

3 - Seeing Colour & Colour Spaces


Prof Peter YK Cheung

Dyson School of Design Engineering

URL: www.ee.ic.ac.uk/pcheung/teaching/DE4_DVS/
E-mail: p.cheung@imperial.ac.uk

In this lecture, we will examine how we can see and interpret colour and explain what is meant by “colour spaces” in the context of designing visual systems.

Photoreceptor cells in Retina



- ◆ Retina have three types of cones that are primarily responsible for seeing red, green and blue colours


Source: American Academy of Ophthalmology

Visible light occupies a small part of the electromagnetic spectrum that spans the wavelength range of 380nm (blue) to 740nm (red).

Our retina has three types of cones that are responsible for us “seeing” colour. They are mostly concentrated in the fovea region.

These are often referred to as RED, GREEN and BLUE cones. Their names are actually very misleading. These photoreceptors are sensitive NOT to single colour light, but a spectrum of frequencies.

Densities of rods and cones




- ◆ Concentration of cones at very high density only in the fovea region, around $\pm 6^\circ$ around the centre of vision.
- ◆ Concentration of rod is spread out.

This diagram shows the density of photoreceptor on the retina from the centre of the fovea. Most of the cones reside within $\pm 6^\circ$ around the fovea region. The rods are spread out throughout most of the retina. Rods do not “see colours” – they are sensitive to a broad spectrum of light. However, rods are both much more plentiful and also more sensitive to photons than cones. Therefore, our low light vision is mostly due to rods. However, cones are sensitive to narrower area of the visible spectrum and are therefore responsible of us seeing colour.

In the area where the optic nerve goes to the brain, the retina has no photoreceptors and therefore there is a hole in the image that the retina can “see”. It is interesting how our brain interpolates our visual in such a way that the hole is generally not perceivable in our day-to-day vision.

How we see colour of objects?

- ◆ The colour that human perceive in an object = the light **reflected** from the object




There are actually two types of light that we normally see: emitted light (e.g. from a computer monitor or TV) and reflected light.

For most objects that we perceive, we are seeing reflected light. The same is true with camera capturing scenes – these are light from a source (e.g. the sun) bouncing off objects. The actual light we see depends on the nature (i.e. spectrum) of the light source, the reflectivity of the object, including its colour and its surface properties (i.e. roughness, texture etc.) and other factors.

Why use Colour in Visual Systems

- ◆ Why use color in the design of visual systems?
 - Color is a powerful descriptor
 - Object identification and extraction
 - eg. Face detection using skin colors
 - Humans can discern thousands of color shades and intensities
 - c.f. Human discern only two dozen shades of grays
- ◆ We even use false colour to characterize greyscale images (e.g. in medicine)





In designing any visual systems, the use of colour is very important. This is partly due to human's ability to distinguish many different shades of colours (depending on individual) while our eyes can only percept perhaps 20 or 30 shades of grays.

As a result, we often even assign colours in greyscale images, such as medical images, to highlight features that may otherwise not be noticed by a clinician. Such false colours are used in many fields such as astronomy (e.g. photo of distant galaxies), remote sensing (e.g. satellite photo of cities) and medical imaging.

Trichromatic Theory

◆ Trichromatic theory (Young and Helmholtz)

- Three types of cones in retina
- Each type has photopigments sensitive to three different wavelengths
- Red (long), Green (medium), Blue (short)



The statement that red cone cells are sensitive to red light is NOT strictly correct. In reality, the three types of cone cells are sensitive not to a single colour light, but to a spectral range where the three colours dominates.

Shown here is a typical sensitivity vs wavelength plot showing the sensitivity characteristics of the three types of cone cells. It is probably more accurate to label them **L, M and S** receptors instead of R, G, B receptors.


How does combining signals from these three types of cones allow us to see all colour? There are two major theories: the trichromatic theory and the opponent process theory.

In trichromatic theory as proposed by Young and Helmholtz, perception of colour is achieved by adding or combining stimuli from these three types of cones. R, G, B and known as primary colours. Mixing them in various proportions can create a very large range of colours. This theory also explains how colour blindness occurs in some individuals when one or more types of cones are defective.

Trichromatic theory of colour is not the only game in town. Some others support the idea that additive effect does not explain everything and other factors resulting in subtractive or inhibitive actions of receptors also affect our colour perception.

In other words, we DO NOT SEE red, green and blue. Instead, we have three types of cones that are activated to different degrees by the incoming light depending on its wavelength, and we **perceive** the colour. It is a psychological things, and NOT a physical thing!

How does our brains "see" blue?




- ◆ Blue light at 440nm wavelength will strongly stimulate the S cone, but not the M and L cones.
- ◆ Our brain interprets this as blue colour.

This shows what happens when light at a wavelength of 440nm falls on the retina. The S (short) cones are strongly stimulated while the M and L cones are barely triggered. The degrees of activation for the three types of cone are shown here. The response is normalized, meaning that the responsivity is relative to the maximum response.

In this case, our brain perceives seeing a deep blue colour.


"Seeing" green and red colours



0.9

1

0



0.4

0.1


0




Here is what happens when light of 540nm and 630nm reach the retina. The brain will interpret these as green and red light respectively. Note that for the 540nm light, both the M and L cones are strongly activated, and yet, we see this as the green colour (not RED+GREEN)!

Don't call them RGB, but LMS!


