

[SQUEAKING]

[RUSTLING]

[CLICKING]

JOSH
MCDERMOTT: So today we're going to finish talking about color perception. I mentioned I did a little bit of reorganization of the ordering of things for logistical reasons. So the sequence of topics on the syllabus has been updated. But today, we'll talk about color, next time motion.

Last time, we started talking about how the world is full of colors, how color is important for lots of things we got to do. Segregating objects, recognizing objects, telling what state objects are in, whether fruits are ripe or not, or whether meat is rare or well done, learning things about the state of the environment around us, and so on and so forth.

And so color depends on the fact that light comes in these different wavelengths. And if you look at light in isolation, it will appear to be colored if you have a particular wavelength that you're looking at. But really, we think the function of color perceptually is to provide information about what objects are made of, pigmentation.

So we talked about the elementary school theory of color, which is selective reflectance. So objects appear colored if they reflect some wavelengths of light better than others. We all learned this a long time ago, and there's some truth to it. So you can see this if you look at the spectral reflectance of different kinds of surfaces. The additional wrinkle, which we're going to talk about today, is that the light that is reflected off of an object is a function not just of the reflectance of the object, but also of the spectrum of the illumination. So those two variables are confounded, just like they were in lightness perception, where we were living in a world of gray.

But in order to perceive color, you have to transduce information about wavelength. And we do that via the three cones. And so we talked about how we have these three types of cones, how they have different spectral sensitivities. We talked about what spectral sensitivity means, how you would measure it, and then spent a while kind of talking about this idea that the cone transduction process you can think of as a mapping of this very high dimensional thing in the world, which is the spectrum of the incident light at a given point on the retina, down to a three dimensional space of cone responses. And that number three is because of the fact that there are these three types of cones. So if we had four cones, it would be four dimensional. If we had two, it would be two dimensional, and so forth.

So this is the spectral image formation process. So there's a light source in the world. That's the illumination of that reflects off of a surface that is characterized by some reflectance. So both of those are functions of wavelength. So this is the spectrum of the illumination. This is the reflectance, which varies with wavelength.

Those two things get multiplied, and that produces the spectrum of the scattered light at a given point in space. That is then projected onto the cone sensitivity functions, and that yields three numbers, the cone responses or absorptions of the long, medium, and short wavelength cones. So those three numbers are what you got to go on for the basis of color perception.

But the fact that we have got three of those cones is critical. If you just had one of them, you would not be able to distinguish different wavelengths. But with three, you have some ability to distinguish different wavelengths, because different wavelengths will, in general, produce different ratios of cone excitations.

But you're still throwing out a lot of information, right? Because you start with this very high dimensional thing that gets multiplied by the sensitivities of the photopigments, and you get three numbers out. So you lose a lot of information when you go from this really high dimensional thing down to three numbers.

And the consequence of that is metamers. So remember, metamers are physically distinct stimuli that appear the same to an observer. And in the case of color vision, they appear the same because of the fact that we've only got these three cone types. So these are just two examples of light sources that have, obviously, really different spectra, but would produce the same responses in the three cones and thus would be indistinguishable.

And so here's some other examples. These are three different light spectra that produce the same cone excitations. And so we talked a little bit about that. And typically, the way to think about metamers and the conditions in which you would experience metamers really is to think about color transduction as linear algebra.

So metamers are, on the one hand this kind of pretty cool thing about perception. This idea that you can have these things that are physically really different, but that are indistinguishable to an organism just because of the way the eye is set up. They're also practically really important, because they're the basis of color reproduction in TVs and so forth.

So what it means is that because you have only three cones, you can reproduce any color experience just with three different color guns. So three different light sources, for instance. You can combine them in different ways and create any particular color. So that turns out to be very convenient.

So then we ended by kind of talking about how a long time ago, before scientists had the ability to go poking around in the eye or in the nervous system and uncover the responses of the components of the eye in the nervous system, people had hypothesized that color vision was trichromatic because of this fact. So scientists discovered that you could take three different light sources and combine them in different ratios to create the experience of all these different colors. So from that, they inferred this idea that there would be three receptor types.

But then there was controversy about this because of these phenomenological problems with trichromatic theory, which is to say that there are relationships between these pairs of colors. And that shows up in colorblindness, where you tend to be red green colorblind or blue yellow colorblind, and not just red colorblind to red or not red blue red colorblind. And there's also these observations that you can make a red look more yellowish or more bluish, but you can't really make it more greenish. It's just not psychologically possible.

And so Hering, who is this other perceptual scientist from back in the day, hypothesized that the receptors had this opponent nature, that they would be measuring the difference between different colors. So nowadays, these two views have been reconciled. And this is where we kind of ended last time, with this idea that you start out with three cone types, the S, M, and L cones, and then combine those in different ways to yield three channels that, in this case, come out of the retina, as it turns out.

So by taking the M and the L cones and computing a difference, you can get what is called a red green opponent channel. By summing the M and L cones, you get an achromatic channel, which people normally think of as luminance. And then by taking the difference between the achromatic channel and the S cones, you get what is referred to as a blue yellow opponent channel.

And so this opponency is set up at the level of the retina. So if you look at ganglion cells, you very often see that there will be a particular type of cone that provides input to the center and then a different type of cone that provides input to the surround. So remember, ganglion cells in the retina typically have center surround receptive fields. They come in these two flavors, on center, off surrounds, and off center, on surround.

And the midget ganglion cells, remember, which project to the parvocellular layers of the LGN, tend to be wavelength sensitive. But they're wavelength sensitive in this particular way. So they tend to be color opponents. So there'll be one cone type that provides input to the center and another to the surround. And they come in these particular pairs. So red and green and then blue and yellow. And in the blue yellow case, it's not one cone type. It would be the combination of the L and the M.

And so the consequence of that color opponency, we think, is that color is organized into a 3D opponent space. So the psychological space of color kind of has these three opponent axes. So one that goes from dark to light and then these two that vary in color from red to green and blue to yellow. And as you transition from blue to yellow, you go through the center of the space where things are neutral or gray.

So you can also describe this in polar coordinates where you have hue, that's like polar angle, and then saturation. That would be kind of the distance from the center. And then value, how light or dark is it. So just different coordinate systems for thinking about color.

