

[SQUEAKING]

[RUSTLING]

[CLICKING]

PROFESSOR: All right. So, so far in this class, we've introduced the basic idea of perception as inferring the structure of the world from sensory input. And we spent about the first third of the class talking about the sense of hearing. So remember, we started out talking about the ear and how the ear is this device to measure sound. So sound is a mechanical vibration. Your ear turns that into electrical signals that get sent to your brain.

And then we talked about how the problem of perception is to take those electrical signals that you're getting from your ear, and then to infer what's happening out there in the world. So the problem of vision is really analogous to that. And so we're going to begin by talking about the eye and the retina, which is the part of your visual system that measures light. Again, the big picture here is that things in the world give off clues to their existence. Your sensory organs detect these clues.

And in vision, the clues come in the form of photons that reflect off of objects and are absorbed by your eye. And the task of perception is to take input from your sensory receptors and infer what is out there in the world. So vision relies on light. Remember, light is electromagnetic radiation. There's a wide range of electromagnetic radiation, as indicated on this diagram. And the part of that spectrum that you can see is a very narrow part where the wavelengths vary from around 400 nanometers to around 700 nanometers. And then on either side there's ultraviolet and infrared, and then lots of other stuff.

All right. So this is a big picture of your visual system, sideways view of the brain. So we've got your eyes here. Light comes in, forms an image on the back of your eye. Electrical signals are carried by the optic nerve to the LGN or lateral geniculate nucleus, which is the part of the thalamus that carries visual signals, and then from there, projects to the visual cortex, which is a large part of your cortex.

All right, so let's take a look at the eye. So cross-sectional view of the eye. Light enters via the pupil. The eye is protected by the cornea, which is this-- it's a very hard part of the eye. But light passes through there, passes through the pupil, which is an aperture. So the iris can change the size of the pupil in order to control the amount of light that gets in. And then you've got the lens, which focuses light into an image on the back of the eye.

So the back of the eye has a few different structures, and the most important one is the retina, which is this layer of cells that does the transduction for vision. And then those electrical signals from the retina are then carried via the optic nerve out of the eye. This is a slice of the retina. So you can see that it's this big matrix of cells. At the back here, we have photoreceptors. And then there's all these other layers of cells. You can see the cell bodies in some of these layers, and then some of the other layers which are primarily containing axons that connect different types of cells.

And then the output layer of the retina are the ganglion cells. Now, one of the curious things about the retina-- and we'll talk more about this in a moment-- is that it's kind of set up backwards relative to what you would think would be a sensible engineering design. So this is the back of the eye. And if you look here, the light is coming in like this. And so naively, you would expect that what you'd want to do is actually put the photoreceptors here, and then have all the other processing after that.

The reason being that, well, in order to get to the photoreceptors, the light has to pass through all this stuff. So it's kind of a funny way to actually set the eye up. As far as we know, it's a fluke of evolution. And we'll talk more about the consequences of that in just a second. But you can see from looking at this that there's a bunch of different types of cells within the retina. And lots of people have devoted their entire career to understanding these cells, and their connections and their function.

For the purposes of this class, the main thing you should be aware of is there's five main classes of cells. So there's the photoreceptors, which absorb photons and generate a voltage. You've got bipolar cells that get input from the photoreceptors, as well as horizontal cells. There's amacrine cells and there's ganglion cells. So one respect in which these types of cells differ is that some of the types don't fire action potentials. They just vary their membrane potential in response to light.

And the main reason we think why that's the case is that the signals don't actually have to travel very far. The connections are primarily local. And remember, one of the things that's very useful about action potentials is it allows you to send electrical signals over long distances. So these cell classes in the retina, they primarily just talk to each other over very local distances. And so presumably, they don't need spikes.

But the amacrine cells and the ganglion cells fire action potentials. The ganglion cells in particular are, as I said, the primary output cell class of the retina, and so they have to send signals all the way up to the thalamus. So this is a picture that shows those five primary cell types. So the processing kind of starts with the receptors. There's two types of receptors, which we'll talk about in a second, and then culminates in the ganglion cells.

And as we discussed, the orientation here is very counterintuitive because light is coming in this way and has to pass through all this stuff. Why is that a problem? Well, you would think that all of these cell bodies would actually cause the light to scatter and would decrease the quality of the image. And so some evidence for this kind comes from species comparisons. So the eye has evolved many, many times. And so if you look across all the different species that exist, you can see things that look like eyes in most of them, but there will be interesting differences and similarities.

