COMP3421

Colour Theory and Exam Preparation

What is colour?

The experience of colour is complex, involving:

- Physics of light,
  - Electromagnetic radiation
- Biology of the eye,
- Neuropsychology of the visual system.

Physics of light

Light is an electromagnetic wave, the same as radio waves, microwaves, X-rays, etc.

The visible spectrum (for humans) consists of waves with wavelength between 400 and 700 nanometers.

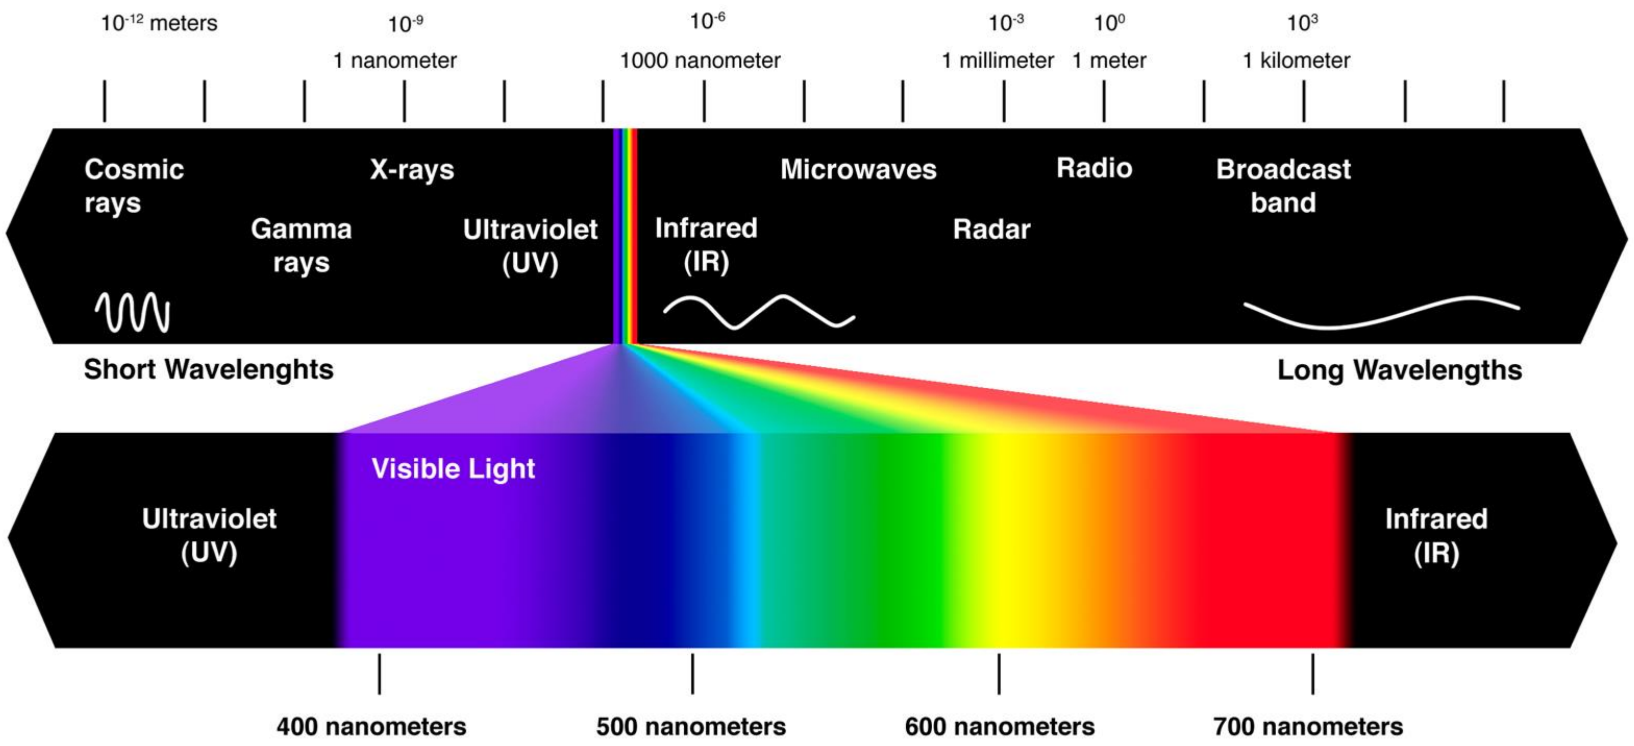

Non-spectral colours

Some light sources, such as lasers, emit light of essentially a single wavelength or "pure spectral" light (red,violet and colors of the rainbow).

Other colours (e.g. white, purple, pink,brown) are non-spectral.

There is no single wavelength for these colours, rather they are mixtures of light of different wavelengths.

Colour perception

The retina (back of the eye) has two different kinds of photoreceptor cells: rods and cones.

Rods are good at handling low-level lighting (e.g. moonlight). They do not detect different colours and are poor at distinguishing detail.

Cones respond better in brighter light levels. They are better at discerning detail and colour.

The Eye

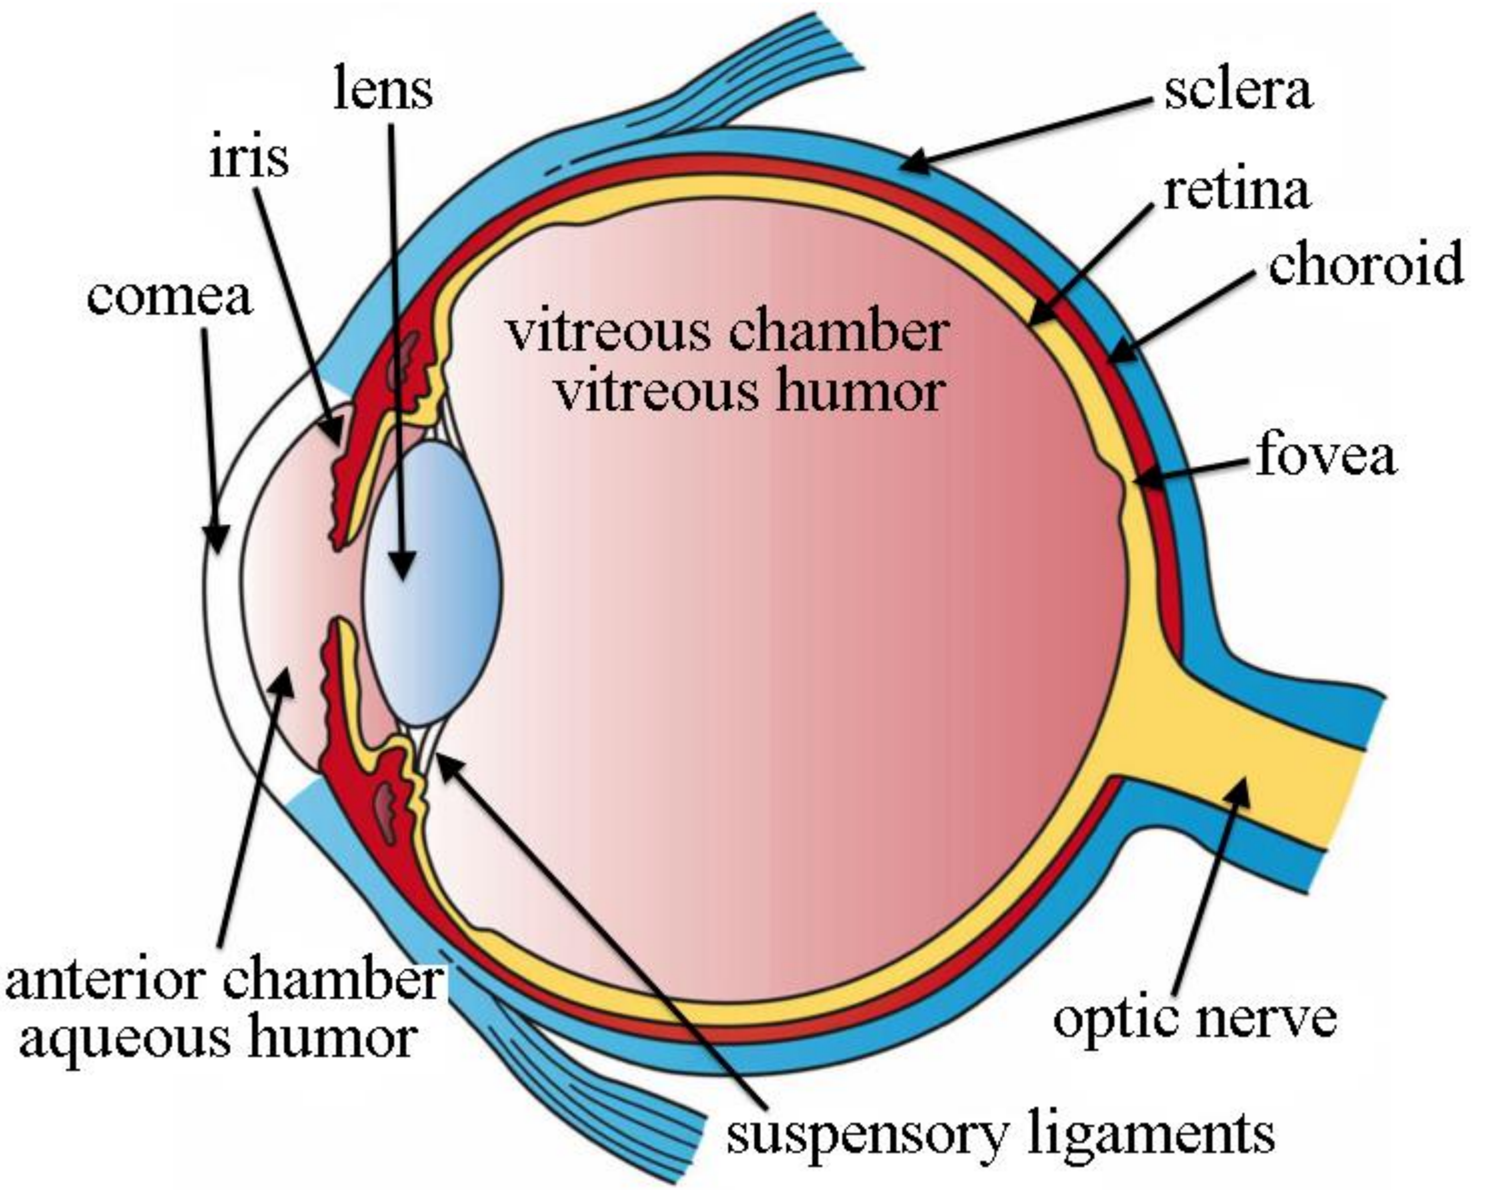

http://open.umich.edu/education/med/resources/second-look-series/materials

Tristimulus Theory

Most people have three different kinds of cones which are sensitive to different wavelengths.

