UCL Institute of Ophthalmology

UCL

UP.GRADE seminar:
Human Vision and Colour Pipelines

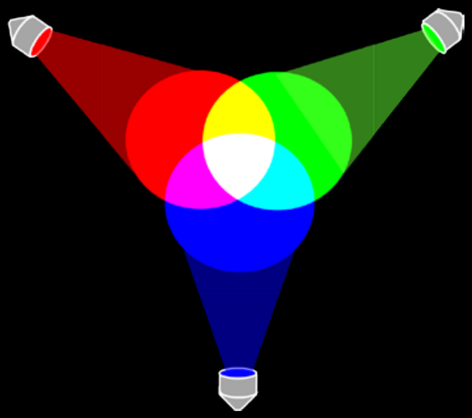

The human colour vision from optics to perception

Andrew Stockman

CAMERIMAGE International Film Festival
Bydgoszcz, 11 - 18 November 2017

CVRL IoO

OUTLINE

1. INTRODUCTION
2. COLOUR VISION AT THE SENSORS
3. COLOUR VISION AFTER THE SENSORS
4. ACHROMATIC AND CHROMATIC VISION
5. COLOUR IN TIME AND SPACE
6. COLOUR IN THE MIND

1. INTRODUCTION

Light

400 - 700 nm is important for vision

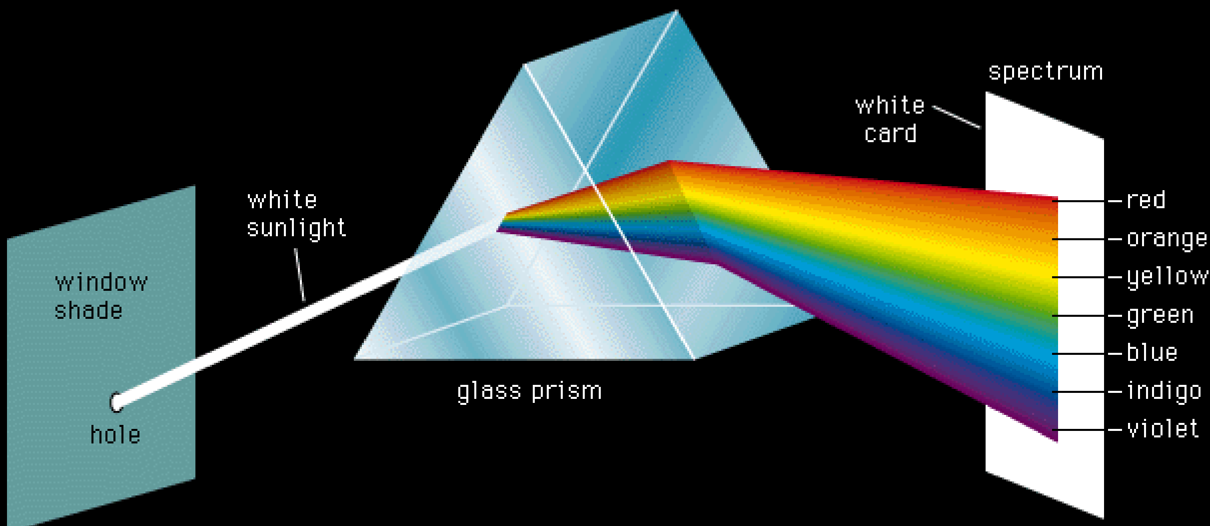

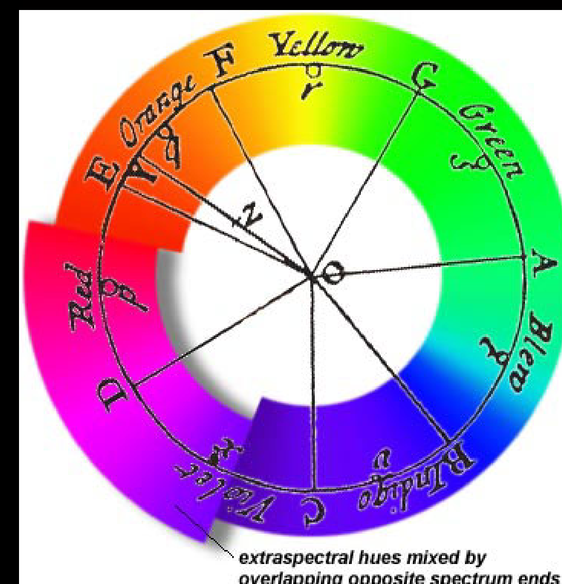

Optical image

The retina is carpeted with light-sensitive rods and cones (the “sensors”)

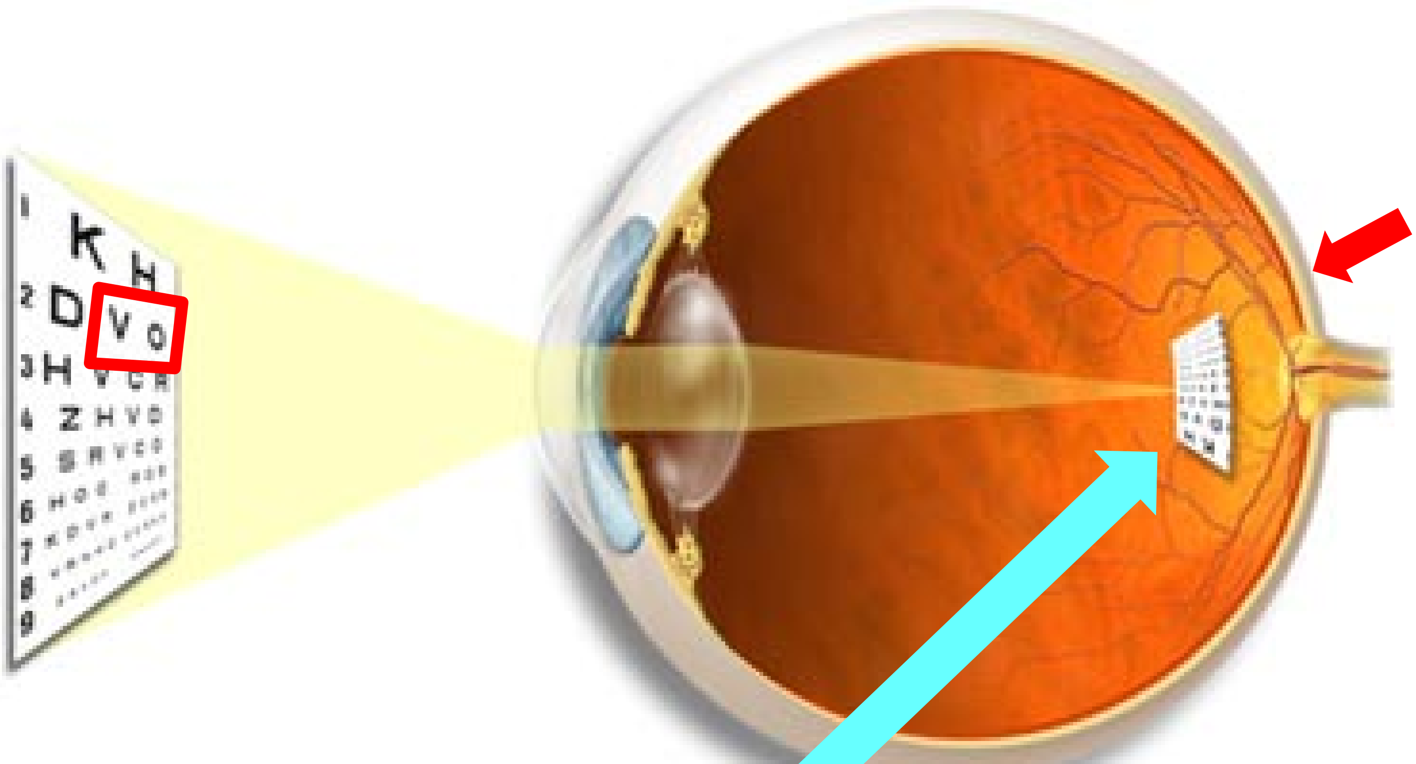

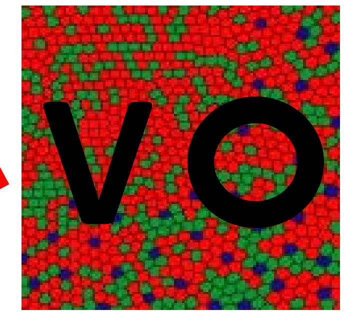

An inverted image is formed on the retina by the cornea and lens.

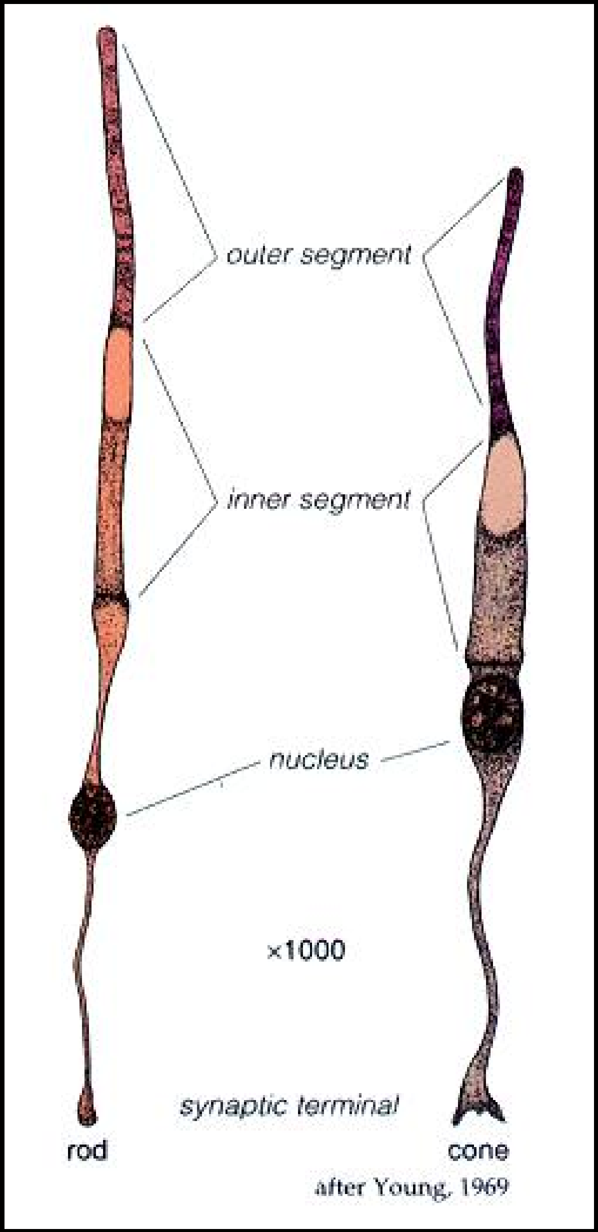

Human photoreceptors (sensors)

- Rods
  - Achromatic night vision
  - 1 type

- Cones
  - Daytime, achromatic *and* chromatic vision
  - 3 types

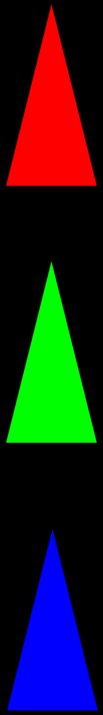

Rod and cone distribution

There are about 120 million rods but only 6 to 7 million cones in the human eye.

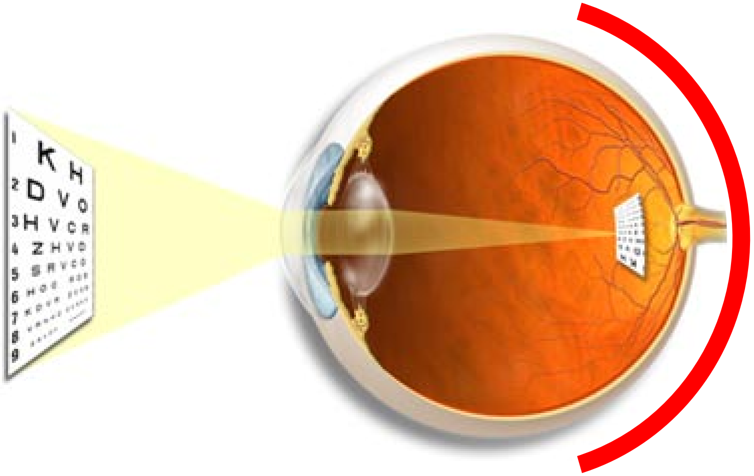

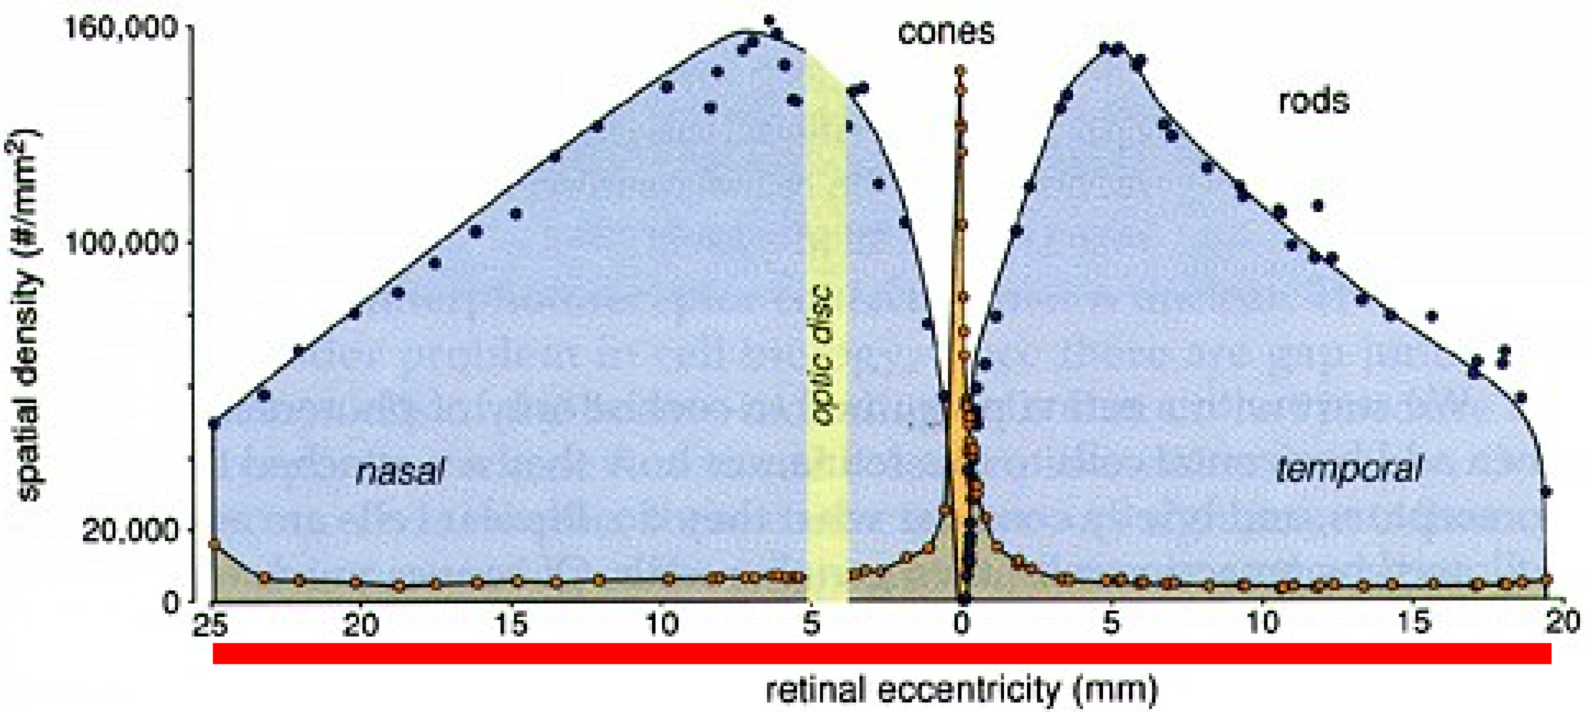

0.3 mm of eccentricity is about 1 deg of visual angle

NOT UNIFORMALLY DISTRIBUTED!

The centre of vision, the “fovea”, is rod-free, and the very centre, the “foveola”, is rod- and S-cone (blue sensor) free.

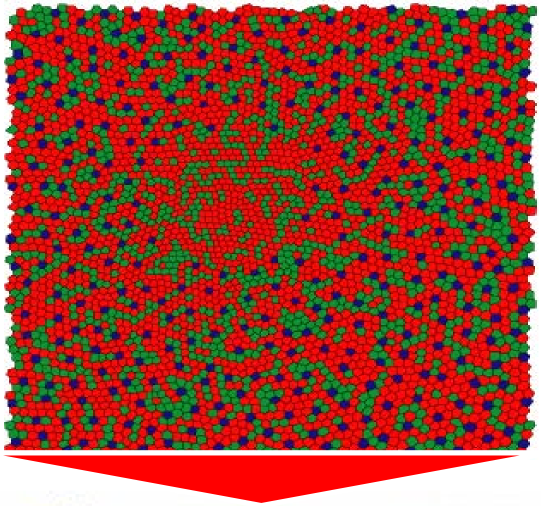

Come back to why the centre is blue-sensor free at the end (if there is time).

