

# BC

## BEFORE COMPUTERS

On Information Technology from  
Writing to the Age of  
Digital Data



Handwritten characters in a cursive script, likely representing an early form of data or code.



Stephen Robertson





<https://www.openbookpublishers.com>

© 2020 Stephen Robertson



This work is licensed under a Creative Commons Attribution 4.0 International license (CC BY 4.0). This license allows you to share, copy, distribute and transmit the text; to adapt the text and to make commercial use of the text providing attribution is made to the authors (but not in any way that suggests that they endorse you or your use of the work).

Some of the material in this book has been reproduced according to the fair use principle which allows use of copyrighted material for scholarly purposes.

Attribution should include the following information:

© Stephen Robertson, *B C, Before Computers: On Information Technology from Writing to the Age of Digital Data*. Cambridge, UK: Open Book Publishers, 2020, <https://doi.org/10.11647/OBP.0225>

In order to access detailed and updated information on the license, please visit <https://www.openbookpublishers.com/product/1232#copyright>

Further details about CC BY licenses are available at <https://creativecommons.org/licenses/by/4.0/>

All external links were active at the time of publication unless otherwise stated and have been archived via the Internet Archive Wayback Machine at <https://archive.org/web>

ISBN Paperback: 978-1-80064-029-0

ISBN Hardback: 978-1-80064-030-6

ISBN Digital (PDF): 978-1-80064-031-3

ISBN Digital ebook (epub): 978-1-80064-104-4

ISBN Digital ebook (mobi): 978-1-80064-105-1

ISBN Digital (XML): 978-1-80064-106-8

DOI: 10.11647/OBP:0225

Cover image: Katsushika Hokusai (1760-1849), *A merchant making up the account*. Wikimedia [https://commons.wikimedia.org/wiki/File:A\\_merchant\\_making\\_up\\_the\\_account.jpg](https://commons.wikimedia.org/wiki/File:A_merchant_making_up_the_account.jpg)  
Public Domain.

Cover design: Anna Gatti.

## 8. On physics and physiology

---

In Isaac Asimov's *Foundation* trilogy, written in the 1940s but about the distant future, he describes a device that plays a recording of a 3D moving image with sound. The display device is a glass cube, in which the viewer sees a human figure talking—something like a talking head on television, but in 3D. Asimov does not explicitly say this, but the impression is that the image inhabits the 3D space inside the cube; people can watch from all around, but the people at the back will see only the back of the figure. Asimov was a biochemist, but one might describe this as a physicist's version of 3D film. This is in complete contrast to Wheatstone's original 1838 stereoscope, and to modern 3D film and television, and indeed to virtually everything tried in between these two dates, which might be described as relying on physiology—on the fact that we perceive depth through our binocular vision.

In this chapter I want to explore this space a little. I will return to 3D vision later, as in the previous chapter, I will start with sound.

### 3D sound

Our sense of the location of the source of a sound depends in part on the fact that we have two ears. The difference between what our ears hear and report to the brain allows us some degree of directional sense of where the sound is coming from. This is the basis for stereo sound systems. Given different sounds from two separated loudspeakers in a room, our ears can have some illusion of sound location.

However, this illusion is not very good. The sounds delivered to our two ears by two loudspeakers (in a room with its own aural character) are only a very rough approximation to what might be heard in a real environment with real sound sources, and of course real echoes from whatever else is in that environment. So what might be a better way?

There are two ways to go. One of them is to have many more loudspeakers, potentially with different signals to each one. ‘Surround sound’ systems, used for example in cinemas, are a move in that direction. But it could go further. I once came upon a public performance of a recording of a 40-part motet. In a large empty hall at the back of a church, there were 40 loudspeakers, each mounted on a stand at head height, distributed in a rough circle around a room. I could wander around and in and out while the music was playing, hearing in different ways, for example concentrating on one or a small group of parts, with the rest in the background. Exactly what I heard at any point depended on which way I was facing as well as my location. In addition to the different relative location of the ears as one turns, our ears are themselves each to some extent directional, and one’s head casts an aural shadow.

That’s a true physical attempt at a solution to the problem. However, it’s not a feasible general approach to hi-fi in the living room!

The other direction would be physiological. We can take much more seriously the idea of delivering different signals to each ear—in fact good headphones make for a much cleaner aural environment, each ear hearing only its own signal, with no interference or cross-over and no echoes. However, in order to do this properly, the recording should be made in a similar fashion. That is, one should use a pair of microphones, each in its own shell-like mount, on either side of a head-shaped object.

This is known as *binaural* or *dummy head* recording, and is quite different from normal stereo recording. It is seriously difficult to do well. For one thing, everyone’s head is a different shape, as are their ears and ear canals. For another, if the listener moves or turns their head while listening, the dummy head that was doing the recording did not move in the same way at the corresponding time during recording, so at this point the listener’s experience will be distorted. Binaural recording cannot be a general solution, any more than multiple speakers can be.