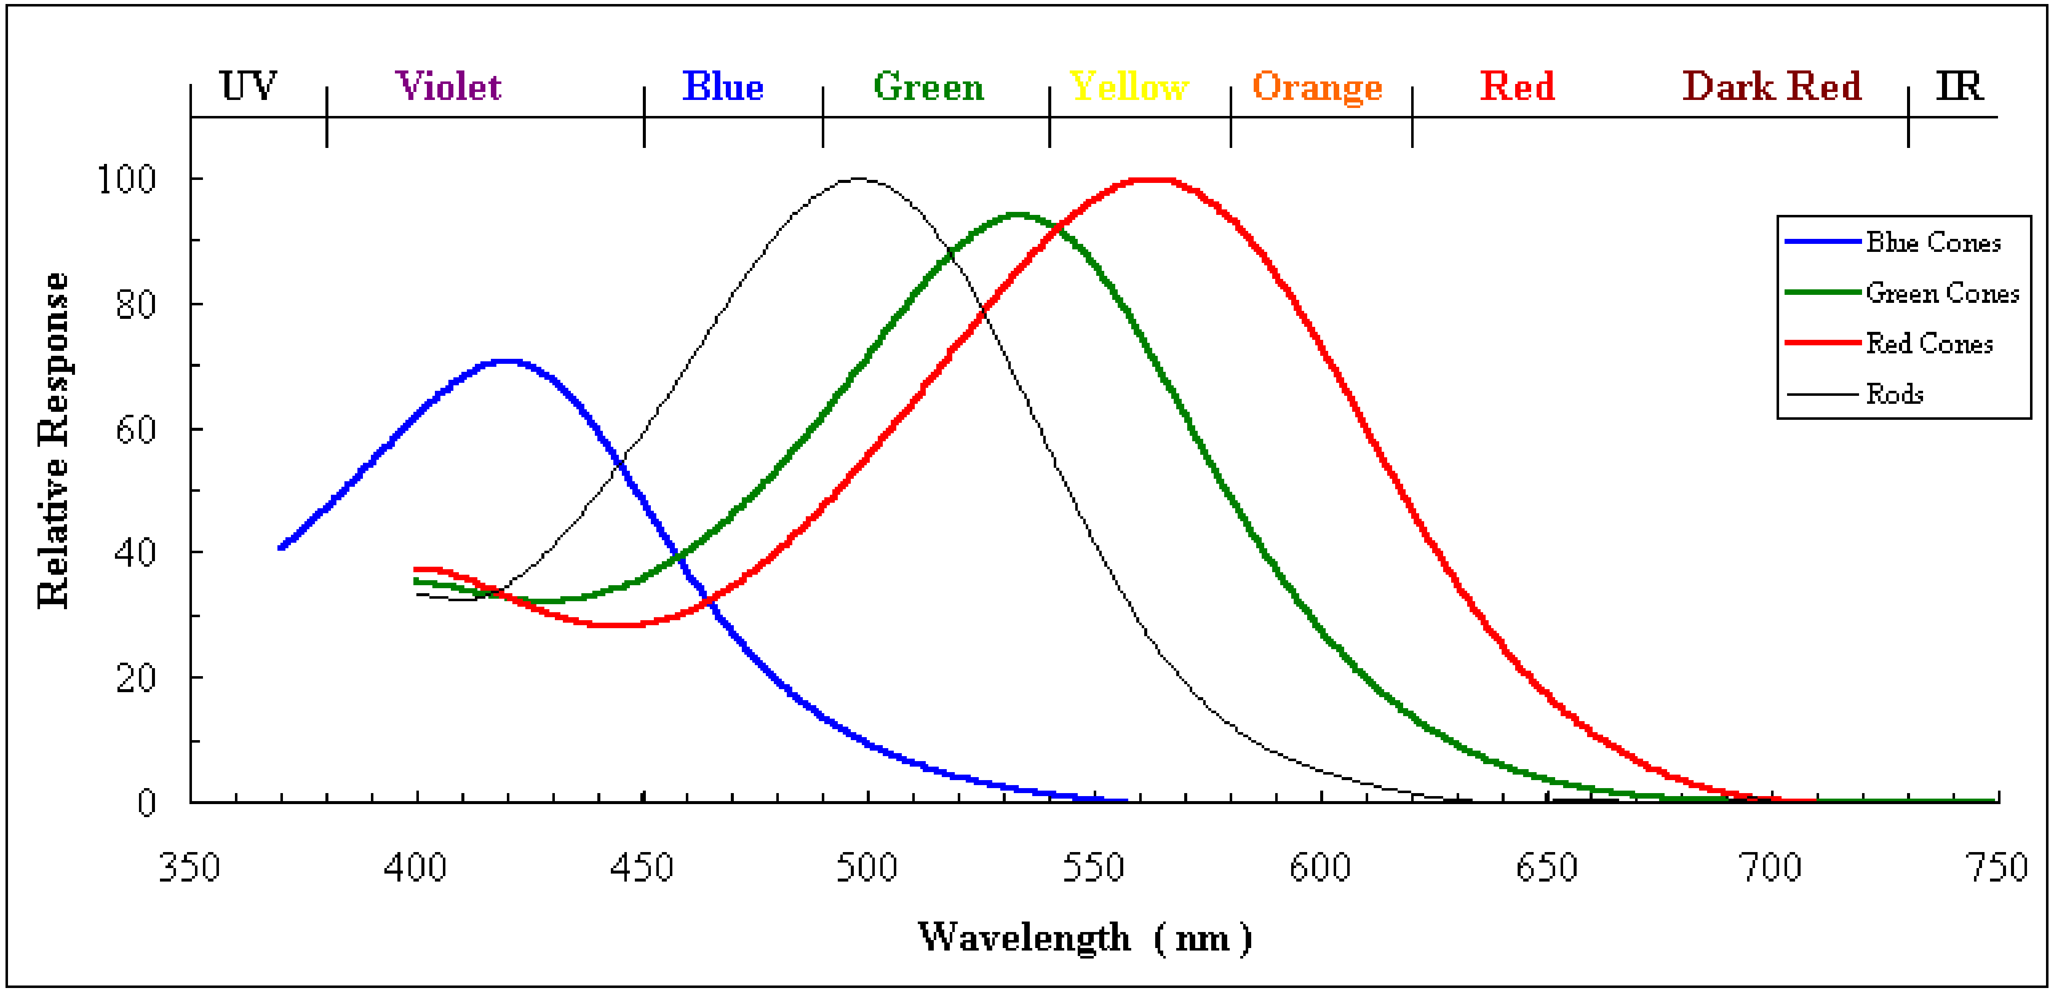

Colour blending

As a result of this, different mixtures of light will appear to have the same colour, because they stimulate the cones in the same way.

For example, a mixture of red and green light will appear to be yellow.

Colour blending

We can take advantage of this in a computer by having monitors with only red, blue and green phosphors in pixels.

Other colours are made by mixing these lights together.

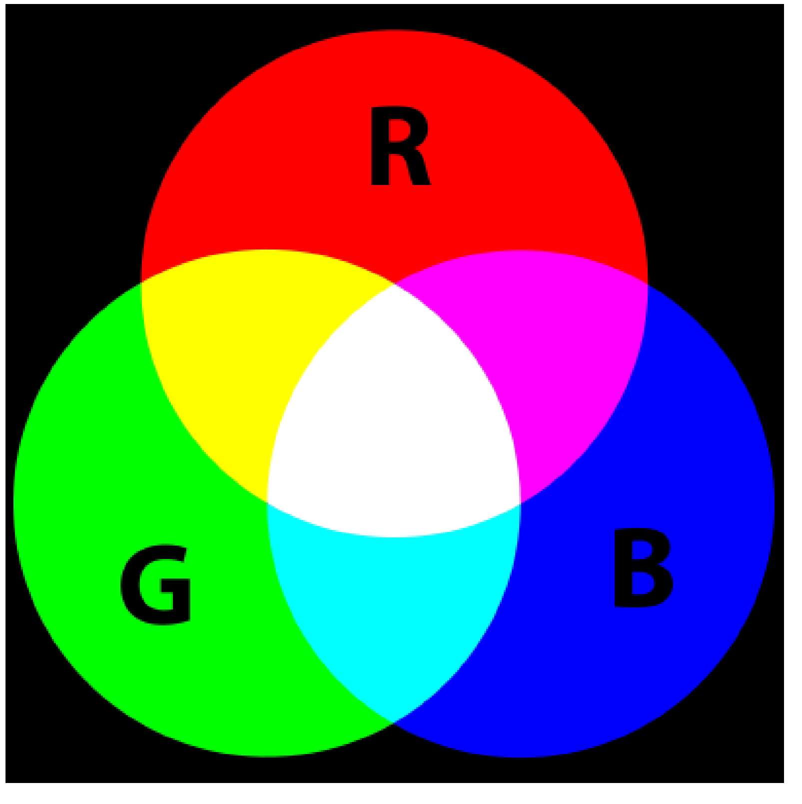

Colour blending

Can we make all colours this way?

No. Some colours require a negative amount of one of the primaries (typically red).

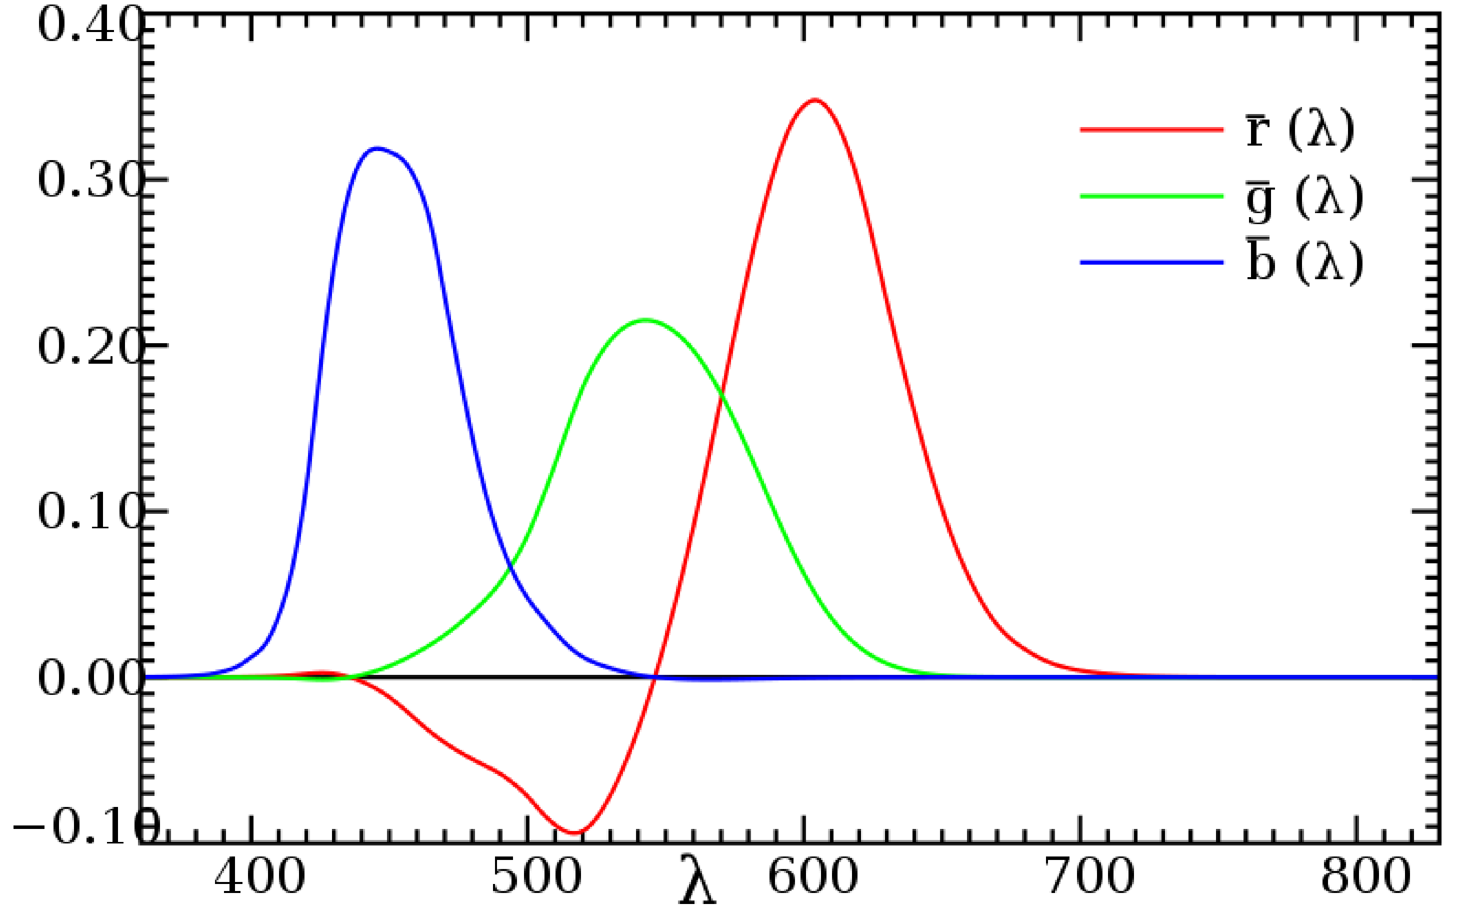

Colour blending

What does this mean?

Algebraically, we write:

$$C = rR + gG + bB$$

to indicate that colour C is equivalent (appears the same as) *r* units of red, *g* units of green and *b* units of blue.

Colour blending

A colour with wavelength 500nm (cyan/teal) has:

$$C = -0.30R + 0.49G + 0.11B$$

We can rearrange this as:

$$C + 0.30R = 0.49G + 0.11B$$

So if we add 0.3 units of red to colour C, it will look the same as the given combination of green and blue.

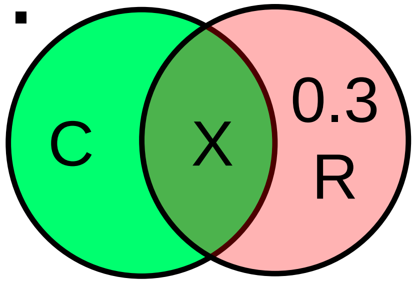

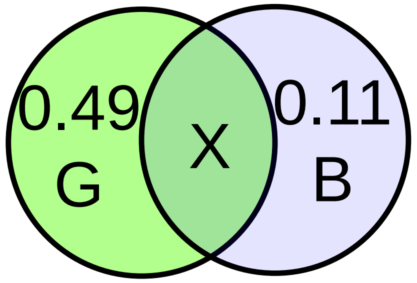

Data source:
http://www.cvrl.org/

Tristimulus Theory and Colour Blending

https://graphics.stanford.edu/courses/cs178/applets/locus.html

https://graphics.stanford.edu/courses/cs178/applets/colormatching.html

Complementary Colors

Colours that add to give white (or at least grey) are called **complementary colours**

eg red and cyan

Retinal fatigue causes complementary colours to be seen in after-images

http://www.animations.physics.unsw.edu.au/jw/light/complementary-colours.htm

Describing colour

We can describe a colour in terms of its:

- Hue - the colour of the dominant wavelength such as red
- Luminance - the total power of the light (related to brightness)
- Saturation - the purity of the light i.e the percentage of the luminance given by the dominant hue (the more grey it is the more unsaturated)