630 nm



540 nm



440 nm



Long cone

~~Red cone~~



Middle cone

~~Green cone~~



Short cone

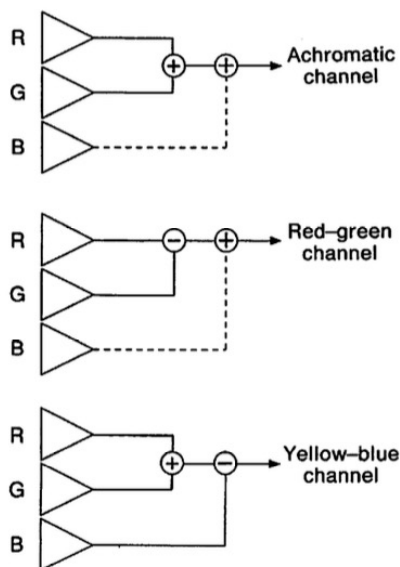
~~Blue cone~~

Therefore, calling these Red, Green and Blue (RGB) cones is actually very misleading.

We shall therefore call them L, M and S cones, respectively responding to long, medium and short wavelength of light.

Nevertheless, it is still useful to label these in the three different colours for each of representation.

Opponent Process Theory (Herring)



- ◆ Three types of mechanisms combine cones signals in both additive and inhibitive ways.
- ◆ The achromatic channel is sensitive to all three cones and detects only the luminance information (i.e. black vs white)
- ◆ The Red-Green channel compares the difference in red and green cone signals (i.e. one inhibits the other)
- ◆ The Yellow-Blue channel compares the difference between blue cone signals with the sum of red and green signals.
- ◆ This theory explains some of the colour perception experiments in human visions

Source: Tovee p45-46

The three additive colour idea is not the only one in town. There is another theory which is widely accepted.

The opponent process theory was proposed by Ewald Herring to explain why we do not perceive some colour combinations such as greenish-red or bluish-yellow. He suggested that our perception of colours involved some form of inhibitive or opponent interaction between the three types of cone signals.

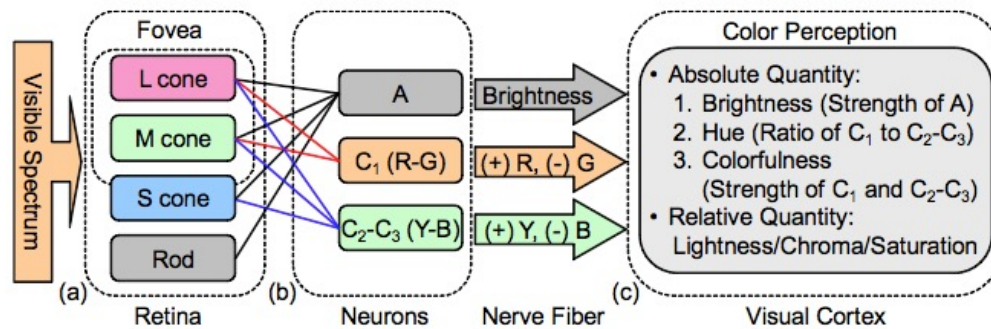
The three opponent mechanisms are shown in the slide here.

The first is responsible for achromatic perception, meaning that it detects the differences in luminance but not colour.

The second compares the difference between the red and the green cone signals. These two colours oppose each other (hence we don't see redish-green or greenish-red).

The third compares the differences between the blue cones and the sum of the red and the green.

Opponent Process Theory of Colour Perception



- ◆ Here is another view of how the opponent process theory explains our colour perception mechanism.

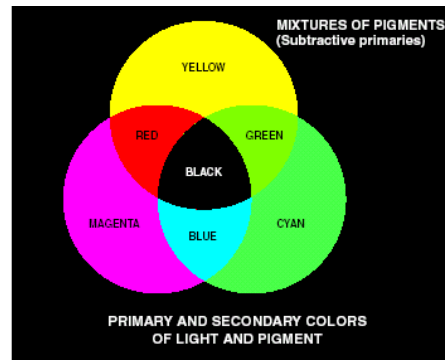
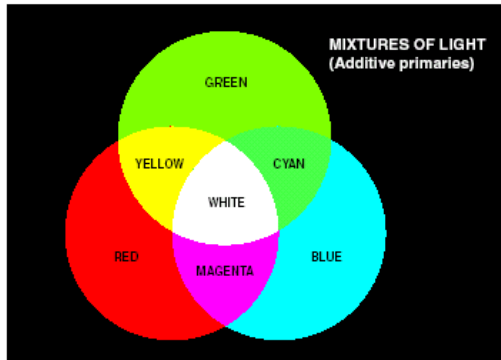
This is another way of depicting the opponent process theory of colour perception.

It is important to note that both trichromatic and opponent process theory are theories – they explain observations of experiments, but they are not absolute. This is because "seeing" colour is a psychological phenomenon, not a physically one.

For example, the colours involved in the opponent process theory have recently become more complicated and not just involving R-G and Y-B.