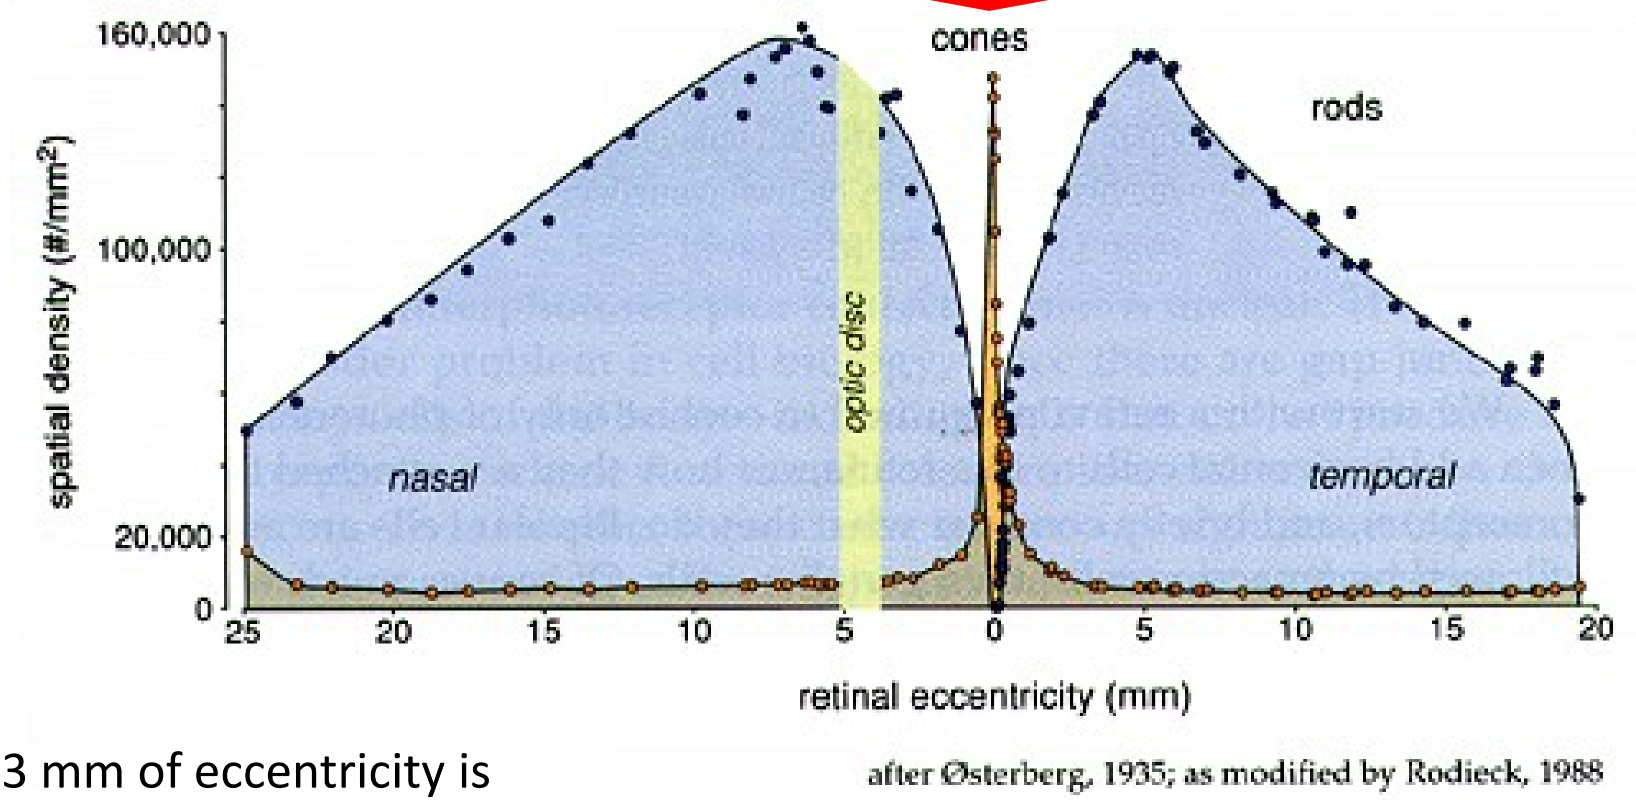

0.3 mm of eccentricity is about 1 deg of visual angle

During the day, you have to look at things directly to see them in detail

Cones peak at the centre of vision in the fovea at 0 deg

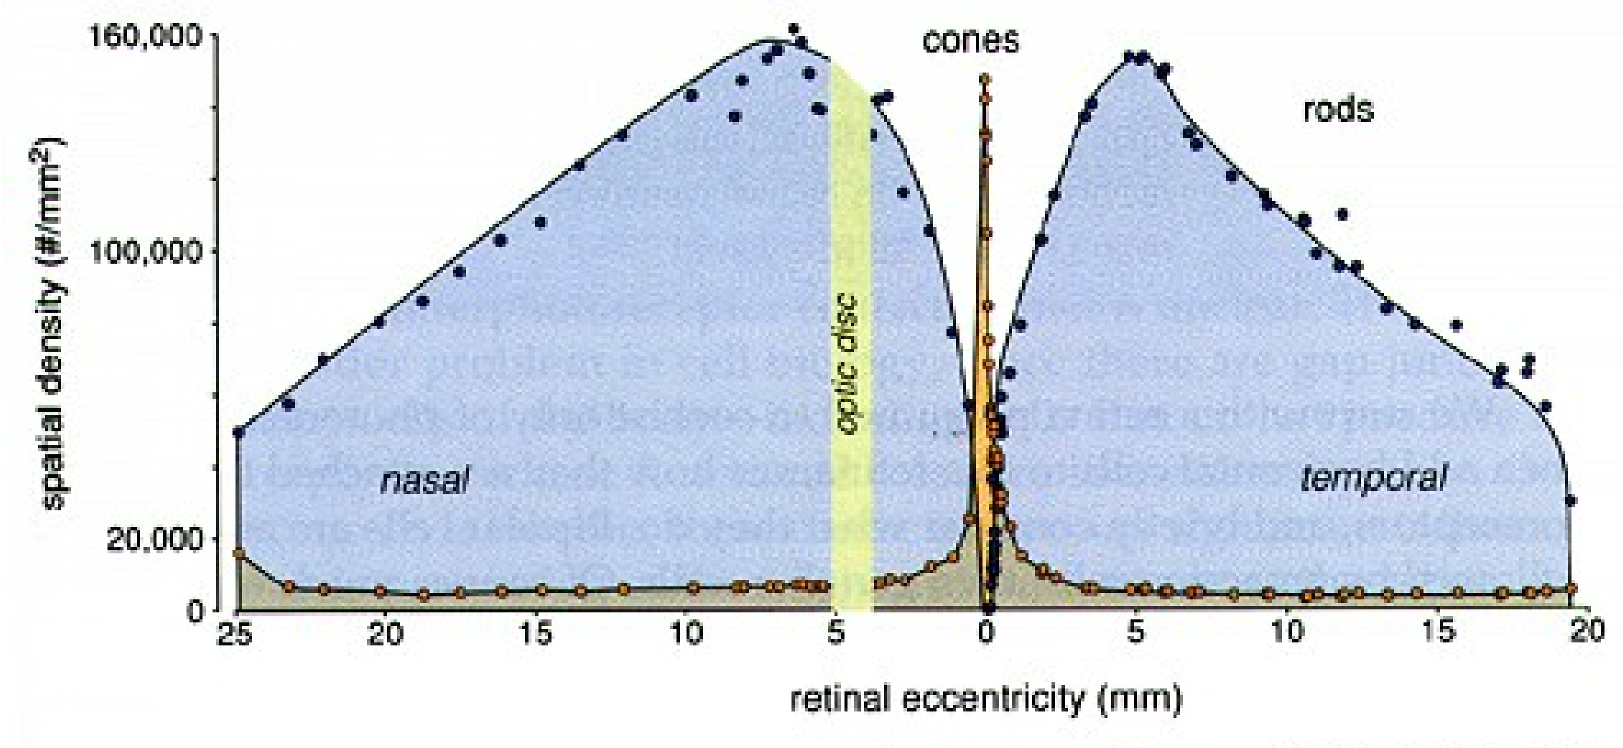

after Østerberg, 1935; as modified by Rodieck, 1988

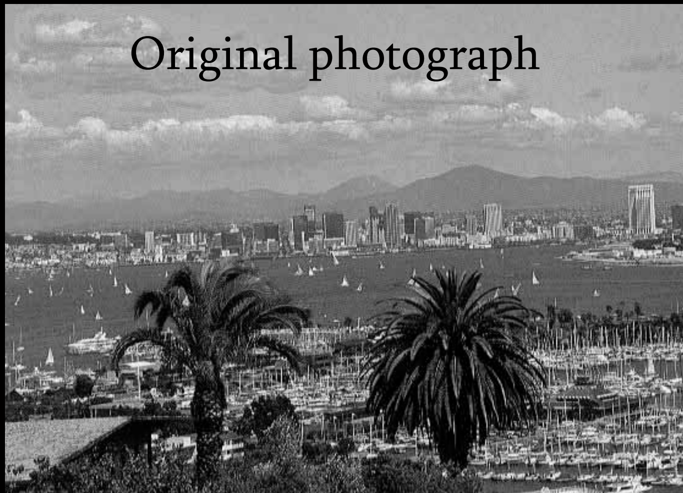

The human visual system is a foveating system

Simulation of what we see when we fixate with cone vision…

Credit: Stuart Anstis, UCSD

The eye is imperfect

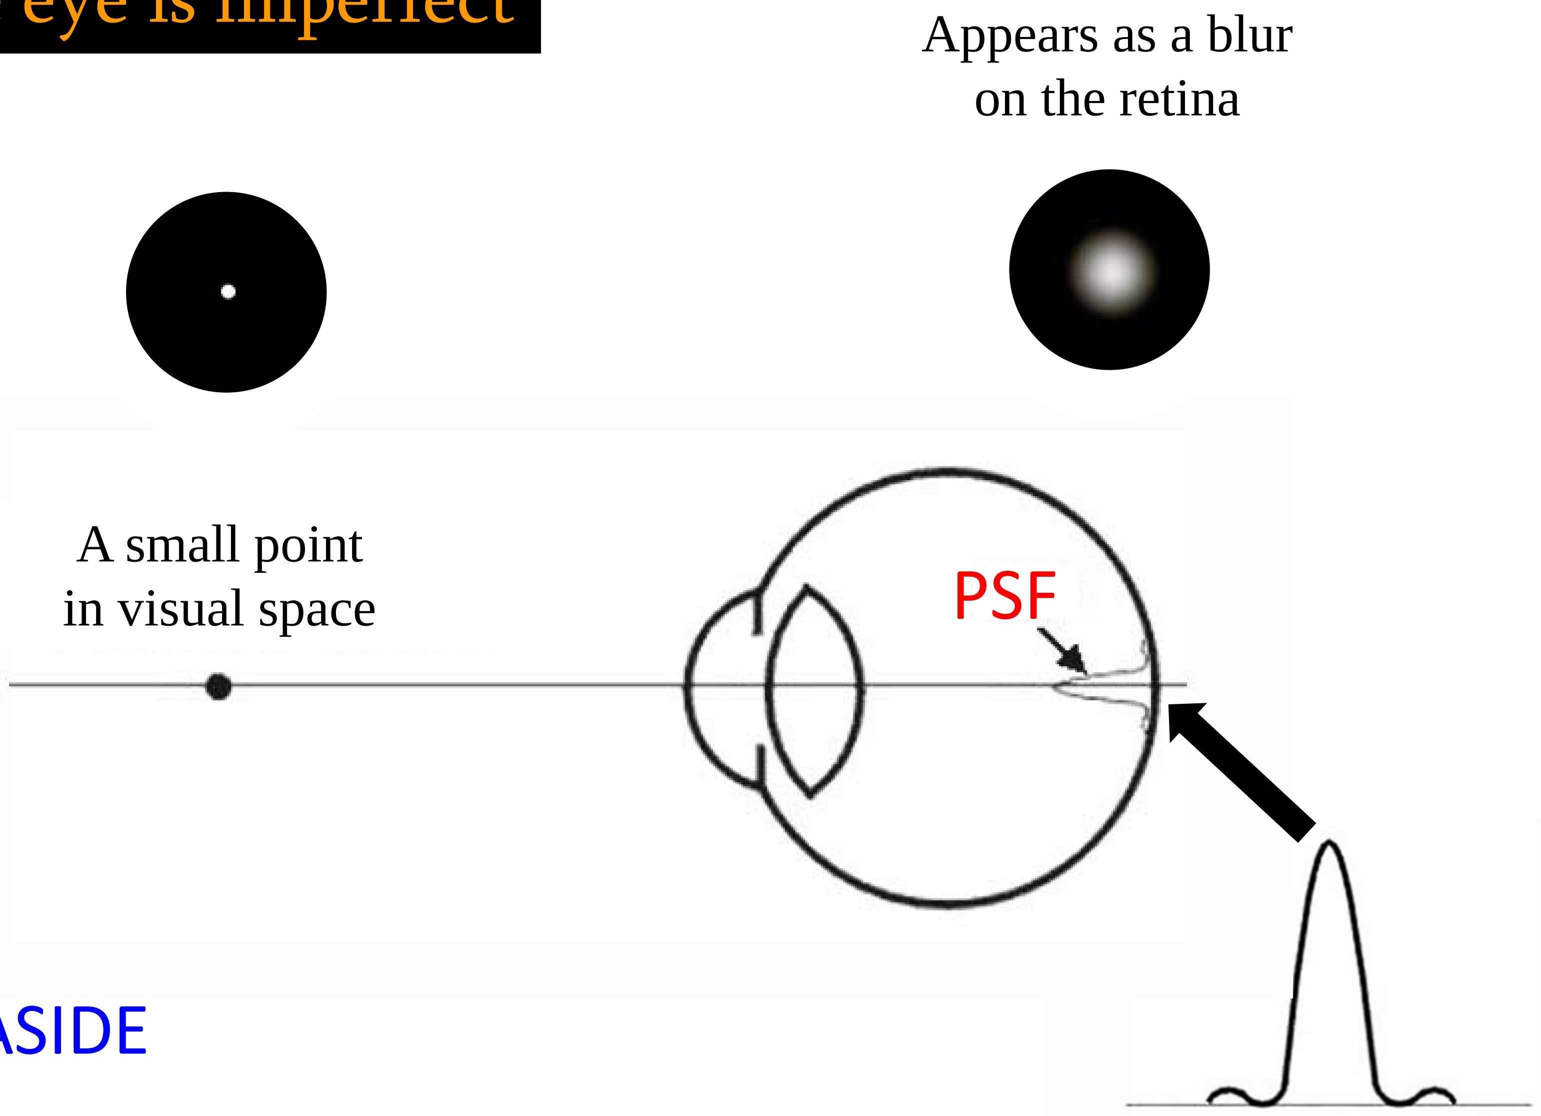

ASIDE

Point spread function (PSF)

From Webvision, Michael Kalloniatis

2. COLOUR AT THE SENSORS

How dependent are
we on colour?

No colour…

Colour…

Colour is important because it helps us to discriminate objects from their surroundings.

But how important is it?

Split the image into…

ACHROMATIC COMPONENTS

CHROMATIC COMPONENTS

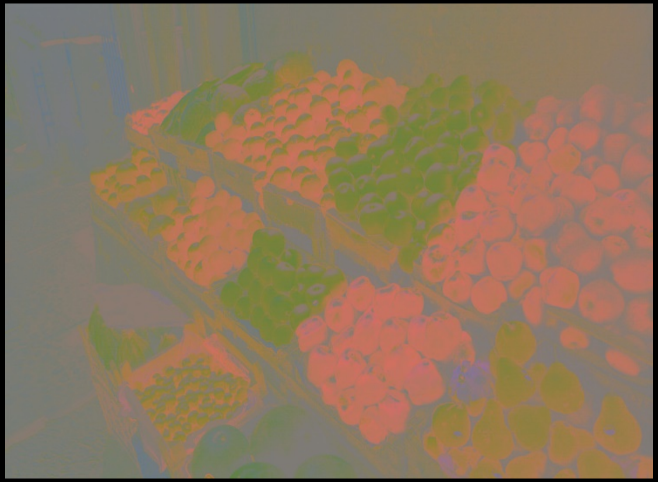

CHROMATIC COMPONENTS

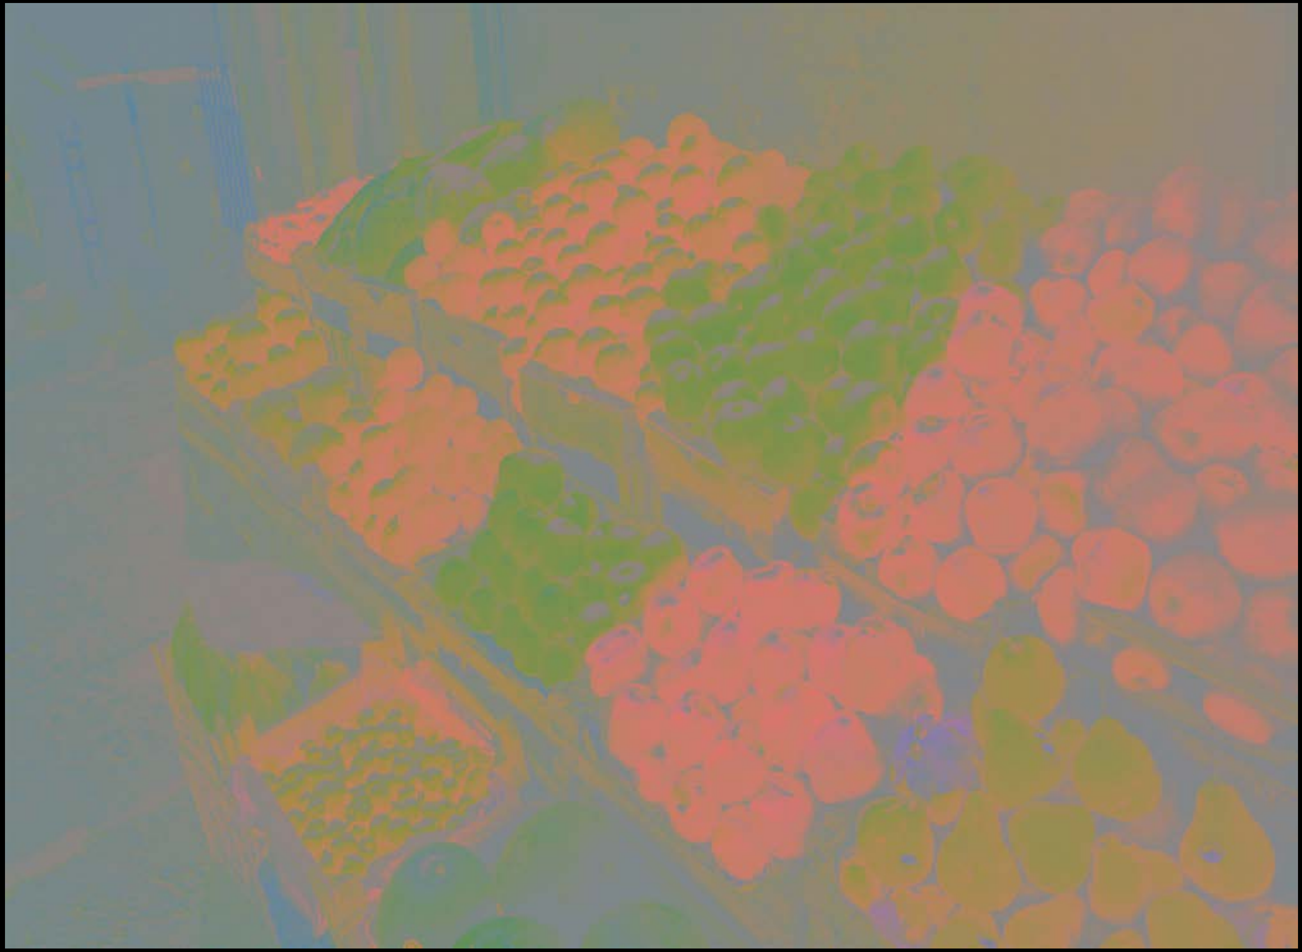

Chromatic information *by itself* provides relatively limited information…

ACHROMATIC COMPONENTS

Achromatic information is important for fine detail in time and space …

Are the colours that we see…

a property of mainly physics or biology?

Colour isn't just about physics. For example:

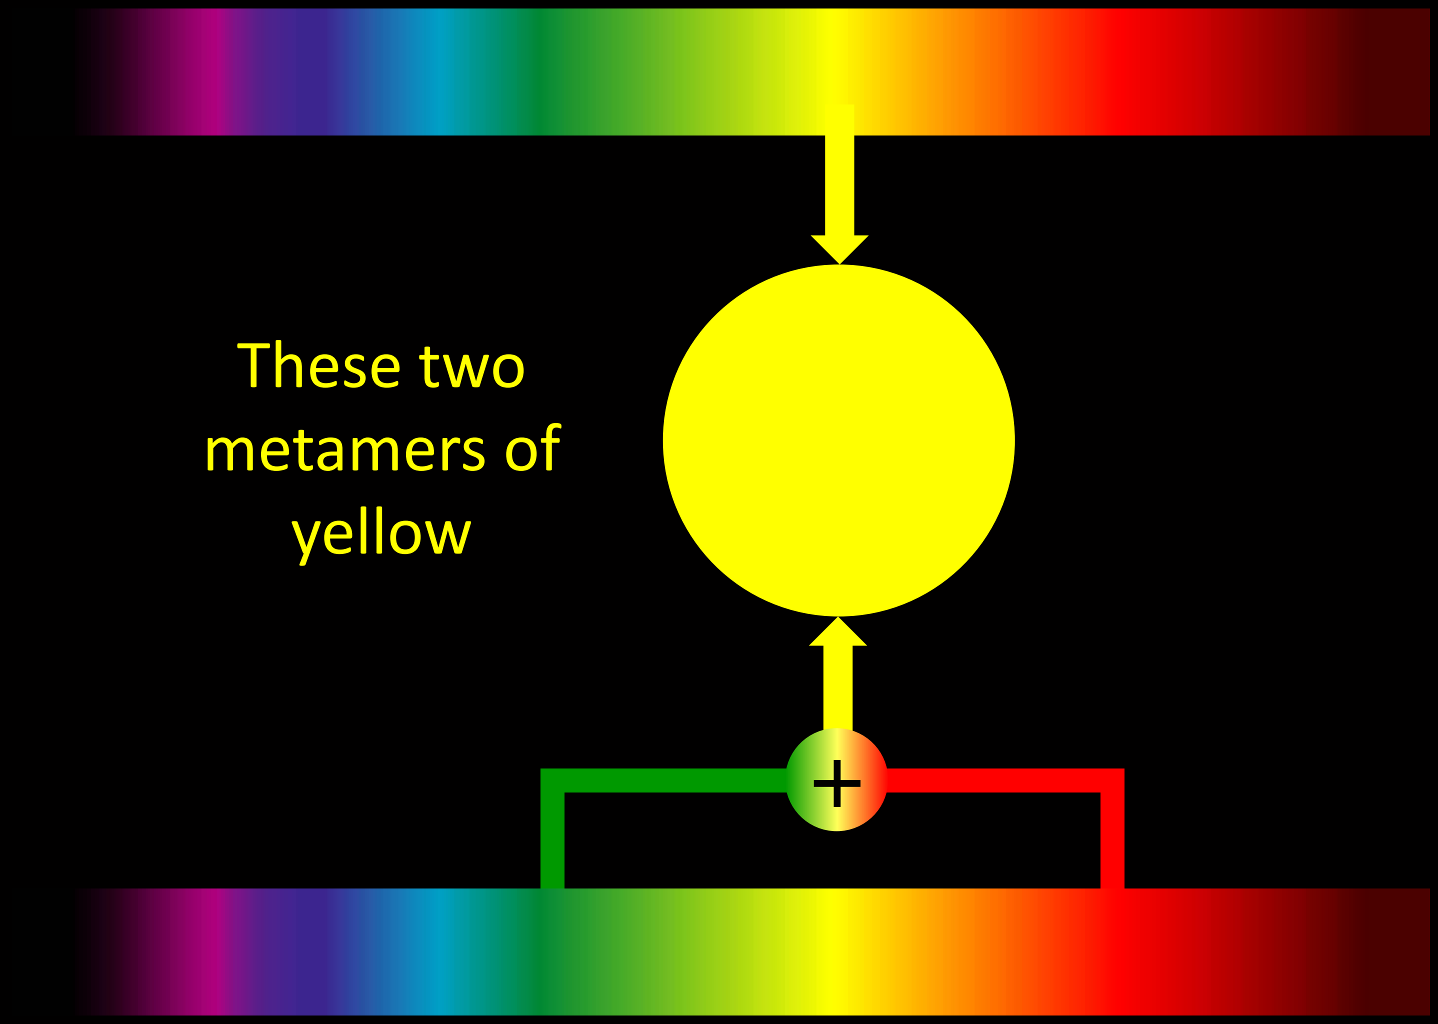

though physically very different, can appear identical.