So in particular, this feature of the receptor layer is kind of being oriented backwards is something that you see in vertebrates, but that you don't see in cephalopods. So cephalopods are things like squids, and octopi and things like that. And so you see this structure that looks like an eye, but you can see that in the octopus, the photoreceptors actually kind come first, followed by other stuff. All right?

Whereas in vertebrates, the photoreceptors are kind of at the back, and then you have these cell bodies that all the light has to travel through. And this has actually major consequences for how the optic nerve exits the eye, which we'll talk about in a second. So as far as we can tell, the fact that the retina is wired up backwards is, as I said, a fluke of evolution. It's some kind of local minimum that was found.

And but there are a couple of adaptations that have been added on to this over the course of evolution, as far as we can tell, that mitigate the consequences of this. And so one of those is that at the fovea, which is the center of your gaze, and as we will discuss later in this lecture, the center of your gaze is where your spatial resolution is highest. So normally, if you want to look at something, you direct your fovea to that something. And then you can get a very high resolution image.

So at the fovea, the cell bodies and the axons are displaced. And this is presumably to maximize the resolution in the part of the eye where the sampling is densest and where you really want to have really good spatial resolution. So you can see that the cell bodies here are moved out of the way, and that means that the light doesn't really get scattered. And I guess the idea is that out here, away from the fovea, the little bit of scattering that kind of happens from the cell bodies is acceptable.

And also, one of the things that kind of mitigates this issue is the fact that the cell bodies are pretty translucent. So it's not like they're absorbing lots of light themselves. Now, the other major consequence of the fact that the retina is hooked up backwards is what's shown here. So the axons of the ganglion cells, which form the optic nerve, those have to exit the eye somewhere. And the problem is that with this being set up the way that it is, the place where they exit the eye is going to be a location where you can't have photoreceptors.

So this is a picture showing the axons from the ganglion cell leaving the eye. So this is a cross-section of the retina. You can see these different cell layers here. The photoreceptors would be back here. The axons all kind of converge and then exit the eye at this particular location. So the consequence of this is that there is a region of the retina where there are no photoreceptors. And the consequence of that is that there is a region of visual space when your eye is directed in a particular place that you can't see, because you don't have photoreceptors there to detect the light.

So we can demonstrate this ourselves. So we're going to do a demonstration of the blind spot that everybody has. And in order to see this, you're going to have to only look with one eye. So I want you to close your left eye and then fixate your left index finger. So I'm going to do this myself. All right? You're then going to place your right index finger next to the left one, and then slowly move your right index finger to the right. So again, close your left eye, and then move that finger away, and then-- yep, there we go.

So you'll reach this place where you can't see the tip of your right finger. Yeah, it's pretty cool. So the tip of your finger should disappear. And this is because it has entered the blind spot of your right eye. So hopefully, that works for everybody. So this is something that's happening all the time. You're stuck with this blind spot. The reason that you most you don't notice this, we think, is partly because the blind spots in the two eyes are kind of in different parts of the visual field.

So usually, you got both of your eyes open. And so there's one eye that is detecting most of the stuff that you can see. And so you have to do these very careful experiments in order to be able to detect this. So again, the blind spot is a consequence, we think, of the retina kind of being set up backwards.

So let's talk a little bit about the photoreceptors. So the photoreceptors are cells that contain photo pigments that absorb photons, and that then produce a change in membrane potential. There's two types of photoreceptors, the rods and the cones. They're called this because the physical shape of the receptor is different. And they're different physiologically in lots of ways. So one respect in which the rods and the cones differ is that their density on the retina varies quite dramatically with what we call eccentricity, or distance from the fovea.

So this is a graph that shows the photoreceptor density plotted separately for rods in blue and cones in red. So zero here corresponds to the fovea. That's the center of gaze. And then you can either move out in the nasal direction towards the nose or the temporal direction. And what this graph is showing is that at the fovea, the density of cones is very, very high, and there's actually no rods. As you move out of the fovea, the density of the cones drops quite a lot, and the density of the rods kind of increases.