Primary and secondary colors

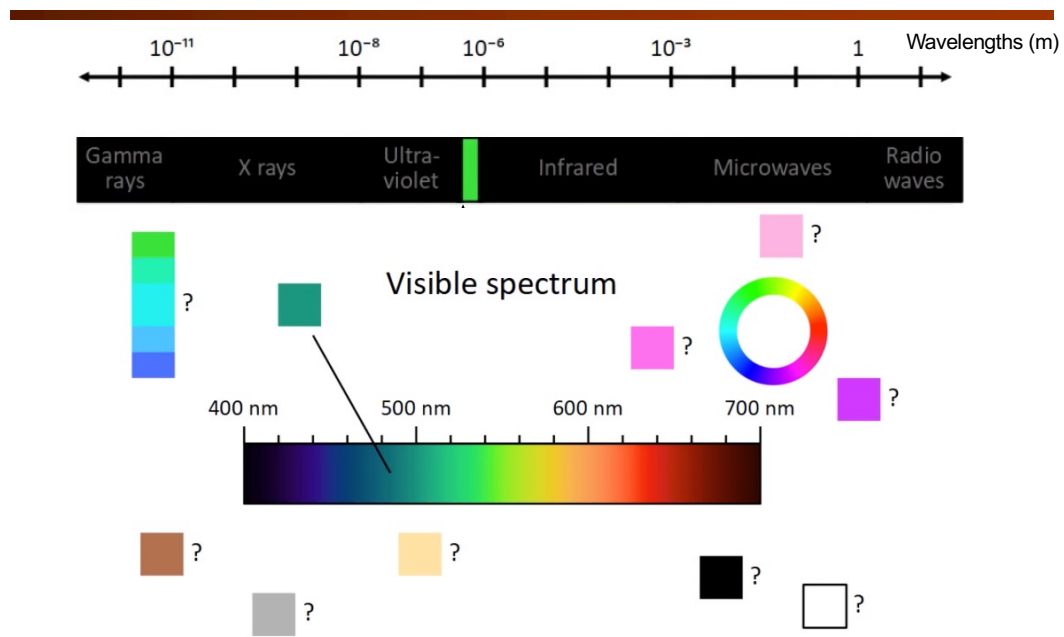
- ◆ In 1931, **CIE**(International Commission on Illumination) defines specific wavelength values to the **primary colours**
 - B = 435.8 nm, G = 546.1 nm, R = 700 nm
 - However, we know that no single color may be called red, green, or blue
- ◆ **Secondary colours**: G+B=Cyan, R+G=Yellow, R+B=Magenta



The CIE also defined R, G, B at respectively these wavelengths: 700nm, 546.1nm and 435.8nm, and called these the PRIMARY COLOURS. Mixing lights from these three primary colours produces yellow, cyan, magenta and white. Mixing different proportion of these colours produces the other possible colours.

They also defined three SECONDARY COLOURS: Cyan, Yellow Magenta. Mixing these produces red, green, blue and black!

Limitation of the visible spectrum



PYKC 19 Jan 2026

DE4 – Design of Visual Systems

Lecture 3 Slide 13

It is important to appreciate two things:


1. Electromagnetic waves outside the very narrow visible spectrum are perceived by us as the colour BLACK. For example, when our eye is presented with infrared or ultra-violet “light”, we see these as black colour.
2. The visible light spectrum shown at the different wavelengths are NOT the only colours we see. This is because light we see are NOT monochromatic. Monochromatic light are those having only light of a single wavelength. But light in general has many wavelengths mixed together. This is called a polychromatic light source.

Light from the sun is an example of a polychromatic source – it contains energy at a very broad spectrum of light including those that are outside our visible range. We perceive sunlight as white which does not exist in the colour bar shown here.

In other words, we cannot use this visible spectrum, a 1D colour bar, to represent colour in general because many colours will be missing! You will not find purple, brown, grey, or even black and white colour, and many other colours in this spectrum.

So how can we describe colour?

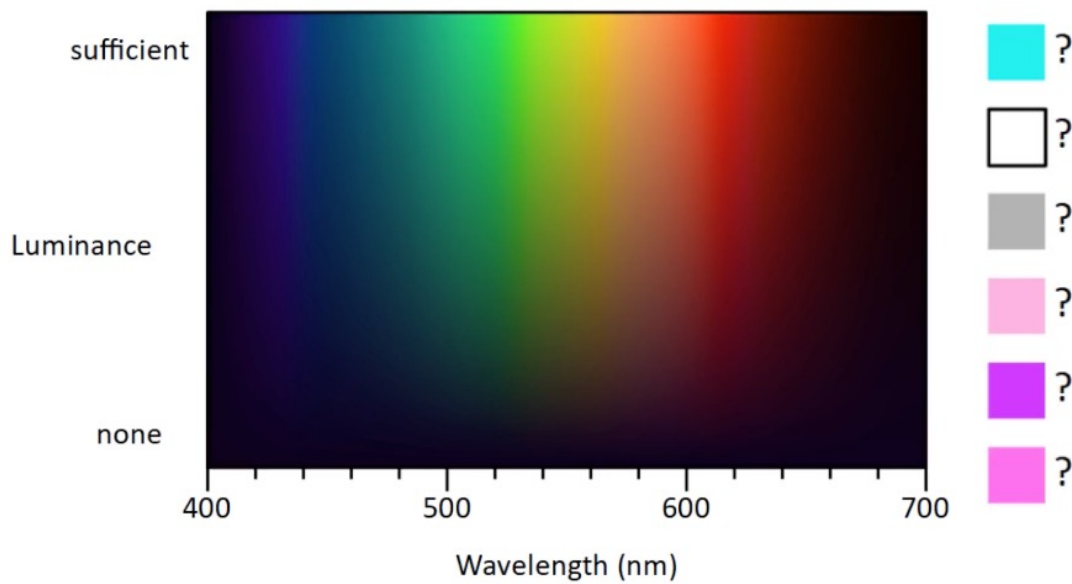
The idea of luminance and chrominance



One way is to first divide colour into two properties: luminance and chrominance (just like rods and cones). Shown here is two examples: left is the case where a very bright green is received by the brain. All the L and M cones are strongly activated (0.9 and 1), but not the S.

When a dim green light reaches the retina, only some of the L and S cones are activated. In that case, we see a much weaker colour and the we perceive a dark green colour. Therefore, our eyes perceive the same green colour but at different intensity values.

Luminance vs Wavelength is not sufficient




We can now use the luminance value of the colour to provide a richer representation as shown in the 2D plot here.