Spectral Density

Hue is the peak or dominant wavelength

Luminance is related to the intensity or area under the entire spectrum

Saturation is the percentage of intensity in the dominant area

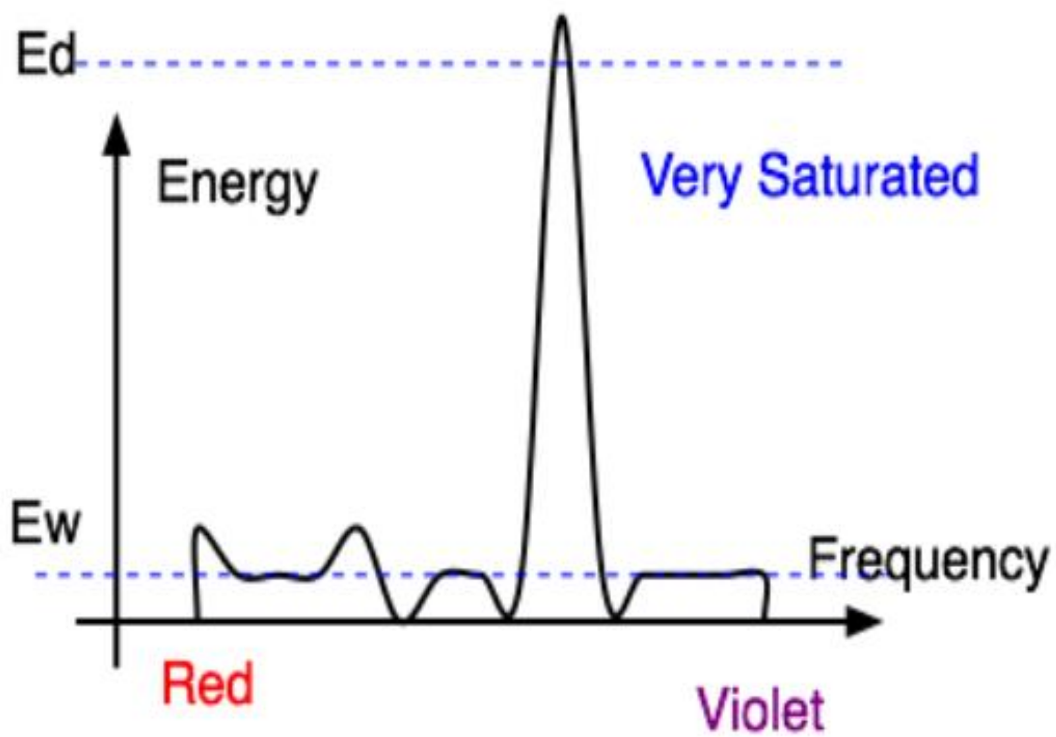

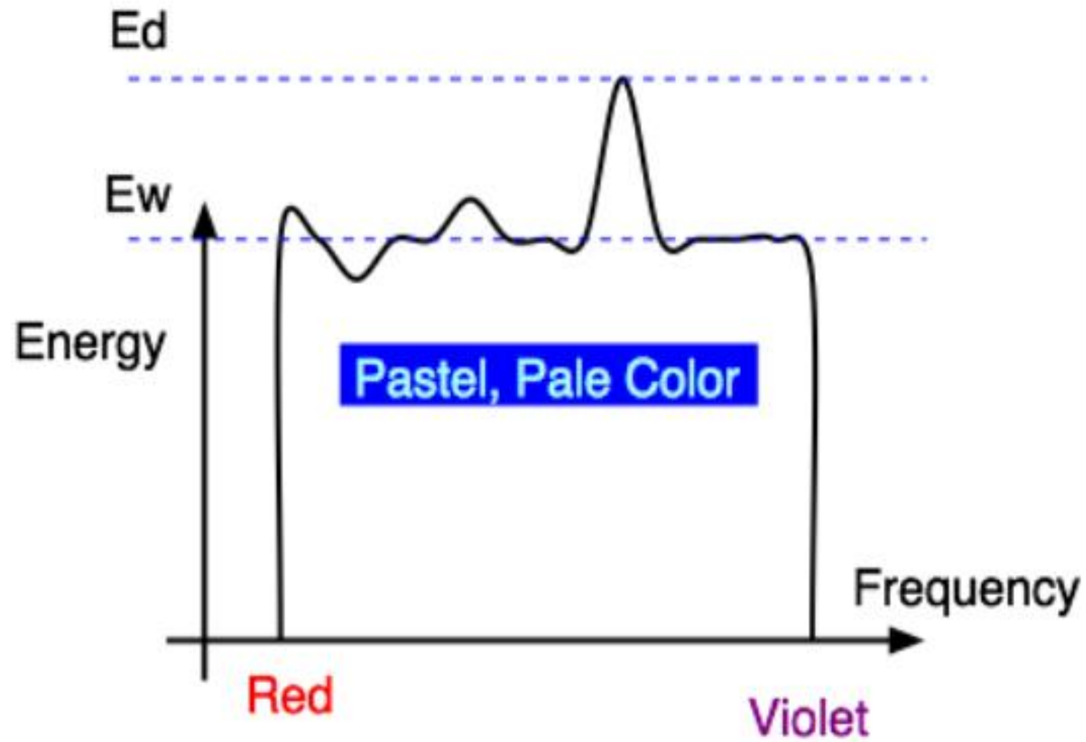

Physics vs Perception

We need to be careful with our language. Physical and perceptual descriptions of light differ.

A red light and a blue light of the same physical intensity will not have the same perceived brightness (the red will appear brighter).

Intensity, Power = physical properties
Luminance, Brightness = perceptual properties

Describing colour

“Computer science offers a few poorer cousins to these perceptual spaces that may also turn up in your software interface, such as HSV and HLS. They are easy mathematical transformations of RGB, and they seem to be perceptual systems because they make use of the hue-lightness/value-saturation terminology. But take a close look; don’t be fooled. Perceptual color dimensions are poorly scaled by the color specifications that are provided in these and some other systems. For example, saturation and lightness are confounded, so a saturation scale may also contain a wide range of lightnesses (for example, it may progress from white to green which is a combination of both lightness and saturation). Likewise, hue and lightness are confounded so, for example, a saturated yellow and saturated blue may be designated as the same ‘lightness’ but have wide differences in perceived lightness. These flaws make the systems difficult to use to control the look of a color scheme in a systematic manner. If much tweaking is required to achieve the desired effect, the system offers little benefit over grappling with raw specifications in RGB or CMY. “

http://www.personal.psu.edu/cab38/ColorSch/ASApaper.html

Standardisation

A problem with describing colours as RGB values is that depends on what wavelengths we define as red, green and blue.

Different displays emit different frequencies, which means the same RGB value will result in slightly different colours.

We need a standard that is independent of the particular display.

The CIE standard

The CIE standard, also known as the XYZ model, is a way of describing colours as a three dimensional vector (x,y,z) with:

$$0 \leq x, y, z \leq 1$$

$$x + y + z = 1$$

X, Y and Z are called imaginary colours.
It is impossible to create pure X, it is just a useful mathematical representation.

The CIE standard

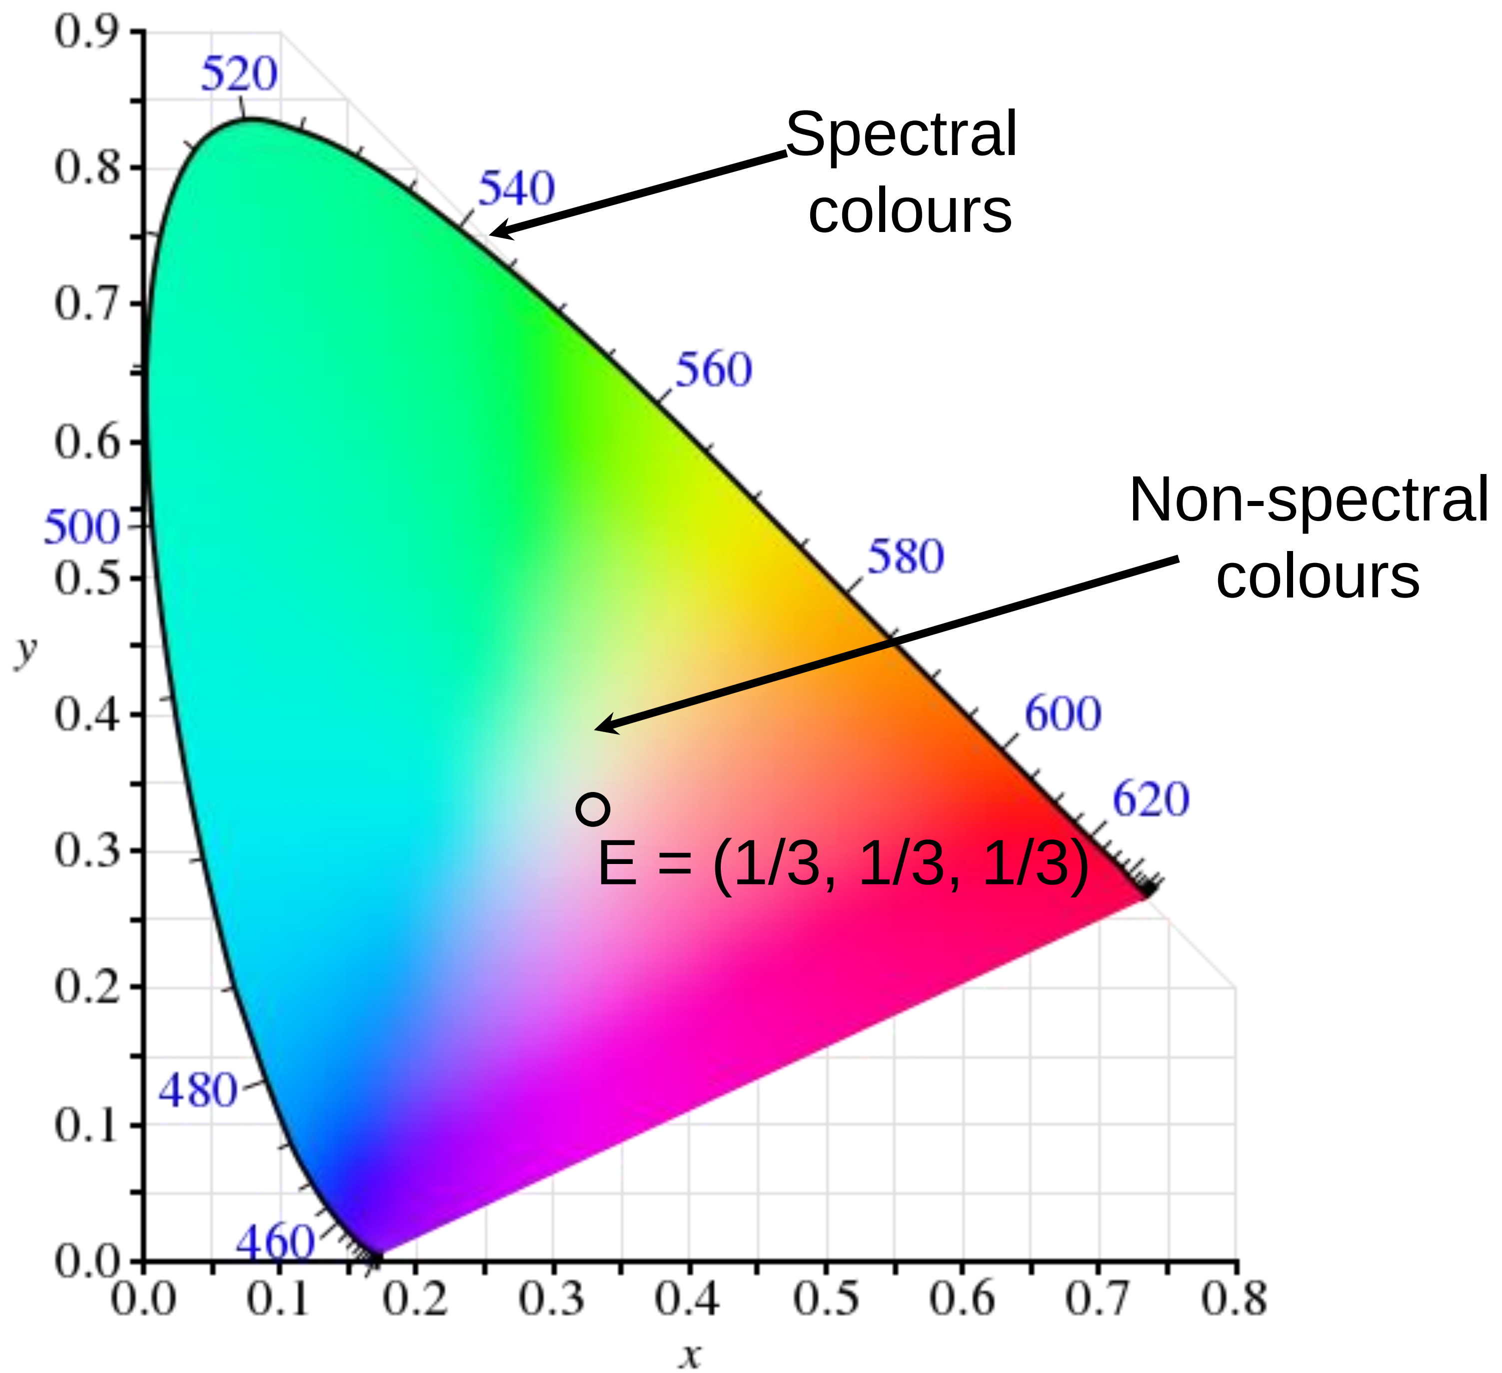

Gamut

Any output device has a certain range of colours it can represent, which we call its gamut.

We can depict this as an area on the CIE chart.

If a monitor has red, green and blue phosphors then the gamut is the interior of the triangle joining those points.