There are many other such metamers or matches…

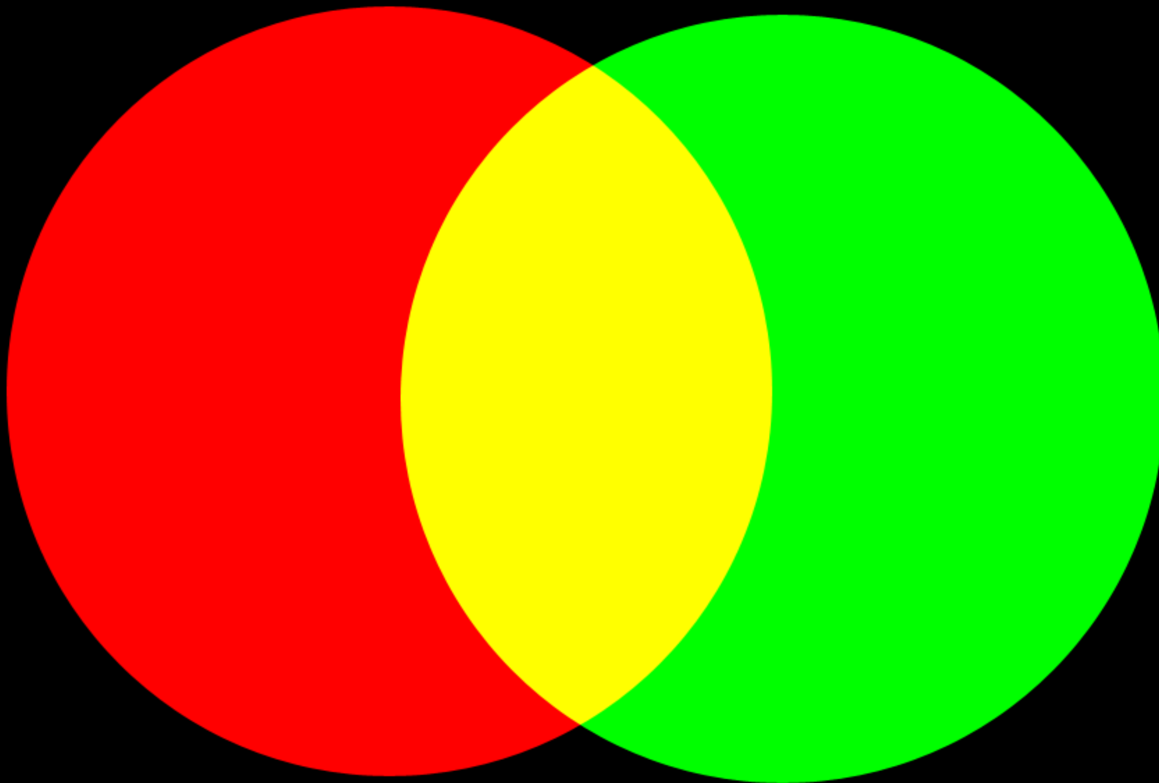

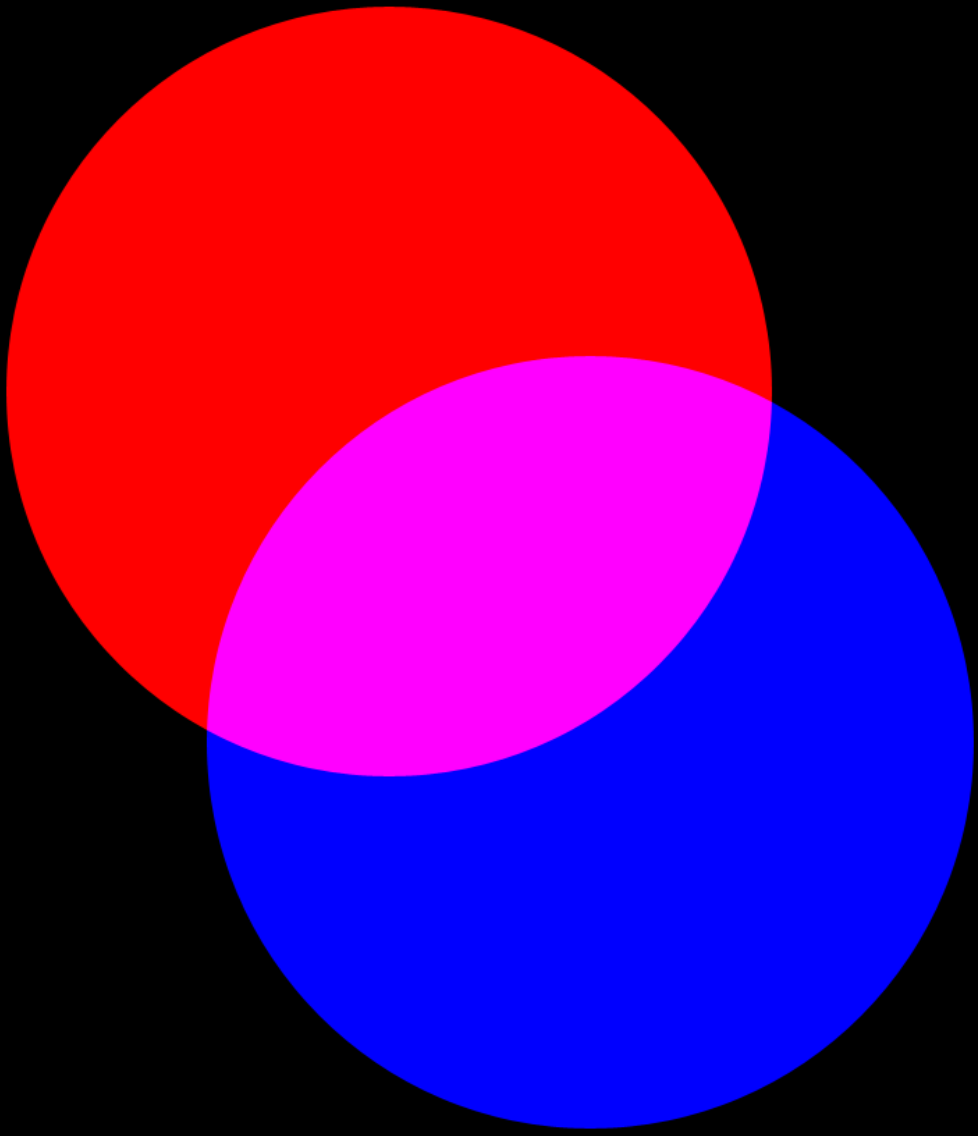

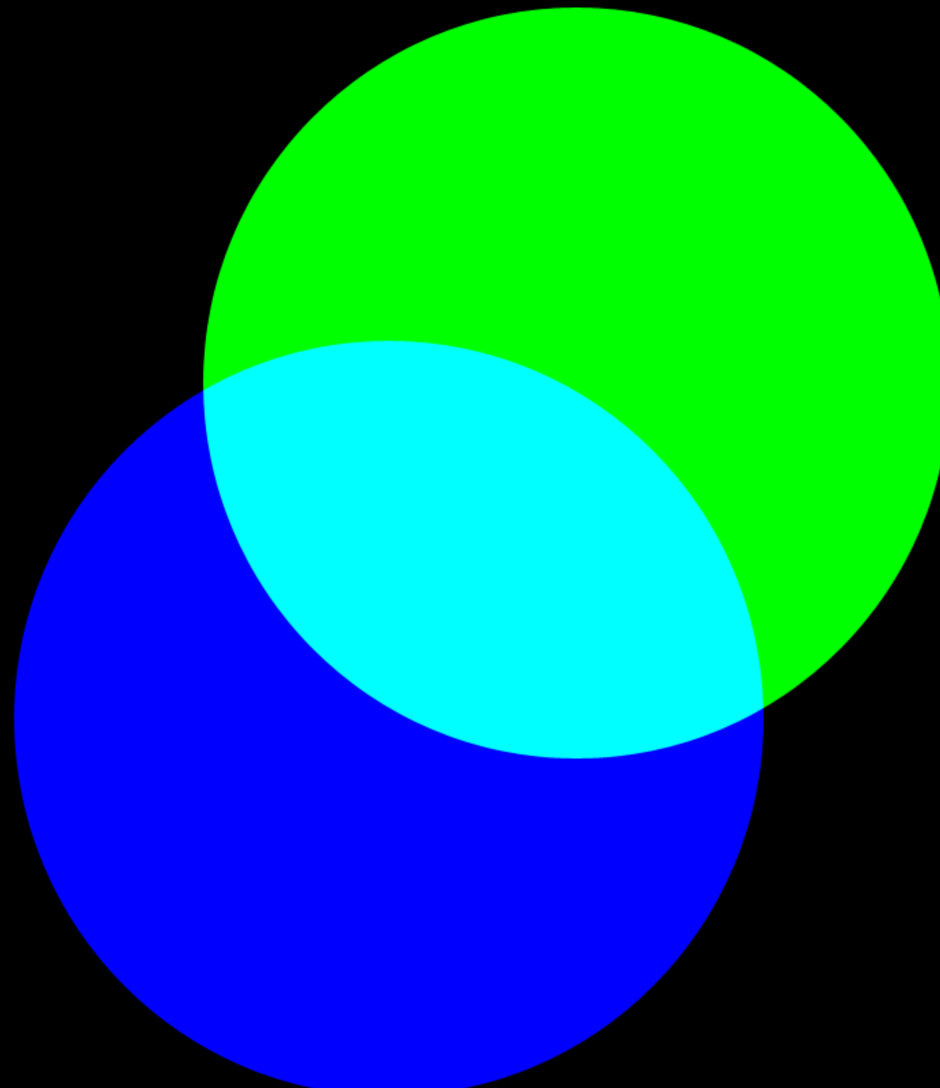

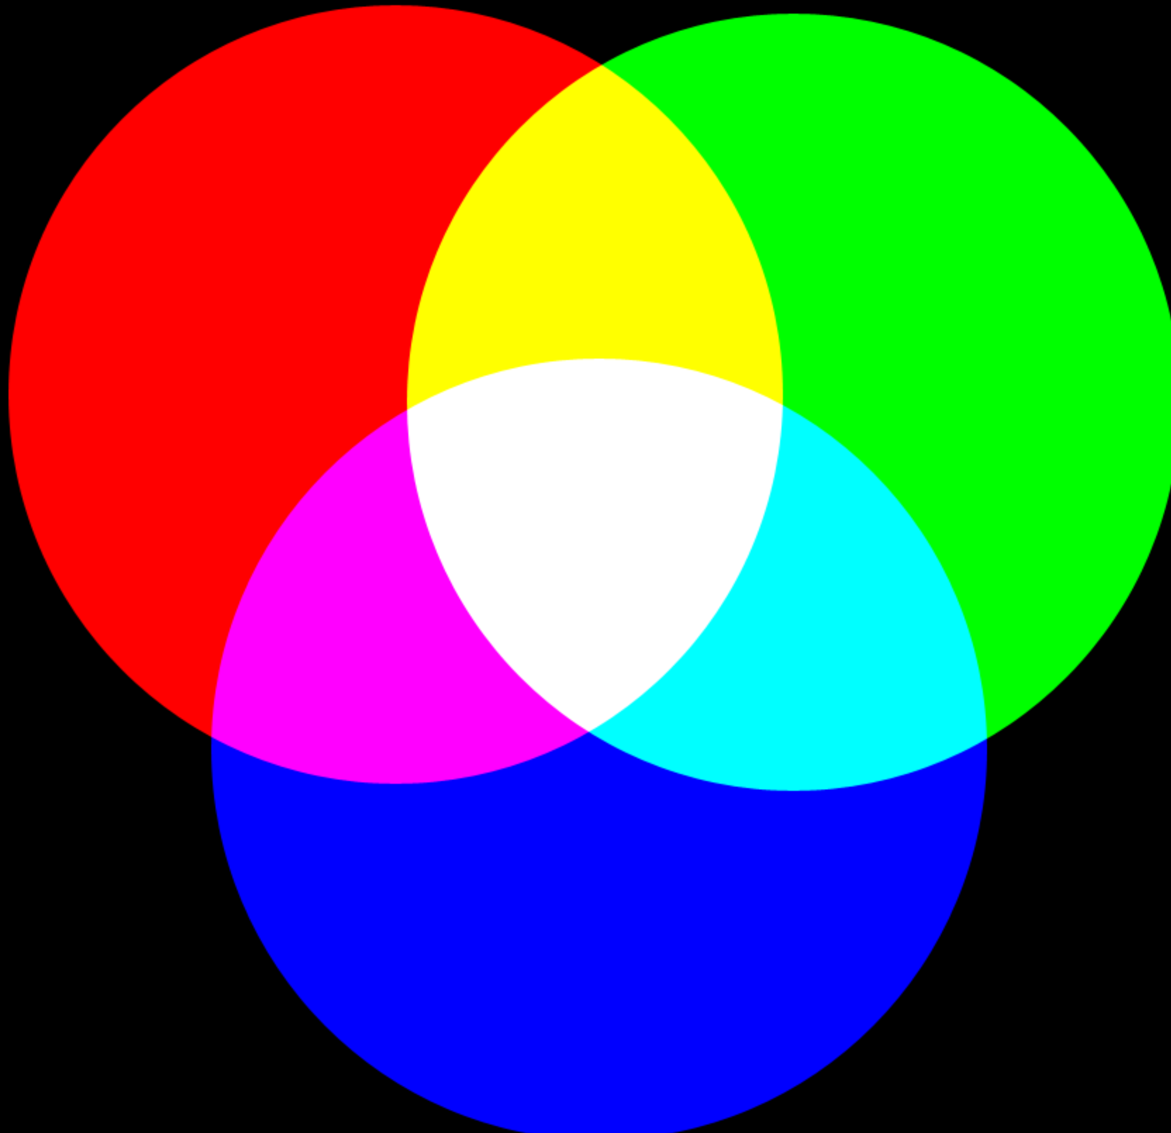

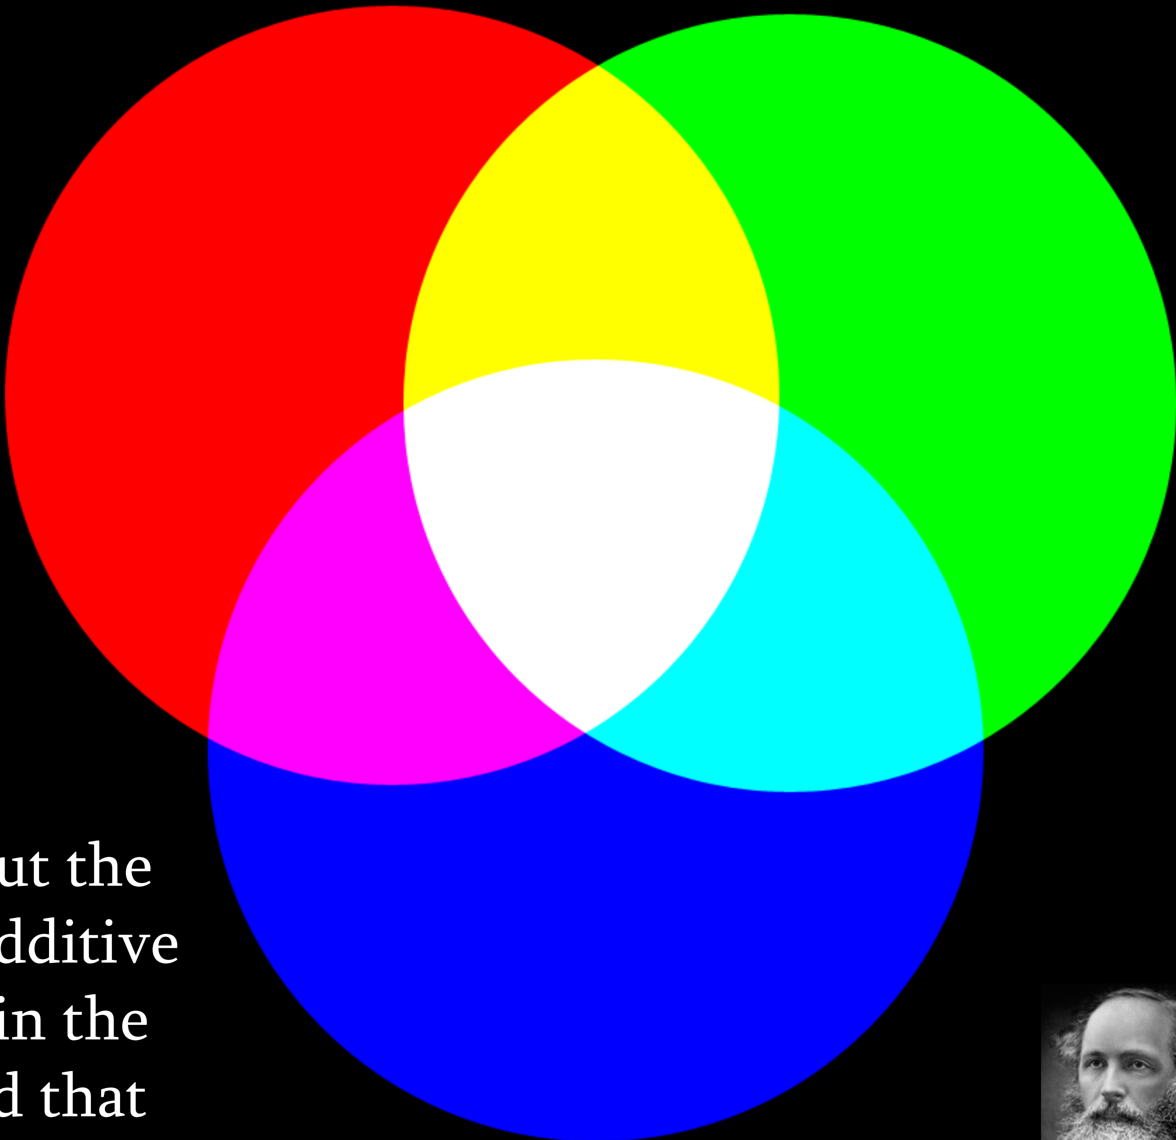

Before we knew about the underlying biology, additive colour mixing done in the 19th century revealed that colour vision was…

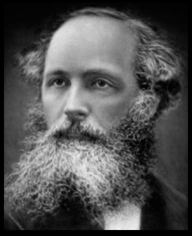

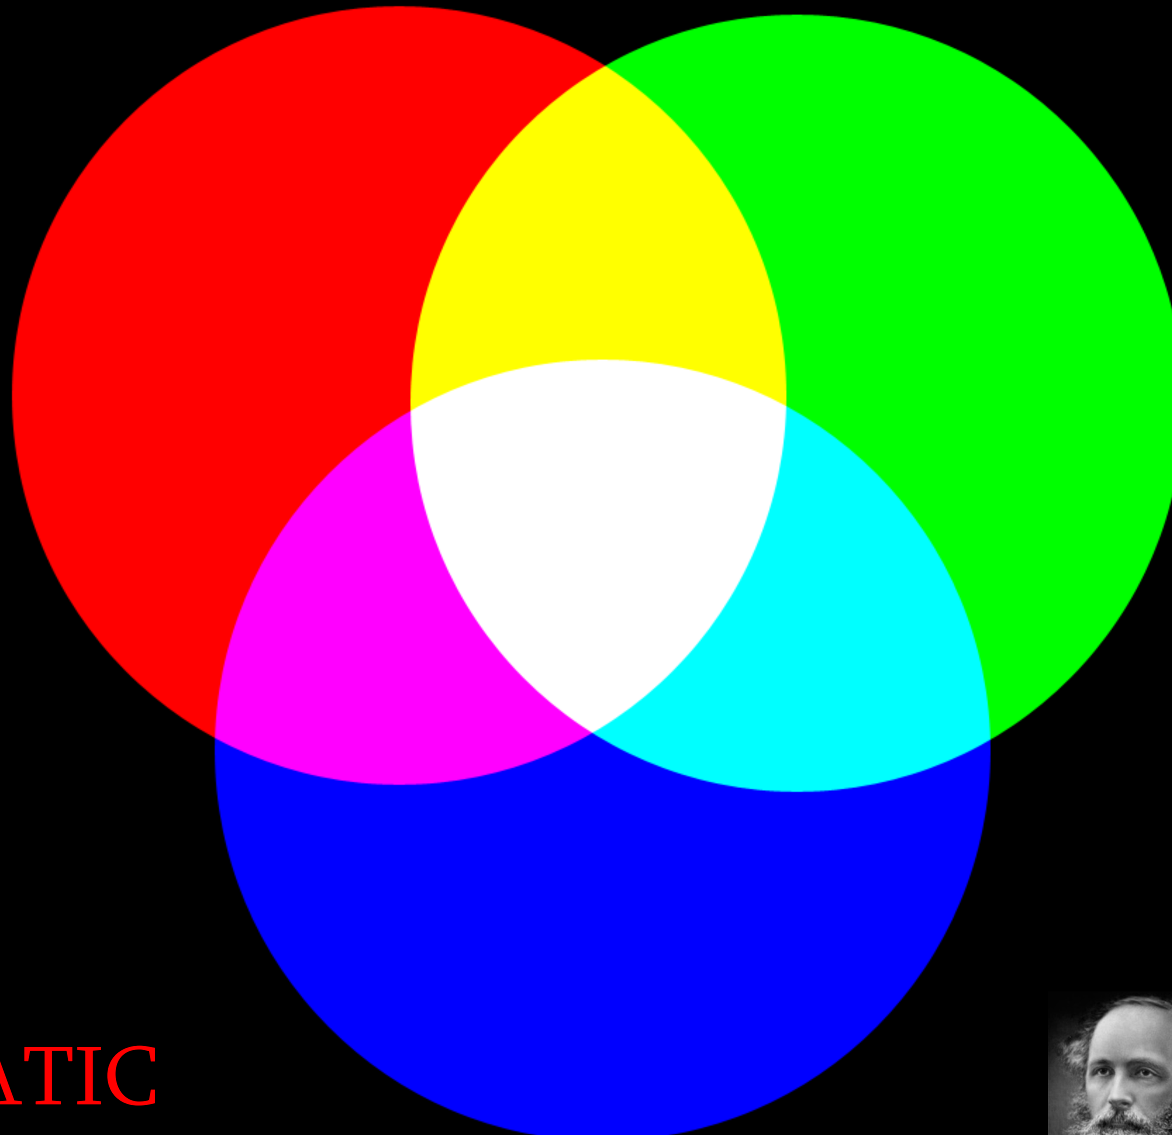

TRICHROMATIC

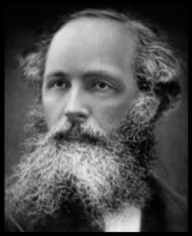

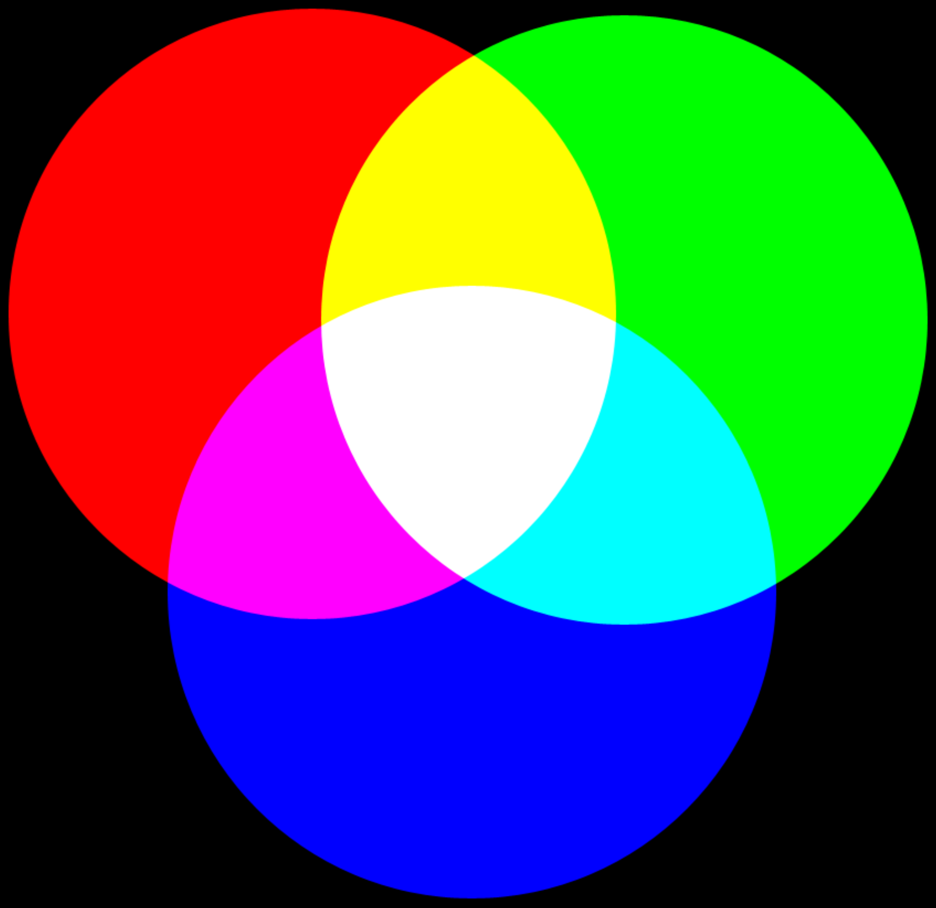

Trichromacy means that colour vision at the input to the visual system is relatively simple.