So here's the interim summary. The wavelength composition of light is represented in the cone responses. So the wavelength composition of light is the basis of color vision transduced by the cones. Metamers result from trichromacy. The cones are then combined to form opponent channels.

But there's this thing that we're going to have to confront, which is that light that is reflected off of a surface is the product of the illumination and the reflectance. We think of color as being the perceptual correlate of reflectance. And the fact that the wavelength that comes off of a surface is also influenced by the illumination is a challenge that our perceptual systems are going to have to deal with.

And the problems are that the cones alone, they just measure the spectrum of light entering the eye. And so if we want to actually disentangle these two causal factors in the world, the illumination and the reflectance of the surfaces that reflect light, we're going to have to do some kind of perceptual inference. And this is another classic example of an ill posed problem.

Now, the fact that we can successfully make this inference most of the time is referred to as color constancy. And this has been appreciated for a long time. So this is actually a quote from Hering back in 1872.

So he wrote, "The approximate constancy of the colors of scene objects, in spite of large quantitative or qualitative changes of the general illumination of the visual field, is one of the most noteworthy and most important facts in the field of physiological optics." That's what they used to call vision. "Without this approximate constancy, a piece of chalk on a cloudy day would manifest the same color as a piece of coal does on a sunny day." So that's like lightness constancy. "And in the course of a single day, it would have to assume all possible colors that lie between black and white."

So here's why this is a big problem. This is three examples of the spectral composition of different light sources. So the y-axis is relative energy. The x-axis is wavelength. So the blue curve plots the spectrum of daylight, some particular time of the day. So that's what you get. The red curve plots the spectrum of a tungsten bulb. So that's the traditional light bulb that used to be in most kinds of lamps. And so that's pretty heavily biased towards long wavelengths. And then that crazy looking green curve is what you get from a fluorescent light.

So these three spectra are just obviously wildly different. And the consequence of that is that the spectrum of light that is reflected from surfaces will vary quite enormously. So the top graph is showing the spectrum of light reflected from surfaces under daylight illumination. So we've got a leaf, an orange, and a tomato. And so again, the relative differences between these is consistent with the elementary school theory of color vision. Because the red object, the tomato, is reflecting more long wavelength light. The green object, the leaf, is reflecting more middle wavelength light. So that kind of makes sense.

But then the bottom graph shows the light that's reflected off of those same three surfaces from a fluorescent light source. This is not fluorescent, but you probably have fluorescent lights maybe at home in your kitchen, or you go to a department store, you might encounter it. And so the wavelength composition of the light that's coming off the surfaces is completely different in those cases. Yet for the most part, the color of these surfaces looks pretty similar.

Now, you're also probably-- you've probably encountered failures of color constancy. So occasionally, you'll be in some department store. And there's something funny about the light source. And you either won't quite be sure of what the color is of some piece of clothing you're thinking about buying. Or maybe it looks one way, and then you get home and it looks different. So occasionally there are these failures. And those are interesting to think about and probably diagnostic of some of the ways in which we're solving this little posed problem. But most of the time, color constancy works pretty well.

So you might say, well, maybe this is some artifact of the modern world, the fact that we got fluorescent lights and tungsten light bulbs and stuff. And that doesn't seem to be the case. So even daylight can vary across the time of the day and season. And so these are two different spectra of daylight at different times of the day. So the relative proportion of short and long wavelength light can change.

And the basis of color vision as we've discussed is the cone responses. And so really what ultimately matters is whether those are going to be affected by this. And they are. And so this is showing three different colors of paper, yellow paper, red paper, blue paper under two different illuminations from a tungsten bulb and the blue sky. And what these graphs are showing is the spectrum of the reflected light. And then up at the top, those bar graphs are the responses of the three cone types. And so the point is that the three cone types for blue paper under a tungsten bulb are pretty different from that of blue paper under a blue sky.

So you can see that for the blue paper under the tungsten bulb, because the tungsten bulb has got so much power at long wavelengths and because the blue paper doesn't reflect very much of those, the overall responses are low. Under a blue sky that's got more short wavelengths, you get a big boost in the response of the S cone. And there are analogous changes in the other colors of paper.

So different illuminations produce different cone responses. And yet the pigmentation of the thing in the world remains the same, and it looks pretty similar to us, suggesting that we're able to somehow compensate for that in some way.

And as is the case with lightness constancy, insight into the nature of color constancy comes often from illusions. So this is a pretty famous one. So these are two images that were kind of constructed so that different squares here that look different colors are actually physically the same.

So in particular, compare the blue checks on the left with the yellow checks on the right. So they look really different. So blue and yellow. But they're actually physically the same shade of gray. So we just kind of have whited out the rest of the image. And you can verify that they're the same.

So the point of this is that physically what's coming off of the page at these points is what we call gray light. So spectrally, pretty neutral. But when you see the entire image, that physically gray light is interpreted either as something that is blue in the world or yellow in the world.

So the high level intuition here really is that this is very much analogous to all of the principles that we talked about in lightness constancy. So remember, in lightness constancy, we saw these illusions where you could have the same gray patch and it would look lighter in one condition and darker in another. And in general, the conditions in which it's lighter are those in which there's evidence that it's in lower illumination. So it's as though you take the same physical stimulus and compensate for the low or the high illumination. And so the same kind of thing is thought to happen here.

So stated again kind of informally, if you add yellow and blue light, you get gray. So if there is gray light that is projecting to your eye and your brain thinks that the illumination is very blue, so biased towards short wavelengths, then the only way that you could have gray light on your eye is if the surface patch is yellow. And the converse holds if the illumination is very yellow.

And so in these kinds of illusions, what's happening here is there's a lot-- in the one on the left, there's a lot of evidence that the illumination is very yellowish. In the one on the right, there's a lot of evidence that the illumination is very bluish, because all this stuff has kind of got a lot of blue tint to it. So again, unconsciously, without you really thinking about it, your brain seems to be compensating for the fact that there's evidence for blue illumination here and yellow illumination here, and then taking that thing that's physically gray and inferring that it's blue and yellow, respectively.

Here's another example. This one is a little bit more subtle, but same essential idea. So the point is that the second card in this person's hand is physically the same shade of light pink. But then the one on the top looks really kind of pink and the one on the bottom looks a lot more white.