RGB Gamut

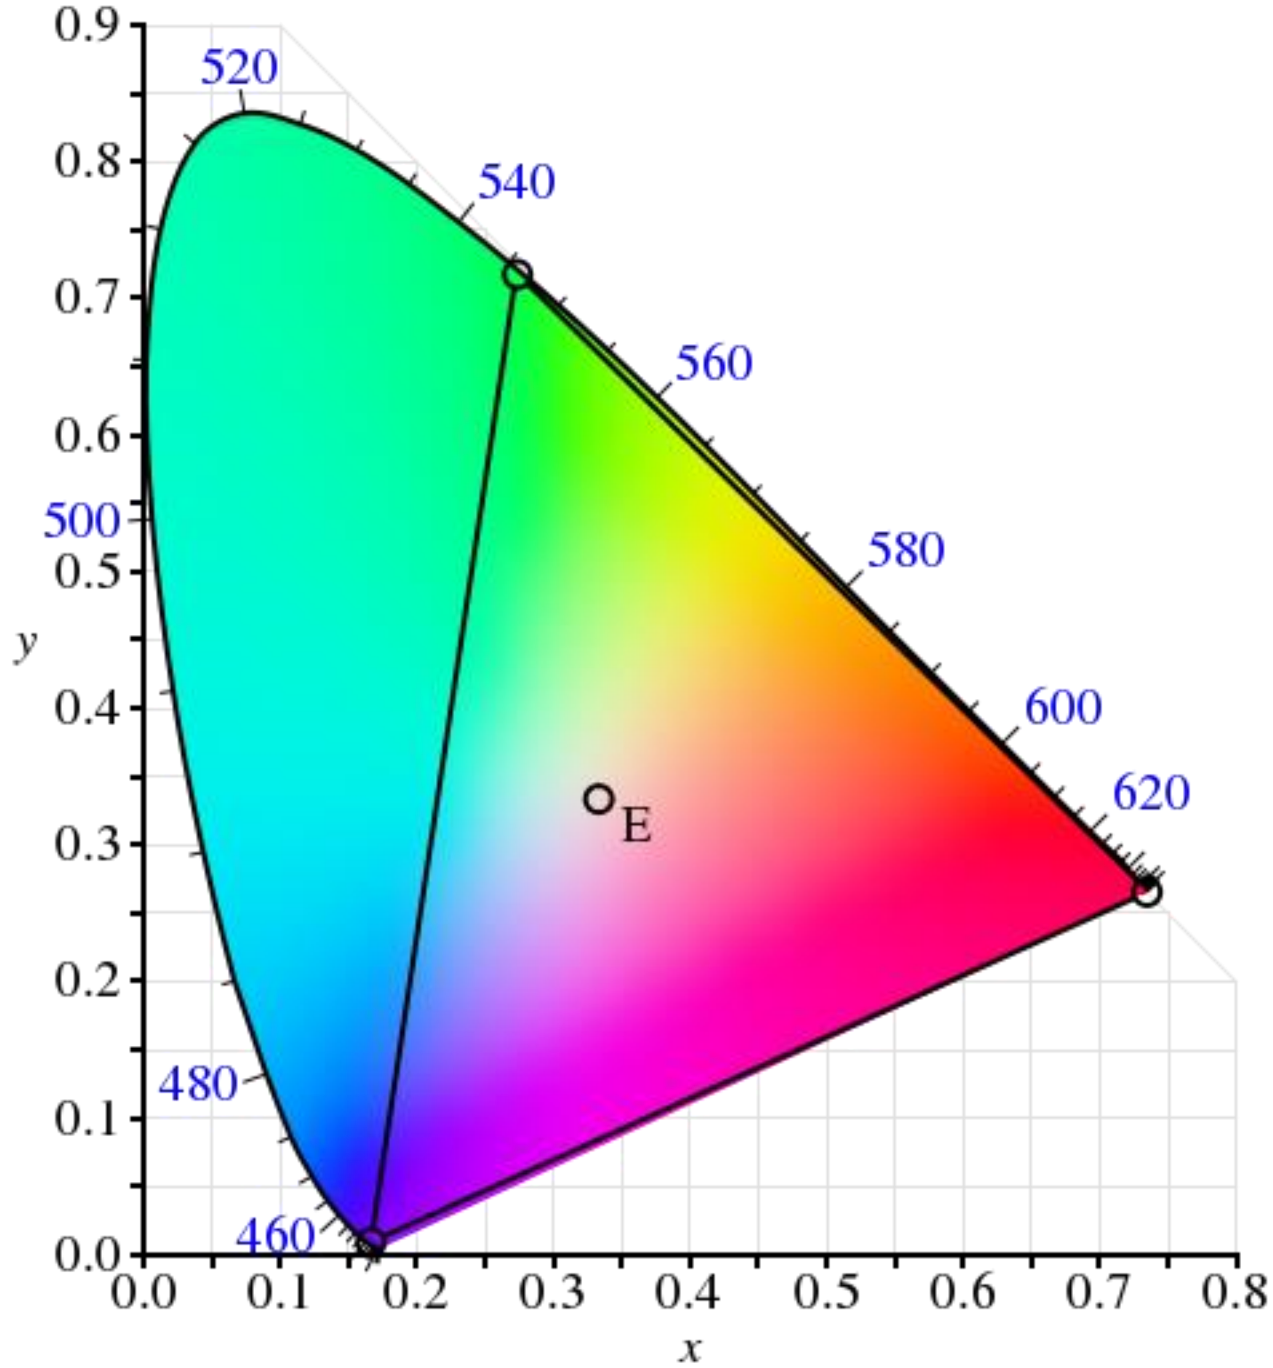

0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0.0
y
520
540
560
580
600
620
500
480
460
E
0.0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
x

RGB Gamut

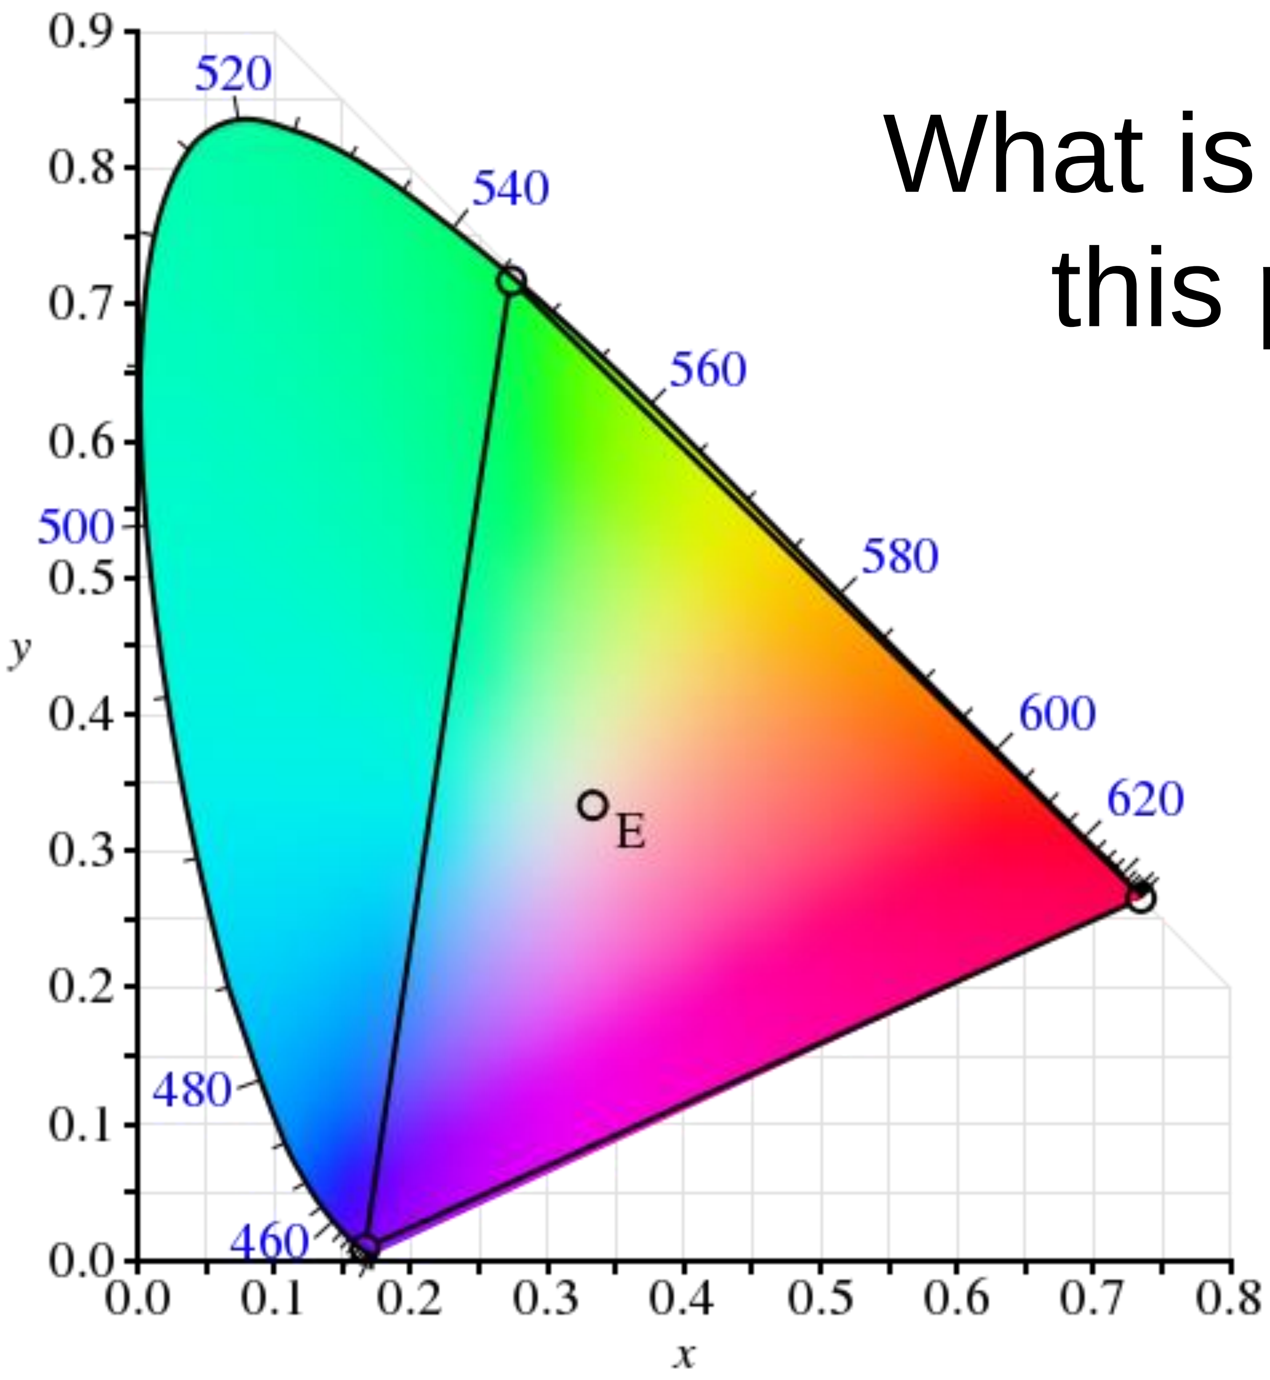

What is wrong with this picture?

Gamut mapping

How do we map an (x, y, z) colour from outside the gamut to a colour we can display?

We want to maintain:

- Approximately the same hue
- Relative saturation to other colours in the image.

Rendering intents

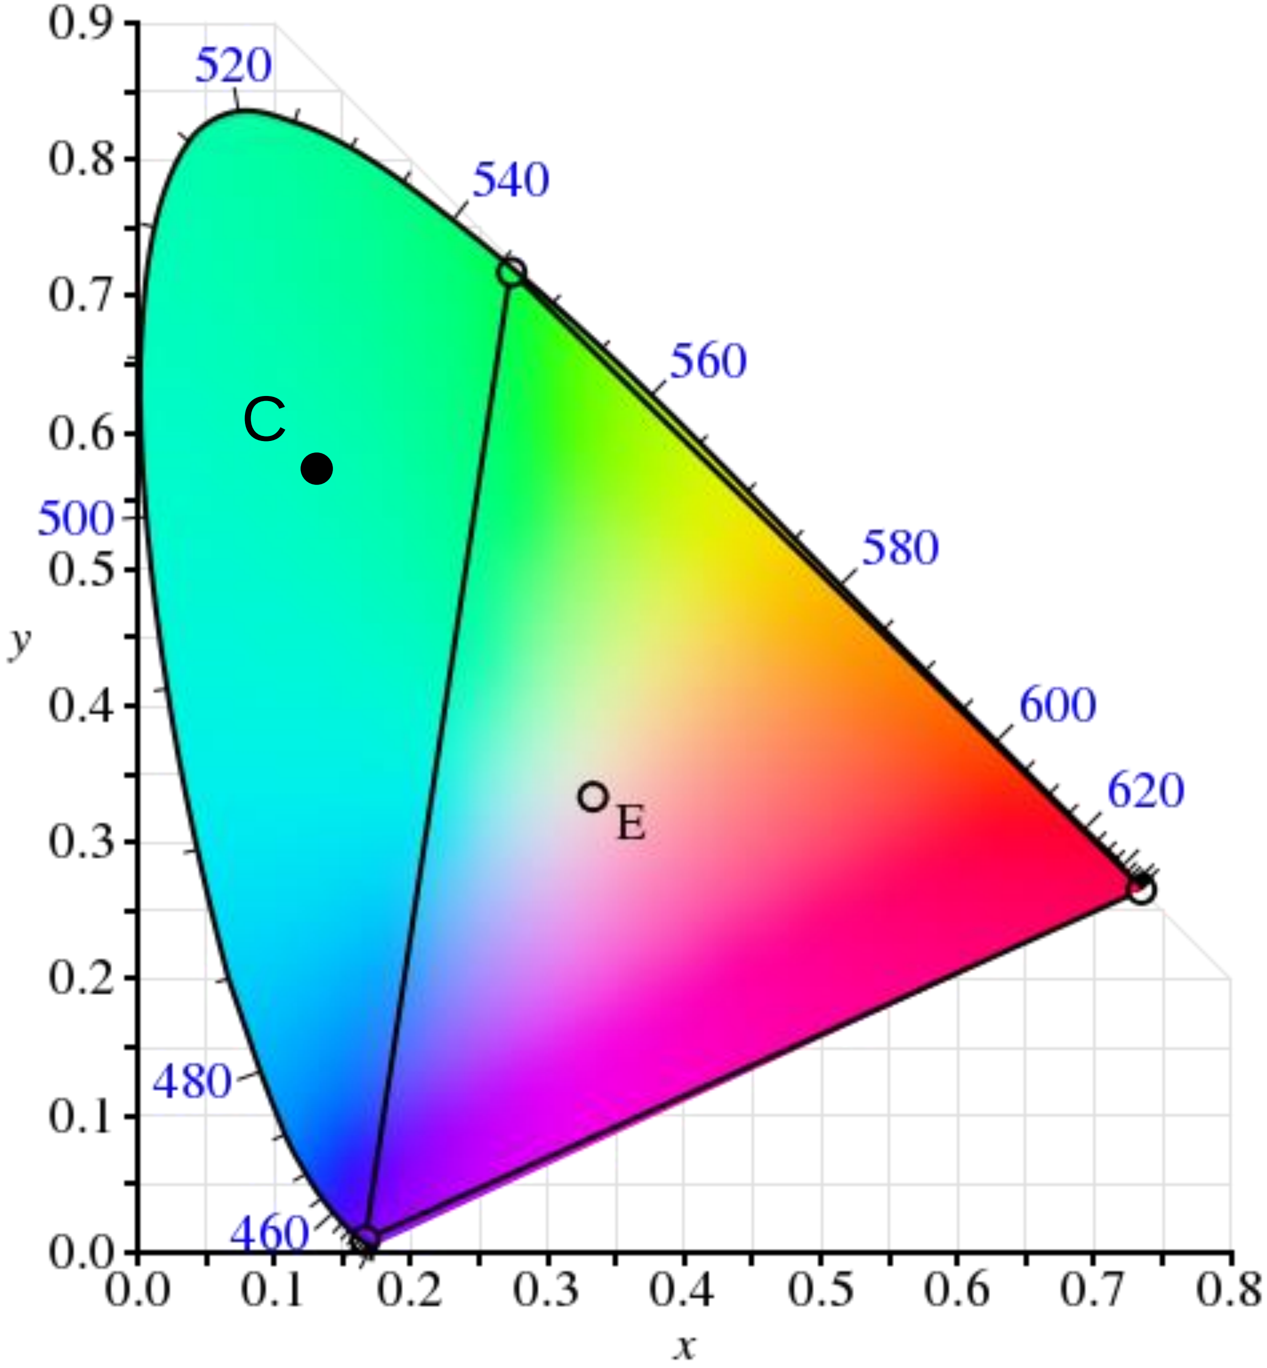

There are four standard rendering intents which describe approaches to gamut mapping.

The definitions are informal.

Implementations vary.

