

Colour

Introduction

Colour is a critical component of the patterns we observe, but it is somewhat orthogonal to the topics of study this year, hence we won't devote a lecture or specific activity to it, but it seems useful to have some background information.

There are some other good sources out there, though they often go into much more detail than we need. For instance see:

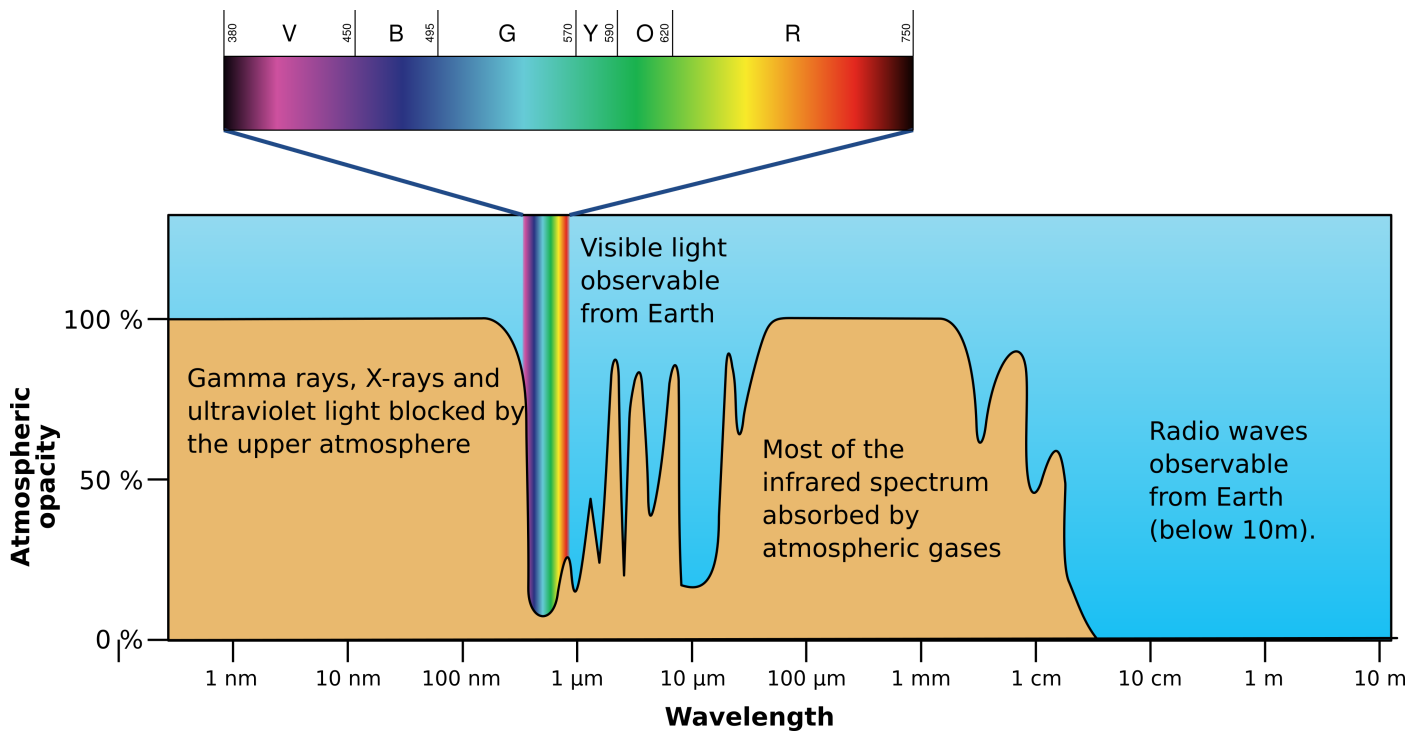
- [Computer Vision, Lecture 11, Rajesh Rao](#)

Light

Any discussion of colour has to start with light. It was first noted by Newton that visible light was made up of a spectrum of "colours". Those colours correspond to bands of frequency of electromagnetic radiation (EMR). They are in fact part of a much wider range of EMR. The visible light spectrum lies has wavelengths from 380 to 740 nm (430–770 THz in frequency). Newton actually equated these frequencies to musical notes!

- https://en.wikipedia.org/wiki/Visible_spectrum
- https://en.wikipedia.org/wiki/File:Atmospheric_electromagnetic_opacity.svg

This band is particularly interesting because our atmosphere is almost transparent in this range, so it is transmitted freely, and hence our eyes have adapted to be able to see it. We call this the *optical window*. The figure below illustrates the optical window (it's a mash up of figures from the links above).



At shorter wavelengths (higher frequencies) we have ultraviolet and on upwards to x-rays. At longer wavelengths we have infra-red ranging into microwaves and radio waves, but the optical window means that this is what we are able to “see”.