The y-axis is the luminance value (from none to high), and the x-axis is a wavelength. We can now see different brightness of the rainbow colour. This is a 2D representation of the 1D colour bar we had before.

This is still not sufficient because all the missing colours are still missing.


LMS Colour Space



Since our brain receives signals from L, M and S cones, we can create a 3D cube, with three axes spanning 0 to 1 each. Each axis represents the degree of activation with the L, M and S cones. In this slides, light at a wavelength of 500nm will active [L M S] by [0.3 0.5 0.1].

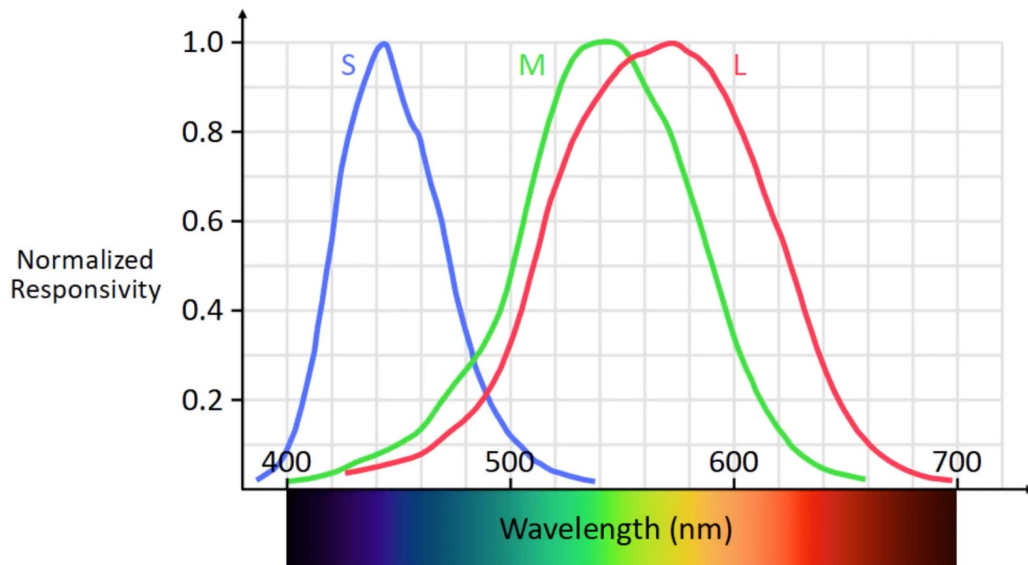
This represents a point in the 3D cube as shown here.

Mapping wavelengths of colour to 3D space




If we now go through all the spectrum wavelengths and map each into the x, y and z axes of the cube, we can see the colour here in 3-D trajectory as shown in the slide.

Normalised Cone Responses



Let us examine the normalized (i.e. scaling maximum to one for all cone responses) in a bit more detail. The shape of the response is “bell shape”, which can be approximated by a Gaussian function. Doing so allows us to provide a mathematical and standardized way of modeling the human visual responses of the three types of cones.

Gaussian Approximation



PYKC 19 Jan 2026

DE4 – Design of Visual Systems

Lecture 3 Slide 19

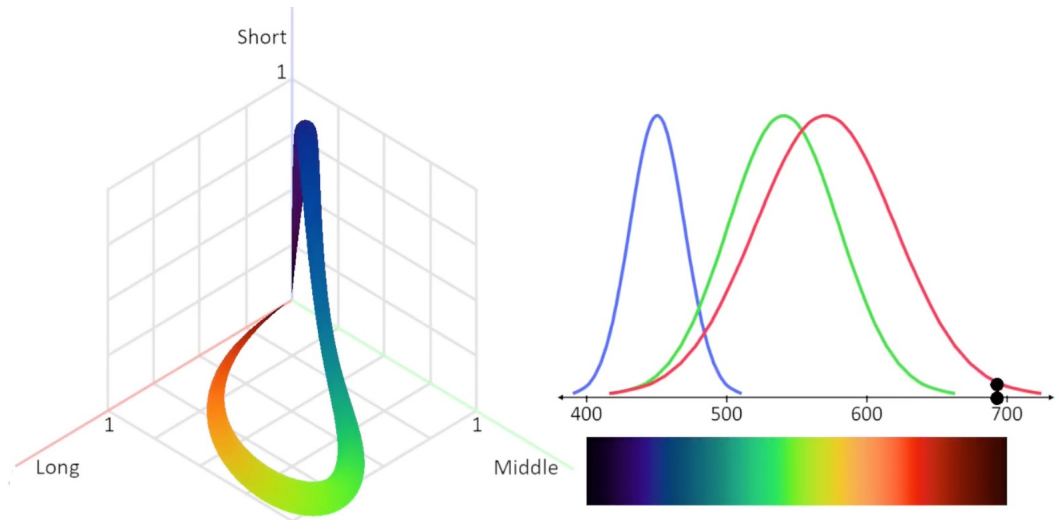
The gaussian distribution function has the form:

$$y = e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}$$

where μ = mean, and σ = standard deviation.