Absolute colormetric

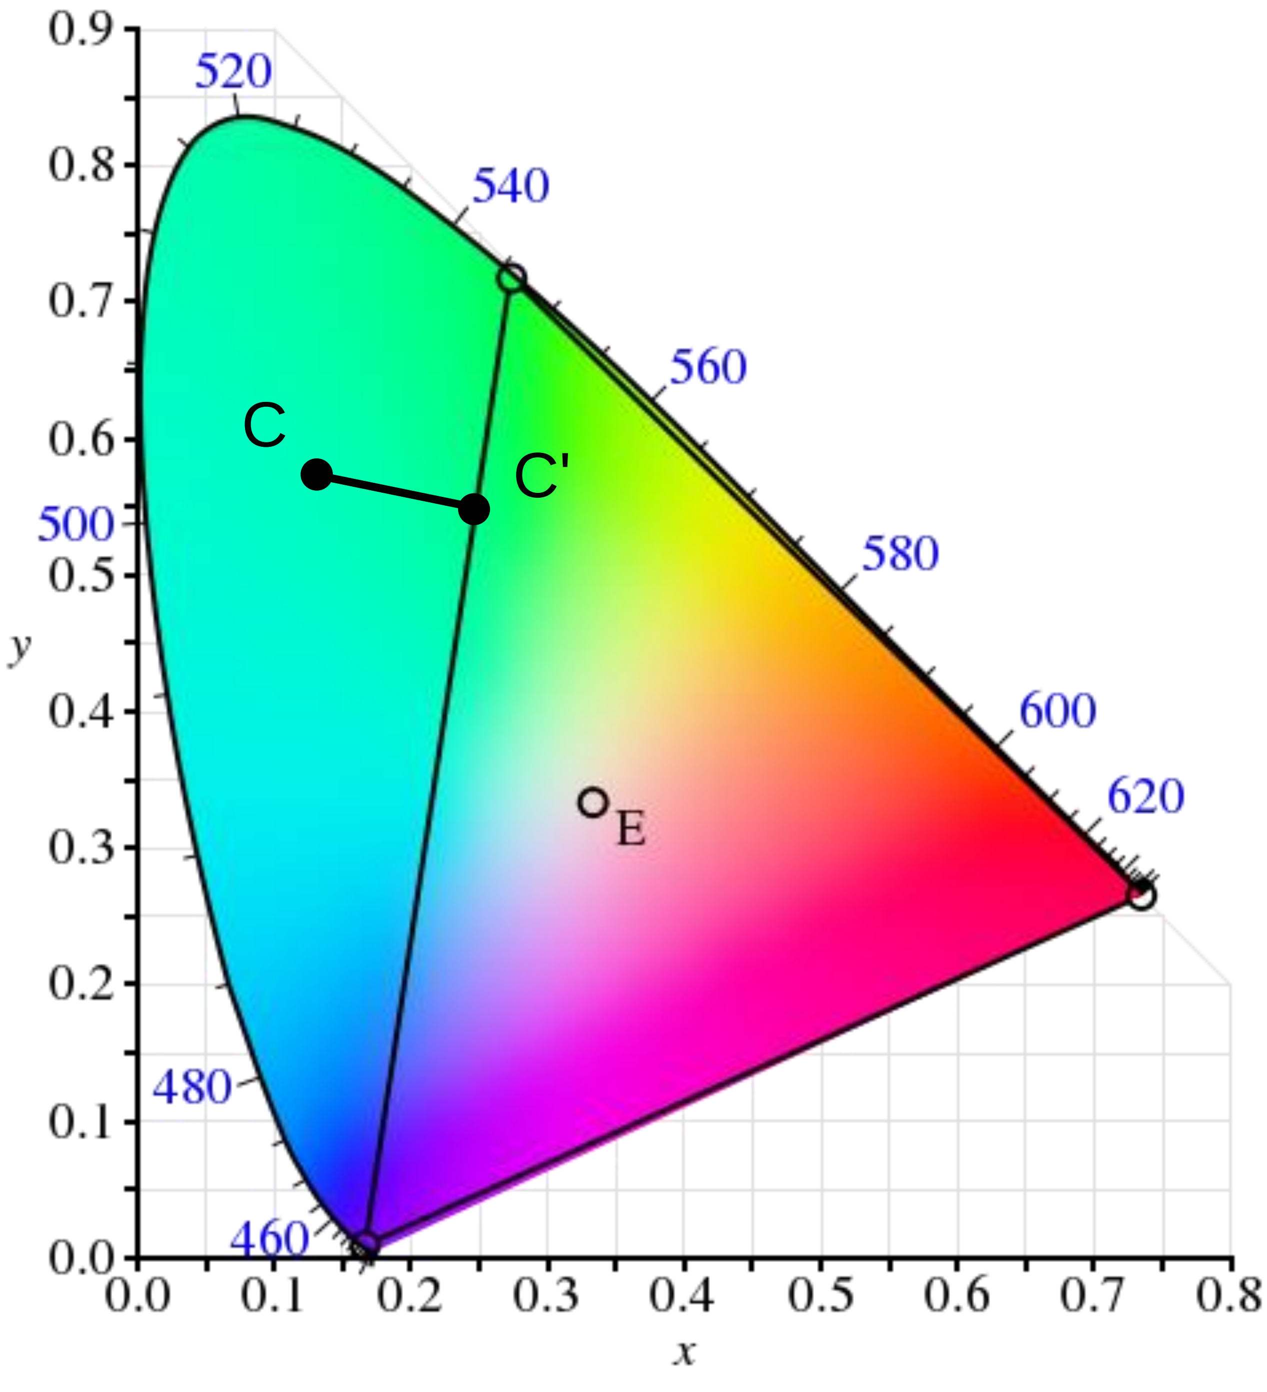

Map C to the nearest point within the gamut.

Distorts hues.

Does not preserve relative saturation.

Relative colormetric

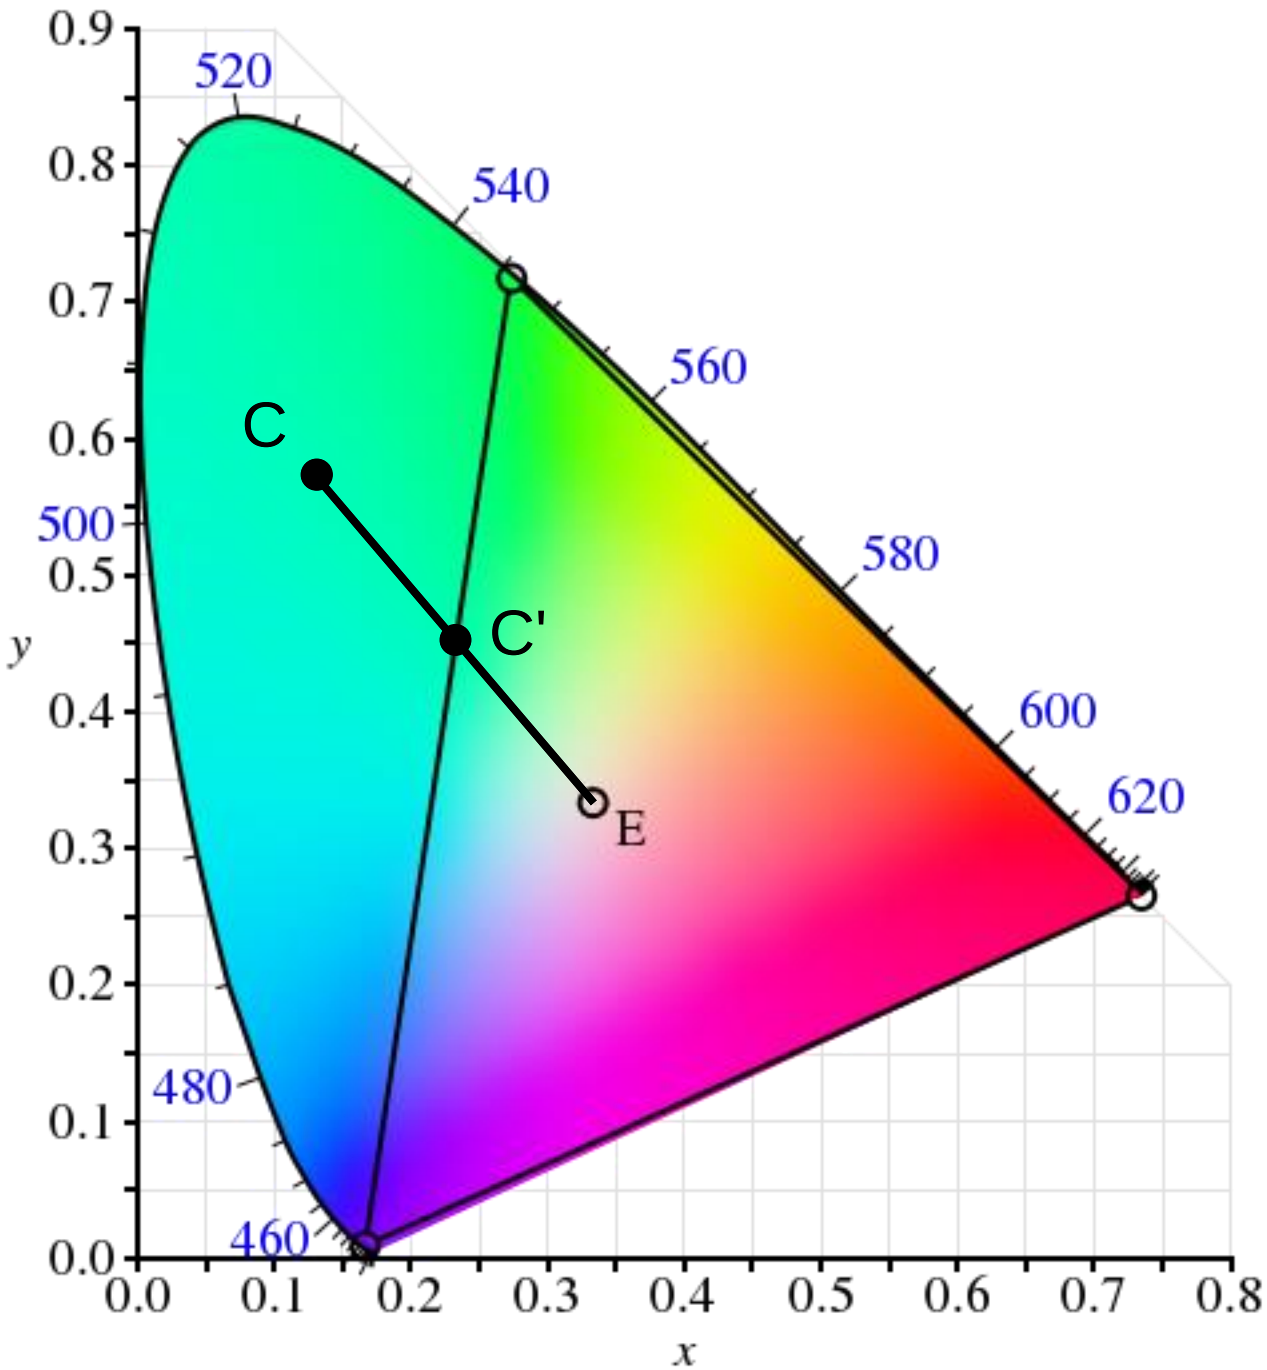

Desaturate C until it lies in the gamut.

Maintains hues more closely.

Does not preserve relative saturation.

Perceptual

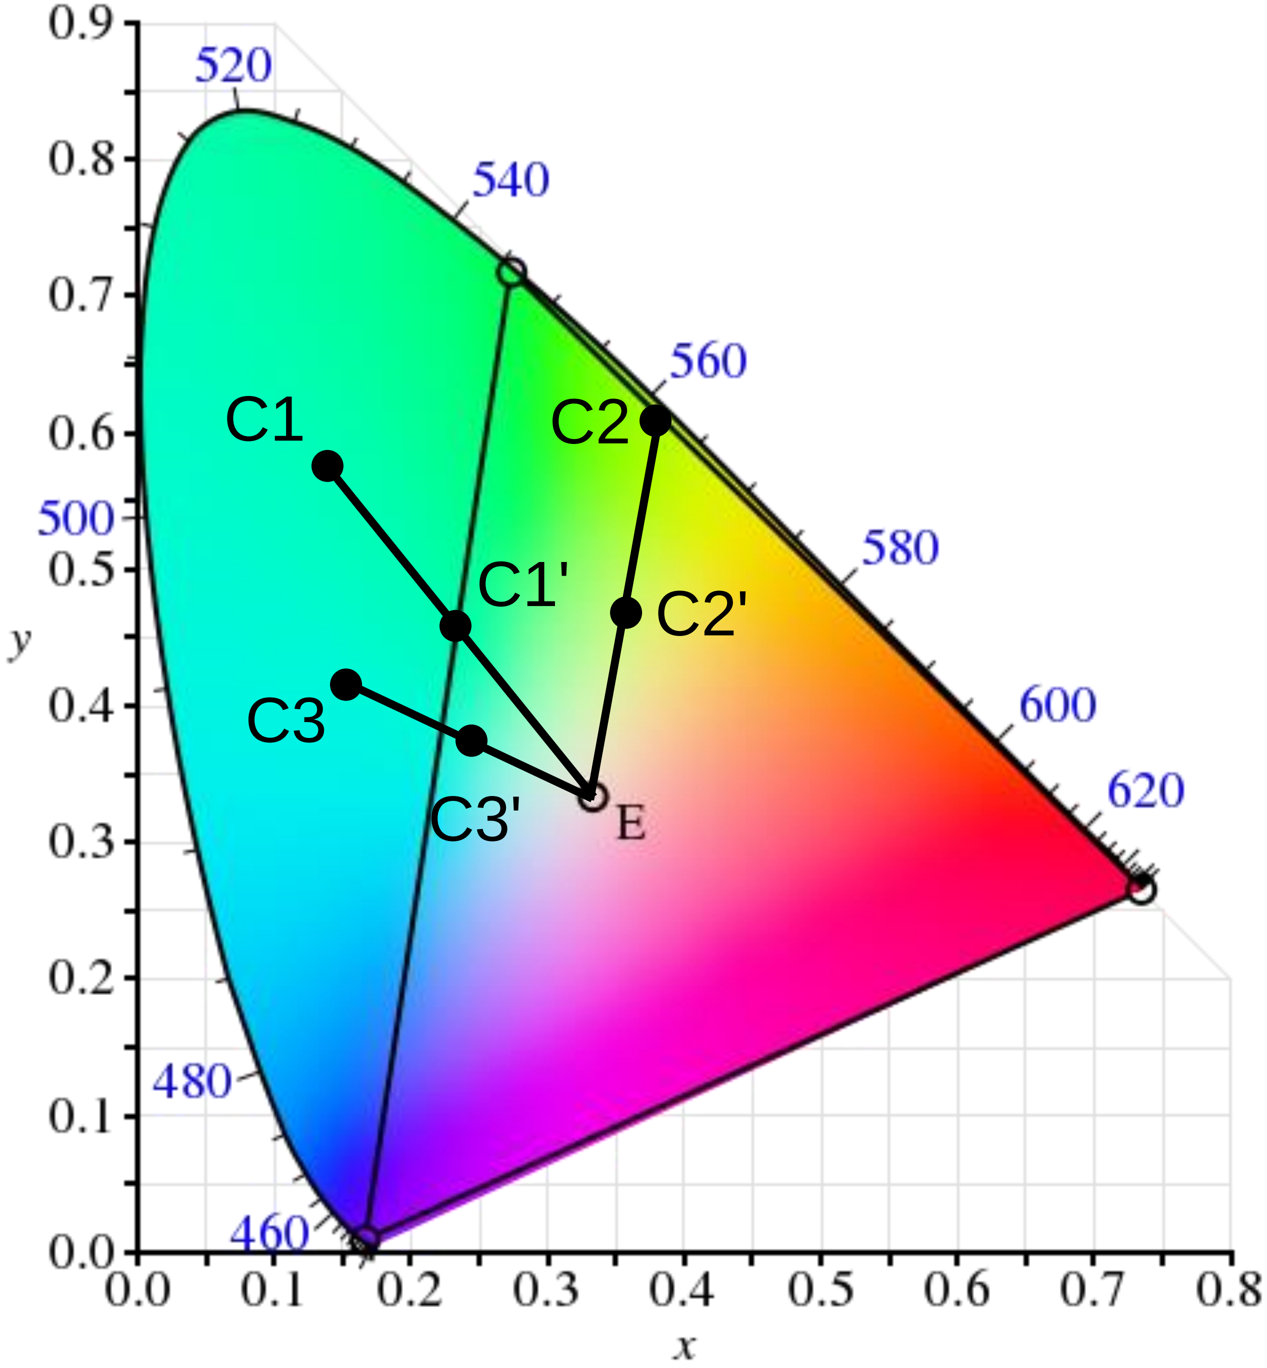

Desaturate all colours until they all lie in the gamut.