It is a 3 variable system…

Colour TV

Trichromacy is exploited in colour reproduction, since the myriad of colours perceived can be produced by mixing together small dots of three colours.

If you look closely at a colour television (or this projector output)…

3-coloured dots

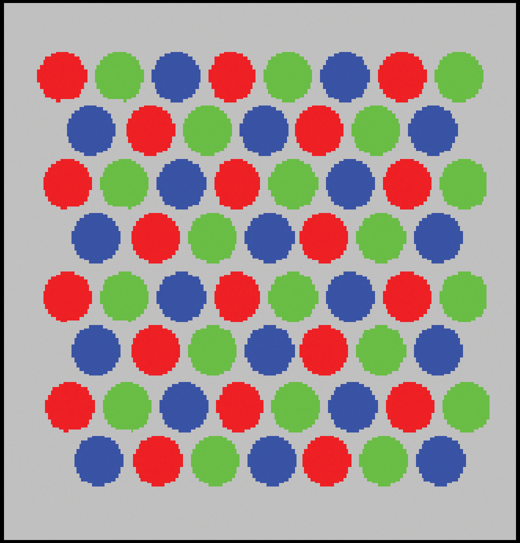

3-coloured bars

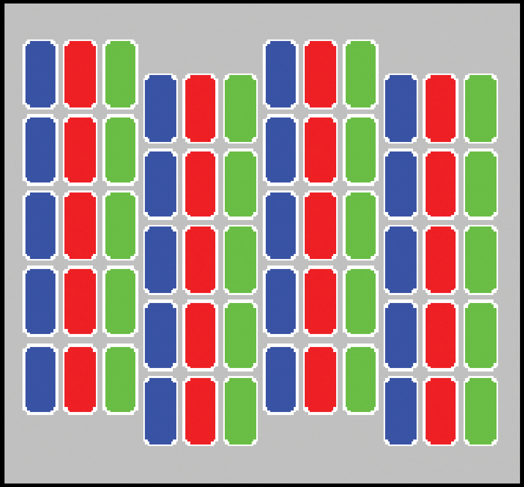

The dots produced by a TV or projector are so small that they are mixed together by the eye and thus appear as uniform patches of colour

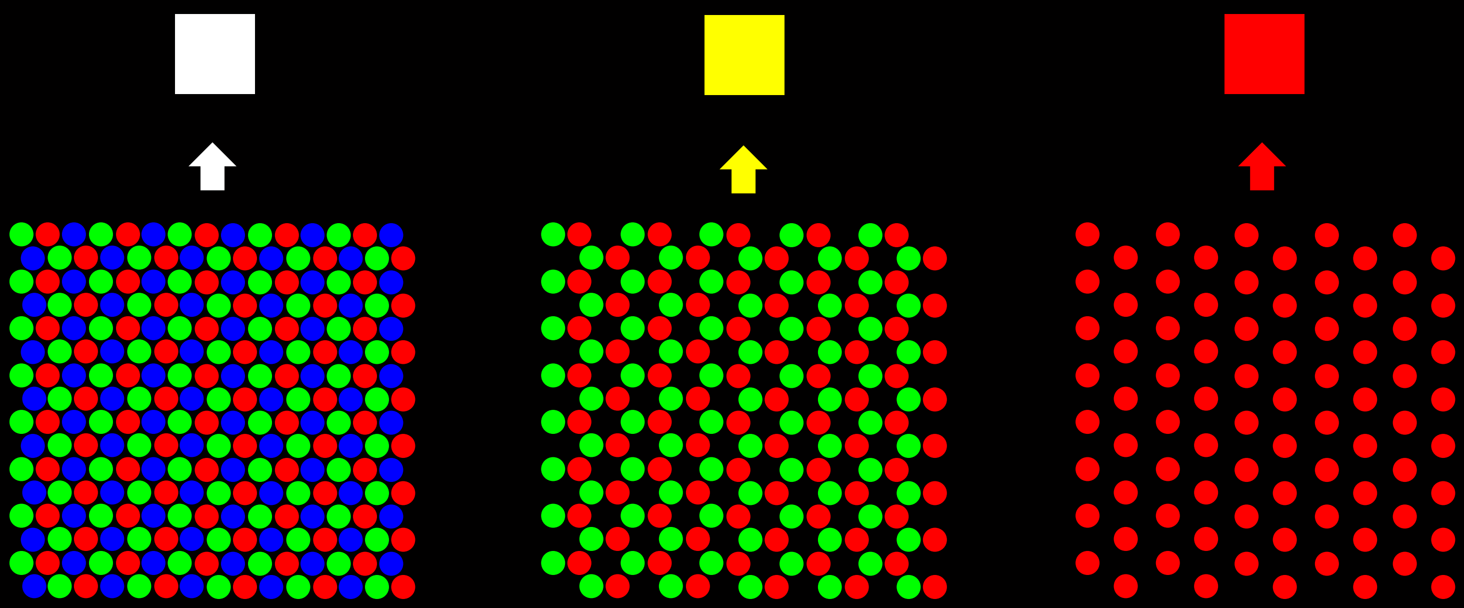

A Sunday afternoon on the island of La Grande Jatte

George Seurat

Why is human vision trichromatic?

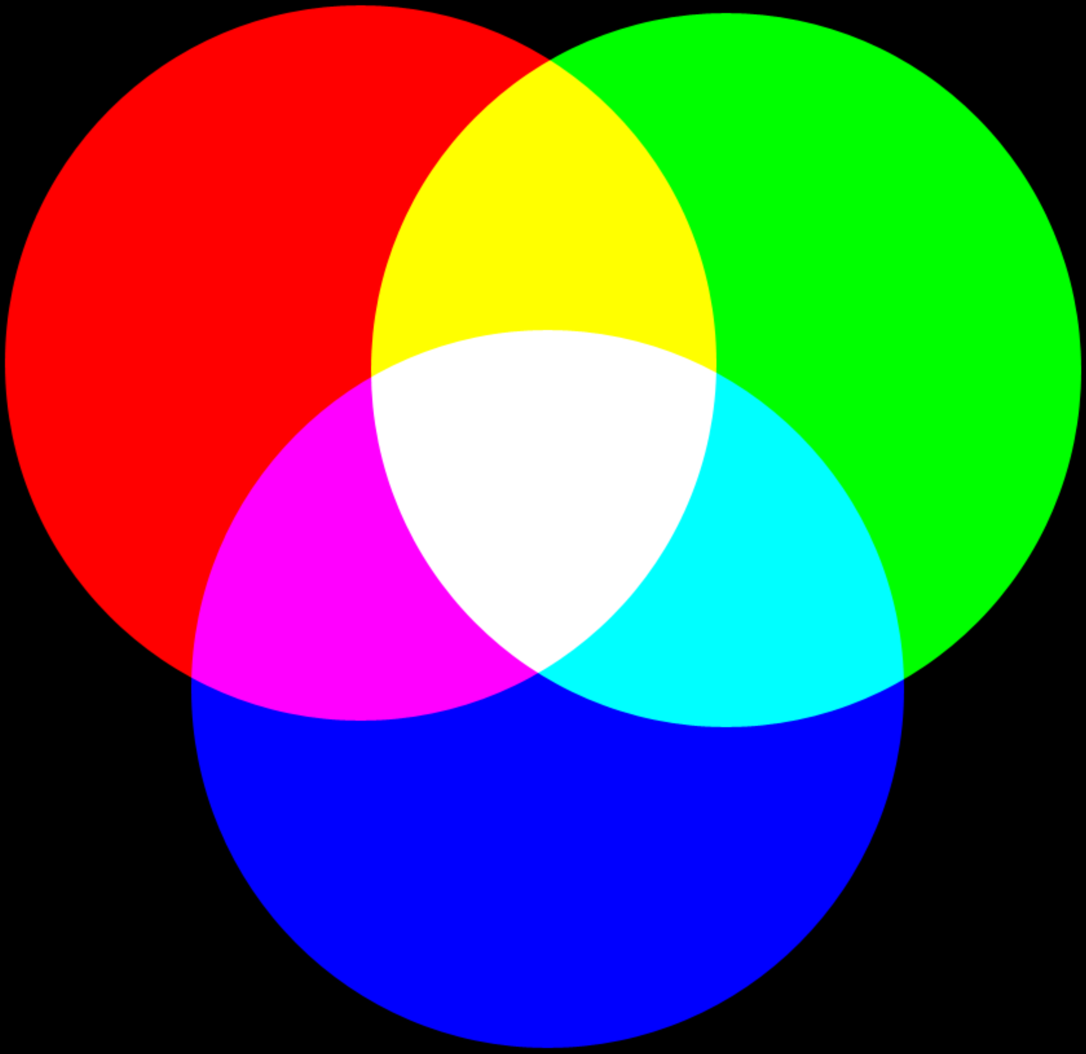

The main reasons are because just THREE cone photo-receptors or sensors are responsible for daytime colour vision:

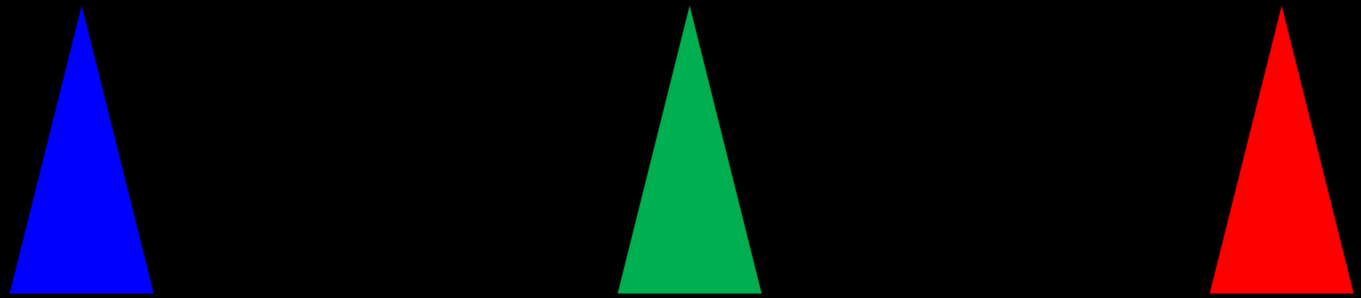

Short-wavelength-sensitive or "blue" sensor

Middle-wavelength-sensitive or "green" sensor

Long-wavelength-sensitive or "red" sensor

And because each sensor produces a UNIVARIANT output.

COLOUR VISION AND VISION, IN GENERAL, IS FUNDAMENTALLY LIMITED BY THE PROPERTIES OF THESE SENSORS…

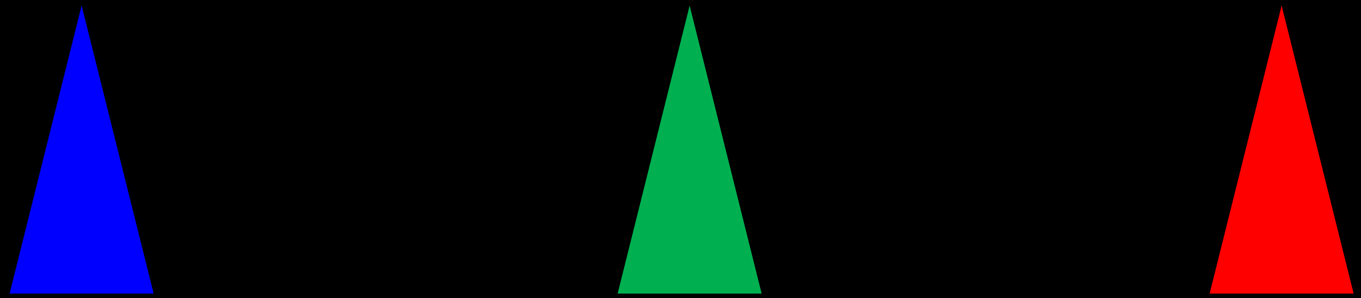

Short-wavelength-sensitive or “blue” sensor

Middle-wavelength-sensitive or “green” sensor

Long-wavelength-sensitive or “red” sensor

What is univariance?

Univariance can be explained simply at the molecular level by the interaction of photons with the photopigment molecules in each photoreceptor…

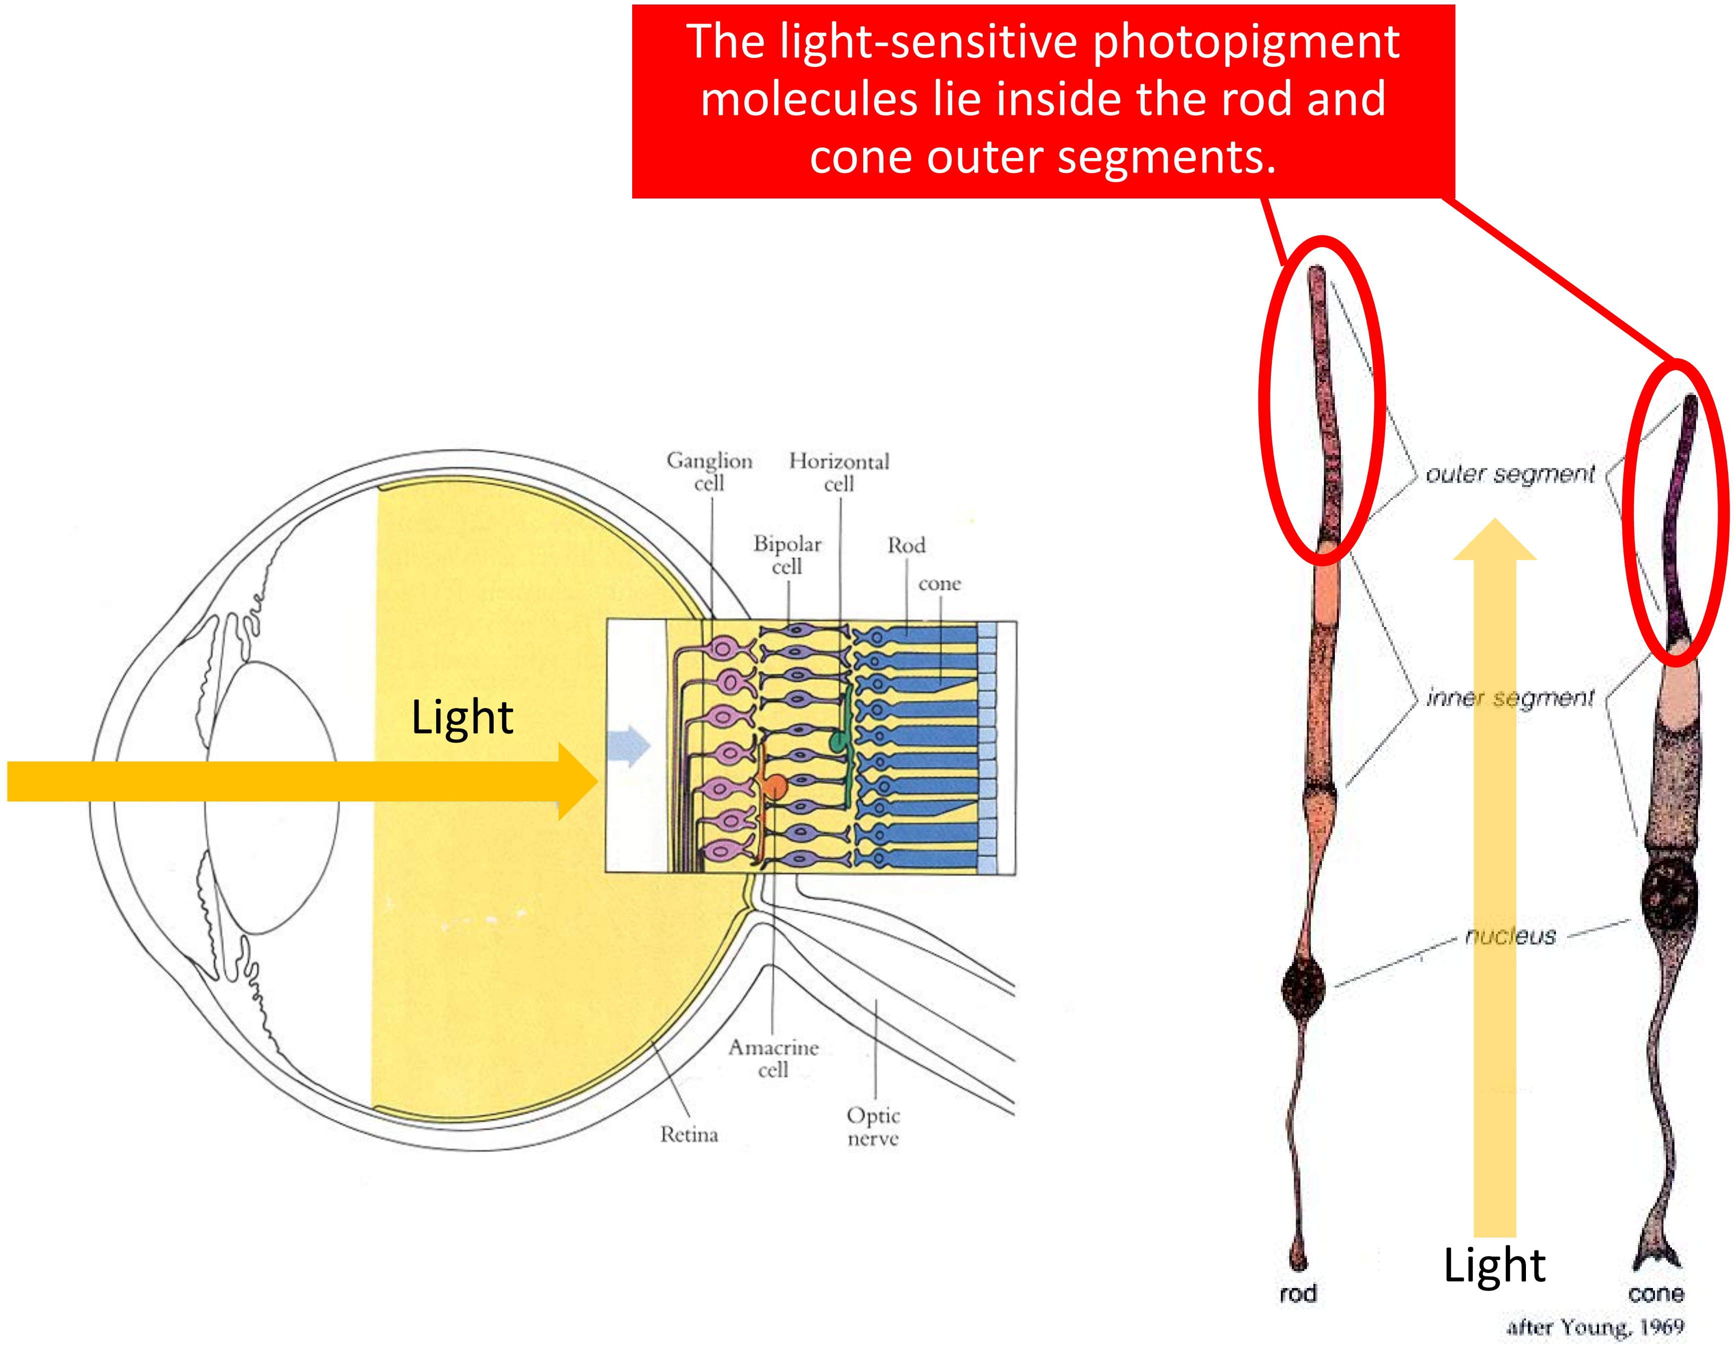

The light-sensitive photopigment molecules lie inside the rod and cone outer segments.
Light
Ganglion cell
Horizontal cell
Bipolar cell
Rod
cone
Amacrine cell
Retina
Optic nerve
outer segment
inner segment
nucleus
Light
rod
cone
after Young, 1969