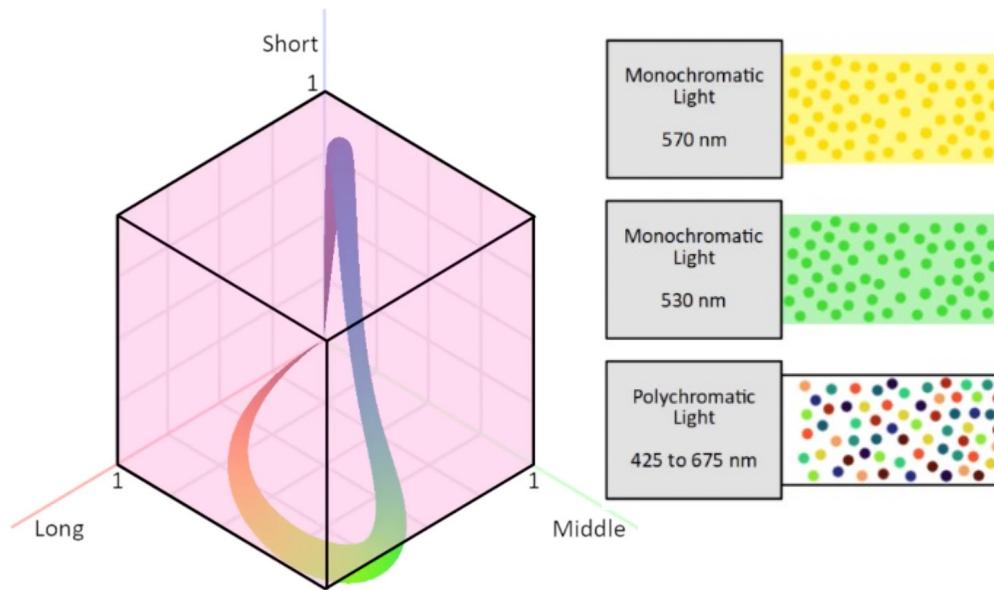
With the values shown, all three responsivity functions can now be represented by a mathematical formula.

Visible Spectrum mapped onto LMS space



By scanning from left to right on the visible spectrum colour bar, and mapping each of the L, M and S values of this gaussian approximation of the cones responsivity curves, we can create this coloured ribbon within the LMS cube.

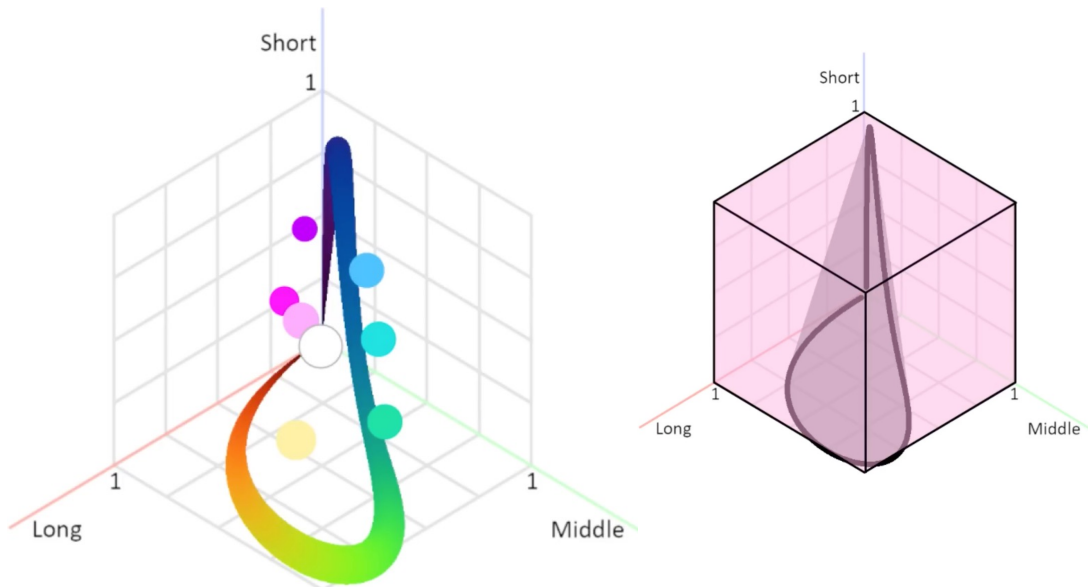
LMS cube can have many more colours



Of course, the entire cube has lots of space outside this ribbon to accommodate many more colours.

In general, a colour source can be polychromatic meaning that it contains lights with many wavelengths. These will fill part of the pink space outside the visible spectrum ribbon.


More colours in the 3D Colour Space



Taking all the possible combination of lights, we can fill the space as shown in the slide. All the perceivable colours occupy an area of a cone shape inside the cube as shown. The coordinate $[1\ 1\ 1]$ is pure white is, and the origin $[0\ 0\ 0]$ is pure black.

Now we can represent colours such as magenta and cyan inside the LMS cube space.

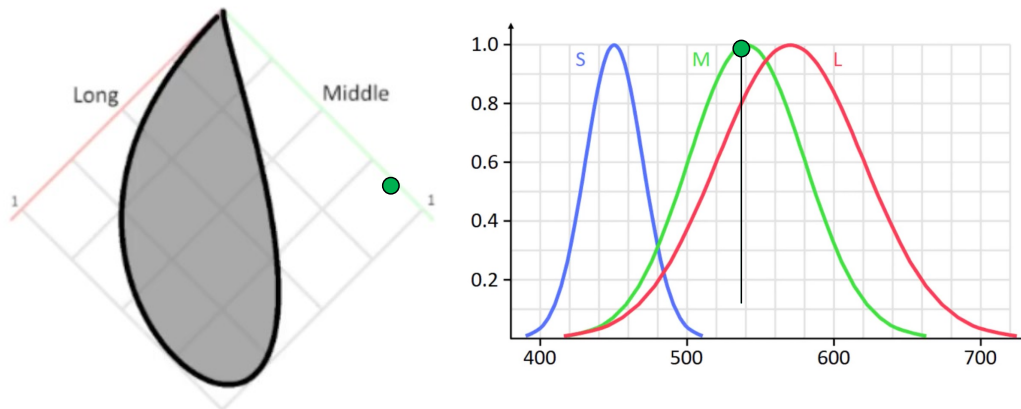
The Visible Gamut



Note that all perceivable colours DO NOT occupy everywhere inside the LMS cube. Only colours inside the the cone are visible by us human. We will examine why in the next slide.