Maintains hues.

Preserves relative saturation.

Removes a lot of saturation.

Saturation

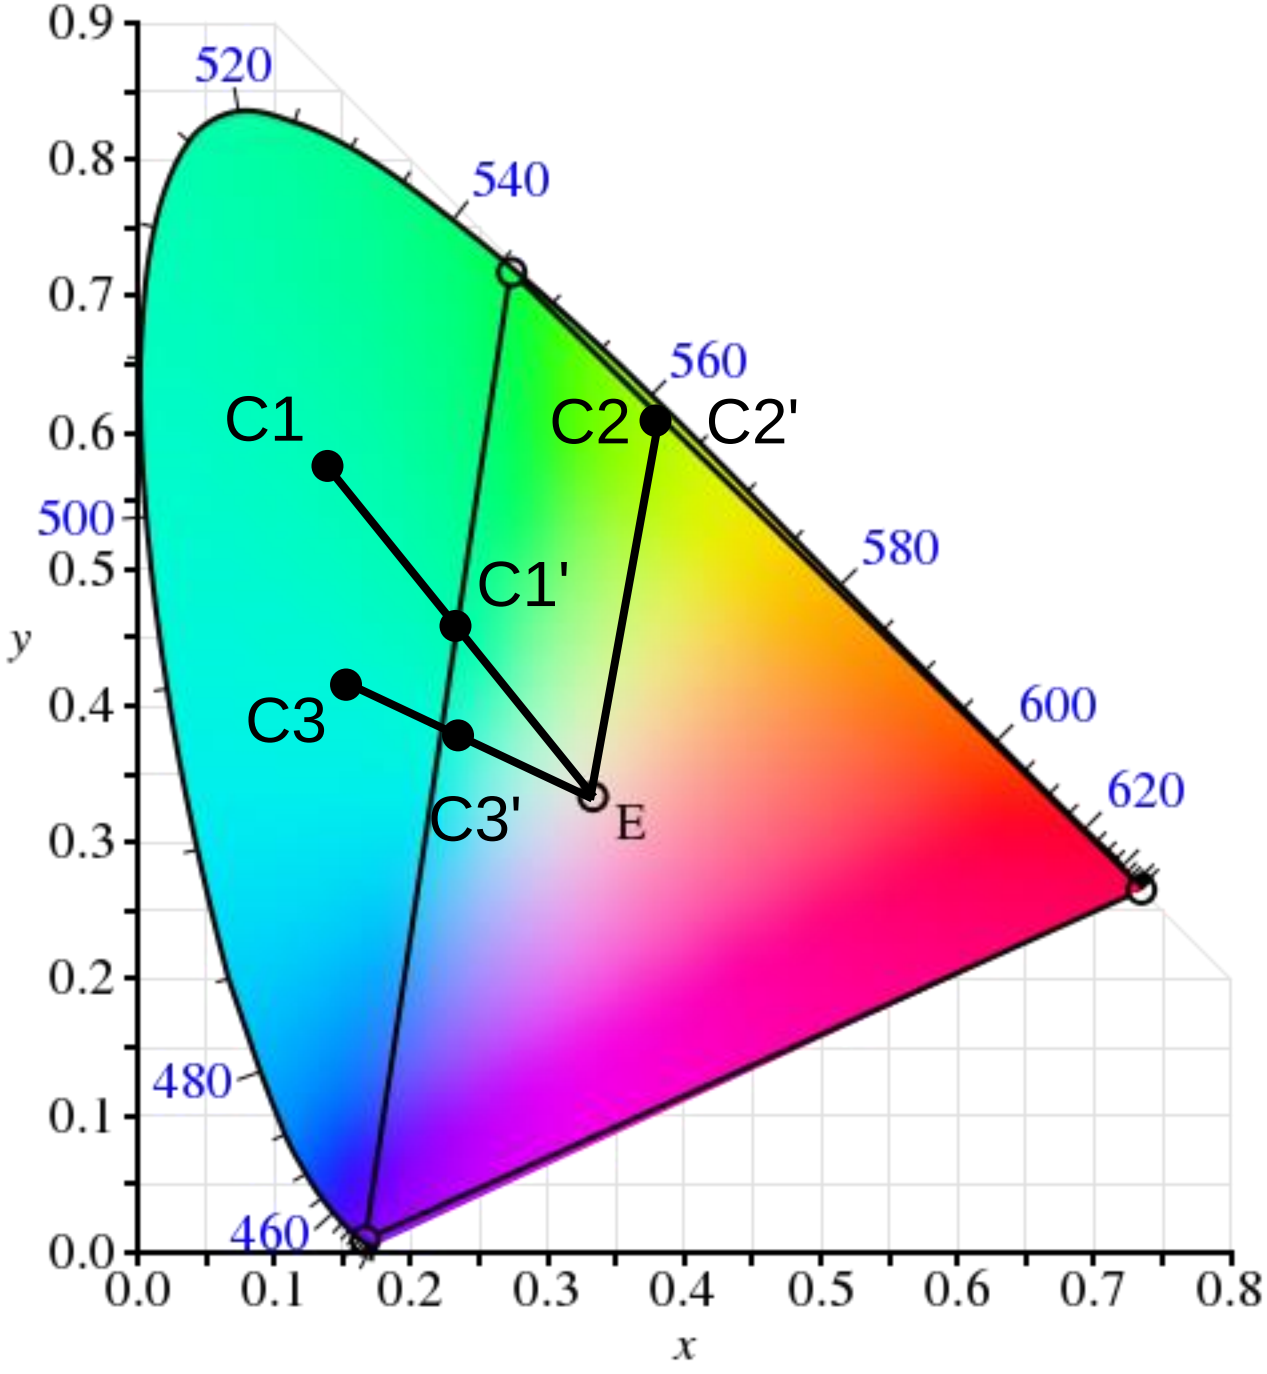

Attempt to maintain saturated colours.

There appears to be no standard algorithmic implementation.

Demo

http://graphics.stanford.edu/courses/cs178/applets/gamutmapping.html

Colour space

Standard colour representations:

- RGB = Red, Green, Blue
- CMYK = Cyan, Magenta, Yellow, Black
- HSV = Hue, Saturation, Value (Brightness)
- HSL = Hue, Saturation, Lightness

RGB

Colour is expressed as the addition of red, green and blue components.

$$C(r,g,b) = rR + gG + bB$$

This is called additive colour mixing. It is the most common model for computer displays.

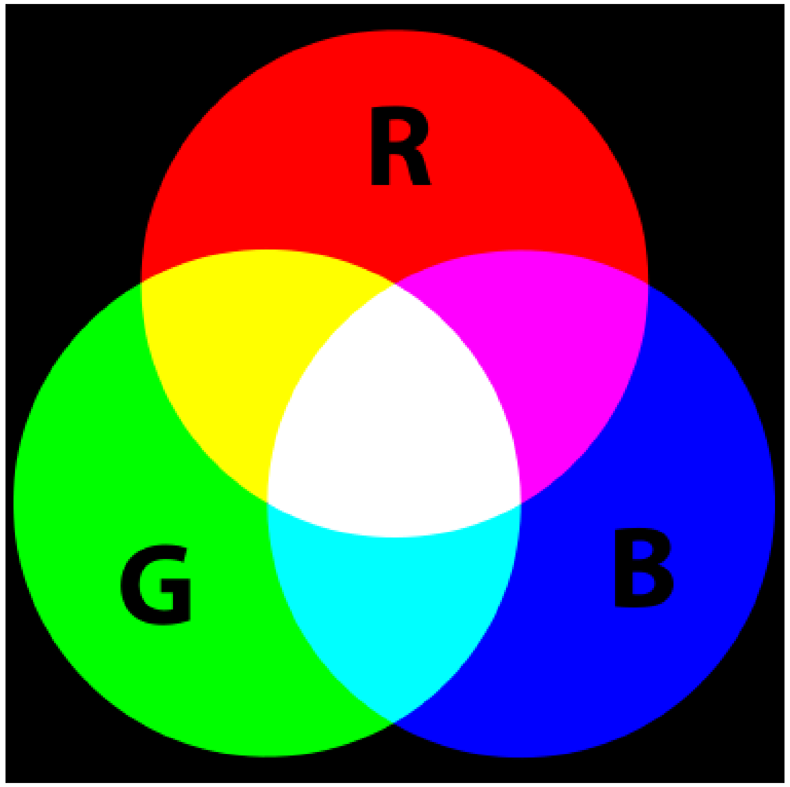

CMY

CMY is a subtractive colour model, typically used in describing printed media.

Cyan, magenta and yellow are the contrasting colours to red, green and blue respectively. I.e.:

$$Cyan = White - Red$$

Cyan pigment/ink absorbs red light.

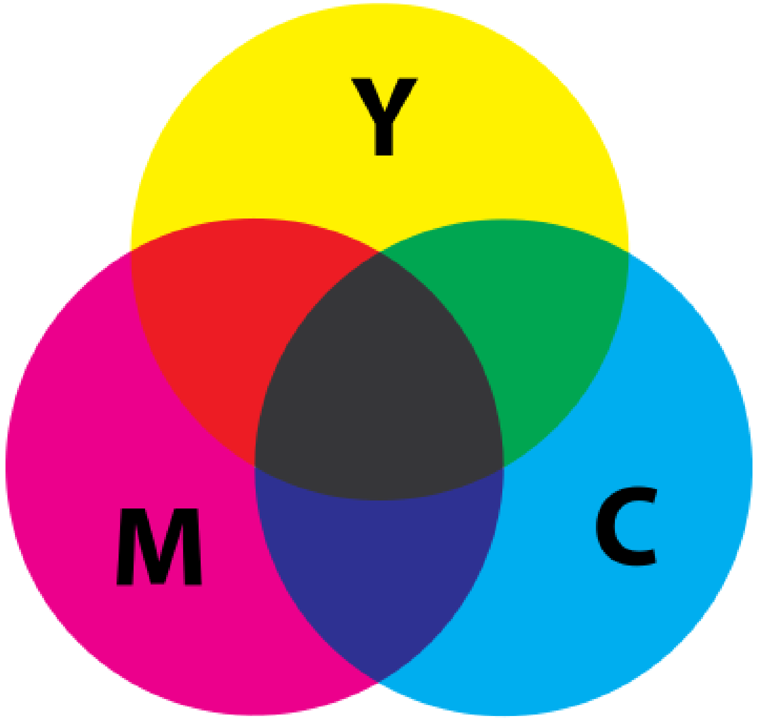

CMYK

Real coloured inks do not absorb light perfectly, so darker colours are achieved by adding black ink to lower the overall brightness.

The K in CMYK stands for "key" and refers to black ink.

HSV

HSV (aka HSB) is an attempt to describe colours in terms that have more perceptual meaning (but see earlier proviso).

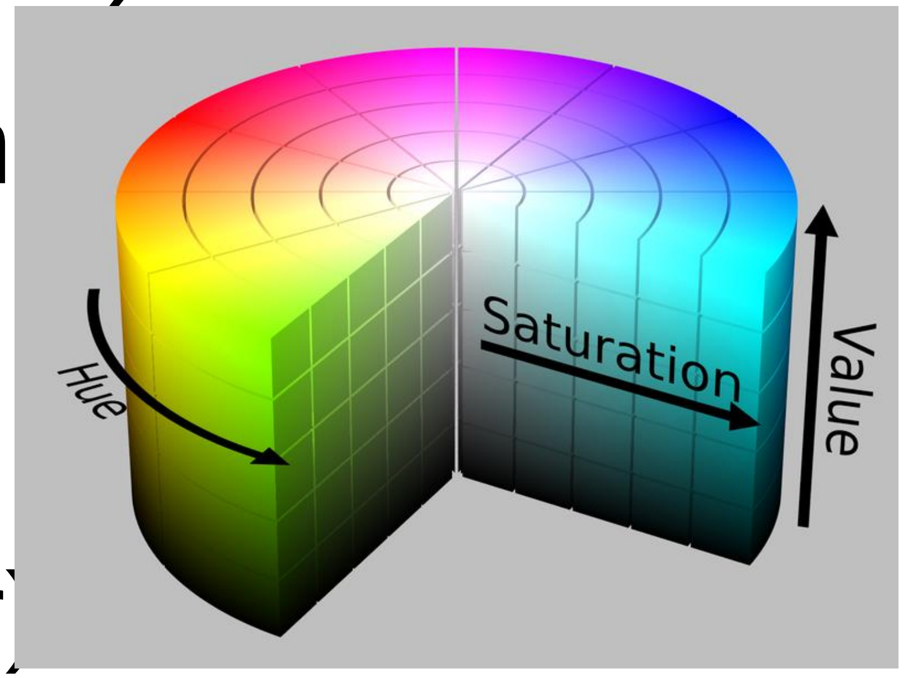

H represents the hue as an an
from 0° (red) to 360° (red)

S represents the saturation
from 0 (grey) to 1 (full colour)

V represents the value/brightness
form 0 (black) to 1 (bright colour).