And again, if you look at what's going on in the image, there's evidence here of bluish illumination up top and pinkish illumination on the bottom. And so you attribute the pink light here to the illumination and thus it looks white. Whereas there, you attribute it to the paper. Any questions about these kinds of illusions?

So these illusions reveal color constancy at work. And again, the logic that's kind of written out in this page, I think, is now fairly familiar, because we have leaned on it pretty extensively at various points throughout the course. The notion here is that the point of vision is to recover properties of the world, not the image.

So the visual system is not designed to just go around measuring light. It's designed to tell you what's out there in the world. The information in the image is used to infer the properties of the world, not the properties of the images in our eyes. And so we call these things illusions, because you're seeing something that at some level is divorced from the physical thing that you would measure from the image. And we think that gives us some understanding of why the colors of objects don't appear to change much when they're in different illumination environments. And it's very much analogous to lightness constancy that we discussed in the previous lecture.

So how does all of this work? So there are similar debates, as with lightness perception, regarding the role of low level filter based mechanisms versus mid and high level mechanisms. And as was the case with lightness perception, I don't really have a totally nailed down theory that can account for all of the data that's available. We have bits and pieces of clues for the ingredients of theories.

And again, some of the pieces of evidence are very much analogous to things in lightness perception. So this is a color analog of simultaneous contrast. So it's a little hard for me to see. Hopefully you can make out that up top, there are these little squares kind of in the middle big squares.

And you can probably tell just by looking at it that the one on the left looks kind of brown and the one on the right looks kind of green. Actually, they're physically the same. So this is a color analog of the simultaneous contrast illusion that we talked a lot about in the context of lightness perception. So the idea is that the surround here causes that little square in the center to look different.

So again, with these color illusions, these things are not happening at the level of the photoreceptors. So the photoreceptors are going to give you the exact same response to those two little squares. So the fact that they look the same is the consequence of some subsequent processing. And some of that, again, could be potentially center surround receptive field type operations that look like taking differences between this. But there's also the possibility that there's kind of higher level things going on.

And this experiment here is kind of a cool example of that that provides some evidence for some fairly sophisticated inferences. And so, Fernanda, do have those things? All right. We can do a little in-class demo of this. So just distribute those evenly. You're going to have to share.

So this is a really cool experiment that was done back in 1999. And the idea is to look for a role for inferred 3D structure in color constancy. And so what some of you are holding in your hands is a folded piece of paper that is magenta on one side and white on the other. And so this is a diagram. Magenta on this side, white on the other.

So if you just hold the thing up and look at it, one side looks magenta and the other side looks white. That's what it is. So the thing that's happening here, though, that you're not aware of is light is coming from a light source and reflecting, to some extent, off of the magenta surface and then off of the white surface and then into your eye.

And so the consequence of those different reflections is that the fact that there's a magenta surface on one side of the card is causing the white surface to actually reflect some light that is tinted towards magenta due to that pattern of reflections.

And so this is actually a physical measurement with a photometer that demonstrates that. And so this is the amount of light that is coming off of the surface at different positions. So this is the magenta side. And so there's a lot of pink light coming off there. But this is the white side. And you can see that there's a fair amount of pink tinted light that's physically coming off of that side.

Now, the remarkable thing is that it doesn't look pink, but you can make it look pink. And the way that you make it look pink is by holding it in your hand, closing one eye, and then waiting. So when you close one eye, the three dimensional structure will become a little bit ambiguous. And it can even become bistable. And so if you hold it and close one eye, the three dimensional structure can reverse. And so you can go from seeing the thing like this to seeing the thing like that. And when you do that, when you get the depth to change, that thing that looked white will start to look pink.

And the inference is that your visual system is making some kind of inference based in part on the 3D geometry that somehow is compensating for the pink light that it knows should be reflected off of the magenta surface onto the white surface. But when you reverse the depth, that reflection is no longer likely to happen, because the surfaces aren't pointed in the right direction for that. And so now that pink light that's actually coming off the surface gets attributed to the surface itself. And so the thing looks pink.

And for me, it even works-- just closing one eye is enough to actually cause the other thing to start to look a lot more pink. That work for you? Yeah. OK. Good. Satisfied customers. That's what we aim for.

And so this is just some data. This is the way that they measure this stuff. So the actual experiment involved showing people these displays and then giving them a big set of colored cards. And they had to pick the one that was the best match to the surface in these two different conditions. And then this is a representation of the color of the card. And the different data points here are the different viewing conditions. So the point is that the different viewing conditions produce different sensations of color.

So this is some evidence that there is a fairly sophisticated element to the inference of color constancy that apparently to some extent involves 3D geometry in a fairly sophisticated way. This idea that you're going to get inter-reflections and certain three dimensional configurations and not in others. What questions you got about this? Yeah.

STUDENT:

This is slightly related to the previous p-set. There was one question that was asking about cortical and subcortical processes. And it talked about the fact that if you close an eye, so if you're [INAUDIBLE] monocular vision, that implied that it was a cortical process. So is color constancy a cortical process rather than a higher level thing?

JOSH

MCDERMOTT:

So the question was whether or not the fact that making you close one eye here relates to the monocular manipulation from the problem set. And so in particular, on the problem set, I think what you were asked to do is to adapt to this display with only one eye open, and then see if the effects of adaptation would transfer to the other eye. Something like that. I think that's right. So that was about transfer of adaptation.

So this is really just about-- so the reason, and I should have said this, the reason I asked you to close one eye is because that eliminates binocular depth cues, which would otherwise make it very difficult for you to see the incorrect three dimensional structure here. So you close one eye and you lose stereoscopic depth. And so depth just becomes a little more ambiguous. And that enables you to see the thing like this, even though it's actually like this. With both eyes open, that would be very hard to achieve. So the purpose of that manipulation is different, in this case. But yeah, it's a good question.

The dress. So everybody has probably seen the dress. This was probably the first viral perceptual phenomena in the era of social media, I think, potentially. So it's this photo. And is there anybody here that hasn't seen this before? So this was viral, like I said.