Thus ordinary stereo and surround sound occupy a slightly uneasy place somewhere in between a true physical solution and a true physiological one. This is not to say that sound recording and playback is necessarily bad—some things come across wonderfully. But it is, necessarily to some degree, a distortion of the original sound.

## The physics of colour

If you pass a bright white light through a prism, onto a white surface, you get a display of the spectrum of colours, as in a rainbow. This phenomenon was studied by Isaac Newton in the seventeenth century, but was not fully understood until the nineteenth. The visible light spectrum is now understood to be a part of a much larger spectrum encompassing all electromagnetic waves, including radio, microwave, x-rays, and gamma rays, as well as those just outside the visible range, called infrared and ultraviolet.

Light and the rest are wave forms, which can be characterised by their wavelengths; the spectrum shows all the different wavelengths. Visible light has wavelengths between approximately 380 nanometres (violet) to 750nm (red). White light normally contains a full range of colours. A surface may reflect light of different wavelengths to different degrees—then the surface may be perceived as coloured. Usually this would be a smear across some range of the spectrum. Also a light source may generate different mixtures of colours. Old-fashioned filament lightbulbs typically produce light that is stronger at the red end of the spectrum than daylight. Modern bulbs can often be made to emulate the old-fashioned ones or to be close to daylight—these are currently referred to as *warm white* and *daylight* respectively, with *cool white* somewhere in between.

One case that will be useful for further discussion is the sodium lamp, often used for streetlights. It's unusual in that the light it produces is (to a close approximation) strictly monochromatic—that is, of only a single wavelength, around 590nm, pretty well at the yellow-orange boundary. If the only light source in a scene is a sodium lamp, it is impossible to distinguish the colour of any surface, because however much or little light is reflected from a surface, all of it is this single colour.

## The physiology of colour perception

The light-sensitive cells in our eyes are of two types, rods and cones. Rods do not distinguish colours; however, cones are further subdivided into three types with different colour sensitivities, which enable us to see colour. These are called *red*, *green* and *blue* cones, which is an approximate way to describe their respective sensitivity to different colours of the spectrum. But actually,

each type responds to a smear of different wavelengths, and these smears overlap considerably.

If our eyes are presented with monochromatic light (such as from a sodium lamp), the response of each type depends on whereabouts in its smeared response the monochromatic light lies. Sodium light lies quite close to the peak of the response-smear of the red cells, but with a significant green-cell response as well (very little blue-cell response). Our colour perception depends on the ratios or proportions of these different responses—the brain says ‘*this* much red, together with *this* much green, but very little blue, looks like a particularly virulent yellow-orange’. It is *only* through these proportions that we perceive colour.

If the light hitting our eyes were not monochromatic, but smeared over a range of wavelengths towards the red end, we might nevertheless get a very similar effect: that is, a very similar proportional response from the three different types of cones. There are actually very many different combinations of the basic wavelengths that our eyes are quite incapable of distinguishing.

### Three-colour theory

Given that our eyes only have the three types of cones to distinguish colours, it seems plausible that we can construct colours using three primaries. That is, it should be possible to fool the eye into thinking it is seeing any particular colour by presenting it with suitable combinations of the three primaries.

What do we need for primary colours? The colour cones suggest something like red, green and blue. Indeed this is what is normally used for what is called *additive* colours. If you start with red, green and blue light sources, you can generate white light and more or less a full range of colours. Exactly this is done in some projection systems, with three separate projectors for the three colours, all focussed onto the same white screen. Something similar also happens in computer and television screens, with closely packed dots of colour. In each case, there is no interference between the colours—if the red light is projected, adding green or blue will not affect the red light itself, and the eye is free to see the mixture.

For printing on white paper, we have a different situation. Here we start with white light, but the printing ink filters some colours out—the more ink we add, the darker the result (this is called *subtractive* colours). For this

purpose it is best to use not red/green/blue but the complementary colours, cyan/magenta/yellow. However, it is much more difficult to get the colours looking right. Most printers also use black ink (because overprinting the three primaries doesn't produce a good black); some do much more complicated adjustments.

For an artist, mixing coloured paints, the situation is different again. Mixing paints is closer to subtractive than to additive colours, but does not work exactly like subtractive printing ink. A more usual set of primaries for this purpose would be red/yellow/blue, but most artists use a much wider range of colours to mix.

## Problems

The three-colour approach to images has proved successful, but it's worth exploring some of the issues around it.