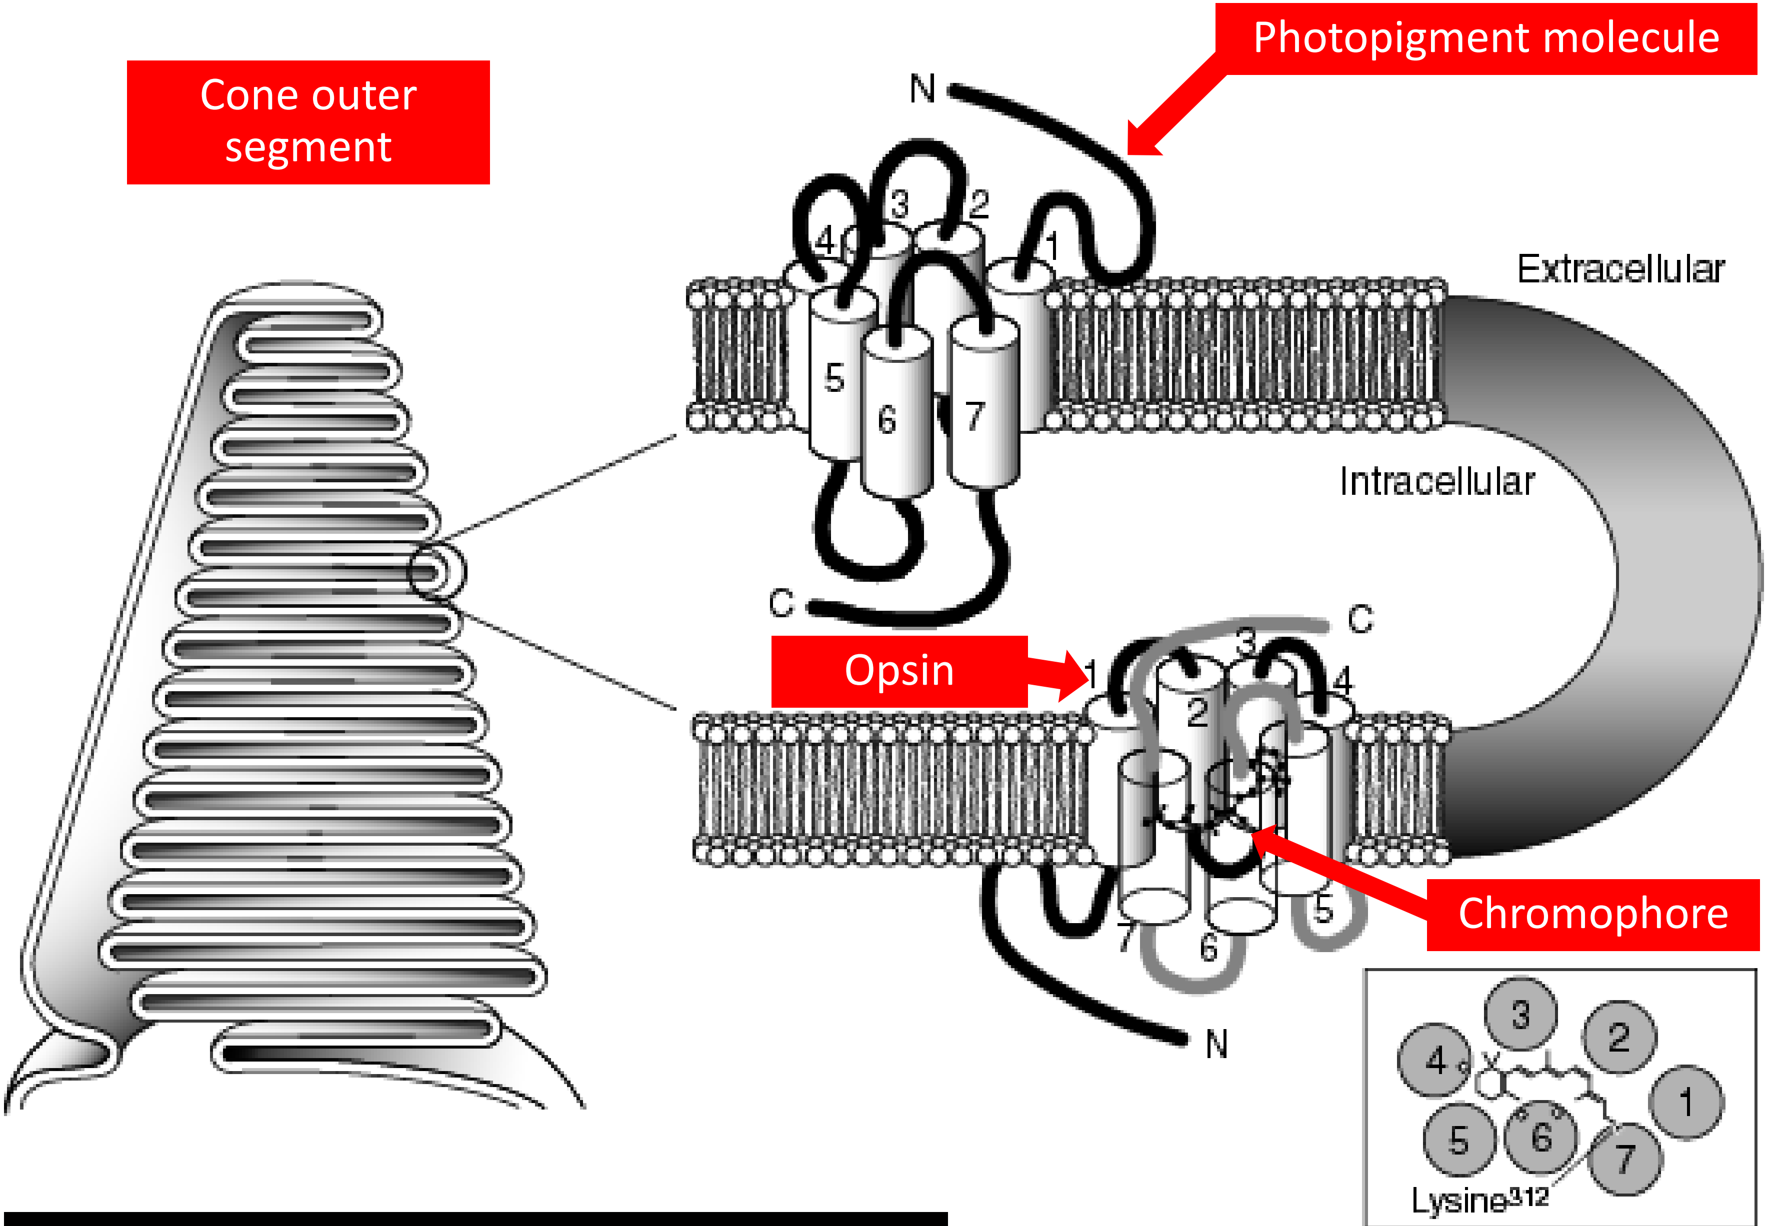

Photopigment molecules (cone)

From Sharpe, Stockman, Jägle & Nathans, 1999

Chromophore

(*chromo-* colour, + *-phore,* producer)
Light-catching portion of any molecule

The molecule is twisted at the 11th carbon.

11-*cis* retinal

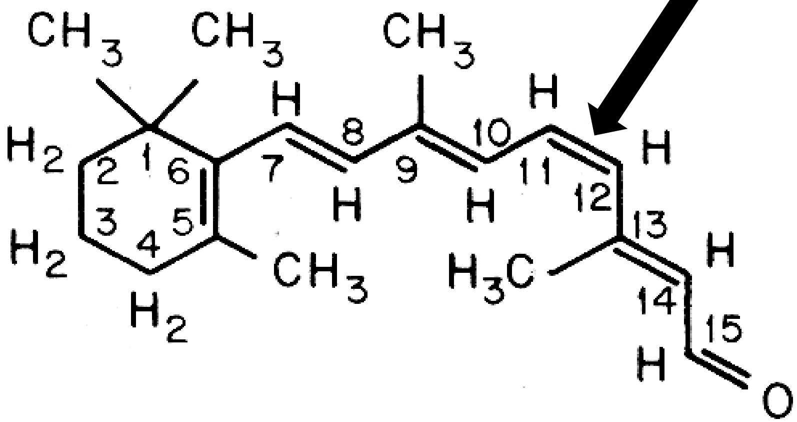

Chromophore

(*chromo-* colour, + *-phore,* producer)
Light-catching portion of any molecule

A photon is absorbed

11-*cis* retinal

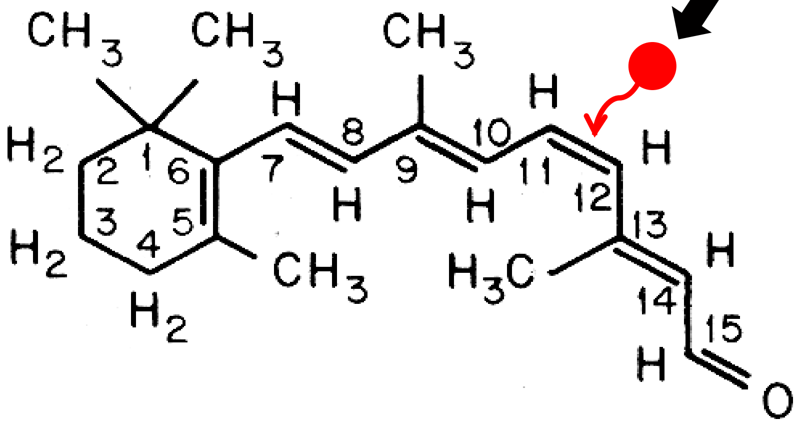

Chromophore

(*chromo*- colour, + *-phore,* producer)
Light-catching portion of any molecule

A photon is absorbed

11-*cis* retinal

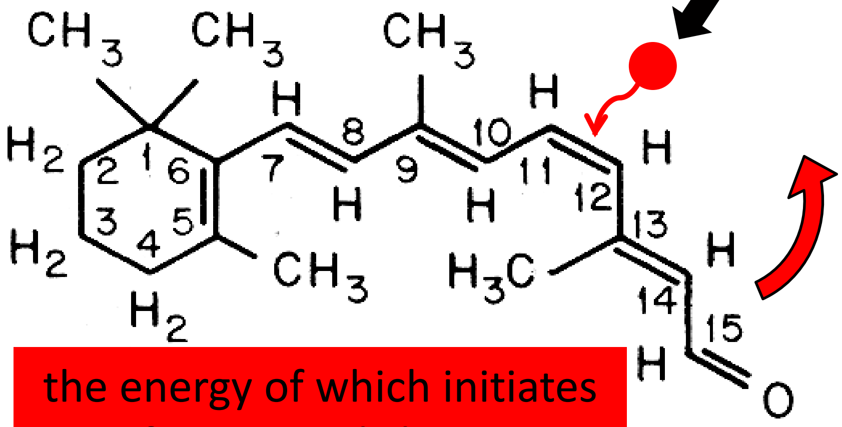

the energy of which initiates a conformational change to...

Chromophore

(*chromo-* colour, + *-phore,* producer)
Light-catching portion of any molecule

A photon is absorbed

11-*cis* retinal

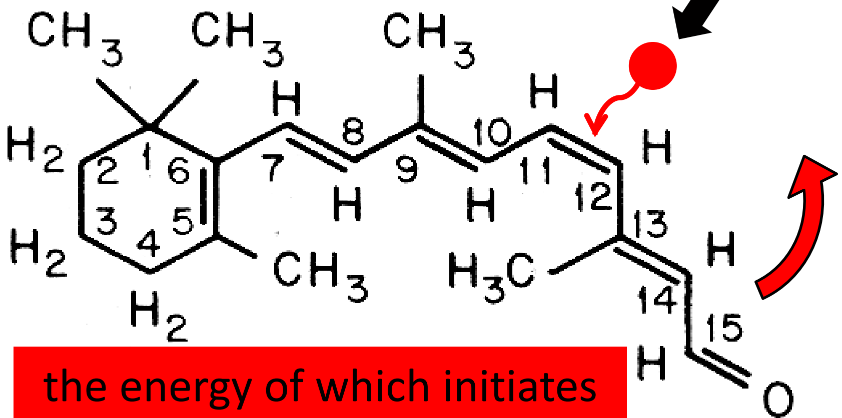

Think of this molecule as a photo-sensitive switch!

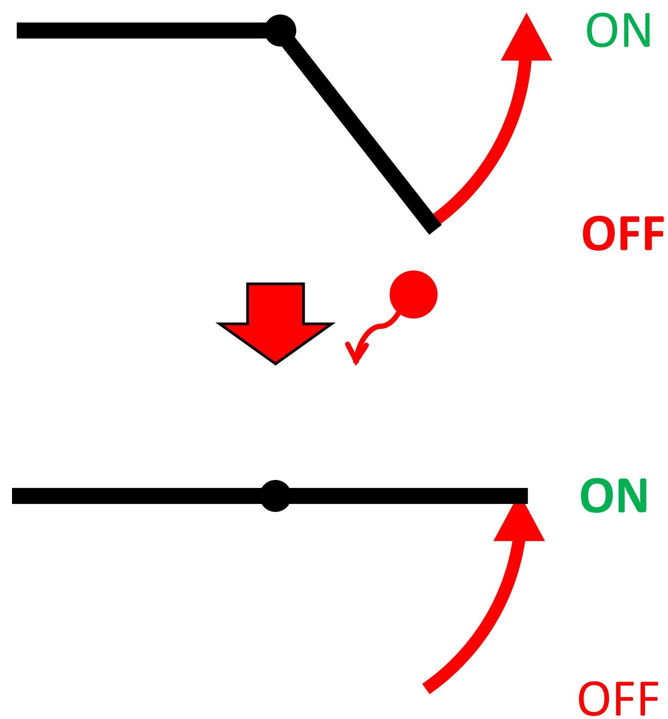

the energy of which initiates a conformational change to...

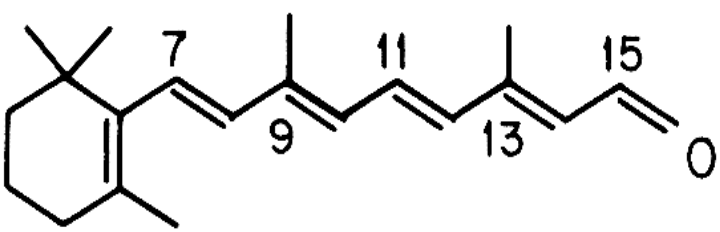

all-*trans* retinal

Chromophore

11-*cis* retinal

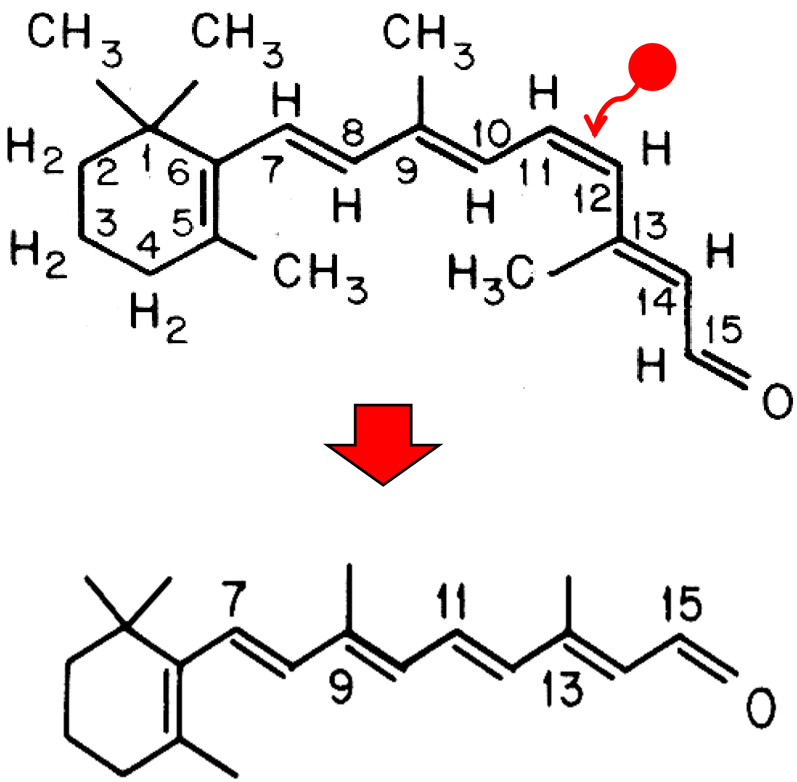

all-*trans* retinal

Crucially, the event is binary or “all or nothing”.

If a photon is absorbed it has the same effect as any other absorbed photon, whatever its wavelength.

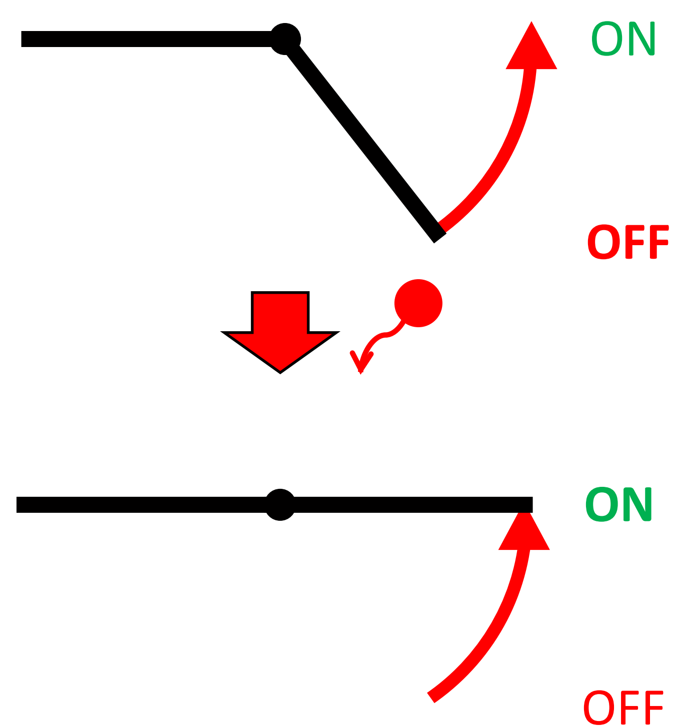

Chromophore

11-*cis* retinal

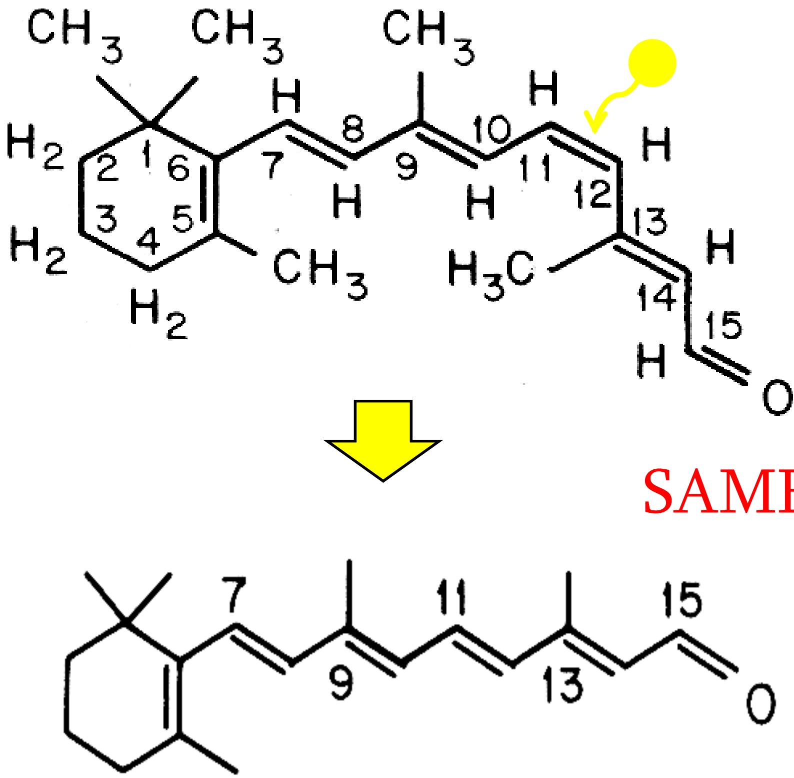

all-*trans* retinal

Crucially, the event is binary or "all or nothing".

If a photon is absorbed it has the same effect as any other absorbed photon, whatever its wavelength.