HSL

HSL (aka HLS) replaces the brightness parameter with a (perhaps) more intuitive lightness value.

H represents the hue as an angle
  from 0° (red) to 360° (red)

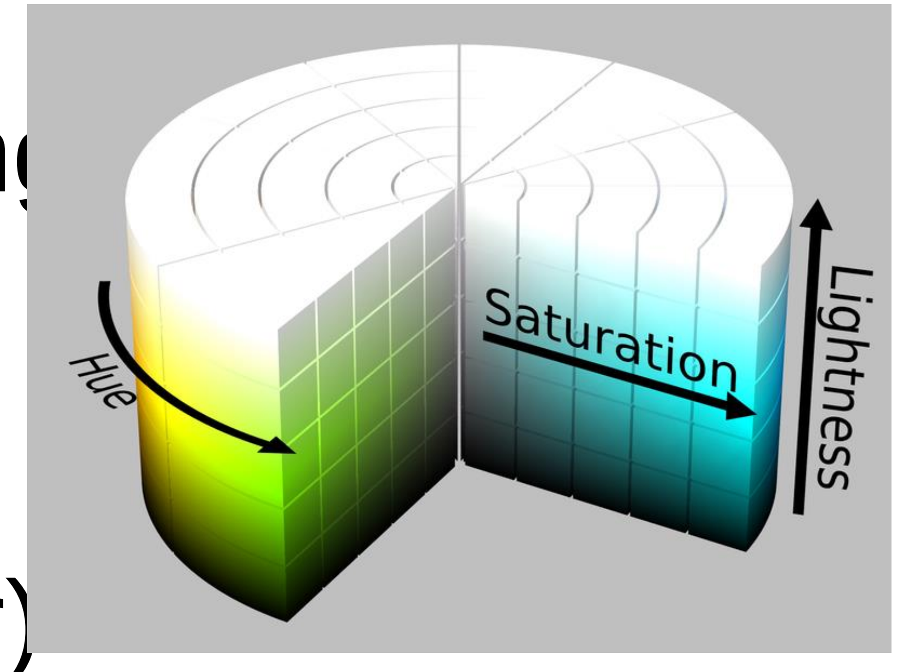

S represents the saturation
  from 0 (grey) to 1 (full colour)

L represents the
  lightness form 0 (black) to 1 (white).