First, let's think again about sodium light, and about taking photographs. If I photograph a sodium lamp, the three primary colour receptors in my camera will respond in a way which is similar to the response of the three types of cones. Then, if I display the resulting photograph on my computer screen (which uses LED technology), the image on the screen will be made up of a combination of red, green and blue LED cells. The challenge of displaying an image that looks good to me is the challenge of reproducing in my eyes roughly the same proportional responses that the original sodium lamp produced. The system might achieve this, though if we think in terms of the spectrum, it is very clear that the smear produced by my screen is hugely different from the monochrome sodium light itself.

Does this matter? Well, it might matter a lot.

For one thing, not all animal species are trichromatic as we are. Some have only two different colour receptors; some have four (specifically birds, reptiles and some fish). A tetrachromatic animal will see colour distinctions that we cannot see. Thus even if the screen image looks good to me, it would fail to satisfy the birds!

One interesting suggestion, not yet demonstrated, is that actually some humans have tetrachromacy—or at least that some of us have four different types of cones, which might give us effective tetrachromacy if we knew how to use them. I say 'us', but actually it's much more likely in women than in

men, for genetic reasons. It may even be the case that some women are able to make use of them, and thus see a wider range of colours than most of us. But even if this does not happen, the responses of individuals may differ.

It is well known, of course, that some people are *less* sensitive to certain colour differences than the majority—this is normally referred to as ‘colour blindness’. But if some people are more sensitive, or even if some people are differently sensitive, this means that something that I see as a good colour match might to these people seem a poor match.

Could there be a physical solution to this problem? Ideally, we might like to represent the full colour spectrum with many different finely graded colours. It would be possible to have more than three ‘primary’ colours, but it’s very unlikely that we could go far in that direction with (say) cameras or display screens. Thus once again, what we have is a compromise.

### Dots, lines, frames, pulses

Most discretisations of smooth variables in the world (but not all, as we have seen) involve dividing the continuity up into very many small steps. This works (when it does) because our perceptions do some of their own smoothing, and thereby restore some smoothness to something that is actually not at all smooth. This process probably involves not only the sense organs themselves, but also the neural processes that follow when sensory input is transmitted to the brain. In some cases, the sensory organ itself generates discrete signals even from smooth input, and these discrete signals must be interpreted smoothly. We have already seen how our eyes generate discrete signals for different colour ranges; it is also the case that brightness (or intensity) is conveyed in the eye-brain by the number of rods or cones that fire; in a given time interval, each one fires or does not, so at some level the internal process is digital anyway.

So some smoothing is natural, and this suggests that there is no problem about presenting data to the senses in discrete lumps, provided they are small enough. But it does raise the question of what ‘small enough’ means, and whether there are any other effects from such discretisation. The pointilliste artists like Seurat had a theory that their method of painting, building up colour shades from small dots of primary colours, actually enhanced our colour perception, making the images seem brighter.

Some recent films have been shot and played at 48 frames per second, rather than the usual 24. Although 24 fps is fast enough that the viewer is not normally aware of flickering, it seems that the smoothness of 48 fps causes some people to feel sick, from something like motion sickness. So there may indeed be effects of a rather oblique kind.

### Three dimensions

Now let's return to 3D display, where we started this chapter.

A physical solution to 3D display would be to create a model image in 3D space, which one could walk around and see from different angles, just as much as if it were real. But this does not seem like a very good solution for (for example) 3D movies. It might work for scenes involving people in a room, but outdoor scenes with buildings would have to be greatly reduced in size, and those with distant vistas would not work at all.

The binocular method pioneered by Wheatstone is a much more plausible solution. As I have mentioned, it depends on the fact that a lot of our depth perception comes from our binocular vision, with the two eyes turning a little inwards to focus on something close. There are other effects—it's also the case that each eye does its own focussing—in something like the way a camera is focussed, by adjustment of the lens. However, this is really only important for very close objects. In the far distance, binocular vision doesn't help much either—one useful clue here on earth is what artists know as *tonal perspective*, where the intervening atmosphere causes distant objects to look hazier and slightly bluer than they would look close to (on the moon, with no atmosphere, it is impossible to tell how far away or how high the mountains are). And of course there is the usual kind of geometrical perspective—because we know what sort of heights humans typically have, one good clue to how far away they are is their perceived size.

In movies and still photographs, both kinds of perspective are present anyway, of course—and indeed, when watching a film, one is normally well aware of the three-dimensionality of the scene. The present 3D film technology does not attempt to adjust monocular focus, but does add the binocular vision component to enhance the illusion of three-dimensionality.

This illusion has some interesting components. For example, suppose you are watching a 3D film, and you see a post in the foreground and some-

one passing behind it at some distance. Someone watching the same film from across the other side of the room will see the same thing—the person and the post will line up with her eyes at the same instant that they line up with yours. This makes no sense geometrically!

Nevertheless, as with sound and colour, the illusion is what is important. A pragmatic mixture of physics and physiology may be quite sufficient to achieve a good illusion.

In the next chapter, I will consider other ways to represent particular kinds of images and sounds.