And then these are images of the retinal mosaic taken from these different points. And so you can see that at the fovea, you only see evidence of only having a single type of receptor. And then you get out here in the periphery and you can see that there's two types of receptors. All right, the other thing that you can see on this diagram is what's called the optic disk. So this is the location on the retina where the optic nerve leaves the eye. And there, there are no photoreceptors.

So you're going to hear me talk a lot about eccentricity. So what is eccentricity? Again, in vision science, eccentricity is distance from the fovea. So we always use the fovea as a reference point. And then typically, we will measure eccentricity in degrees of visual angle from the fovea. So a good rule of thumb is that one degree of visual angle is approximately the width of a fingertip at arm's length.

The full moon, which is about the same size as the sun, when considered as an image on your retina, is about half of a degree, and the blind spot is about 5 degrees, so pretty big. So there's these big differences in the distribution of receptors in different parts of the eye, so that's just an empirical fact. And there's also big differences in the responses of the different types of receptors. And one place that this shows up is in dark adaptation.

And so the big picture here is that the rods are very sensitive to light and are primarily used in low light conditions, so at night. The cones are less sensitive to light and are mostly active during the day. So if you go outside on a bright day, the cones will be active and responsible for your vision, and the rods won't be doing anything. And then you go inside to a dark room, and then the rods will take over.

But this process is not instantaneous, and it involves a process of dark adaptation. And that's because both types of receptors adapt, to some extent, to the ongoing light. So this is the result of a classic experiment that measures the detection threshold for light. So your task is just to say whether a light is present or not. You vary the amount of light that's present and figure out the threshold for telling whether there's a light that's present or not.

And so that threshold in this experiment is measured as a function of the amount of time that you've spent in a dark room. So the idea is, you're out on the street, you go into this dark room like a movie theater, or in this case, this would have been a room in a lab. And then you measure your detection threshold. So everybody has had the informal experience of walking into a movie theater and initially not being able to see very much. And then you sit there, and then over time, you can see the seats around you, and the people around you and so forth.

So that's the process of dark adaptation. So something's changing in your visual system to make it easier to see in low-light conditions. And part of that are changes that are happening at the level of the photoreceptor. Now, if you do this experiment and you measure the threshold as a function of time, you're going to get the purple curve that's shown here. So what happens is that, initially, your threshold is high. And then as time goes on, the threshold drops. And then you get to this point, and it kind of starts to drop, again, a little more aggressively.

So the curve has got of funny shape. And we believe-- in fact, there's now lots of evidence-- that that funny shape comes from the contributions of two types of photoreceptors, the cones and the rods. All right? So as I said initially, in high light levels, your cones are the ones that are active, and your rods are saturated. And so you walk into this dark room, and the cones are the photoreceptor class that determines your light sensitivity.

Both types of receptors start adapting, but the sensitivity of the rods is, for this initial period, still going to be worse than that of the cones. So the cones adapt a little bit, but then they kind of cap out at this level of sensitivity. That's the best they can do. So then the rods continue to adapt, and their sensitivity starts to exceed that of the cones. And the idea is that your behavioral threshold will be determined by whichever class of photoreceptors is more sensitive at that moment in time. And so here, that'll be the cones, and then here, that will be the rods.

All right. So there are conditions in which you are primarily dependent on rods and conditions where you are primarily dependent on cones. And so-called rod and cone vision have a number of differences. So in general, rod vision is more sensitive than cone vision. So when you're reliant on your rods, you're better able to detect light. And so some of this comes from the fact that individual rod photoreceptors are more sensitive to light than individual cone photoreceptors.

But part of this also derives from the fact that there is higher convergence from the rods to the ganglion cells. So remember, photoreceptors of transduced light send their signals to this set of different cells, culminating in the ganglion cells. And so with the rods, there's a high degree of convergence. What that means is that a single ganglion cell will get input from roughly 120 rods. So they're pooling signals from a lot of different rods, and that makes them very sensitive, of course, at the cost of spatial resolution.

By comparison, the convergence from the cones to the ganglion cells is much less. So in the fovea, it's very often 1 to 1. So a single ganglion cell will get input from a single photoreceptor. And again, we believe that's in order to maximize spatial resolution in conditions where the light levels are high. All right, so rod vision is more sensitive than cone vision, but this comes at the cost of acuity, like we said. So rod vision has lower acuity than cone vision.

