




Colours of the World

by **David Morgan-Mar**

The world around us is colourful. What's more, it looks colourful to our eyes. It didn't have to be that way. In fact, there is a lot of subtlety and nuance in the coloured world that we humans are missing out on. To see why, we need to understand how our eyes work.

Our eyes detect light — that much is fairly obvious. But the details of exactly how they do it are important. The front parts of our eyes are essentially equivalent to the lens assembly on a camera — something to focus the incoming light into an image. Our eyeballs have a **lens near the front**. Rather than focus by moving the lens back and forth as a camera does, we focus our eyes by stretching the shape of the lens with tiny muscles.

These muscles are largely involuntary and respond automatically to focus our eyes on whatever we happen to be looking at. If you move your gaze between objects nearby and objects far away, you can see the effect of your focusing muscles. In some people, the lens shape is slightly wrong, or the muscles don't work as well as they should, resulting in an inability to focus on objects at certain distances. This can be



Layers of the retina. Light enters from the left; the light-sensitive cells are on the right. Based on a **Creative Commons Attribution-ShareAlike** image by **Wikimedia Commons** user **chris**.

corrected by glasses or contact lenses. As we get older, the lens **stiffens up and the muscles weaken**, so we may end up needing glasses then too.

Immediately in front of the lens is the **iris**, the coloured part of your eye. This opens and closes around a hole in the middle, called the pupil, to let in different amounts of light: wide open at night time, and narrow in daylight. This gives us the ability to see over a wide range of lighting conditions. The iris performs the same job as the **aperture iris** on a camera lens, adjusting the amount of light that enters.



Eye, showing the coloured iris.

Creative Commons Attribution-ShareAlike image by **Sarah Cartwright**.

The light then passes through a **clear fluid that fills your eyeball**, and lands on the back surface, which is called the **retina**. Because of the focusing of the lens, the light forms an image that resembles what is in the outside world. Special nerve cells in your retina are sensitive to light, in a similar way to how special nerve cells in your fingers are sensitive to touch, or ones in your nose are sensitive to smell. The physical principles of how all these special nerve cells work differs, since they detect very different stimuli, but the overall result is the same: they detect the presence of something and send an electrical impulse into connecting nerves, which conduct the signal to your brain. And so your fingers tell you when you touch something, your nose tells you when you smell something, and your eyes tell you when you see something.

The special nerve cells in your eyes come in four different types. The types are called rods, L cones, M cones, and S cones. The **rod cells** are shaped like long, thin cylinders, or rods. The three types of **cone cells** are tapered to one end, making them conical. The L, M, and S stand for long, medium, and short, though this does not describe the length of the cone cells - we'll get to what that's all about in a little while.

An interesting aside about the structure of the eye is that all of these light-sensitive cells need a supply of oxygen, which is delivered by the bloodstream. Arteries lead into the back of your

eyeballs to deliver this blood, then fan out into a network of tiny capillaries to carry it to all the rods and cones scattered across your retina. They also need to send their signals out to the brain, via nerves. The most sensible arrangement for all this would be to have the blood vessels and nerves connect to the back of the rods and cones, so they don't get in the way of the incoming light that the rods and cones detect. But this isn't how our eyes are structured. Instead, the main artery and the optic nerve both punch right through the retina and then spread out into broad networks on the surface, in front of the light-sensitive rods and cones. So there is a considerable fraction of your retina covered by blood vessels and nerves, where you are not sensitive to light.



Retina of my right eye.

You can see the blood vessels on your retina for yourself if you look at a dark background, shine a (not too bright) light into your eye from the side, and wiggle the light around a bit. (The nerves are smaller and not as easily visible.) **These articles** describe **how to do it** in more detail. Give it a try! So our eyes clearly sense the shadows caused by this network of blood vessels on the front of our retinas. It's just that our nervous systems have adapted by editing them out,

presenting a cleaned up ("Photoshopped") image of what our eyes really see to our brains. Although our brains **fill in the gaps for us**, the end result is that our eyes are less sensitive, both to light and to fine detail, than they could be if the blood vessels and nerves were behind the rods and cones. You can also detect the hole in your retina where the optic nerve and the main artery enter your eye, which produces a **distinct blind spot** in your vision.