So everybody knows that the reason this went viral is that people disagree on what colors the dress appears to be. So some people, and in general, people tend to either see the thing as being white and gold or as being blue and black. And this is the result of a big online experiment with 8,000 people where they asked people to label the colors that they see. And so you can see 60% of the participants say that the dress is white and gold, and then almost 30% of the participants say blue and black. And then there's assorted others that see other things. But most people see it as either white or gold or blue or black.

And so it went viral because people are very confident that it's either white or gold or blue or black and disagree about what it looks like. And so the explanation of this, I think at this point, is pretty well settled. So this is a quote from one of several papers that came out shortly after this. And it's a really nice example of color constancy and the ill posed nature of the problem.

So here's the idea. So as it says, those who assume a relatively long wavelength illuminant, such as from incandescent light, could be expected to see the dress as black and blue, as they would subtract long wavelengths from the image in order to arrive at a relatively shorter wavelength percept of the object. Conversely, those who assume a relatively short wavelength illuminant, such as daylight, could be expected to see the dress as white and gold, as they would subtract short wavelengths from the image in order to arrive at a relatively longer wavelength percept of the object.

So that's kind of the essence of how perceptual scientists think about this. The idea is that different people are making different assumptions about the spectral content of the illumination. And you can also see, if you look at this, the picture is quite ambiguous. You could actually imagine that this might be outside. You could also imagine that it might be inside, taken in a way that there's not a ton of unambiguous cues to the illumination. So that's what people think is going on.

Here's just a graphical depiction of the idea. And so we've got these two different dresses, one that is actually blue and black and one that is actually white and gold. And then those two different dresses are under two different illuminations. So this is kind of a yellower long wavelength illumination, and this is a shorter wavelength, bluer illumination. And the point is that those same two dresses under these two different illuminations create physically identical images.

And that's what this movie is supposed to show. So you can see this little patch here getting taken out. So two different combinations of pigmentation and illumination create the same image.

So what's the evidence for this kind of explanation? So this is the result of an experiment in which people were queried both about what color they believed the dress to be, as well as whether they thought the illumination was natural or artificial. And so the y-axis here is-- and so people are grouped into three different groups here.

So this is the group. This is the people who said they thought the lighting was natural. These are the people who thought the light was artificial. And these were the people who were unsure. And the y-axis is the proportion of each of these groups that report the white gold percept. So what this is showing is that the people who think that the illumination is natural are substantially more likely to see the dress as white and gold than the people who think the illumination is artificial.

Now, again, I mean, this is assuming that people are actually accurate in reporting the representation of the illumination in their heads, and they might not be so. But it's nonetheless in the direction that you would expect.

Another thing that's pretty interesting is sleep schedule seems to influence assumptions about illumination. So in the same study, they asked people to characterize their circadian type. So an owl is somebody who likes to stay up at night and sleep in late in the morning. A lark is somebody who wakes up really early and goes to bed early.

And so the idea is that somebody who's an owl is going to have more experience with artificial light, less experience with daylight, and thus more likely to implicitly assume artificial lighting, because that's what they experience more of the time. That's the idea. Whereas a lark is going to experience more daylight in their lives and will be plausibly more likely to assume natural daylight as being the illumination.

And so this is showing for each of these four self-categorizations the percent that people report the white gold percept. And so remember, the white gold percept is the one that you would arrive at if you assumed a short wavelength dominated illumination, namely daylight. And so you can see that the people that self-categorize as larks are more likely to see the white gold percept than people that self-categorize as owls.

So again, it's kind of consistent with the idea that the way you live your life affects the distribution of illumination that you experience, that that has been internalized into priors, and that that influences your inferences about color.

Another consistency check. This is the same self-categorization into owls and larks. And this is the percent of people in each group that say they think there's artificial lighting. And so the people that are owls are a fair bit more likely to say that they think there's artificial lighting than people who say they're larks.

So that's the dress. The basic story is that this is a really interesting case where there are individual differences in the implicit assumptions that people are making about, in this case, illumination. And you have to make some assumptions about illumination in some cases in order to resolve the ill-posedness of color perception.

So this is one kind of case where there are these individual differences. And this is actually pretty cool, just because it's an interesting case where priors can a little bit be linked to everyday experience and differences in the way that people live their lives.

So a lot of cases, we think that priors are everywhere in perception. We typically don't really have hard evidence as to the extent to which those are hardwired and derived from an appearance on an evolutionary time scale, or whether they are learned over the course of development. But this is one case where there seems to be some role for the experience over the course of development, maybe even into adulthood. So that's kind of cool. Any questions about the dress? Yeah.

STUDENT: Are people able to switch which one they see?

JOSH
MCDERMOTT: I think possibly with enough practice, you might be able to get yourself to. I think the reason this kind of blew up the way that it did is it's actually really not that easy to do that for most people. And so people are just like, well, that's definitely white and gold or that's definitely black and blue. And that's what made it interesting is they couldn't understand how someone could see it differently.

But I think possibly with practice, you might be able to do some of that. A lot of times this stuff is-- it's not very cognitively penetrable. Your perceptual system does its thing. You get the result, and it's hard to access the internal computations that give rise to it.

So how is color perception implemented in the brain? There's a variety of pieces of evidence for some degree of functional specialization, namely the idea that there are particular parts of the visual system that mediate color perception. One really interesting phenomena that provides evidence for this is the phenomenon of achromatopsia.

So this is a phenomenon that can result from certain types of brain damage in which people selectively lose the perception of color. And so the idea is that if there is some bit of the brain that's mediating some particular function, if you have an unfortunate accident, typically a stroke, part of the brain is deprived of oxygen. That part of the brain dies. And you could get some selective damage to a particular type of function.

So this is a quote from a case study. This was a person who was an artist who then had this stroke and lost the ability to perceive color. "So Mr. I arrived at his studio expecting that the horrible mist would be gone. But as soon as he entered, he found his entire studio, which was hung with brilliantly colored paintings, now utterly gray and void of color.

He saw people's flesh, his wife's flesh, his own flesh as an abhorrent gray. He found foods disgusting in their grayish dead appearance and had to close his eyes to eat, but this did not help very much, for the mental image of a tomato was as black as its appearance." Yeah, sounds tough.