In this slide, we show the cone volume in three different views: top, left and right. The unshaded areas are the illegal space where colours occupying these spaces cannot be visible by our eyes. Why?

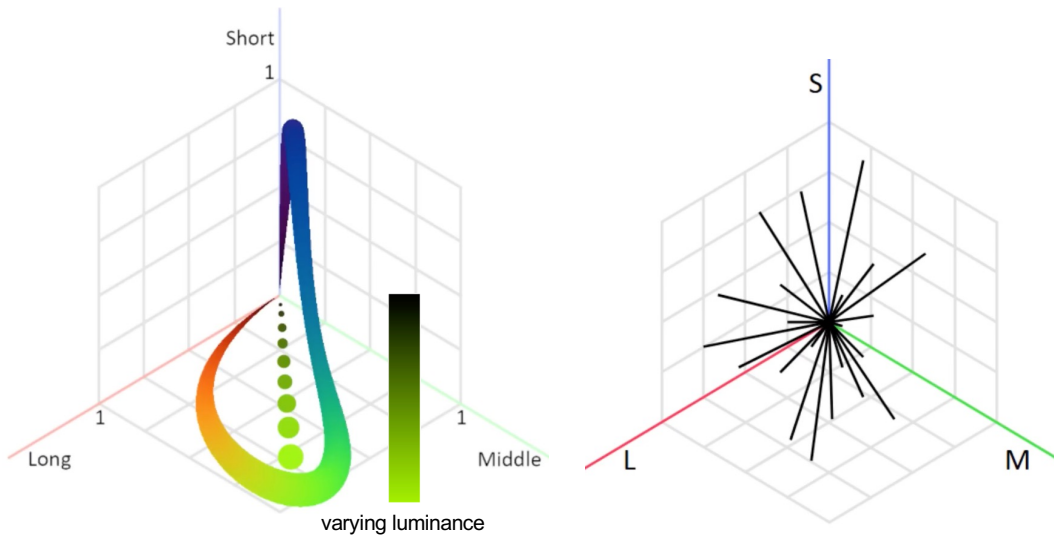
Why are some colour impossible in 3D Space?



Suppose we wish to have a colour light source that activates the M cones strongly, but not the L cones. That's where the dot in the TOP view of the space is placed.

Looking at the responsivity curves, there is NO WAY that one can activate the M cones strongly WITHOUT also activating the L cones quite strongly. This is shown on the curves on the right. Therefore, such a colour does not exist that produces such a response in the retina. So, this colour in the space IS NOT possible. It is called an impossible colour.



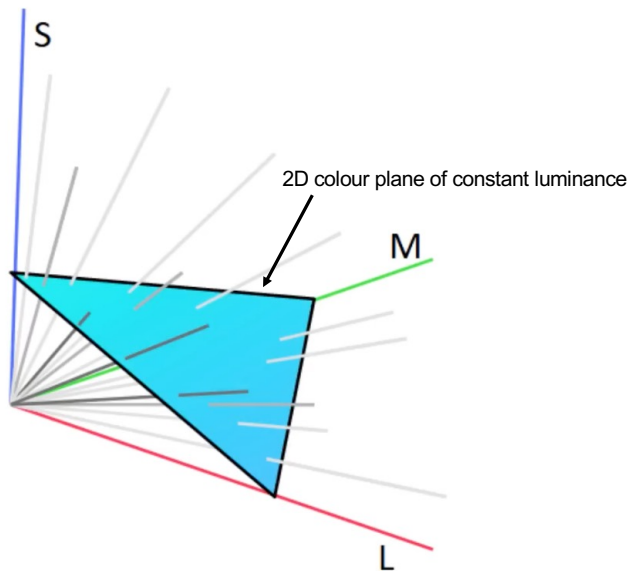
Constant Colour (Chromaticity)



Viewing and perceiving colour in a 3D cube is difficult. It would be very nice and convenient if we can somehow map colours onto a 2D space – then we can draw them on a sheet of paper or display on the computer screen.

To understand how this can be achieved, let us consider where the green colour with different level of luminance sits in the 3D LMS cube. This is shown as the dots here. All the green colours of different luminance value follow a straight line. In general, the trajectory of a constant colour (chromaticity) with varying luminance are always straight lines emitting from the origin (which is BLACK).