The frequencies of visible light are broken into the traditional spectrum — ROYGBIV — based on our perception of those frequencies. However it is critical to note is that any (visible) light is made up of a spectrum of frequencies of light ranging from 380 to 740 nm wavelength. The amount of the various frequencies leads to the apparent colour of the light. But the transformation into perceived colour is not at all simple. A red coloured object might be reflecting “red” frequencies, but more commonly it is reflecting a variety of frequencies that we perceive as red overall.


Perception

The eye is, of course, critical for detecting light. Its structure and organisation are all interesting, but the critical information for consideration of colour is that it has two specialised cells called *rods* and *cones* that absorb and respond to light. Rods respond to light intensity, and so are “colour blind” but there are more of them, so they give us more sensitive black and white vision. Cones are responsible for detecting colour.

- <http://hyperphysics.phy-astr.gsu.edu/hbase/vision/colcon.html>
- https://en.wikipedia.org/wiki/Cone_cell

There are three types of cones, and each has a different response to the possible spectrum of visible light. Roughly, they correspond to detecting red, green, and blue light, but it is simplistic to think of them as responding to a narrow band. In fact that

bands overlap. A lot. So scientists refer to the three types as long (L), medium (M) and short (S) corresponding to the peak of wavelength to which they respond. The figure below illustrates the strength of response of each type to the different frequencies of light (again it's a mash up from the links).



Any light *signal* composed of a set of frequencies will be absorbed by these cones in proportion to the cones responsivity at the respective frequencies, and our brain translates this into a colour. So although we might think of a colour as being composed of red, green and blue components it is much more complicated.

A *metamer* is a group of frequencies of EMR that we percive as a single colour. There can be many. For instance, a green laser's light is made up of a single frequency, but we can combine a group of different frequencies of light into the same green colour.

We call this typical situation *trichromatic* vision, but if you are colour blind you might be missing one type of cone, or having a different response from one of the cones. It is reported that some people have four types of cone, and thus have *tettrichromatic* colour vision, and can see colours no-one else can see.

Colour Models

So colour is incredibly complex, but for most people a colour can be modelled by 3 numbers: the expected response of the red, green and blue cones to the light. Thus we often represent colour spaces by a 3D space. Commonly, you will hear of RGB triples specifying the amount of red, green and blue in a colour. Matlab uses this model (by default) for giving the colour of lines or regions.

However, this is still an oversimplification, and understanding the oversimplifications involved can be important in some contexts.

There is a *gamut* of colours we can perceive. The goal of a colour model is a mathematical abstraction of color, meant to simply describe and create each of these colours. The difficulty is that any achievable combination of three components may overrepresent some parts of the gamut, and not reach every corner.

Most colour models use three numbers (e.g. RGB triples) but others involve 4 numbers or more. Some are related to a model of perception (e.g., RGB) others are based on other conceptions of color. For example, HSV = (Hue, Saturation, Value) is based on the idea that the 'color' or *hue* is a 1D value (a position on the ROYGBIV spectrum), which is then modified by its value (how bright it is) and its saturation (how intense the colour is: low saturation corresponds to washed out, grey colors, and high saturation to vivid, primary colours). Various colour-picking tools will let you choose colors according to their RGB or HSV values. It is often possible to transform between colour models via simple linear transformation, but some models are more complex.

- https://en.wikipedia.org/wiki/Color_model
- <https://www.designersinsights.com/designer-resources/understanding-color-models/>
- https://en.wikipedia.org/wiki/CIELAB_color_space
- <https://www.rapidtables.com/convert/color/rgb-to-hsv.html>
- https://en.wikipedia.org/wiki/HSL_and_HSV
- http://coecsl.ece.illinois.edu/ge423/spring05/group8/finalproject/hsv_writeup.pdf

Which model is best? That's hard to answer. Some models are essentially equivalent, in that they can produce the same gamut, but not all models are such. Also, most color "numbers" the R,G and B for instance, are given with limited resolution, often only 8 bits, allowing only 256 values. That limits the accuracy with which parts of the colour spectrum can be reproduced. Some are better at, for instance, reproducing human skin tones. Model are also highly related to how the colours are actually reproduced, which will we talk about more below.

Complicating colour perception is the fact that our optical system is *highly* adaptive. Our colour perception is altered by adjacent colours, and many optical illusions are possible based on colour perception. But we won't go into those details here.

Hex

Colour values (e.g., RGB triples) are often given in hexadecimal values or base-16 values. Hexadecimal (or just hex) values are written using the 'digits', 0,1,2,3,4,5,6,7,8,9,a,b,c,d,e,f, representing the numbers 0-15.

Sometimes hexadecimal numbers are written starting with a '#'.