The next question is why are our eyes built in this sub-optimal way? If they'd been engineered in the same way as we build modern electronic sensor chips for digital cameras, all of the necessary circuitry (the blood vessels and nerves) would be behind the light sensitive parts where it wouldn't get in the way.¹ But evolution doesn't always come up with the best solutions for problems. We inherited the layout of our eyes from generations of ancestors stretching back millions of years. We share the basic anatomy of our eyes, including the backwards structure of our retinas, not only with all mammals, but also birds, reptiles, and fish.

This anatomy developed with the very first vertebrates, and has been with us ever since. It arose because the skin of pre-vertebrates without a sense of vision was, quite naturally, covered with a network of capillaries to supply the cells there and nerves to provide a sense of touch. At some point, a mutation let some nerve cells below the skin become sensitive to light. This is all you need to distinguish night from day, or shade from sunlight, for example. Common earthworms have this ability, from light-sensitive cells distributed throughout their skin. Once you can detect light, the sense becomes better if the light-sensing cells are clustered on the inside of a pucker in the skin. The walls of the pucker give some

directionality to the incoming light that can be sensed, and so the animal can now tell roughly which direction the light is coming from. This directionality gets more precise the deeper the pucker, until it's very good when the light-sensitive cells are at the bottom of a pit in the skin.²

If the pit then partially closes over, leaving a small hole for the light to come in, it then effectively becomes a pinhole camera (which I've **discussed before**). The light-sensitive cells on the inside of the cavity can then detect an image of the outside world. This is clearly better than mere direction sensitivity, and is in fact the structure of the eyes of the **nautilus** - these animals' eyes have a small hole in the front and are filled with sea water. The next step is to close the hole with a transparent membrane, and the cavity can then be filled with a clear body fluid, which provides structural advantages. The transparent membrane can then thicken and become shaped to better focus the incoming light, allowing the eye to grow larger and become more sensitive to low light levels.

And thus we have arrived at the vertebrate eye, by following this evolutionary process in which a series of small changes, produced by random mutation between generations, each lead to a specific advantage over the preceding organs. The problem is that evolution is not directed. It would have been sensible, from an engineering viewpoint, to have a mutation somewhere along the line reverse the layering of the blood vessels being above the light-sensitive cells. If this mutation had happened anywhere along the evolutionary chain from the first vertebrates to modern mammals, then some or even all vertebrates today would have more efficient eyes. But



such a mutation didn't happen, and so we're stuck with an inside-out retina anatomy.³

Okay, aside over, let's get back to the light-sensitive cells themselves: the rods and cones. The rod cells detect light, and they're not fussy about it. They detect light of all wavelengths, which is great if you're in a dark place and need all the sensitivity you can get. The problem is, all a rod cell can do is send a nerve signal saying, "I've detected some light!" It can't tell which colour it is. You'll notice this whenever you're in a dark room, with just enough light to see shapes and not bump into things. It's your rods doing all the work, and you won't be able to detect colours of objects.





Figure 1. Responses of cone cells to different wavelengths of light. Based on a **Creative Commons Attribution-ShareAlike** image by **Wikimedia Commons user Vanessa Ezekowitz**.

Colour vision is all about the cone cells. Here's where the three types of cone cells, L, M, and S are explained. The L cones detect long wavelength light, that is, towards the red end of the spectrum. The S cones detect short wavelength light, or towards the blue/violet end of the spectrum. And the M cones detect medium wavelength light, which is the light around green. But it's not quite as simple as that. In fact the colours of light that the three different types of cones detect overlaps quite a lot, as shown in Figure 1. And it's this overlap that gives us the ability to detect what we perceive as millions of different colours, despite the fact that we only have three different types of colour-sensitive cells.