SAME EFFECT

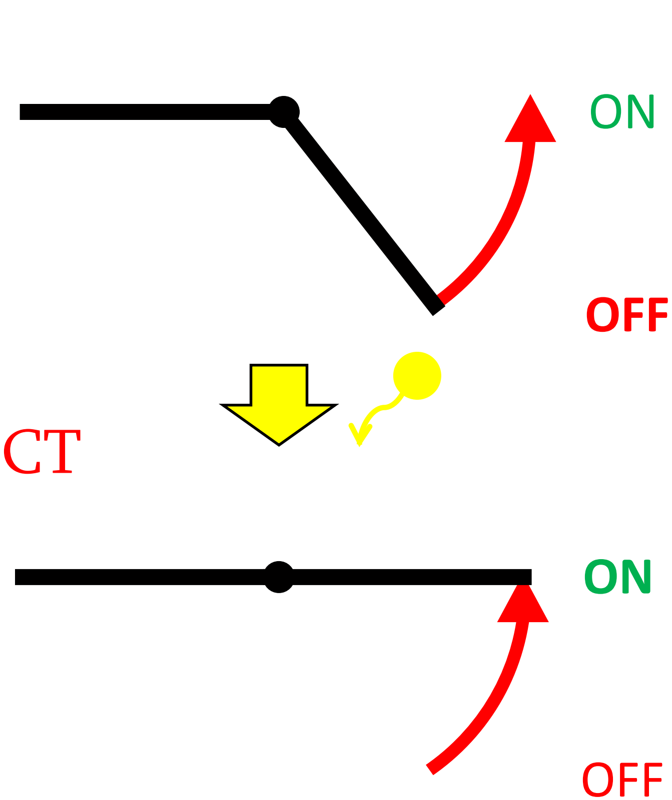

Chromophore

Crucially, the event is binary or “all or nothing”.

If a photon is absorbed it has the same effect as any other absorbed photon, whatever its wavelength.

11-*cis* retinal

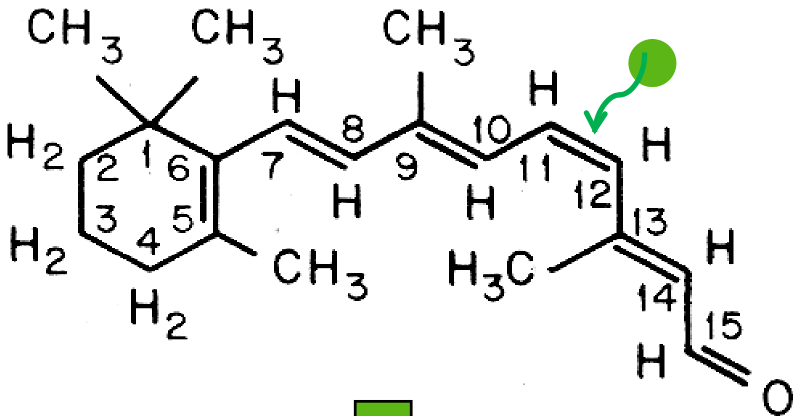

SAME EFFECT

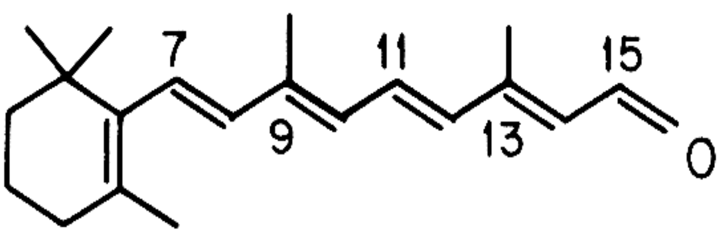

all-*trans* retinal

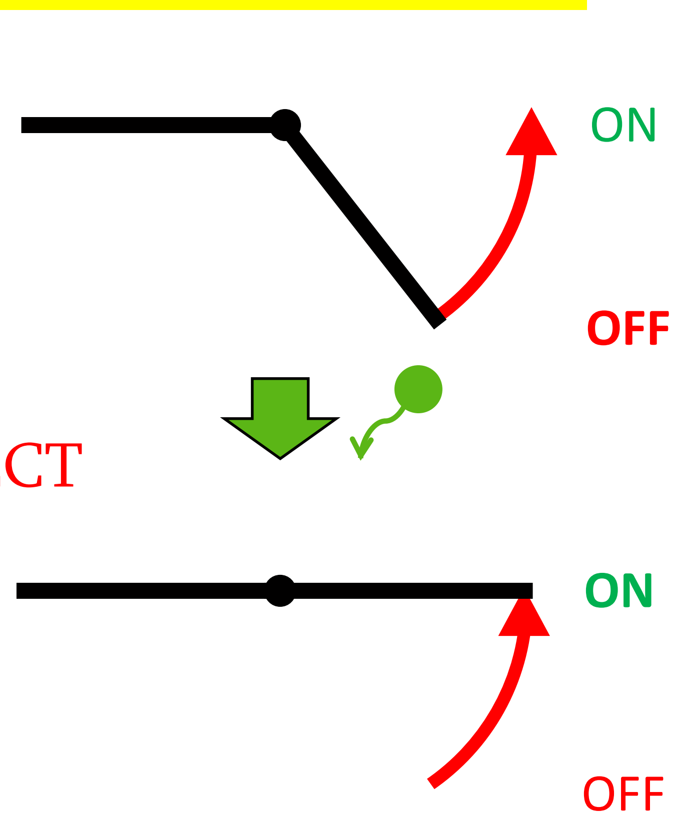

Chromophore

Crucially, the event is binary or "all or nothing".

If a photon is absorbed it has the same effect as any other absorbed photon, whatever its wavelength.

11-*cis* retinal

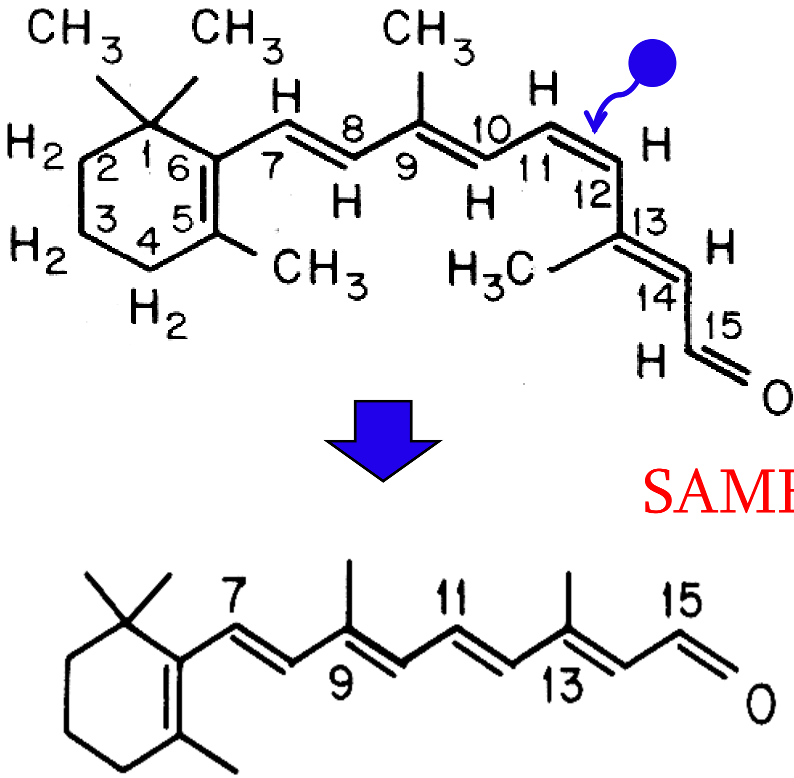

all-*trans* retinal

SAME EFFECT

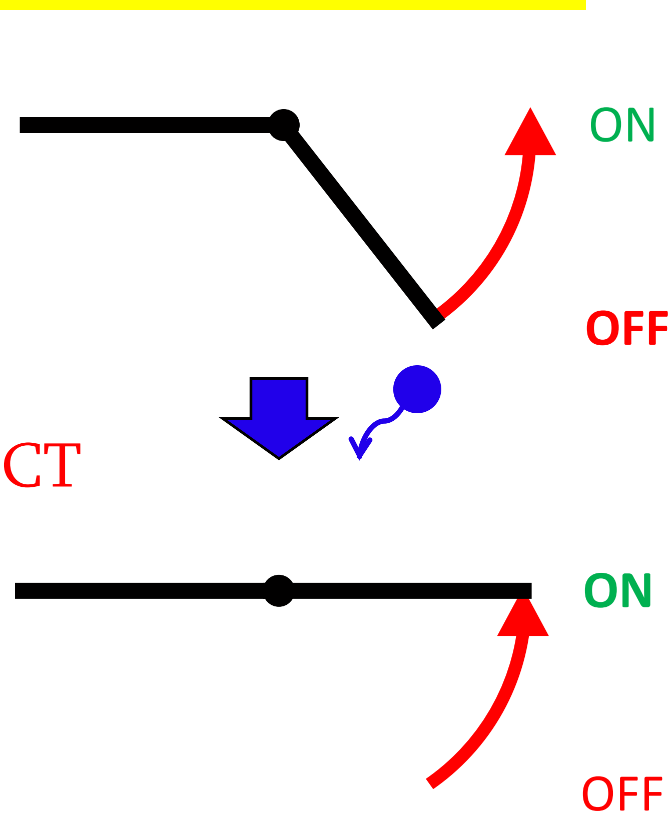

Chromophore

This process cannot encode wavelength (colour)!

It is “UNIVARIANT”

11-*cis* retinal

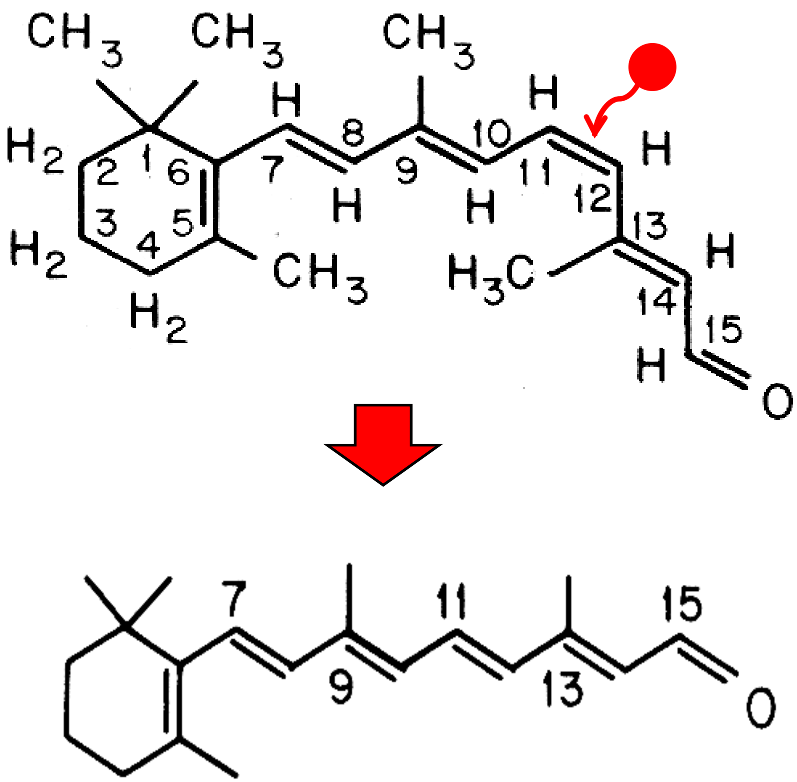

all-*trans* retinal

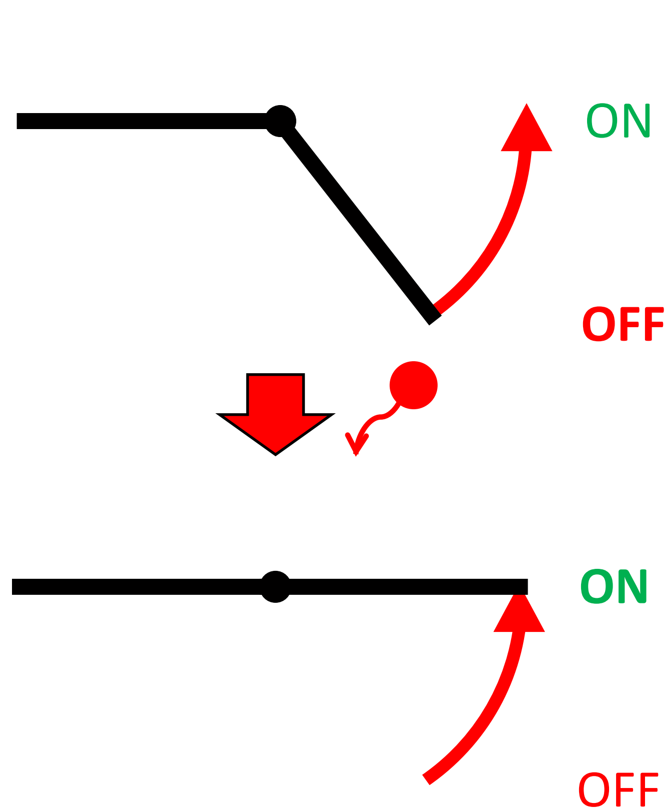

UNIVARIANCE

Once absorbed, all these photons…

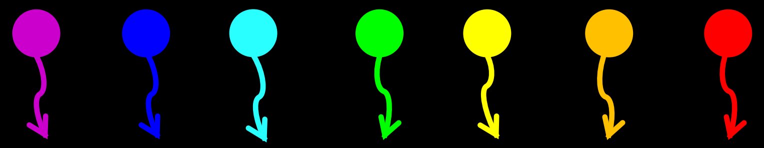

have the same effect.

UNIVARIANCE

What does vary with wavelength is the probability that a photon will be absorbed.

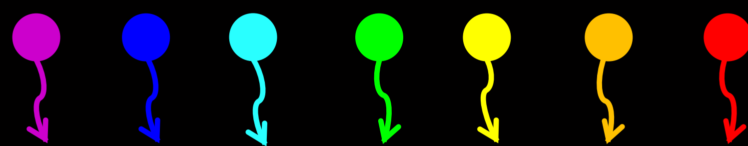

This is reflected in what is called the cone “spectral sensitivity function”, an example of which is the middle-wavelength-sensitive (M-) cone function…

IN TERMS OF A CAMERA, THIS IS A SENSOR SPECTRAL SENSITIVITY.

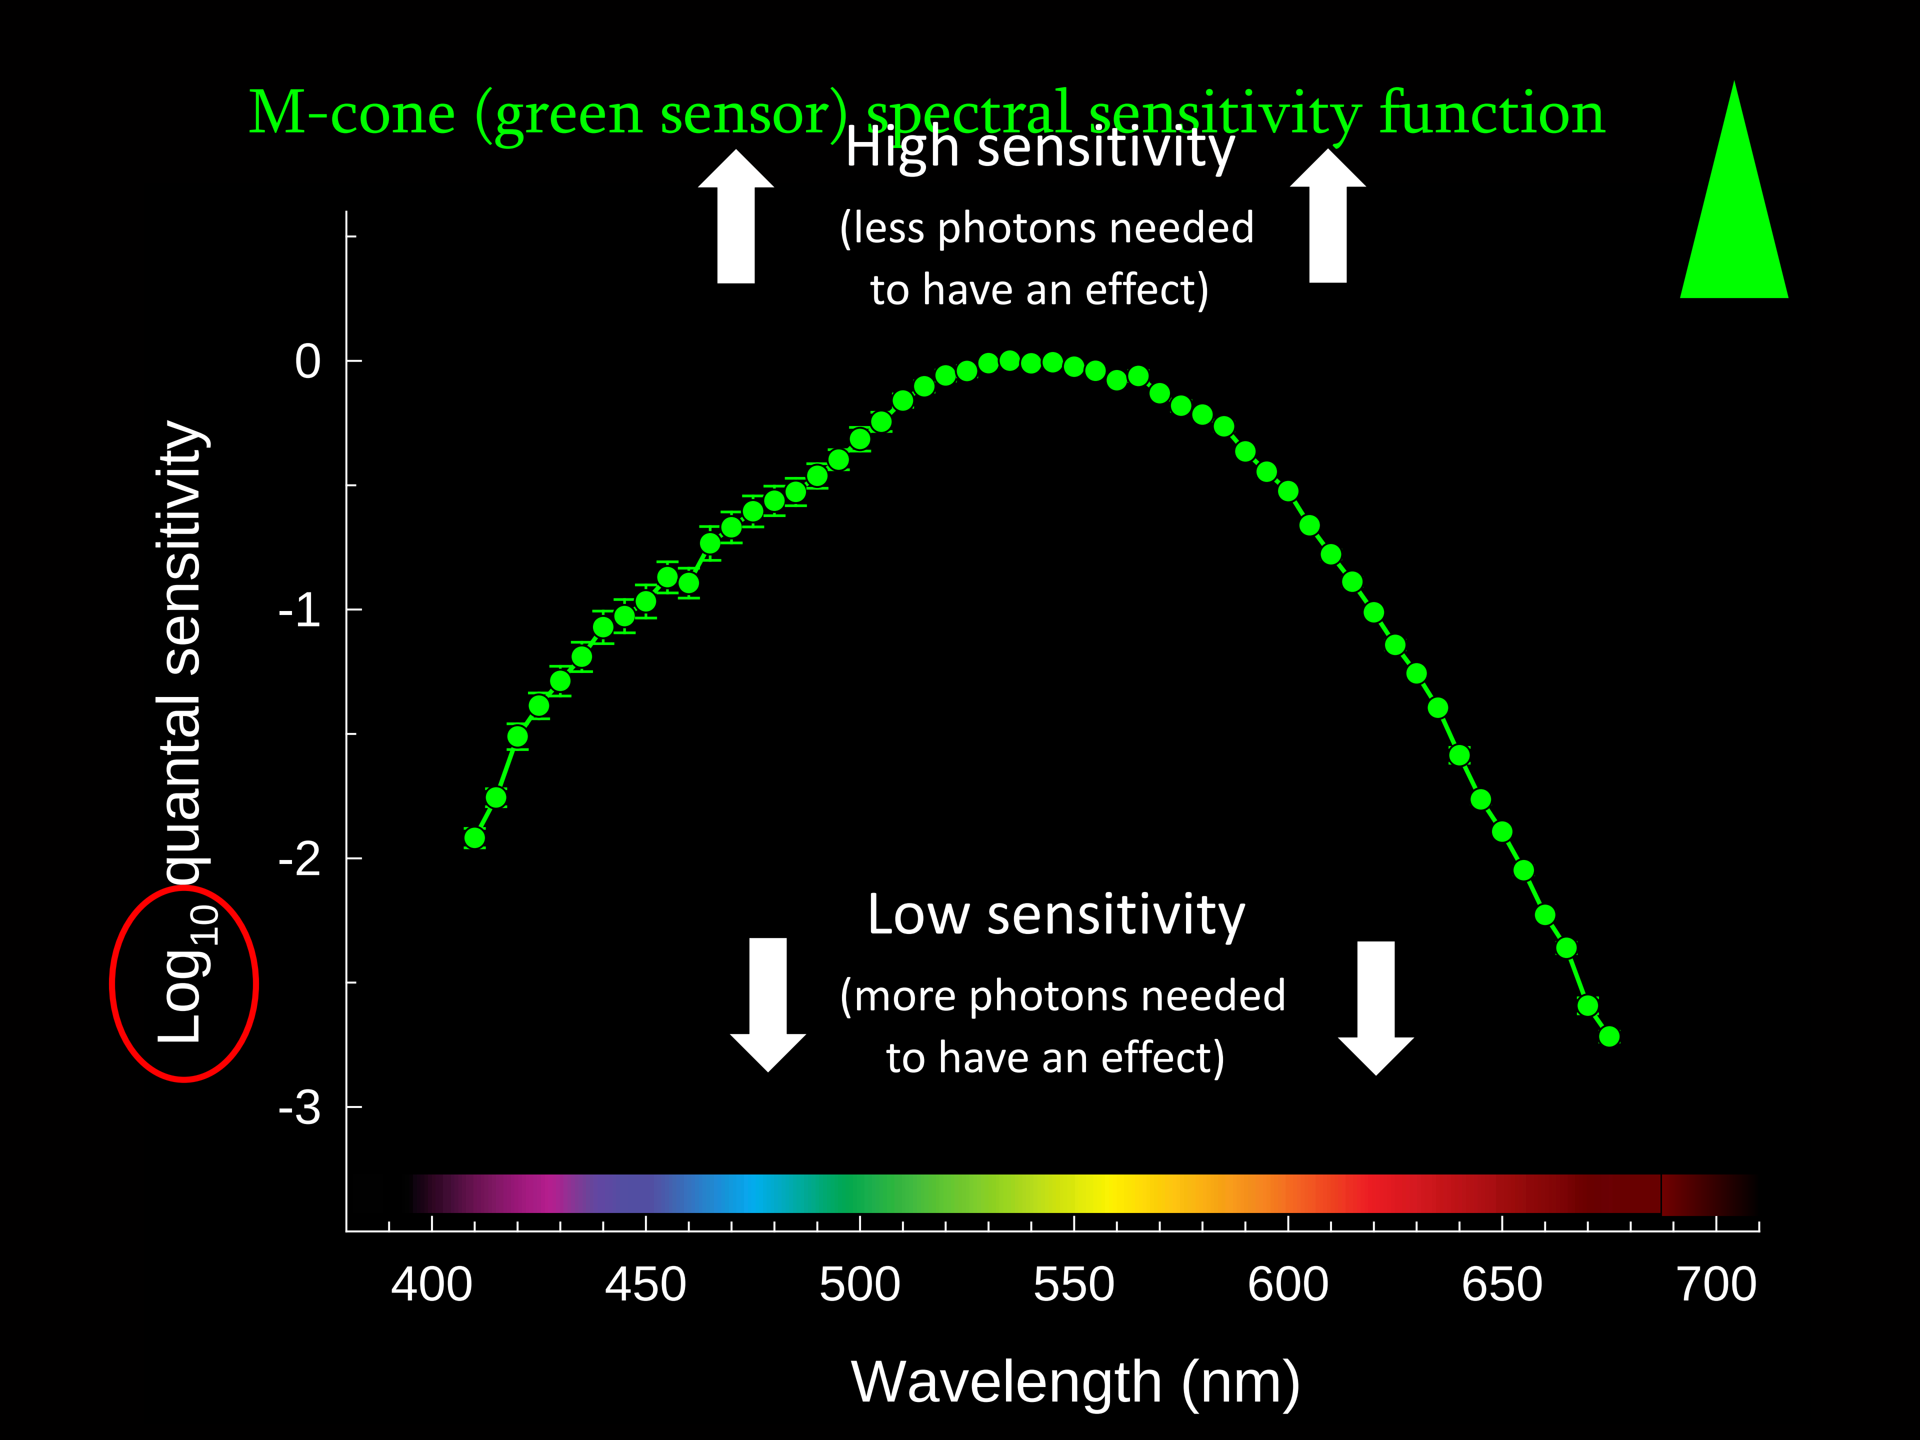
M-cone (green sensor) spectral sensitivity function
High sensitivity
(less photons needed
to have an effect)
Low sensitivity
(more photons needed
to have an effect)
Log_{10} quantal sensitivity
0
-1
-2
-3
400
450
500
550
600
650
700
Wavelength (nm)