But scientifically, really interesting, because it suggests that you can lose the ability to see color while preserving most other aspects of vision. A guy can still recognize objects and navigate around and so forth.

STUDENT: So it seems that his memory was also black and white?

JOSH
MCDERMOTT: Saying the image was black and white.

STUDENT: Like if he imagined the tomato, [INAUDIBLE].

JOSH
MCDERMOTT: Yes.

STUDENT: And if he remembered stuff, like tomato in the past or something ?

JOSH Well, the question is, is the memory different? I mean, in order to access a visual memory, you often have to
MCDERMOTT: form a mental image. And so we often think that the way that imagery works is you activate the visual representation that would result from actually looking at something. And so it could be that the memory is in some sense normal, but you don't have the part of the visual representation that represents color anymore. And so all that's left is the representation of all the other stuff. So you can't represent the color. So that's one interpretation.

STUDENT: But blind people can see images. Like, when they remember stuff, they can't see it? It's just like, different?

JOSH I think so. Yeah, so someone who loses sight later in life. Yeah, I mean, I think they can remember what things
MCDERMOTT: looked like.

STUDENT: Probably depends on how they lose their sight. If you just lose the actual, I mean--

JOSH Oh, absolutely. Yeah, sure.

MCDERMOTT:

STUDENT: [INAUDIBLE] Then you probably wouldn't remember.

JOSH Yeah, and that's probably true. So I was imagining that there's an issue with the eyes. So if you had some
MCDERMOTT: massive stroke, major lesion of the occipital cortex, yeah, I don't think you would have imagery anymore. And in fact, I believe it's the case that if you get a big lesion on, say, the right hemisphere, left hemisphere is intact. So imagery of things that is on the right would be normal, but things on the left would not. That's a whole field which you can check out. It's pretty interesting.

So this is the location of the lesions that typically result in achromatopsia. So they're non-random. So this is a figure that took a whole bunch of cases of people that reported achromatopsia following brain damage. And this is an overlap of the lesions. And so there's particular places that are much more commonly damaged. So that suggests some localization of function.

So in people with normal visual systems or with monkeys with normal visual systems, there have been experiments to try to find regions for representing color. So this is an fMRI experiment that is showing brain regions that respond more to images that vary in color than to those that just vary in luminance. And so you see these little places that are more active to color than to images that don't have color.

And for reasons that I think are unclear, there are multiple such regions, and they tend to be right next to regions of the ventral pathway that are very responsive to faces. So these color biased regions here in the temporal lobe, and they're right next to these so-called face patches where you see lots of neurons that are responsive to faces. So that's just a feature of the visual system. I don't think we have a very clear explanation for that, but it is what it is.

So there's been lots of interest in the relationship of color to other aspects of vision. But first we got a question.

STUDENT: So it wasn't that the man with the condition was only using his rods. His vision didn't get blurry or anything either?

JOSH

MCDERMOTT:

Based on the case report, I believe that to be true. I mean, yes, but the retina is fine. So again, you have a stroke and the eye is not going to be affected. So it's just the part of the brain that is deprived of blood flow for that period of time that would be damaged.

So what is the relationship of color to shape processing? So there's a bunch of interesting demonstrations here. So one kind of empirical fact is that the spatial resolution of color processing is more limited than the spatial resolution of luminance processing. So this is just a really simple demonstration where there's a checkerboard. The middle part of the checkerboard is blue and yellow.

And you can see that as this is scaled down, the very small versions, you lose the sense of color variation. But you can still see the checks. So the black and white variation is pretty easy to discern, even in the smallest one. But the blue yellow variation, you probably lose. So this will depend to some extent on how far back you're sitting in the room.

And some of this probably reflects the receptor lattice. So for instance, the short wavelength cones are just kind of spaced pretty far apart in the cone mosaic. So that's probably part of it. But there's just also a lot of other pieces of evidence that color behaves a little bit differently than luminance variation. So one such piece of evidence is that you can come up with images that can purely contain chromatic variation.

And I'm not going to tell you exactly how you come up with these things, but you can tell from looking at these things that they look kind of weird. And in particular, it's a lot harder to perceive shape in images that are so-called isoluminant. It's also much harder to actually perceive motion. So if you take a isoluminant image and you and you set it in motion, the motion is often really hard to perceive.

And this is another example of the poor spatial resolution of color, where you've got this gray cloud and then an analogous yellow cloud. And it's really pretty easy to see the fuzziness of the border here. This is actually much harder to see. And then you superimpose it on the black circle, and it looks just kind of a yellow, that the black circle is yellow in the middle. When in fact, you have this scruffy border all around.

So there's this kind of intuition that form processing, the processing of shape, may be mostly luminance based, with color kind of filled in around the form. And there's a pretty amazing illusion that I will show you that is consistent with this. So let's check this out. There are these illusion contests and this is the winning illusion from some year.

And so what you're supposed to do here is fixate on that little dot in the middle. And the illusion here consists of this colored pattern that kind of pops up. And then after the color pattern pops up, you see an outline of one kind of part of the pattern. And alternately, you will see one part and then the other part. And they'll be different on the two halves of the screen.

And so the really amazing thing here is that the color that you perceive is really dependent on which of the two outlines you're looking at. And so the thing kind of flips from being blue to pink and vice versa. And what's amazing about this is that the color afterimage, so the adaptation in the cones should be exactly the same, irrespective of which outline you're looking at. So you see that color image. There's a slight amount of adaptation in the cones. So you get an afterimage. But then what you perceive, what you see, is a function of which outline you get.

And so it's kind of consistent with this idea that a lot of color is filled in based on boundaries of objects that are kind of derived from the luminance system. And again, this is all a little bit hand-wavy. We don't have rigorous accounts of this, but there's these very suggestive and powerful illusions that are kind of consistent with that general thing. Yeah.

STUDENT: [INAUDIBLE] the opposite of what it is that [INAUDIBLE]?

JOSH MCDERMOTT: Yeah. So typically, if you generate a color afterimage, you'll adapt to red, and then you'll see something that will be more greenish. So the fact that it looks like the opposite color, that's kind of a fairly standard thing. What was new and remarkable about this was the dependence on the outlines. Any other questions about this? So this was the illusion of the year in 2008. It's pretty good.