Plane of Constant Luminance

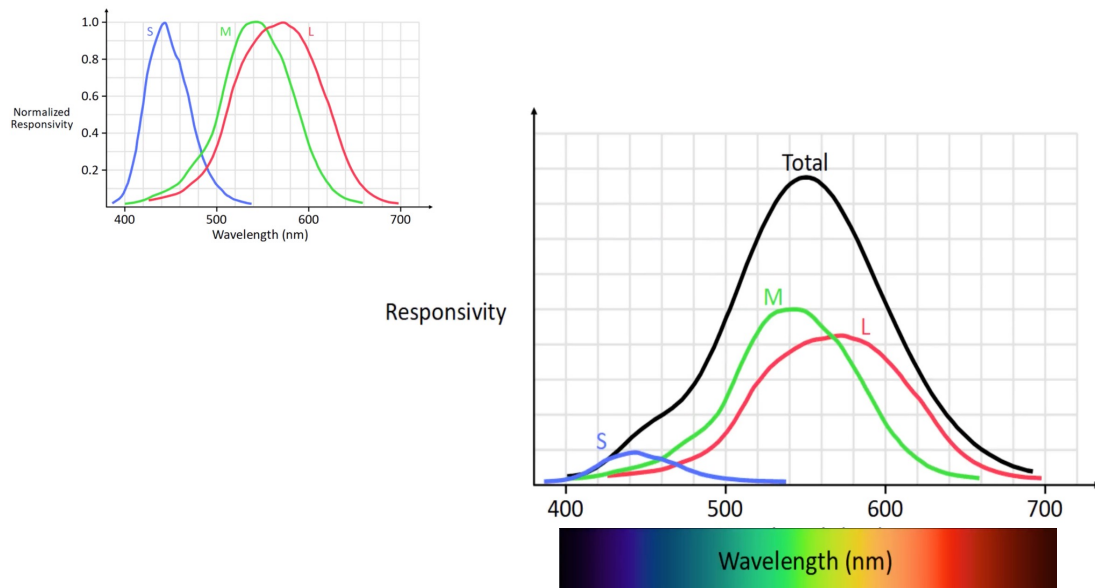


On these lines of constant colour (chromaticity), we can find the point where the luminance are the same. These points fall onto a 2D plane.

In other words, constant chromaticity forms lines radiating from the origin. Constant luminance forms a plane that intersects all these lines.

Such plane can be represented either as a colour circle or a colour triangle – a commonly used method that allows a user to pick a colour to use for painting or drawing.

A Complication - Actual Cone Responsivity



PYKC 19 Jan 2026

DE4 – Design of Visual Systems

Lecture 3 Slide 27

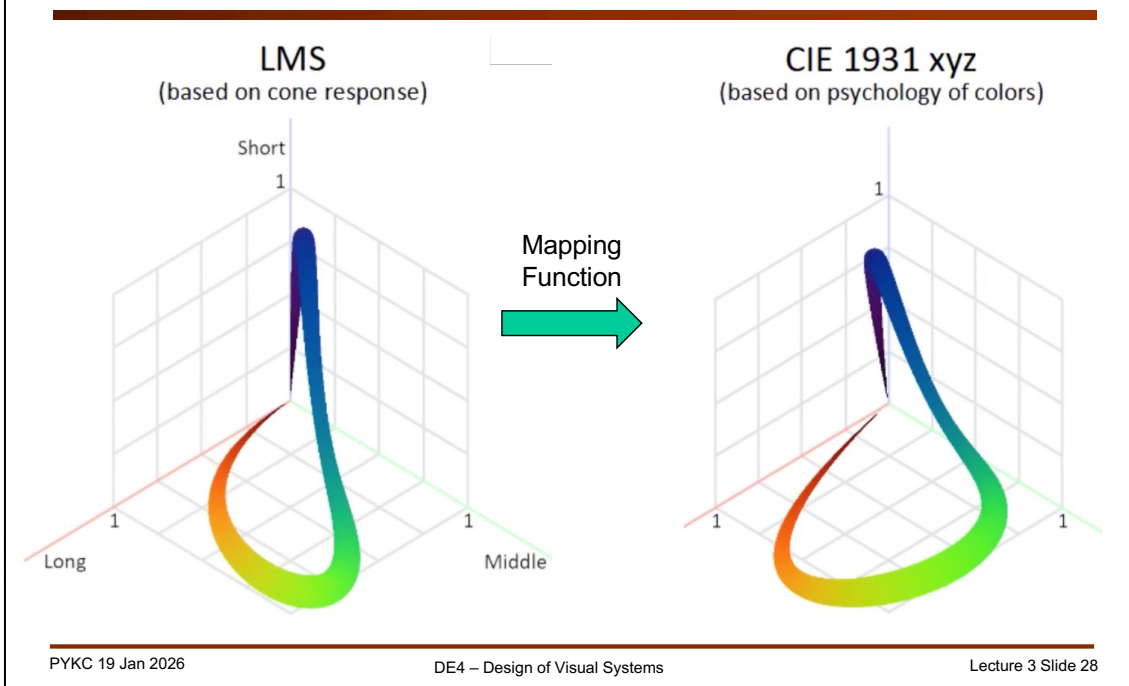
So, LMS cube appears to be a good way to represent all visible (and impossible) colours. Further, within the cube, we could even separate luminance and chrominance.

However, the LMS model we have used so far is NOT actual response of the human eye. Why? The responsivity curves were all normalized. In reality the S cones are far less responsive than the L and M cones as shown in the slide.

If we add all three responses together, we get the total response curve in black. This suggests that the different colours produce different levels of activation to the brain. The peak level is for the yellow wavelength. That's why yellow colour appears brighter than the blue colour.

So, we need to adjust the luminosity to the human perception.

LMS vs CIE 1930 xyz Colour Spaces



Because of the way human perceive colour is not exactly the same as that based on normalized cone responses, an organization known as the **International Commission on Illumination** (CIE) got together to define a new colour space known as the **CIE 1931 xyz** colour space (or just xyz for short). The LMS cube is now mapped to an xyz cube taking into account the psychology of human vision.


The mapping is linear and obeys the following two rules:

1. A straight line in LMS (RGB) space also maps to a straight line in xyz space.
2. The origin stays the same, therefore black is the origin in both spaces.

The xyz space is generally used for many UI applications.


We normally refer to this CIE space as lower-case xyz instead of XYZ, because the parameters are “normalized” such that $x+y+z=1$.