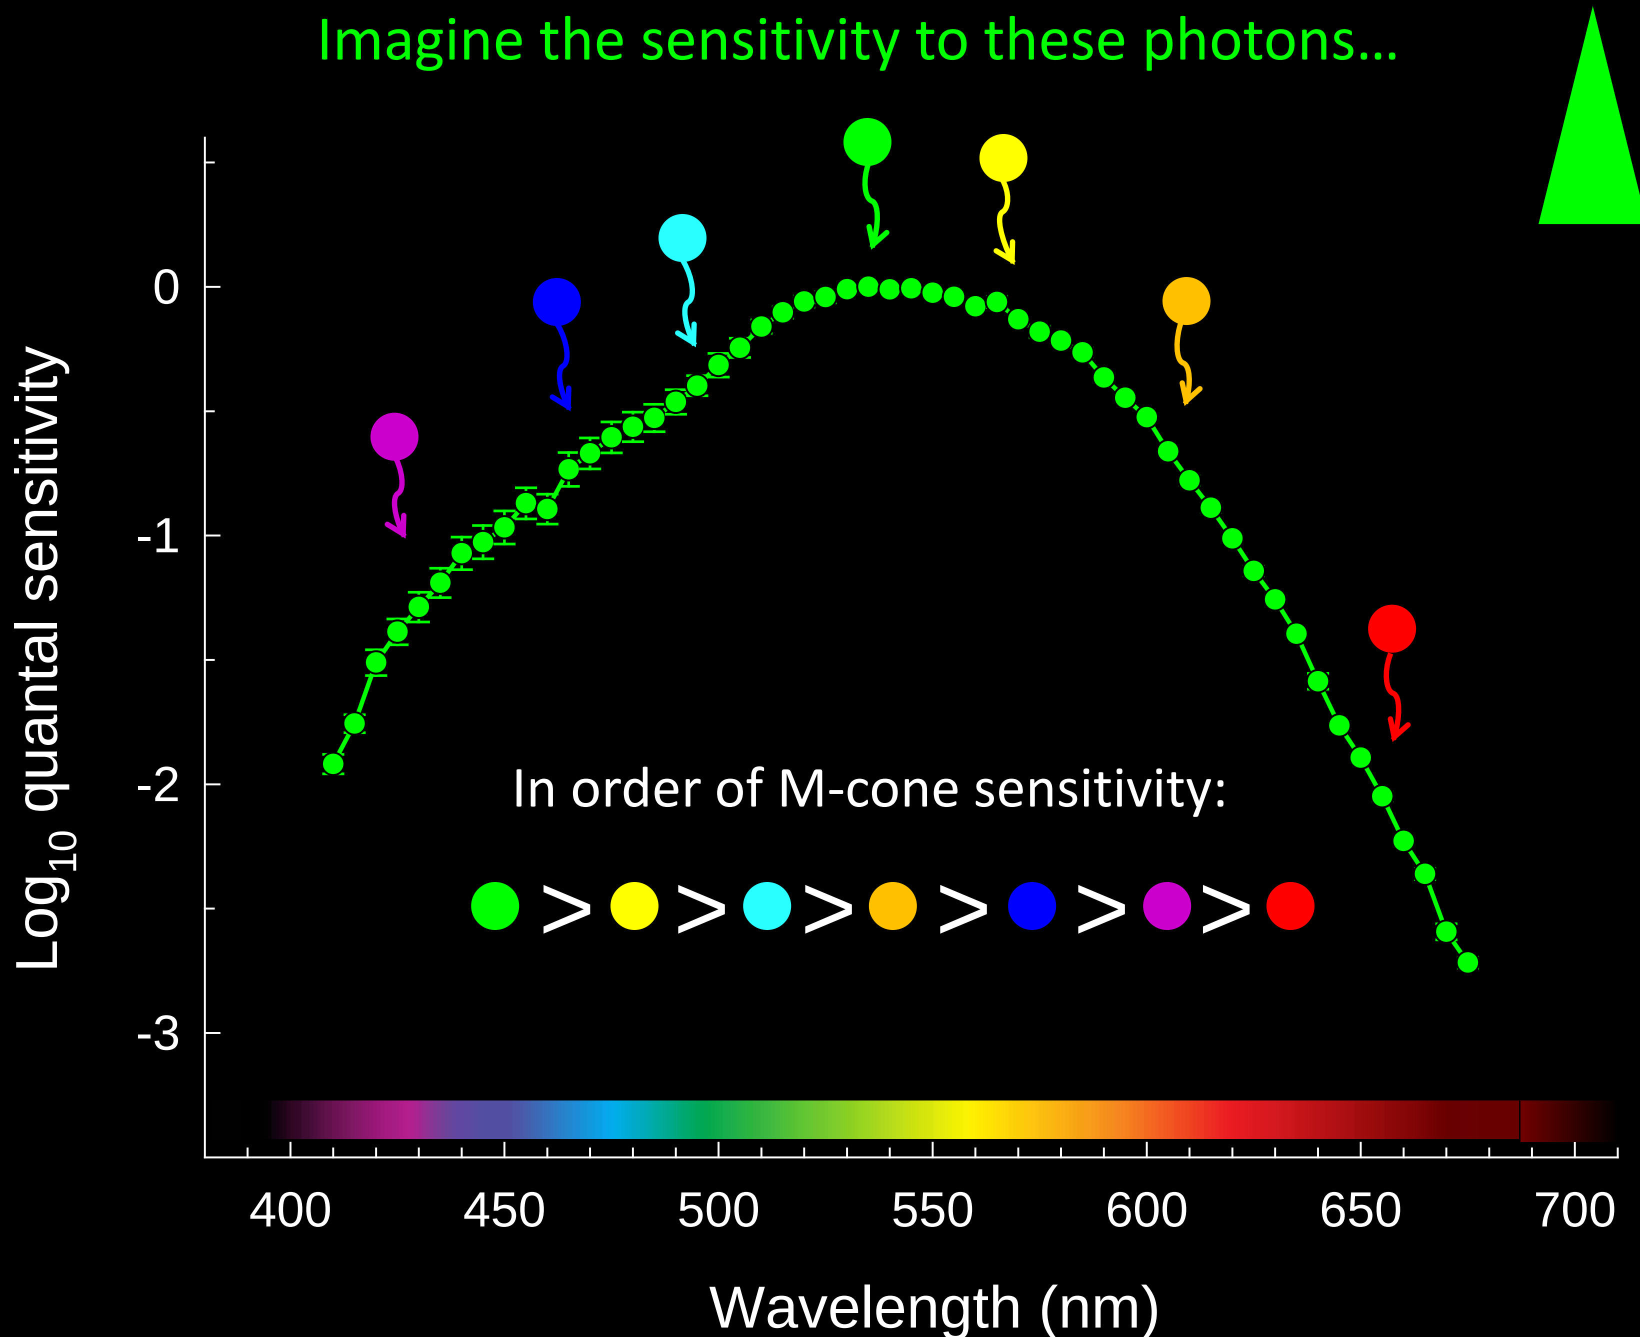

Imagine the sensitivity to these photons...
0
-1
-2
-3
Log10 quantal sensitivity
400
450
500
550
600
650
700
Wavelength (nm)
In order of M-cone sensitivity:
> > > > > >

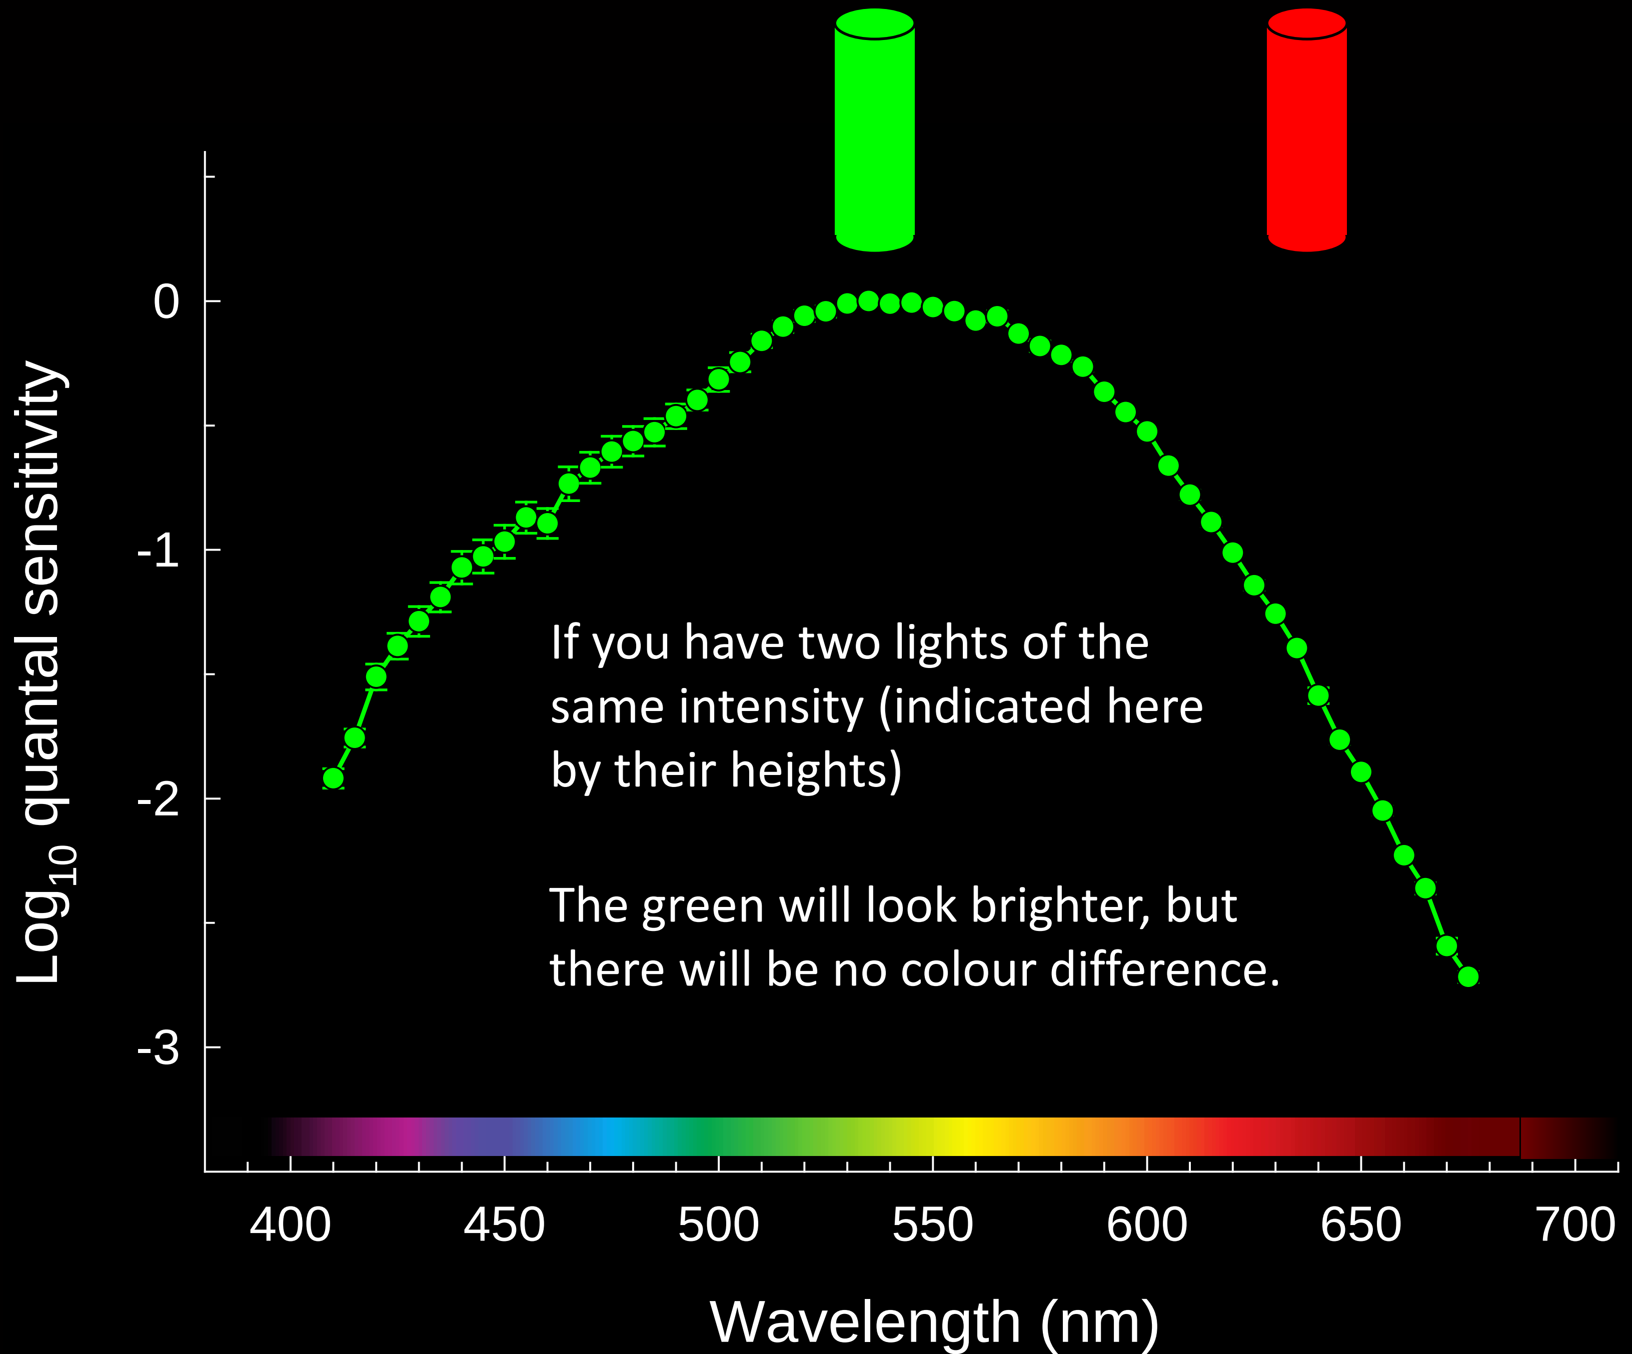

0
-1
-2
-3
Log10 quantal sensitivity
400
450
500
550
600
650
700
Wavelength (nm)
If you have two lights of the same intensity (indicated here by their heights)
The green will look brighter, but there will be no colour difference.

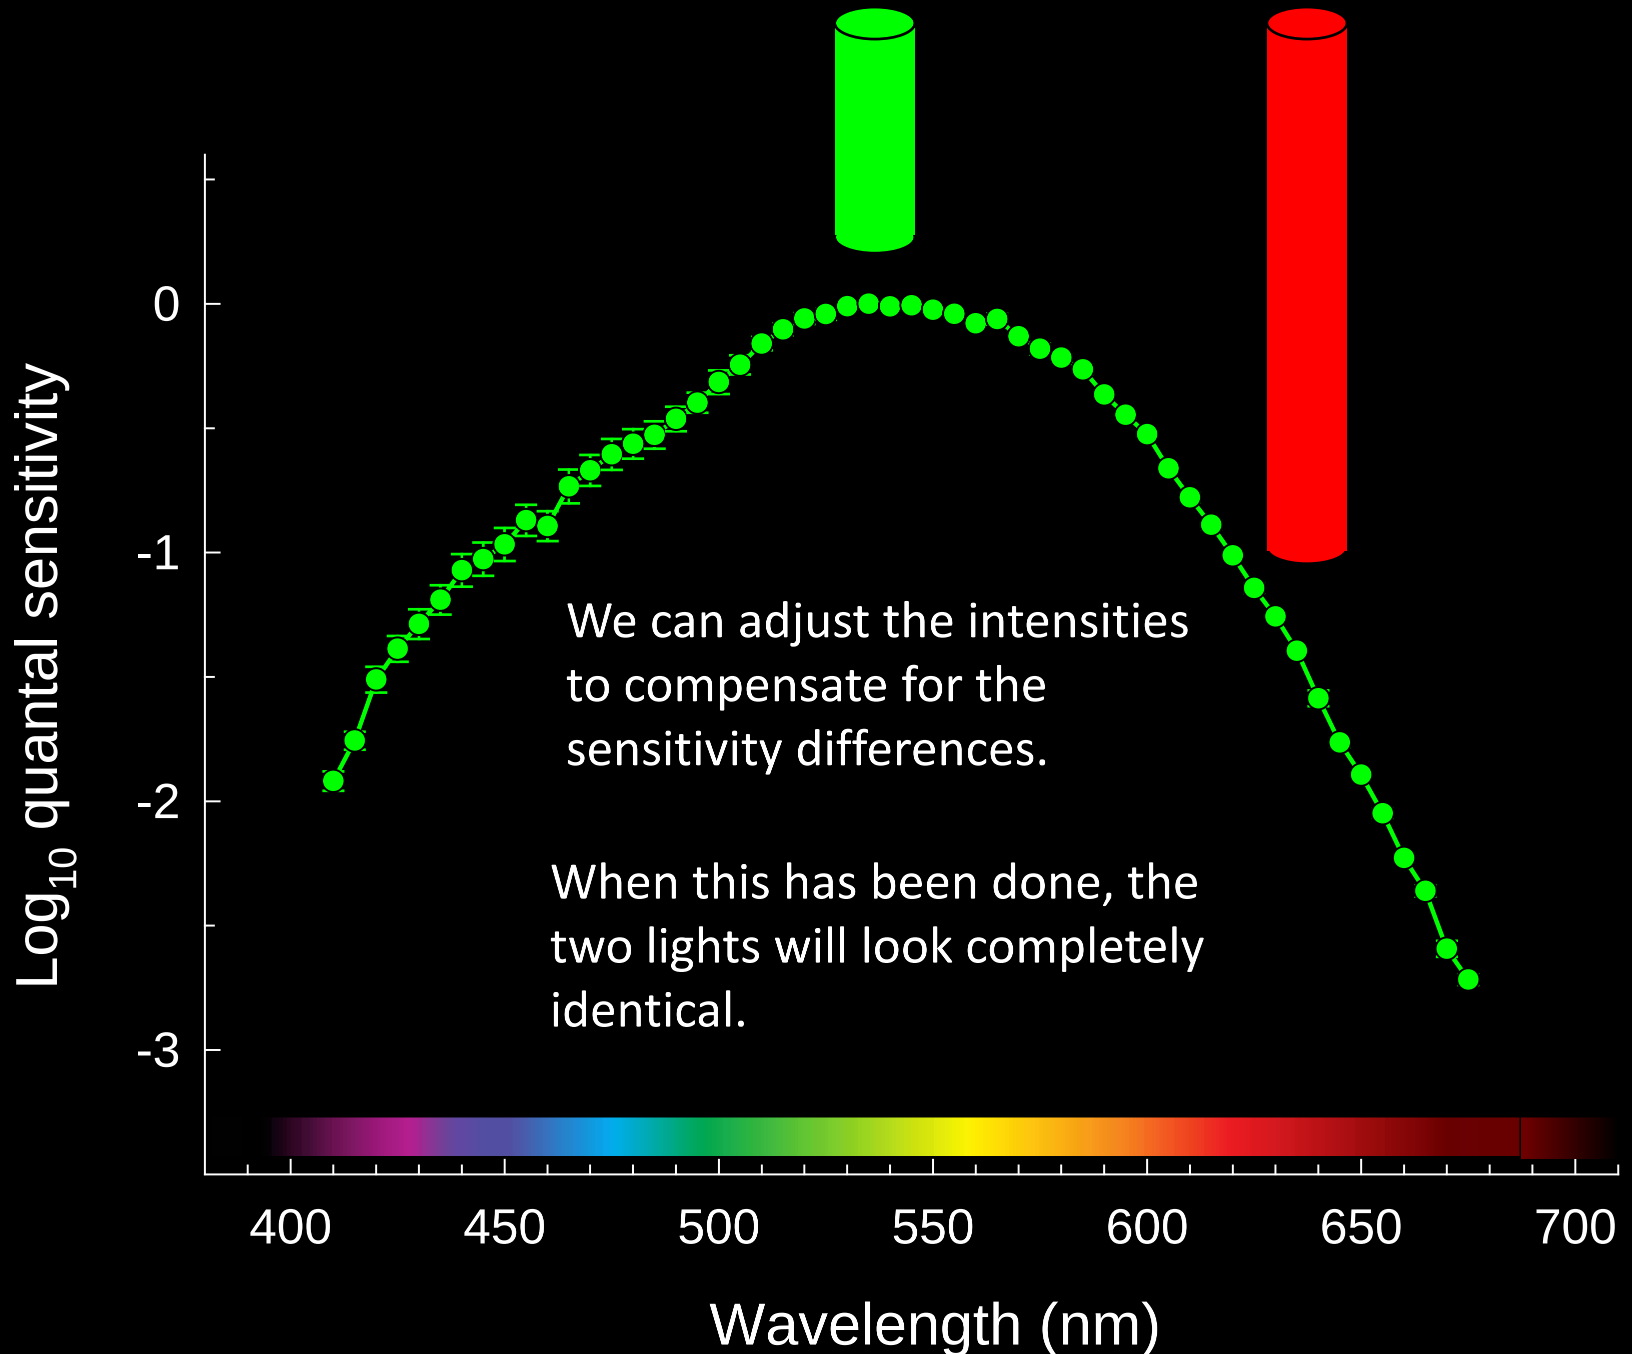

0
-1
-2
-3
Log10 quantal sensitivity
400
450
500
550
600
650
700
Wavelength (nm)
We can adjust the intensities to compensate for the sensitivity differences.
When this has been done, the two lights will look completely identical.