So there's this evidence that there's some separation in some sense between color and form in the sense that color is not very good at supplying form. It tends to follow the luminance structure in some of hand-wavy way. But these things are not completely independent.

So if you look in the visual system, you can find orientation selective neurons that are also wavelength selective. So this is an example of a receptive field in V1 where they separately map the receptive field for L cone stimulation and M cone stimulation. And you can see that there's opponency here that's related to the orientation selectivity here. So you get excitatory input from the M cones in one lobe and in the other lobe for the L cones.

So this was a physiological discovery. But what predated that was this very well-known and very powerful after effect called the McCullough after effect. And so what I want you to just do here is start looking at this. So look alternately between the green and the pink stimuli. We're going to do this for about a minute. Just get started on it while I talk to you about it.

So what you're doing while you look at this is you're adapting to these combinations of color and orientation, but you're doing it at the same time. So you're adapting your cones simultaneously to green and to pink. So you might think that they would cancel out. But at the same time, the color is paired with a particular orientation. And as you will see, this will produce a color orientation contingent aftereffect. It's called the McCullough effect, named after the person who discovered it.

This was initially taken as evidence that there are neurons that are tuned to conjunctions of color and orientation, because that's the traditional logic of after effects. So if you get an after effect for something, that usually is an indication that you have neurons that are tuned for it, that you're adapting out when you adapt. But as we will see, the story may be a little bit more complicated than that.

So you're adapted out. Now look at this. And what you should see-- so first look at the one on the left and then look at the one on the right. And these gratings should look slightly tinted, either pink or green. But the tint will depend on the orientation. So if you on the one on the right, you'll see one particular tint in the middle. And then you look at the one on the left, and you'll see a different tint in the middle.

So this is not happening at the retina. So it's very different from the photoreceptor adaptation that we've talked about. It's dependent on the orientation. So it's something more central.

So the traditional kind of naive interpretation of this is that there are neurons that are tuned for combinations of orientation and color and that you just adapt those neurons. Now, the reason that people are generally skeptical of that kind of explanation is that this particular after effect, for reasons that are unclear, lasts for a very long time.

And so we'll actually look at it again in a few minutes, and you'll probably still be able to see it. If you go home tonight and you look at this again, you'll probably still be able to see it. It will eventually go away, but it lasts a really long time. And so people think that it may have a bit more to do with perceptual learning or something than a traditional after effect, but it's a pretty remarkable effect. It's worth knowing about.

A few other phenomena. I should also just stop and say that in this lecture and the next few lectures, I'm going to show you a bunch of phenomena. And part of the reason I'm showing you these things, well, first of all, they're kind of fun, but also because these are potential food for thought for the illusion lab that is coming up. So I don't expect you to memorize all of these things. I just want to introduce you to a few of these in case they give you inspiration for your illusion labs. And they also just provide a little bit more context.

So these are a couple of other interesting cases where color interacts with shape. So it contributes to contour grouping. So we saw some stimuli like this a couple lectures ago, where you can detect the kind of continuous path of the Gabors. But if they alternate between being chromatic and achromatic, like they do here, on the right, it's a lot harder to see. So that kind of suggests that color is contributing to contour grouping in some way.

It also interacts with something called shape from shading. We haven't talked about shape from shading yet. We're going to talk about it in a fair bit of detail in about a week or two. But shape from shading is this phenomenon whereby if a surface is Lambertian, then the amount of light that is reflected will be a function of the orientation of the surface. And so if you have something that is curved, there will be a gradient in the luminance as the shape changes.

And so this is a pretty interesting demonstration where you get a sine wave grating. So you get this variation in luminance here. And maybe if you are feeling really good, you can see that as a corrugated surface. But then if we take this color grating and we just superimpose it on the luminance grating, now you get a pretty compelling sense of a corrugated surface.

And I think that the idea here is that you almost never essentially would get color variation due to shape from shading. And so if there is color variation in an image and then you get luminance variation, that's kind of like uncorrelated with that or in this case orthogonal with that, it's very likely that that luminance variation is not due to reflectance. And so you would interpret that as actually being more likely due to shape. So the color here is actually kind of enhancing shape from shading.

And we'll talk a lot more about shape from shading when we get to that. This is just another example of color interfacing with shape. So although there's a bunch of things about color that are kind of weird in the sense that isoluminant images don't do a very good job of supporting shape, you tend to fill in color based on luminance boundaries, there are also these complicated interactions with shape as well.

Last thing we're going to talk about is color blindness, scientifically usually called color deficiency. So most color deficiencies are due to one or more cone types being either missing or abnormal. So the most common thing is dichromacy. So that is when one cone is missing. Because there are three types of cones. There are three different varieties of dichromacy. Deuteranopia is the scientific term for when you're missing the M cone. Protanopia means you're missing the L cone. Tritanopia means you're missing the S cone.

Colorblindness is typically diagnosed with these color plates. How many people have had a doctor show these things to them? I remember my daughter had a checkup when she was three or four or something over here at MIT Medical. And they brought out these colored plates and they were asking her what the numbers were, and she didn't know her numbers very well at that point in time.

And so the test depends on the person saying the correct number. And so she kept saying the wrong number, and I felt really embarrassed. But she's not actually colorblind. So that was my experience with these color plates. But that's how you diagnose colorblindness by introducing differences that somebody who's missing one of the cones would not be able to see.

There's also anomalous trichromacy. So that's when one of the cones has a slightly abnormal spectral sensitivity relative to the normal population. That typically occurs in the M or the L cones. Again, deuteranomaly and protanomaly. And these are genetic in origin. So most color, they call them defects here or abnormalities, they involve the M and the L cones.

And those are coded on the X chromosome. And so they're much more common in men, and actually pretty common in men. So about 8% of men have red green color blindness because of missing the M or the L cones. You probably know people who are like this. So these are like the incidence of these different types of color abnormalities. And you can see that they're basically all way more common in males than females.

So there's a couple other kinds of things that can happen. So this is very, very rare, but some people are rod monochromats. This means they don't have any cones at all. So as it says here, we're all rod monochromats at low light levels. Because remember, at low light levels, the cones don't operate. So you're just reliant on your rods. That means that you don't have color vision. So rod monochromats typically would have to wear sunglasses because under normal light levels, the rods would be saturated.