Some of this is due to the convergence. And then another difference is that the rods are slower, so they also have a longer integration time. And again, this is a trade off between, in this case, temporal resolution and sensitivity. So integrating over time maybe increases the likelihood, the ability to detect some photons by pooling over time at the cost of being able to distinguish temporal patterns of stimulation.

Another big difference between rods and cones is that rods do not support color vision. And the reason for this is that there's only one type of rod. By comparison, there are three types of cones, and the cones thus allow you to distinguish different wavelengths of light. So we'll talk a lot about this when we talk about color vision in a few weeks. Another big difference is that the rods are absent from the fovea.

And this has the consequence of something called that is known as Arago's phenomena. So if you're outside on a very dark night, so in conditions where you're relying on your rods, and you're looking up at the stars, if you look directly at a star, you won't be able to see it. And that's because you're using your rods for vision, and there are no rods directly in the fovea. So you'll have to look a little bit away from the star in order to be able to see it.

So if you measure sensitivity to light as a function of wavelength, you'll see that there are some differences. And the two words that we're going to use here to refer to conditions in which you are reliant on rods and cones are scotopic and photopic. So scotopic refers to conditions where the light level is low, so kind of during the night most of the time, if you're outside at least, and when you're mostly reliant on rods. Photopic refers to high light levels, so daytime vision where you're using your cones.

So these graphs here show sensitivity as a function of wavelength. So sensitivity is 1 over the threshold. So if you're very, very sensitive, that means your detection threshold is low. And so in photopic conditions where you're reliant on your cones, peak sensitivity is around 550 nanometers. Scotopic vision, by comparison, has peak sensitivity at 505 nanometers. And you can see that your sensitivity to red light is basically nonexistent. So by red light, I mean wavelengths that are close to around 700 nanometers, so longer wavelength light.

And we call that red light, because if you look at that in isolation, all other things being equal, it'll appear red. And this is just an empirical fact about the rod sensitivity. So as we mentioned, there are three types of cones. The three types of cones are distinguished by their spectral sensitivity. So this is a graph that plots sensitivity, and it's been normalized here. So the peak sensitivity is 1 as a function of wavelength. And there are three curves here, each for the three types of cone photoreceptors.

And so those will often be referred to as the S cones, the M cones, and the L cones. And so S stands for short wavelength, M for medium, and L for long wavelength. And so you can see that the S cones have their peak sensitivity down here at around 450 nanometers. The M cones, it's around 550, and the L cones, it's a little bit higher.

So an important idea around photoreceptors, which we will return to when we talk about color vision, is what is called the principle of univariance. And the idea is that the receptor response doesn't contain any information about the wavelength of light. So these curves, they tell you how likely a receptor is to absorb a photon of a given wavelength. All right? And so when your sensitivity is very high, that effectively means that the photopigment in the receptor is very likely to absorb that particular wavelength.

And that means that if you are firing some number of photons per unit time at that photoreceptor, at this particular wavelength, you're going to get a large response, whereas at this particular wavelength, you'll get a lower response. But the concept here is that the response of the photoreceptor is a scalar. It can just vary its membrane potential. And the effect of the wavelength on the photoreceptor is just to vary the probability of absorption, which varies the magnitude of the response.

And the consequence of that is that you can trade off between wavelength and intensity. So in other words, I can take a low-intensity light that is at 450 nanometers and present that to the S cone and get an identical response than if I have a 425 nanometer light that is at a higher intensity. So the consequence of that is that a single photoreceptor on its own doesn't really directly give you any information about wavelength. And that's why you actually need more than one receptor type if you actually want to be able to distinguish different wavelengths.

So that's the principle of univariance. There's just one dimension along which the response of the receptor varies. All right. So one interesting empirical fact is that there's a fair amount of individual variation in the nature of the receptor mosaic on the retina. So these are pictures that were generated with a method called adaptive optics, whereby you can look inside someone's eye and image the cone mosaic. So this was introduced back in 1999 by David Williams at Rochester.

And these are two different images of the cone mosaic from two different individuals. And the point is just that they look kind of different. The distribution of the receptors is not identical. And this probably indicates that there would be individual differences in the nature of color vision. But exactly what those would be is, I think, not yet clear. So those are photoreceptors. The other kind of class of cells that we're going to talk a fair bit about today are ganglion cells. So remember, this is the type of cell that provides the output from the retina.