Let's say we have something yellow, a nice ripe banana. If you look at Figure 1, you'll see that both L and M cones are sensitive to yellow light, so both of these cone cells will send signals to your brain. Effectively they say, "this object is both red and green". Your brain collects this information, combines it, and creates the colour sensation you know as "yellow". If you look at an orange instead, the L cones will send a stronger signal and the M cones a weaker signal. They tell your brain, "this object is strongly red and weakly green", which your brain interprets as the colour orange. And so on.

In short, any object which emits/reflects light of a single particular wavelength (or colour) stimulates the three different types of cones in your eyes to different extents, then your brain interprets the unique mixture of the signals from the cones as a unique colour. And here's where things get really interesting. If a certain shade of yellow light strikes an area on your retina, it stimulates both the L and M cones equally, and your brain senses the stimulus as "yellow". On the other hand, if an equal mixture of red and green light, with no yellow light in it whatsoever, strikes your retina, it also stimulates both the L and M cones equally, and your brain again senses the stimulus "yellow". Your brain cannot tell the difference between these two very different mixtures of light, because your cone cells respond in exactly the same way.



Red, green, and blue subpixels on a TV screen.

And this is why colour televisions and computer screens work the way they do. They have **pixels made up of tiny red, green, and blue components**. Those are the only three colours the screen can produce. Neither your TV nor computer screen can produce yellow light at all.

But by producing equal amounts of red and green light so close together that they hit your retina in essentially the same spot, they can stimulate your L and M cones at the same time, and make your brain interpret the colour as yellow. The same happens for every other colour that you perceive on a display screen. It is generated in your brain by a precisely controlled mixture of different amounts of red, green, and blue light, which stimulate your cone cells in the same proportions of actual light of the desired colour.

Furthermore, this standard **RGB (red-green-blue) system** for colour displays only works to mimic other colours because it is carefully tuned to match the human visual system. You can imagine an animal with four types of cone cells, which detect, for example, red, green, blue, and yellow. For such an animal yellow light would stimulate the yellow cones mostly, and perhaps the red and green ones a bit. And that would be interpreted by the animal as "yellow". But if you show it "yellow" on a display screen, this is composed of red and green light only. This mixture would stimulate the red and green cones in the same way as yellow light, but it would not stimulate the yellow cones as much. In other words, the animal's brain would not interpret this as "yellow", but would see it as some other colour. (A colour that we have great difficulty in imagining, because that's patently not how our brains are wired.)

So it's possible to have a TV or monitor show a photo of a scene that appears essentially perfect to human eyes, but if you showed it to our hypothetical four-coned animal, it would see completely the wrong colours. And this is not so hypothetical, either. Most bird species have **four different colour-responsive types of cone cells**. The extra type isn't sensitive to yellow,



Eye of a bird.

though, it provides sensitivity to violet and ultraviolet light. It doesn't stop there either. Pigeons and a few other birds, and also butterflies, have **five different types** of colour sensitive cells. **Mantis shrimps** have twelve! It's also possible that through mutation, a few humans have more than three types of cones. There is some anecdotal evidence of this in some cases, but it has not been reliably proven.

On the other hand, it's reasonably common for humans to have impaired cone cells or to be missing one of the types of cone cells completely. This condition is inherited via genetic mutations passed on by parents, and causes what we call **colour blindness**. The most common form of colour blindness is impairment of the M cones, which results in shades of green providing stimulus mainly to the red-detecting L cones. To such a person, greens and darker shades of red appear to be very similar colours, and this is called red-green colour

blindness. About 8% of males have this condition; it is much rarer in females. A more severe form is the lack of M cones altogether, which renders reds and greens completely indistinguishable. Similar impairments can affect the L cones, which makes reds look similar to dark greens. Very rarely, the S cones can be impaired or missing, resulting in problems at the blue end of the spectrum, which makes blues and yellows indistinguishable, and is called blue-yellow colour blindness. This form is more common in females, but only affects about 0.03% of the population. In extremely rare cases, two types, or even all of the cones be dysfunctional, leaving the person with no ability to perceive colour differences at all.