Colour codes are often 8 bit numbers, so, we might commonly see values in the form

```
#00 = 0
#01 = 1
#0f = 15
#10 = 16
#11 = 17
...
#ff = 255
```

Often the three number are put together, so a highly saturated red might be written

```
#ff0000
```

Many common colours used in web pages can be seen at

https://en.wikipedia.org/wiki/Web_colors

Colour Creation

Colour models are often created to support a particular colour creation system, by which we mean a system whereby colours are created from simple elements. Typical examples include

- a colour printer
- a TV or computer monitor


These are only two examples, but they cover the two most common cases, and show one of the main dichotomies in color reproduction.

1. A TV (or any other device that projects its own coloured light) typically uses an *additive* process. It adds different colors together by adding their values in the RGB space. Many such devices use red, green and blue sources of colour, these days often from Light Emitting Diodes (LEDs) directly linking them to the RGB model.

2. The colour on a printed page is created by reflection of ambient light from the page. When mixing colours in this case typically a *subtractive* model is used whereby we describe how much Cyan, Magenta and Yellow is present. These are how much of these inks to put on the page. The ink reduces the amount of reflected light, making the colour darker.


The two are often illustrated with Venn diagrams:

Additive model (RGB)



Transmission tech., e.g., TV or computer screens

Subtractive model (CMY)



Pigment mixing, e.g., printers

It's hard to create pure blacks by mixing pigments, so often it is augmented to be a CMYK model, where the K allows a separate amount of black to be added.

In many printing processes it is difficult to (i) mix inks before printing, and (ii) difficult to accurately specify the intensity of the ink used. Thus, printing often uses (very) small dots. The pattern of dots used is called a *dither* pattern, and the process called *halftoning*. If less of a particular colour is needed, smaller, more sparsely spaced dots are placed on the page.

- <https://en.wikipedia.org/wiki/Halftone>
- <https://en.wikipedia.org/wiki/Dither>
- https://en.wikipedia.org/wiki/CMYK_color_model

Here is a simple greyscale/black and white example of dithering.



Color Choice in Procedural Generation

One approach to generating colours in a pattern is to try to mimic reality. For reasons given above, and others, it isn't quite that easy.

However, there are many simple approaches that can work in some settings depending on the applications.

Reproducing nature

Reproducing colours could be as simple as measuring them (in an image, for instance) and reproducing the observations, or as complex as modelling the way the colours are created (e.g., modelling the chemicals causing colours — chlorophyll — the natural light, and the reflection process). That might seem impossibly complex, but it's almost what seems to happen in modern computer-generated images in games and movies. We will talk, in this course, about some of the underlying models that can be used.

Data visualisation

Colour choice is an important part of designing a good data visualisation. This is not formally part of this course, but good colour design is so important it's worth mentioning.

In this case, it is not about single colours it is about how the colours work together. Good contrast is important so that the data being displayed is obvious. But also relations between colours can draw the eye to relationships.

Matlab's default colours are very familiar but not always a good choice.

There are some good tools available choosing colours, for instance:

- <http://paletton.com>
- <http://colorbrewer2.org>
- <https://coolers.co/9e0031-8e0045-770058-600047-44001a>

And more advice

- <https://blog.hubspot.com/marketing/color-theory-design>

Art

In art, the goal is an aesthetic set of colours. As in visualisation, choices of colours for art are not about single colours, but rather about palletes.

No-one can completely quantify such a process, but there are many attempts to generate aesthetic colours. One could do an entire course on the history of colours in art. However, there are some simple ideas that can help. For instance models like HSV can be used to generate a set of colours with the same hue but different brightness or saturation. Or to generate a set of colours with some other similar characteristic.

- Pastel colors tend to be desaturated (low saturation)
[https://en.wikipedia.org/wiki/Pastel_\(color\)](https://en.wikipedia.org/wiki/Pastel_(color))
<https://www.kentor.dev/posts/generating-pastel-colors-for-css/>
- Saturated colors (colours with high S) are primary-looking colours.

So we can use these ideas to generate colours in a particular model by keeping some values constant, and allowing others to be randomised.

Another approach is to have colours that are randomly drawn from another image whose colours are already appealing.

One interesting possibility, however, is that the colour-generation mechanism becomes a subject of the art itself. For example, Roy Lichtenstein, is famous for making dithering patterns a very obvious component of his art, or for choosing colours in fine art that were derived from colours imposed in common printing processes. The contrast between comic-book-like art and more sophisticated topics draws the eye in an engaging way.

https://www.google.com.au/search?q=roy+lichtenstein&source=lnms&tbm=isch&sa=X&ved=0ahUKEwiSz8Do8JfWAhWGJZQKHaiQCvoQ_AUICigB&biw=1036&bih=799

For more general advice see

<http://devmag.org.za/2012/07/29/how-to-choose-colours-procedurally-algorithms/>