And there are two main types of ganglion cells that are defined anatomically, and they're called midget cells and parasol cells. So this is a picture on the bottom of the dendritic field of each type of ganglion cell, so an example of each. So remember, the dendrites in a neuron are the place where they collect input right. So axon is the kind of output-- provides the output of the neuron. The dendrites collect the input. And so the dendrites in this case are distributed over space in order to get input from photoreceptors that are at different parts of space.

They're not usually directly connected to photoreceptors. They're getting input from other types of cells that are, in turn, getting input from photoreceptors. But the spatial distribution reflects the region of the retinal image from which they're getting input. And you can see that the midget cell has a small dendritic field, indicating that it is collecting input from a very small region of the image. And the parasol cell has a large dendritic field, indicating it's collecting input from a larger region of the image. All right.

So parasol cells and midget cells differ in a bunch of their response properties. So parasol cells are what are called achromatic, so they tend to respond equally to all wavelengths, because they pool responses across cones. They tend to prefer first stimuli that change a lot over time. Midget cells are often sensitive to both luminance and color. That's often because they'll receive input from particular types of cones, and to be less sensitive to temperature changes and stimuli.

So these two classes of ganglion cells project to the lateral geniculate nucleus, or LGN, which is the part of the thalamus that's involved in vision. And the lateral geniculate nucleus itself contains two main classes of cells. We now know there are some other classes, as well, but there's two primary ones. All right. And the classes of cells in the LGN are called parvocellular and magnocellular.

Now, the nomenclature here is quite unfortunate and makes this a little bit hard to remember because the parasol cells in the ganglion cell layer of the retina project to the magnocellular layers of the LGN. The midget cells in the retina project to the parvocellular layers of the LGN. So it's kind of the opposite of what would be convenient to remember. Now, for both cell types, the size of the dendritic field, and thus what we call the receptive field, so the region of space to which the neuron responds, increases with eccentricity.

Remember, eccentricity is distance from the fovea. And so this is a diagram that shows this. So what this is plotting is the receptive field's diameter-- so it's the region of space to which the neuron responds-- as a function of retinal eccentricity. So this is the result of an experiment where you stick an electrode in the retina. You're recording the responses of a ganglion cell presenting light stimuli, and then measuring the region of the image that evokes a response in the neuron. And so that's the quantity that's plotted on the y-axis. And then you do this at different retinal eccentricities.

And so the point here, there's two main things that you should note from this. One is that when you're close to the fovea, the receptive field size is small. And that's consistent with this general idea that the fovea is where you have really high-resolution vision. So remember, there's lots of cones there. You keep the receptive field small so that you can distinguish different fine grained patterns there.

As you move from the fovea out to the periphery, so increasing eccentricity, the receptive field size tends to increase. But you can also see that the receptive fields, they fall in these two lines. And these two lines correspond to the midget cells-- that's the lower of the two lines-- and the parasol cell. That's the upper of the two lines. So the parasol cells at a given eccentricity have bigger receptive fields than the midget cells. But in both cases, the receptive field grows with eccentricity.

And so this phenomenon here, whereby receptive field size increases with retinal eccentricity, this is a very general organizing principle within the visual system. So receptive fields tend to be small in the fovea and then larger in the periphery. So this is an eye chart that was designed by Stuart Anstis, a vision scientist, that is intended to compensate for changes in resolution with eccentricity. And so the idea of this eye chart is that if you look at the center dot, then all of the letters should be equally recognizable.

And the reason for this is that the size of the letter has been scaled up with eccentricity. So the ones that are closer to the fovea, those are smaller. The ones that are out in the periphery are larger. So the resolution falls. So the spatial resolution of your vision falls in proportion to the distance of your fovea. And this is some evidence that those phenomena that we're seeing in receptive fields actually have consequences for what you can see.

So major organizing principle of the visual system is the foveal periphery distinction. And it's an interesting question as to why the visual system is set up the way that it is. And so as we have discussed at several earlier points in this class, one of the ways that we can answer these "why" questions is by building computational models that are optimized for different conditions, and investigating the conditions in which this optimized system ends up finding some of the same solutions that you see in biology. And so there's a paper that I like a lot that came out a few years ago that did this for the problem of retinal image sampling.