CIE 1931 xy Chromaticity Diagram - monochromatic



Keeping the luminance at a constant value, we can trace in the x and y axes their corresponding values after the transformation for the visible spectrum at single wavelength as shown in this animated slide. This is now sufficient to represent all visible colours in a 2-D plane as before.

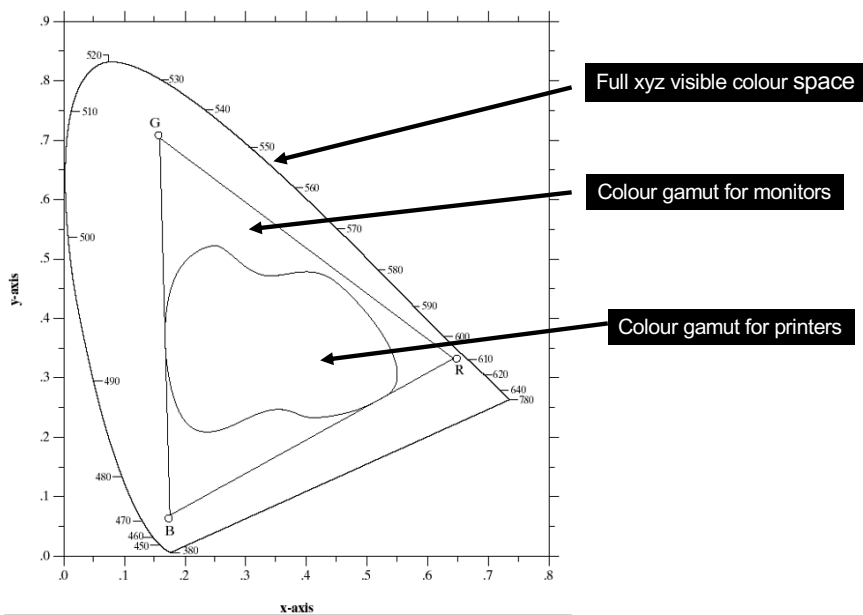
CIE 1931 xy Chromaticity Diagram



If we now fill in all the possible colours, we form the full CIE 1931 xy chromaticity diagram. The edge of this shape corresponds to colour at different wavelength.

The colour inside the shape is for mixture of different colour light with the centre being white!

Restricted colours in printers and monitors





While our eyes can perceive all the colours in the cone region, reproducing this full range of colour in monitors or on printed document is not possible.

This 2D plot shows the region of inside the triangle which encloses all the colours that an RGB monitor can produce. The region for a printer is even more restrictive as shown here.


HSV Colour Space



0 to 360
~ chromaticity
angle



0 to 100
~ chromaticity
distance
from center



0 to 100
~ luminance



So far, we have considered the LMS space based on the cone responses, the xyz space based on human responses. There is yet another colour space that are commonly used, the HSV (or HSI) colour space. This can be derived from the xyz space as follows.

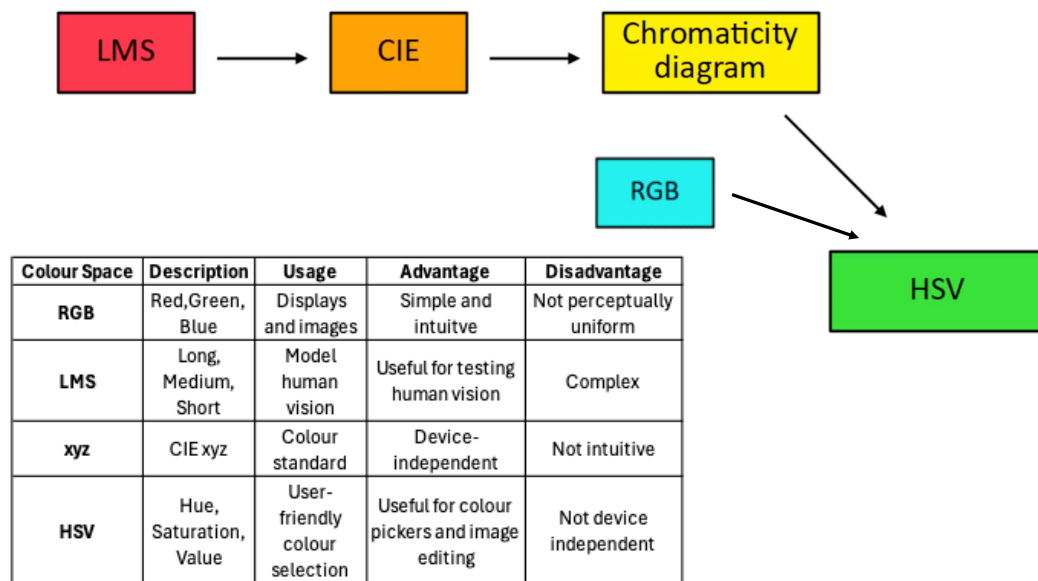
The V (value) parameter measures the luminance level or the intensity, and is between 0 to 100.

The S (saturation) and H (hue) parameters measure one of two aspects of chromaticity.

S (saturation) measure the amount of saturation of the colour, and It is the straight-line distance from the centre or white colour of the CIE xyz colour cone and is normalized to 0 to 100. Therefore, it characterizes the shades of a colour (e.g. shades of blue).

H (hue) measures the direction of the colour radiating from the white colour centre. This is called the chromaticity angle. H and S together defines exactly the coordinate of the colour in the plane.

Summary of Colour Spaces



This is summary of the various colour spaces we have discussed in this lecture.

There are in fact very large number of different colour spaces being proposed and used. They are all linearly transformable from one to another. We will experiment with representing images in a few commonly used colour spaces in visual systems during Lab 2.