And then as we've discussed, you can also get cortical color blindness. So that's achromatopsia. So that's if there's brain damage to the cortex that causes you to not be able to see color. So two very different sources of individual differences in color. Any questions about color? Yeah?

STUDENT: Would people with the first one see a little bit blurrier and would glasses be able to fix that?

JOSH
MCDERMOTT: Yeah, the resolution would be worse because the resolution of the rod system is a lot worse. No, glasses are not going to help that. There would not really be-- I don't know of any intervention that's currently available that would be able to address that. Yeah. I mean, there's a lot of interest in retinal implants. And that, in principle, might be able to help with that.

So how many people still see an aftereffect here? So it lasts a long time. It's pretty crazy. So it will eventually go away.

So high level summary. Color is the perceptual correlate of selective reflectance. It is constrained by the three cones. Metamers result from trichromacy. So the fact that you take this very high dimensional thing, the spectrum of the light that's coming into the eye, and turn that into three numbers at every point in space. The cones are subsequently combined to form opponent channels. We haven't really talked about why that is. It's mostly just I've treated it as a fact, and it has consequences for perception with the opponent relationships between different colors.

Humans have to discount the illuminant to achieve color constancy. The mechanisms of color constancy and a formal computational level understanding, that remains a work in progress. It pretty clearly involves priors over illumination that can differ between individuals, as we saw with the dress. And there's some evidence for specialized brain regions for color processing, although lots of different pieces of evidence that color kind of interacts with other parts of vision, shape, and so forth, as we saw with some demonstrations.

So motion occurs when things in the world change position over time. And today we're going to talk about how the visual system detects this change.

So here's an example. So this is a square of random dots. And there is a square of dots in the center that will move left or right between frame one and frame two. And it's very easy for you to see that motion. And when it stops, you lose the sense of the square.

So people are very sensitive to when things move. We believe that there are neurons pretty early in the visual system that are selective for motion. And as we've just been discussing, one way that we can probe for the presence of neurons that are selective for a stimulus property is to test for aftereffects from adaptation. So we're going to do this here.

So there's a link here for an opportunity for you to try this out yourself. We're going to attempt to do this live. This is an umbrella. It's a really old umbrella. And it was painted a very long time ago with a white spiral. Unfortunately, the paint has slowly deteriorated over time. But what I'm going to do here is I'm going to rotate this for about a minute. And you all have to stare at the center. And so you're going to adapt yourself.

So just start doing this, and hopefully you'll get a sense of-- I don't know if it's expanding or contracting. One of the two. So we're just going to do this for about a minute. So just stare at the center here. And this is going to do some motion aftereffect. The motion aftereffect is one of the most powerful aftereffects in all of perception.

So keep looking at this. And hopefully all those little imperfections in the paint are not going to totally mess this up, but I don't know. Just keep staring at the middle. You're getting sleepy. Keep staring at the middle. It's kind of hard to turn this thing very consistently, but I'm doing my best. We're going to do this for about 10 more seconds.

And then what you're going to do is you're going to look at my face. And I'm not going to look the same. Look at me. Did my face distort? You can also look at the back of your hand. This effect doesn't last that long. This goes away on the time scale of normal visual aftereffects. You just missed a great demo. It's OK.

So that's the motion aftereffect. And you can get little tops that have these spiral things on them. So you spin the top and you look at it and then you can get a really nice aftereffect.

So we typically think of these aftereffects in terms of population codes. So remember the tilt aftereffect. The idea that you've got this population of neurons with different orientation tuning. You adapt. The gain of some subset of neurons is reduced, and thus the inference about the orientation is altered. So the same thing applies to motion. The idea is that you have a population of neurons that are tuned to different directions. You adapt to a particular direction. That subset of the population has a reduced response, and then the inference about the motion is altered. So you can think through the logic there.

So we got evidence that there are these motion detectors early in the visual system. One way to think about detecting image motion is the Reichardt detector. So you can construct something that would be responsive to image motion by wiring together two neurons that have spatially displaced receptive fields here and here with a time delay for one of them.

And so the consequence of the time delay is that the input from this receptive field and the input from this receptive field would reach this at the same time for one direction and not for the other. So that's the essence of the idea behind the Reichardt detector. And if you suppose that the downstream neuron was kind of performing an AND function or a coincidence detector, that's something that neurons often do, you could get something that would be responsive to motion.

So you could also take a neuron that is direction selective and actually look at the receptive field. So remember the spike triggered average. This is the stimulus that, on average, precedes the spike. So we can take some random movie, play it to a neuron, record the frames, the sequence of frames that precedes the spike and then average them, and that gives us a movie. So it's the movie that was the average thing that preceded an action potential by the neuron. So it's one way to think about the receptive field.

And so these are the frames of the movie. 50 milliseconds, 100 milliseconds, 150 milliseconds, and 200 milliseconds before the spike. And if you look closely, you see this orientation selective thing at every frame of the movie, but you can see that the phase is different in the different frames.

So the exact spatial pattern that the thing is most responsive to at a given point in time changes over time. And so this thing, if you played it, would actually look like a stimulus that was moving.

So another way to think about motion detection is that motion is like orientation in space time. And spatiotemporally oriented filters can be used to detect and measure it. So that's kind of weird to think about, because we haven't really been thinking about filters in that way. But time is just another dimension. And if you've got something that is moving through space, so in this case, this is one frame of a movie.

So we just have a vertical bar. So this is the x dimension of the image and the y dimension of the image. And now we introduce time. And if we let that bar translate, in this case in the x direction over time, we get this thing that is oriented in space time. So motion is orientation in space time.

And so if we look at the spike triggered average, we can actually see this. So again, this is a movie. It's three dimensions. There's the x dimension, the y dimension, and the t dimension. So now we're just projecting down onto the xt plane. And we see this oriented thing. So here it is rotated to make it a little bit easier. So orientation here.

So two different ways to think about motion detection. The more of implementation level explanation is this Reichardt detector kind of thing where you can imagine if there was a way to put in a delay in the connection between one neuron and another. You could kind of set something up like this. It's maybe a little bit less obvious to you how you would create something that would do this, but this is a more computational level description of what the detector ought to be doing. And if you actually look at this filter kind of characterization of the receptive fields, you see these things that are oriented in space and time. So there's now lots of evidence that you find these things in early stages of the visual system, in particular V1.