There are a couple of other interesting things about human colour vision I want to mention briefly today. Firstly, if you check Figure 1, you can see that pure green light stimulates not only the "green detecting" M cones, but also the "red detecting" L cones. In fact, it's impossible to stimulate only the M cones with any particular colour of light. Any light that stimulates the M cones will also stimulate the L and/or the S cones. So any time you see what your brain thinks of as "pure green", it's interpreting the stimulation of multiple types of cones. If there was a way to stimulate only the M cones, then you could sense a colour that is in some sense a more intense and pure green than anything your brain has ever experienced. You can't easily imagine such a colour because you've never seen it. But you can sort of get a feel for it by staring at a bright red object for a minute or so, to adapt your L cones, then switching your gaze suddenly to a bright green object. Your L cones will provide a weaker response than normal for a few seconds, and you can briefly experience this über-green sensation.



David Morgan-Mar

David Morgan-Mar is a research engineer living in Sydney, Australia. Currently working for **Canon Information Systems Research Australia** on image processing projects, he also finds time to write webcomics and role-playing game supplements, photograph at a professional level, follow cricket, travel the world, and be a drummer in a band.

How does he find the time to do all this?

I have extremely little spare time. I am always lamenting how I don't have enough time to do all of the stuff I want to do. What I do have is a creative urge. Ideas. The desire to make things, and do things, and learn things. What I have is a list of ideas for things I want to do, or make, or places I want to go. A big list. A really, really big list. I can't possibly do them all.

What I also have is the burning desire to make sure I damn well do at least some of the things on that list. I can't sit still in front of the TV. I'm always thinking about what cool thing I could be doing instead. So I'll run off in the ad breaks and fiddle with my photos in Photoshop, or write snippets of dialogue for comics, or bake some banana muffins. Despite not having enough spare time, I make the time to create things, because I can't bear the thought of not creating things.

People who are going out of their way to find the time to be creative and to make new things are taking steps to make something concrete out of the ideas and projects and creative desires locked inside their heads that other people would otherwise never get to see. They are making the most of their time. Go out and make the most of yours.

Secondly, again looking at Figure 1, you can notice that there is no single colour of light that can stimulate both the L and S cones without also stimulating the M cones. You can, however, stimulate both the L and S cones without also stimulating the M cones (much) if you look at a mixture of red light and blue-violet light. This mixture produces the colour sensation we know as magenta, or hot pink, or purple, depending on the proportions of red and violet. These colours around magenta and purple are known as **non-spectral colours**, because they are not single-wavelength colours that appear in the familiar rainbow spectrum. They are only constructions of our brain when stimulated with light of at least two different wavelengths. This means, in practical terms, that it is impossible to build a magenta laser! (Since lasers pretty much by definition consist of light of a single wavelength.)

So now, if you ever see a magenta, pink, or suspiciously purple laser beam in some science fiction, you can call it out as unrealistic.

Human vision is fascinating and amazing. And it's the whole reason why RGB displays work. If our eyes were built differently, our TVs and computer monitors and digital cameras and all sorts of other things would need to be built differently too.

Notes

1. By “modern” sensors I mean the new style **back-illuminated sensors**, which are engineered this way, with the circuitry safely behind the light-detecting elements. Older sensor chips were actually built with the circuitry in front of the light-detecting elements, because it was easier to build them that way. But now that we have the capability, we're switching over to the more efficient design.
2. This is the stage at which some snakes such as pit vipers have sense organs adapted to detect infra-red radiation (which is really just a form of light, with a longer wavelength than visible red light). These snakes have pits or dimples, at the back of which are cells sensitive to infra-red light, and which give them a good directional sense for heat sources.
3. Interestingly, the eyes of cephalopods, including the aforementioned nautilus, are built “the right way around”, with all the blood vessels and nerve conduits connecting to the back of the light sensitive cells, rather than the front. This is because molluscs evolved eyes independently of vertebrates and somewhere along the way this advantageous mutation appeared for them where it failed to appear for us.