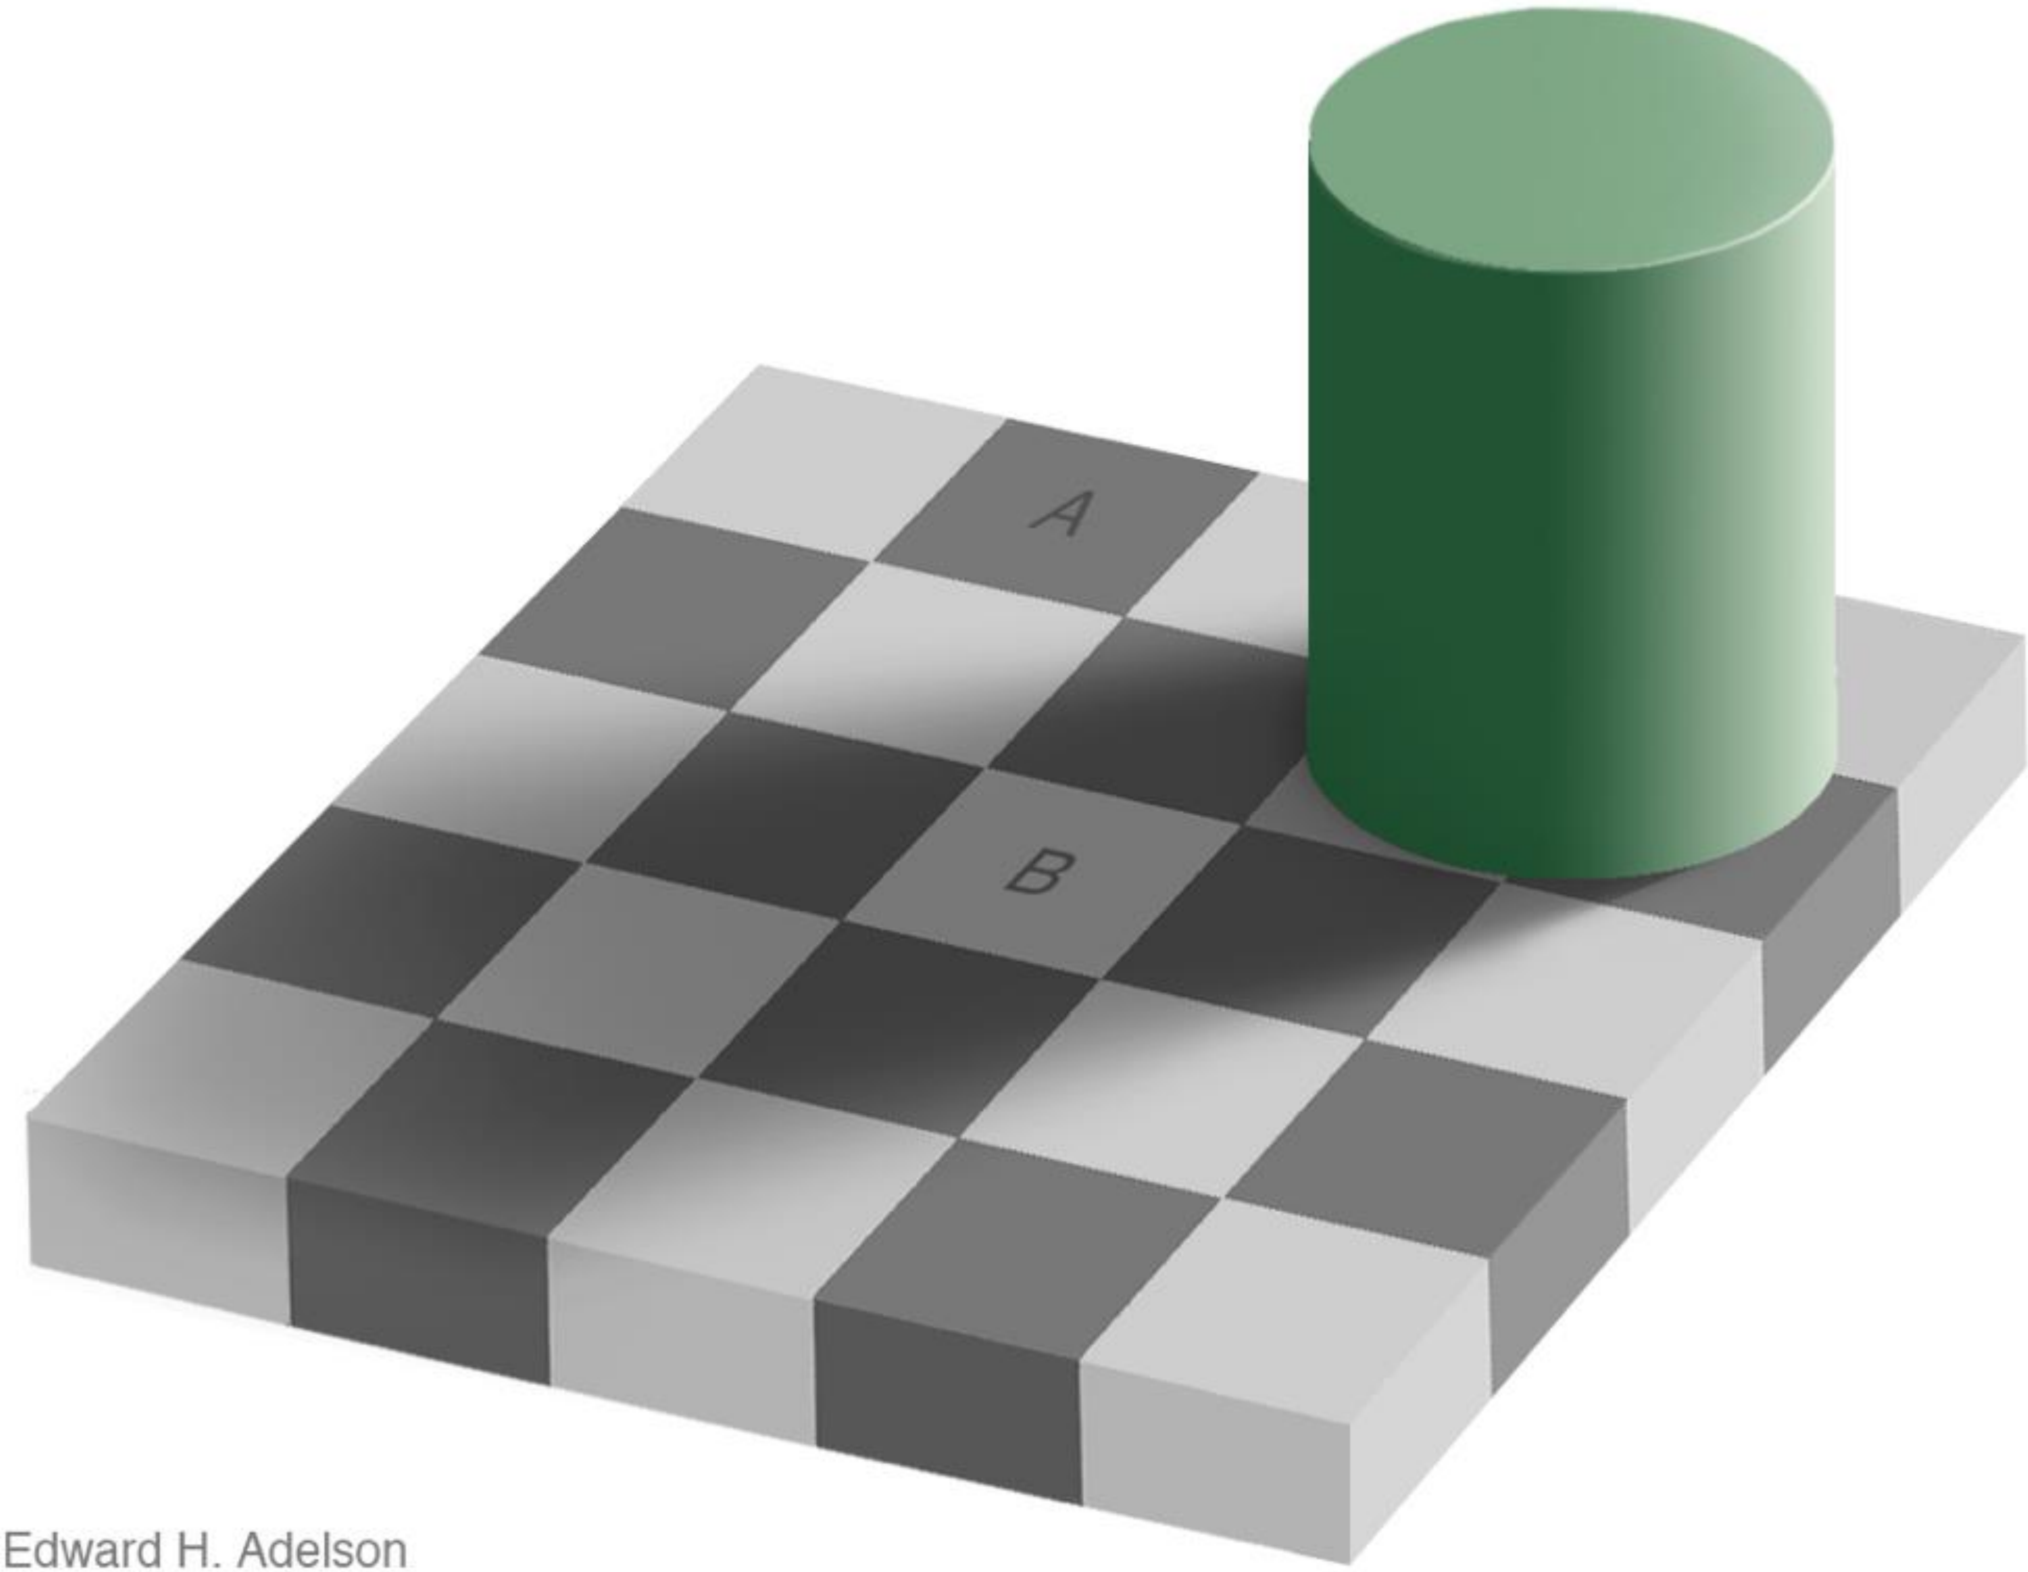

Edward H. Adelson

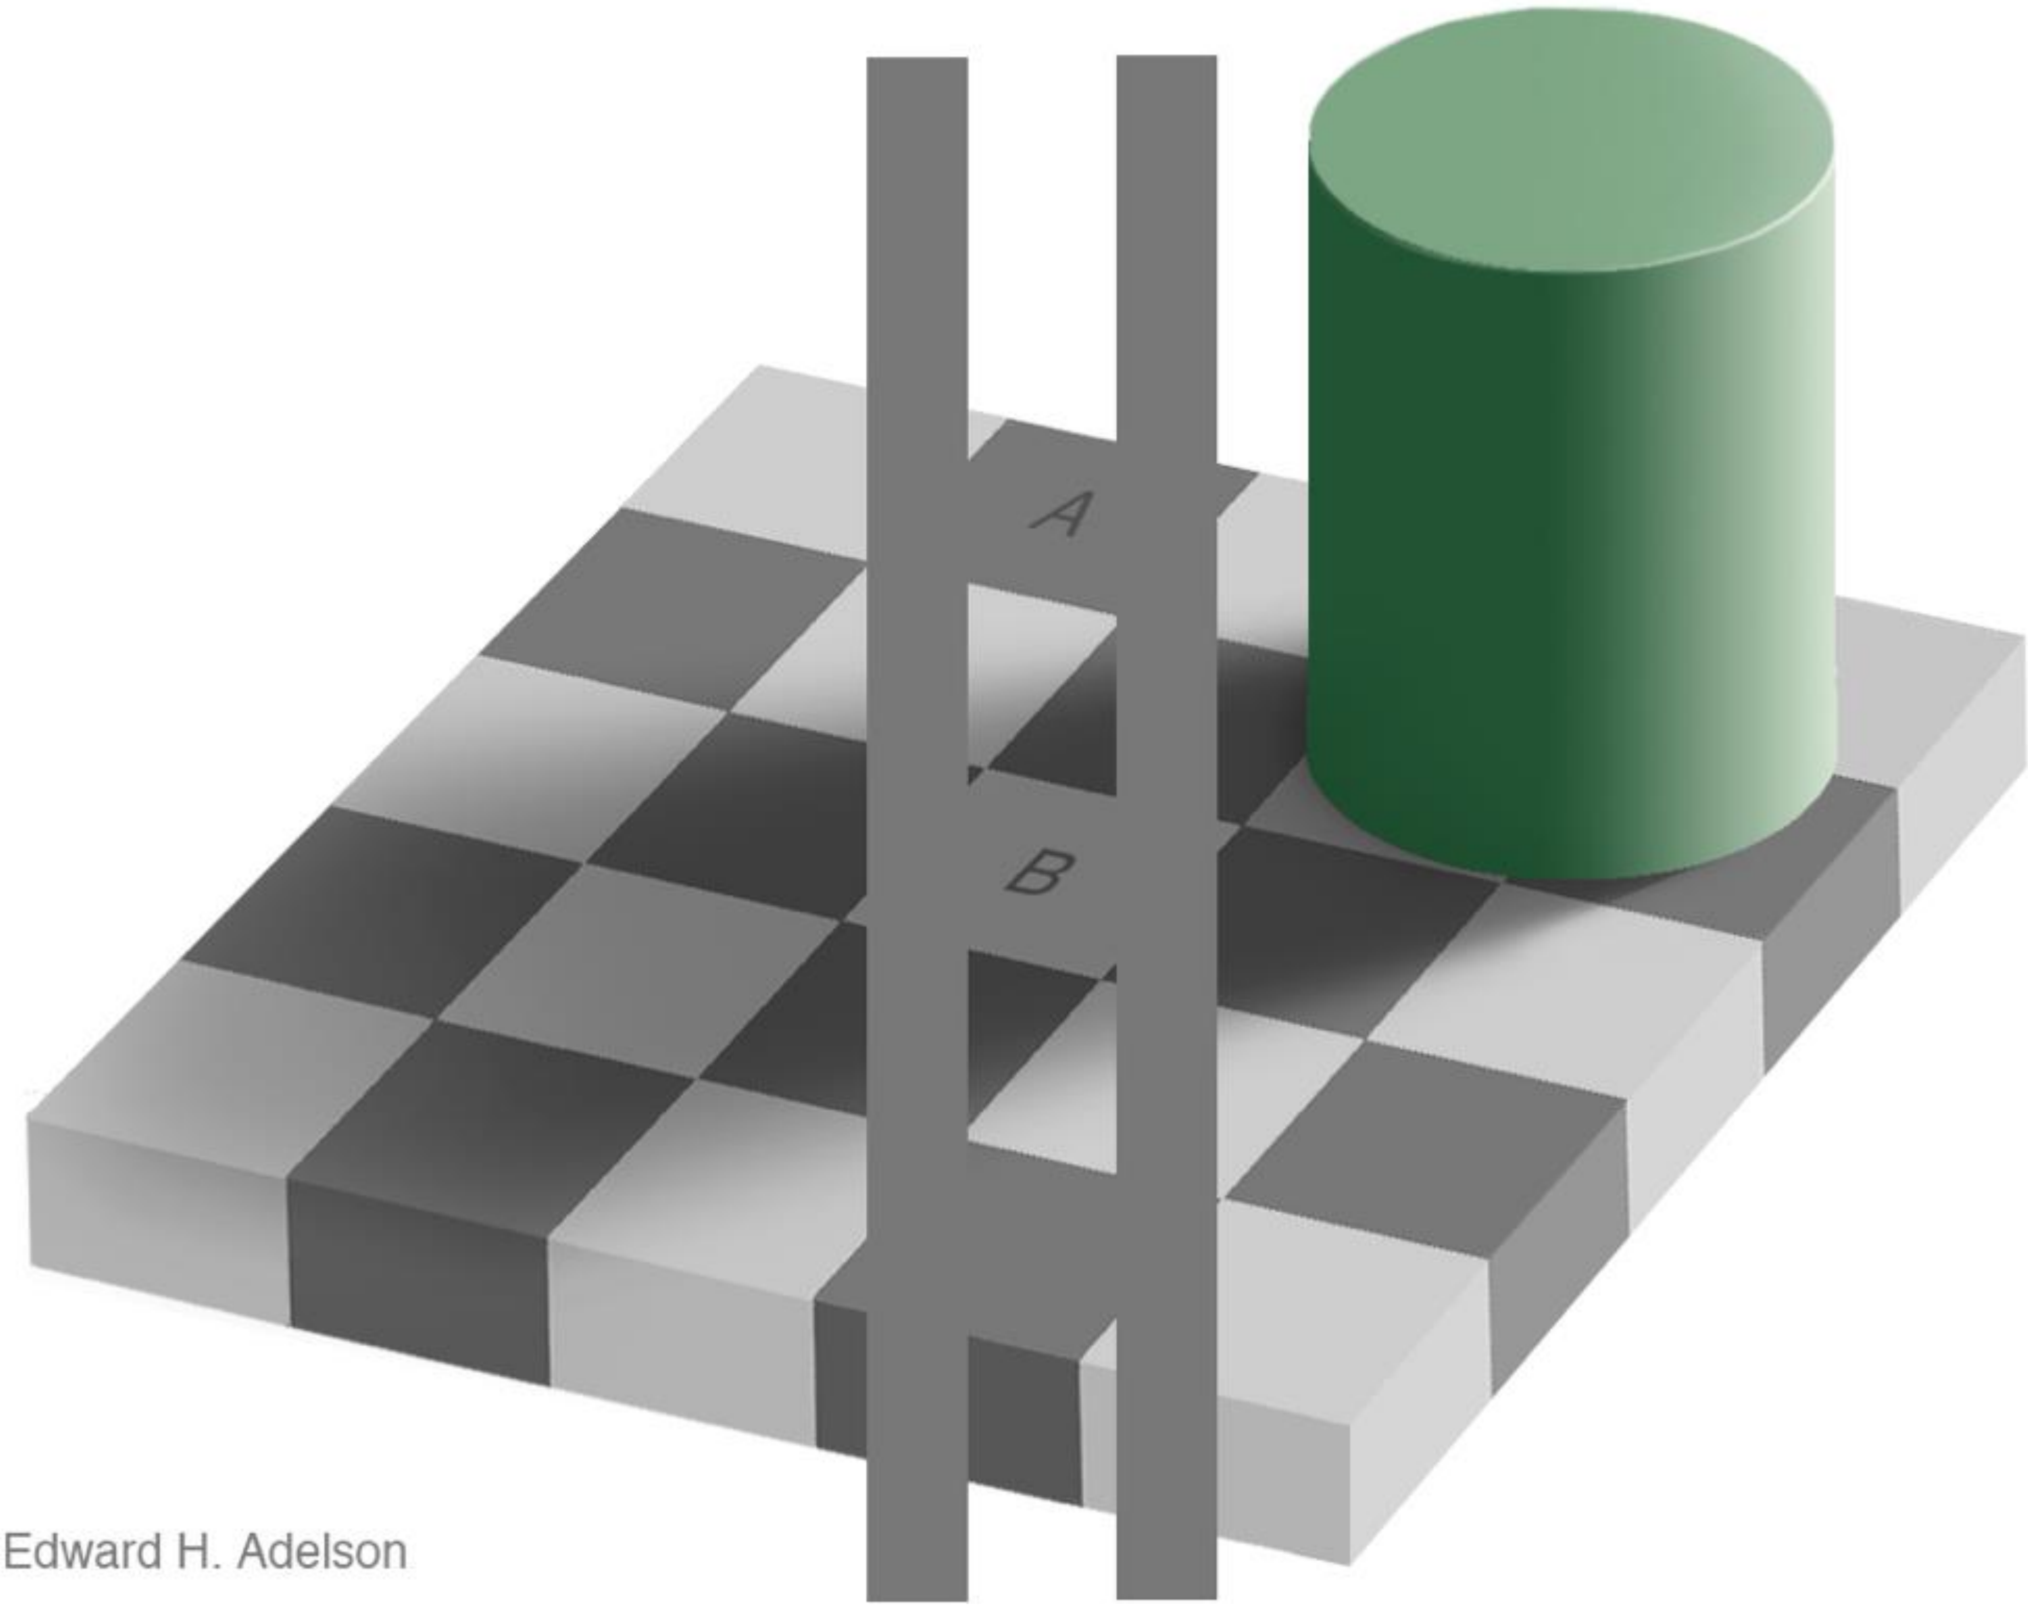

Edward H. Adelson

Video

https://www.youtube.com/watch?v=z9Sen1HTu5o

Exam

- 2 hours
- 15 questions
- 60% of your final mark
- Open book
- Calculators allowed
- Bring a ruler and pencils/eraser

Part A – Algorithms + code

Demonstrate use of an algorithm. Similar to many tutorial questions.

"In the scene shown, the camera is at (2,3,1) in world coordinates, and is rotated to point straight down. Calculate the view matrix for this situation. "

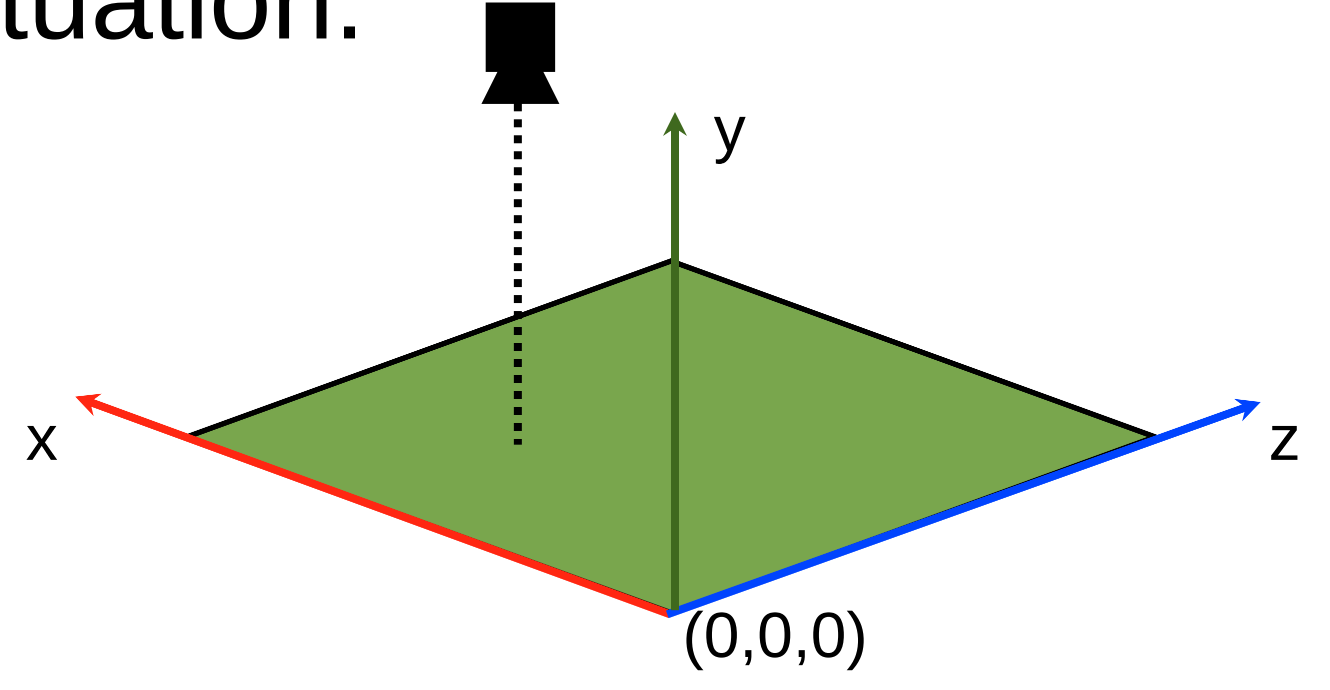

Solution

To move an object to pos(2,3,1) and then rotate by 90 to face the ground we would need to do

translated(2,3,1)

rotated(-90,1,0,0); //around x axis

But since this is the camera we need to do the inverse…so

Solution

rotated(90,1,0,0);

translated(-3,-2,-1);

$$\begin{vmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{vmatrix} \begin{vmatrix} 1 & 0 & 0 & -3 \\ 0 & 1 & 0 & -2 \\ 0 & 0 & 1 & -1 \\ 0 & 0 & 0 & 1 \end{vmatrix} = \begin{vmatrix} 1 & 0 & 0 & -3 \\ 0 & 0 & -1 & 1 \\ 0 & 1 & 0 & -2 \\ 0 & 0 & 0 & 1 \end{vmatrix}$$

Part B- Definitions

Definitions and advantages/disadvantages of different approaches.

"What is gamut mapping? Explain the difference between absolute colormetric and relative colormetric mapping?"

Sample Solution

Any output device has a certain range of colours it can represent, which we call its gamut. Gamut mapping is a way of mapping a colour from one device that can’t be represented by another device. Absolute colormetric mapping maps to the closest color in the gamut. This can distort hues and does not preserve relative saturation. Relative colormetric mapping desaturates the colour until it lies in the gamut and maintains hues more closely.

.

Part C- Applications

Questions that present a particular scenario and ask what methods you would use and why.

"For a 2D game you want to generate a variety of realistic images of apple trees in winter, like the one shown.

What kind of algorithm would you use to do this? What considerations affect your choice?"

Sample Solution

L-Systems are useful for producing realistic botanical models. I would use a stochastic L-System to incorporate randomness to provide a variety of trees.

Because it is a 2d game we would be able to render them fairly quickly. (If we were in 3d, we would have to be careful about run-time which can get high with complicated realistic 3d models)

Week 13

Demos are on. Locations will be posted on course website. Do not be late. Bring laptop you want to demonstrate on, or you can use cse machine.

Additional exam period consultation times will be scheduled. See the course website.

Week 13 Demo

Have some sample worlds that show off your extensions as well as your normal behaviour ready to go.

Have any code you want to show off ready.

This is your chance to shine ☺