, using these kind of projections or using these kind of filters, you can help solve the problem. You can at least get some help in telling the difference between colors for people with colorblindness.

So, [INAUDIBLE] go over some of the limitations. There's diffraction artifacts. Because you have a pinhole, you need a reasonably small f number in order to get a large number or even some different color bands. And that's basically the limitations.

And so future work, one of which you just mentioned is actually using a mask over here. And I'll go into that just a little bit. And one of interesting cubed application, there's actually a company which does that they use this RGB wavelength multiplexing in order to do red and light-- the red and right eye separation. So, they project using--

PROFESSOR: [INAUDIBLE]

ANKIT MOHAN: Left and right. Left and right. And so they use these projectors which have-- when you view them with naked eye, both this and this appear blue because they're both in the blue wavelength. But they're actually different parts of the wavelength that's being occupied by those two projections. And if you put a color filter in front of your eyes, a different for each eye, you can have the right eye only see this part and the right eye see with that part.

So, when you view it visually, it would appear just like a normal projection. But if you put on the eyeglasses, you can actually distinguish between the two and you have multiplexing in wavelength rather than in time or in polarization. So far, we've talked about what's called the spectral light field. But it turns out, you can actually combine the concept of a light field with the spectral light.

So, what we have here is a traditional [INAUDIBLE], now it's x and θ . And you've placed a prism in front of this. So, what this prism has done is in a direction that depends on the distance between the x plane and the prism plane, you're actually going out and smearing it and getting all the wavelengths in that direction. So, it gives you what looks kind of like that or something like that.

So, you have a whole bunch of overlapping spectral light fields on a spatial light field or you can think of it the other way. You have a whole bunch of overlapping spatial light fields of a spectral light field. So, I mean, you can either put x θ here and λ here or you can replace it and put λ here and θ here. And it's really not very different.

But what becomes interesting is the case where this plane is diffused in front of parallel. In which case, each of these rays are exactly the same intensity because it's a [INAUDIBLE] or diffuse scene. So in this case, what you get-- the light field that you get over here is the special light field that we call a blurred light field.

It's because you have-- you're taking the standard spectral light and then blurring it in this direction. So, what you could do is what Kevin was pointing out is use-- if this blurring is not just a box function, but a special coded function, you could then invert that and you could get the full spectral light field out of it without having to use a pinhole. So, this is what that design looks like.

So you can put-- so, before I get into that, one of the things I didn't talk about is so you can-- by putting a pinhole and putting a light field camera over here, you can simply by using a standard lightweight camera, you can capture a multi-spectral image. So you can build a multi-spectral image by taking a light field camera and putting a prism, a lens, and a pinhole in front of it. So, what you're essentially doing here is replacing the-- is getting rid of the θ component of the light field and replacing it with the λ component of the light field.

So, once again, you had this light field and now you put a light field camera in front of it, you can sense that data. And this is initial results that we had in that where each-- it's basically a microlens [INAUDIBLE] that's placed next to the sensor and so you can see-- you still have the Bayer mosaic. So, you have this red to blue back with green in between, which is what it's sensing.

But another thing you can do is put what's called a linearly variable interference filter. So, this filter allows only red wavelengths to come through at one end green in the middle and blue at the other end. And this is a very spectrally pure filter. It's basically built out of multiple layers of coating and using interference. And what this does is that it replaces the θ component of the light field with the λ component of the spectral light field.

And now if you put a light field camera in front of it, you can again capture the complete light the complete multi-spectral image. The trade over here is that any one point on this filter allows only one very narrow wavelength to go through. So, it's replacing a pinhole in space with a pinhole in wavelength. And so again, you're going to get very little light.

But what you can also do is place instead of that filter, you can place one of these carefully designed masks, which are similar to the mask we use in the [INAUDIBLE] approach, but I guess we haven't talked about that.

PROFESSOR: No, we haven't.

ANKIT MOHAN: So, you can do this thing and then use deconvolution in order to recover the multispectral image.

PROFESSOR: Let's take a short break and then we can--

ANKIT MOHAN: I just have this thermography stuff that you said--

PROFESSOR: Yeah, we'll do it after the break.

ANKIT MOHAN: OK, cool.

PROFESSOR: Cool, excellent.