So the paper is posing this question of why we have this foveal image sampling strategy. This is work by Brian Chung, and Eric Weiss and Bruno Olshausen. So I'll just read this little excerpt from the paper. So they write, "The human retina contains 1.5 million ganglion cells, whose axons form the sole output of the retina. These essentially constitute about 300,000 distinct samples of the image, due to the multiplicity of cell types coding different aspects, such as on versus off channels," which we're going to talk about in a moment.

"If these were packed uniformly at highest resolution, they would subtend an image area spanning just 5 by 5 degrees. Thus, we would have high resolution but essentially tunnel vision. Alternatively, if they were spread out uniformly over the entire monocular visual field, spanning roughly 150 degrees, we would have wide field coverage, but with very blurry vision, with each sample subtending a quarter of a degree, which would make even the largest letters on a Snellen eye chart--" that's the kind of standard eye chart for diagnosing vision-- "illegible.

Thus, the primate solution, which is to have a fovea, makes intuitive sense as a way to achieve the best of both of these worlds. However, we are still lacking a quantitative demonstration that such a sampling strategy emerges as the optimal design for subserving some set of visual tasks. We explore whether what is the optimal retinal sampling lattice for an overt attentional system, performing a simple visual search task requiring the classification of an object. We propose a learnable retinal sampling lattice to explore what properties are best suited for this task.

While evolutionary pressure has tuned the retinal configurations found in the primate retina, we instead utilize gradient descent optimization for in-silico model by constructing a fully differentiable, dynamically controllable model of attention." So here's the idea. They've got this system that has a receptor lattice. So there's some number of receptors that can sample an image. And each receptor has a particular position and a particular width. So there's this kernel that has these learnable parameters, and there's a whole bunch of these.

And the idea is that they are going to optimize the position and width of the receptors on this retina for a vision task. And in this case, the vision task is to find and identify the digit on a cluttered background. So these are examples of stimuli from this task. So you can see this one has a four. This has a one. There's a nine. So they're in different places, and there's some other junk that's kind of added in to make the task hard.

So the system has to solve this problem. And they're training the system up to solve this problem, but allowing the system to optimize its receptor lattice. And they additionally introduced these other constraints on the problem. So in one version of the system that they trained up, the system was allowed to translate its retina, so essentially, to make eye movements, to position the receptor lattice at different places within the image.

And then in another condition, it was allowed to translate its retina, but it could also zoom the retinal lattice. This is like the kind of thing you can do with your phone, but which the biological visual system doesn't actually do. We don't have a zoom function. We can just translate. And so what they found was that in the condition where the retina could be translated around, eccentricity dependence of the sampling in the retina emerged as an optimal solution.

So these are showing the settings of the receptor lattice over different amounts of training, so optimization for this problem. So this is like the initialization. They initialize these things as uniformly spaced. And then you can see that over the course of the optimization, what happens is that some of the receptors get moved kind of close together in this one part of the sampling lattice, and they also get very small right. So that's like the system is kind of evolving this strategy where there's a whole bunch of very small receptive fields that are packed into one particular part of the sampling lattice.

And so they quantified this by plotting the receptive field size as a function of eccentricity, and so in cases where the model could translate its retina. So these are two solutions that it found, and you see this kind of eccentricity dependence. Whereas in the alternative problem constraint of being able to zoom in addition to translate, the solution they found was different. There was much less of a fovea, and you see much less of this eccentricity dependence.

So here are examples of the model actually performing the task. So this is one trial of this task, and it's just showing you where the sampling lattice is kind of positioned relative to the digit. And the point is that over time, the sampling lattice kind of gets moved so that the fovea, in this case, is kind of centered on the digit, so that it can optimally or best recognize the digit. And this bottom row is what happens when the thing can zoom its retina, as well. And there you can see this uniform receptor lattice, and it just kind of shrinks it down to fit onto the digit.

So what do we learn from all of this? So the inference that you should make here is that this is evidence that a fovea is a useful way to attain high resolution vision, given, one, a limit on the number of samples that can be transmitted to the brain from the optic nerve, and number two, the ability to move the eyes to position the fovea at different locations. And so it seems that biological vision has evolved this strategy that's just a good sampling strategy for being able to solve tasks.

So in other words, you have the ability to move your eyes. And then in concert with that, the retina has the receptors packed in this very, very non-uniform way. So there's this one part of the retina where you have really good sampling. And if you want to see stuff, you typically direct your fovea to that part of the retina.