So one of the problems that you run into is that the local motion detectors that you find early in the visual system on their own can't determine the true direction of motion. And the essential idea is shown here. So the motion detector is generated with these oriented inputs, as we saw in the previous slide. So if you have two of them that are at different spatial positions, one with a time delay, you wire them together, you get your motion detector.

But these two stimuli will both elicit the same response from the neuron, because you end up with this stimulus that is overlapping the excitatory lobe here at this point. And then it moves over here. And if it reaches this with the correct time offset, you're going to get a response.

But this is a different direction of motion that would produce a similar response, because it's got the correct spatial offset in the right direction. So two different directions of motion. They're both going to elicit a response in this neuron.

So motion detectors that have these oriented inputs are not able to determine the specific motion direction and speed of a 2D signal. There's a sense in which they can only see the component at their orientation. So this raises this question of how then does the visual system detect the motion of 2D features that are two dimensional?

And so one kind of important idea is that although the responses of one of these local motion detectors are ambiguous, they're nonetheless giving you a lot of information about motion. And in particular, they're giving you a constraint line in velocity space. So if we think about the stimulus in velocity space, so this is the x component of the velocity and the y component of the velocity, a response in one of these detectors tells you something about a component of the velocity.

And it leaves the other component kind of unconstrained. But that one component gives you a line in the space of velocity, which is all the velocities that could have caused that detector to respond. And it just kind comes back to this idea that really what matters is the spatial offset in one direction. This direction doesn't matter for the response of the neuron. And so there's this line in the space of velocities, which is caused by all the different y components that you could add to a particular x component.

So any single detector with a particular orientation in its input gives a constraint line in velocity space. And so it follows from that if there were two different detectors that were both responding, they would each give you a constraint line. And if you combined those two constraint lines, the velocity could be determined from their intersection.

So the idea is that if two detectors are responding, one detector is telling you the stimulus is somewhere on this line. The other detector is telling you the stimulus is somewhere on this line. There's only one point that's consistent with both, and that would yield the velocity. So this is called the intersection of constraints model of motion detection. This was popularized by two people back in the early '80s, Ted Adelson and Tony Movshon.

And so one way to think about this in terms of receptive fields and in terms of spatial frequency components is remember how we talked about how every image can be represented as a sum of sine wave gratings? This is just Fourier's theorem. So every translating pattern can be decomposed as a sum of 1D sinusoidal components. So you got that random 2D pattern. You could generate that by adding together a whole bunch of different 1D components. And you can think of any neuron in V1 as kind of tuned for one of those components. It's tuned to an orientation and a spatial frequency.

And so if you wanted to build a detector for a particular 2D velocity, all you really need to do is to add up the responses of the V1 neurons whose preferred orientation and speed, namely a constraint line, is consistent with that 2D velocity. So let's suppose we wanted to build a detector for something that's moving right with this speed. So then we just take all of these different constraint lines that intersect this point. And each one of those corresponds to a particular elementary motion detector. And we just have those kind of converge and project to the same neuron.

So this was a very influential idea that suggested that motion processing might involve two stages. So there's this initial stage where the stimulus is decomposed into a set of 1D components. So you can think of that as what happens in V1 where you get Gabor-like receptive fields. And then a second stage where these 1D components are combined using a form of these intersection of the intersection of constraints algorithm.

So this was initially something that was a theoretical idea. There was then some psychophysical evidence for it, which we'll discuss, and then some physiology experiments that provided some confirmation that this occurs. And this was very influential because it was a case where there was this theoretical idea that there might be these two stages of processing. And then people went and looked in the brain, and those two stages of processing seem to map onto two different visual areas in a relatively clean way. So this was a very impactful, early demonstration of how computational stages might be associated with stages of the visual system.

And the story here involves an area called MT. So MT is a visual area that's midway up the dorsal or parietal pathway. Remember, we can think of the visual system is kind of coarsely structured into these two pathways, one that kind of starts out with the midget cells and then the parvocellular layers of the LGN all the way up to the inferotemporal cortex. That's often called the ventral stream. And then the dorsal stream predominantly gets input from the parasol cells and the magnocellular layers of the LGN. And there are these distinct subregions of V1 and V2 that then project to area MT. And this is where all this stuff is in the brain.

So MT is very noteworthy in that the responses are very distinctive. So most of the neurons in this particular visual area are selective for the direction of motion. So this is a graph that shows the proportion of neurons that are selective for different stimulus attributes. So there's one bar for MT. There's another bar for area V4, which is at the same level of hierarchy but in the ventral stream.

So those two regions are often compared. And the point is that the black bar here, which is for MT, is 95% or something. So almost all the neurons are selective for the direction of motion. By contrast, almost none are selective for color. And so with V4, you see the opposite. So some segregation of function. It's also the case that if you lesion MT, like in, say, a monkey, that produces deficits in motion perception.

So the basic idea of these experiments that I'll tell you about are to probe for this type of processing using stimuli that are composed of two sine wave gratings. And these are stimuli that came to be known as plaids. And so the idea is that each sine wave grating will stimulate a different elementary motion detector.

And so this is what the sine wave gratings look like on their own. So these are two different gratings, one vertical, one horizontal. They're drifting. It's not very exciting right. They just move. But we're going to add these two things together, and that creates the plaid.

And so the first interesting observation is that even though we think that this stimulus, when it's represented in V1, different populations of neurons will be picking out this component and picking out this component. That's what those receptive fields are doing. They're orientation selective and tuned for spatial frequency. So the stimulus is going to get decomposed into these different components, in the same way that the cochlea takes some sound that's got lots of frequencies and decomposes it into different frequencies. So we think that in the early visual system, this thing gets decomposed.

But when you look at it, you'll see a single thing that's kind of moving in a different direction than either of the two gratings. So drifting downwards and to the left. So that's a plaid. So that suggests that even though early in your visual system, this thing gets taken apart into these different components, those things are then reassembled or combined in some way to infer the direction of motion in the pattern. And so when we come back, we're going to talk about physiological evidence for this happening in area MT. All right. Let's break there.