Thus, if we had only one photoreceptor, we would be colour-blind…

Examples: night vision (rod achromatopsia)

With two, we
are dichromatic:

Protanopia (missing red sensor)

Tritanopia (missing blue sensor)

Deuteranopia (missing green sensor)

Simulations from Sharpe, Stockman, Jägle & Nathans, 1999

And with three, we enjoy trichromacy:

So, if each photoreceptor or sensor is colour-blind (univariant), how do we see colour?

Or to put it another way: How is colour encoded at the input to the visual system?

Colour is encoded by the relative cone outputs

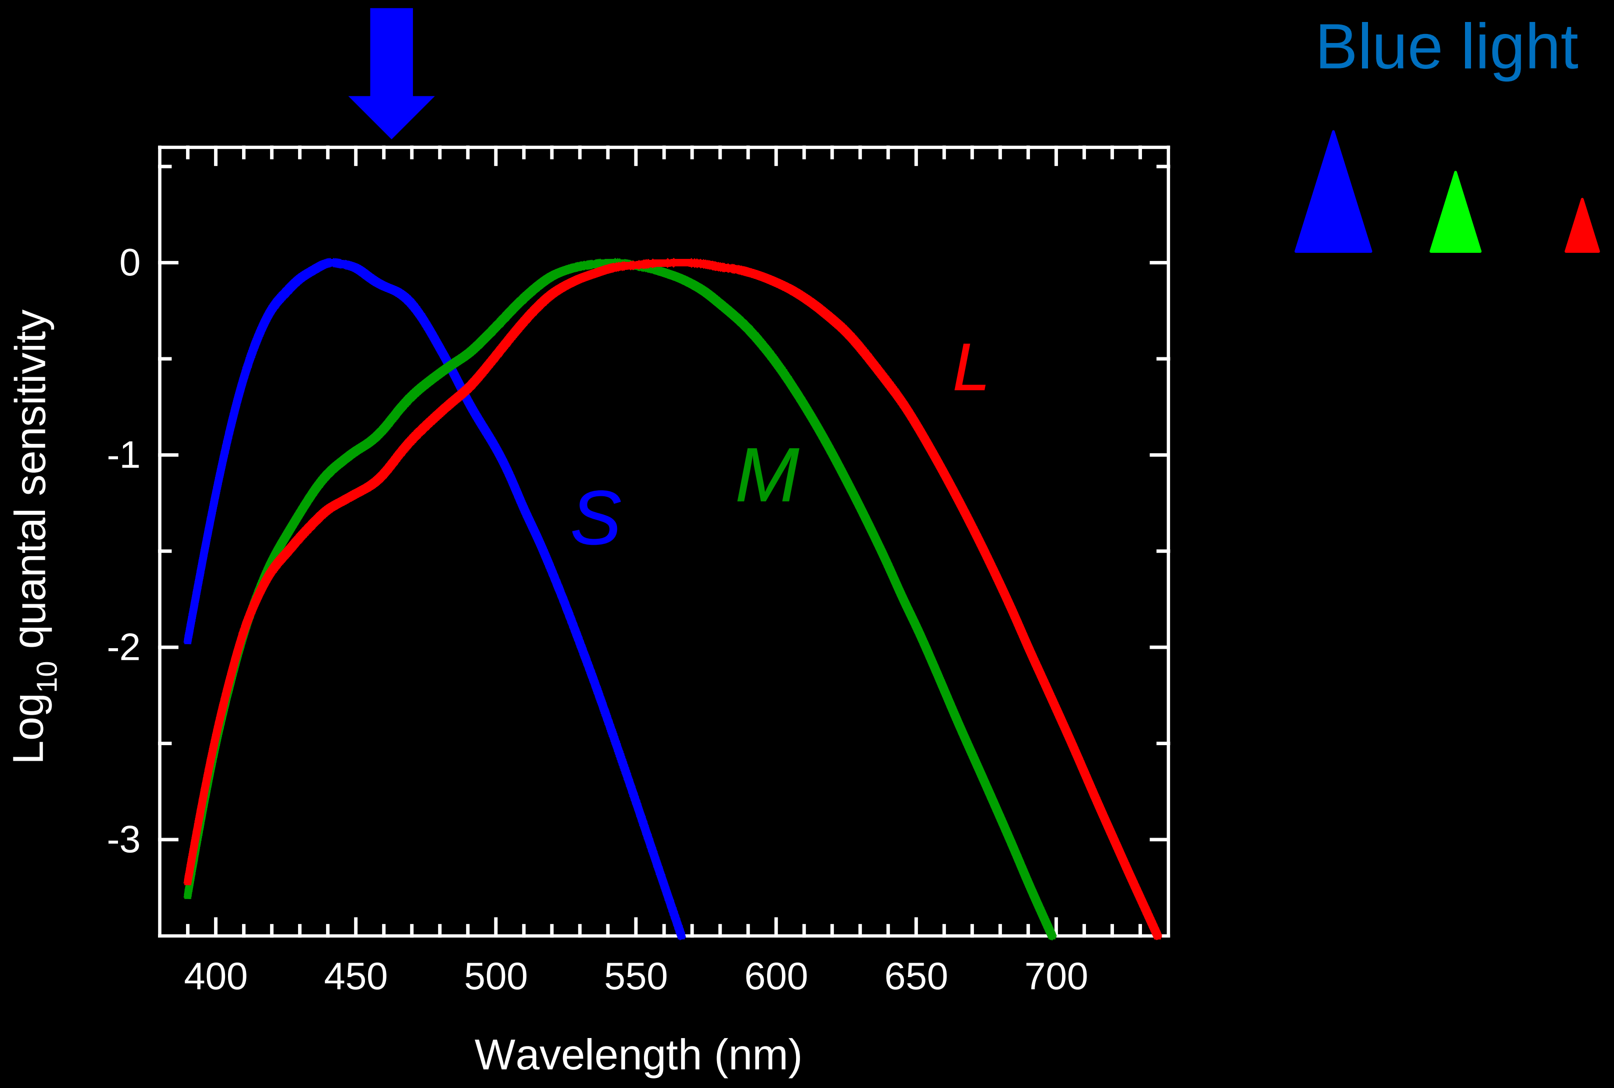

Colour is encoded by the relative cone outputs

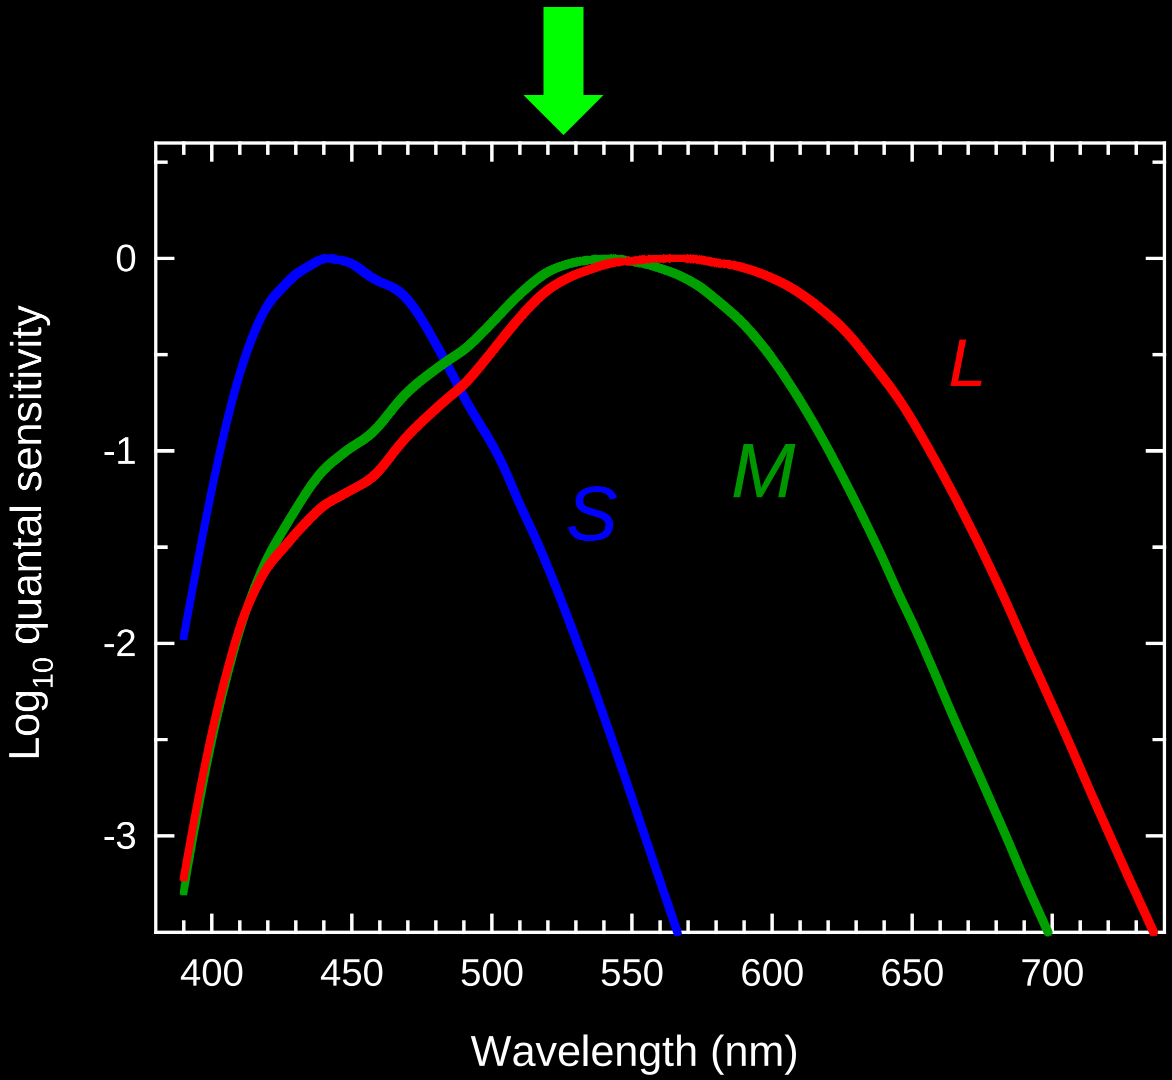

Blue light

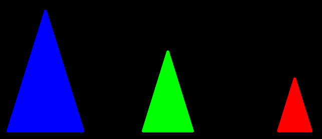

Green light

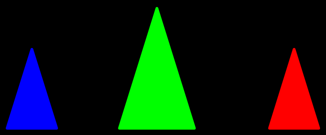

Colour is encoded by the relative cone outputs

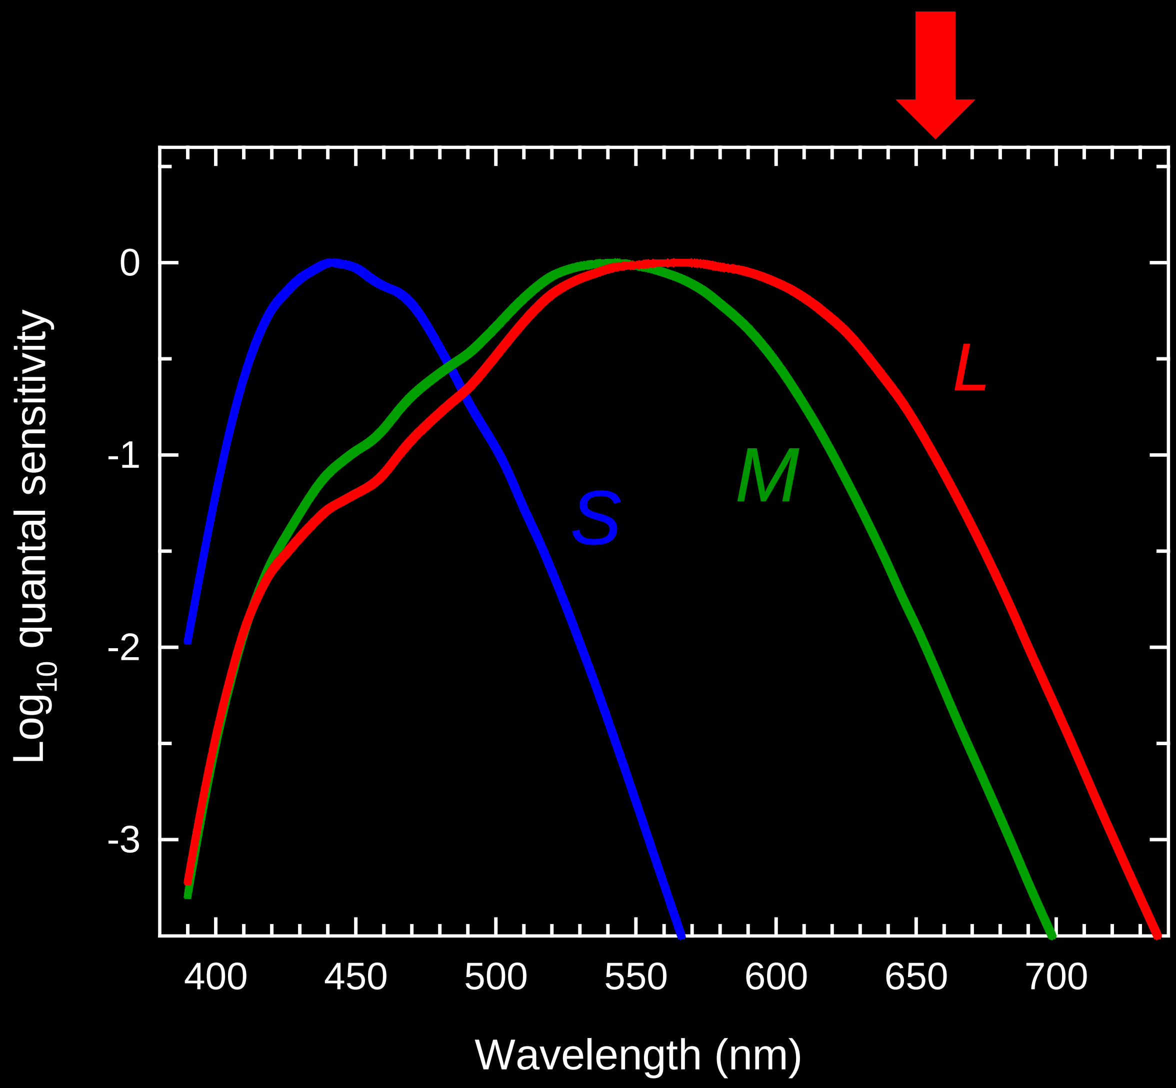

Blue light

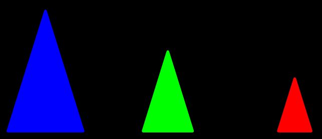

Green light

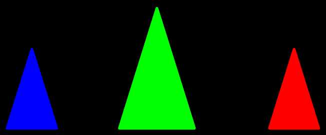

Red light

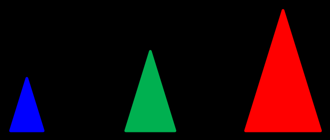

Colour is encoded by the relative cone outputs

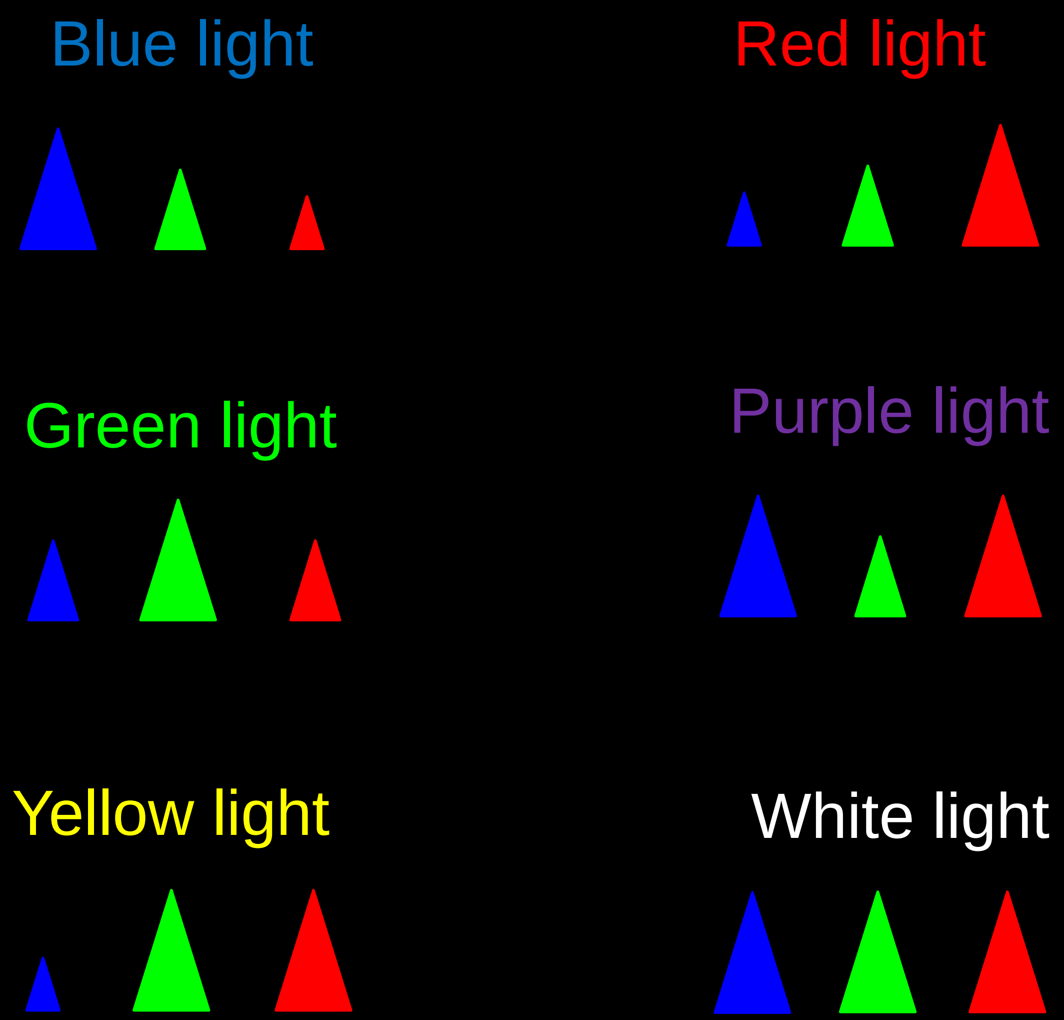

TRICHROMACY

A change in colour from green to red causes a relative increase in the L-cone output but causes a decrease in the M-cone output.

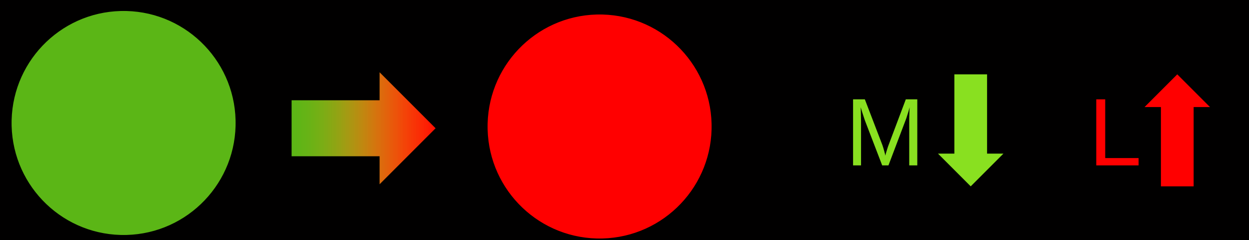

A change in colour from red to green causes a relative increase in the M-cone output but causes a decrease in the L-cone output.

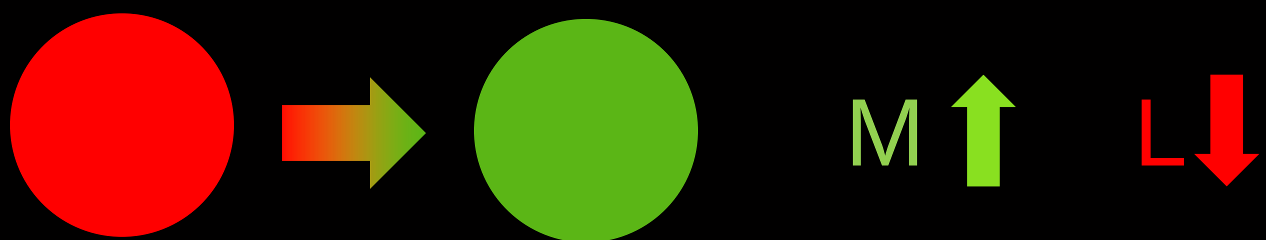

Thus, colour is encoded by *comparing* the outputs of different cone sensors…

The Stockman & Sharpe (2000) cone sensor spectral sensitivities:

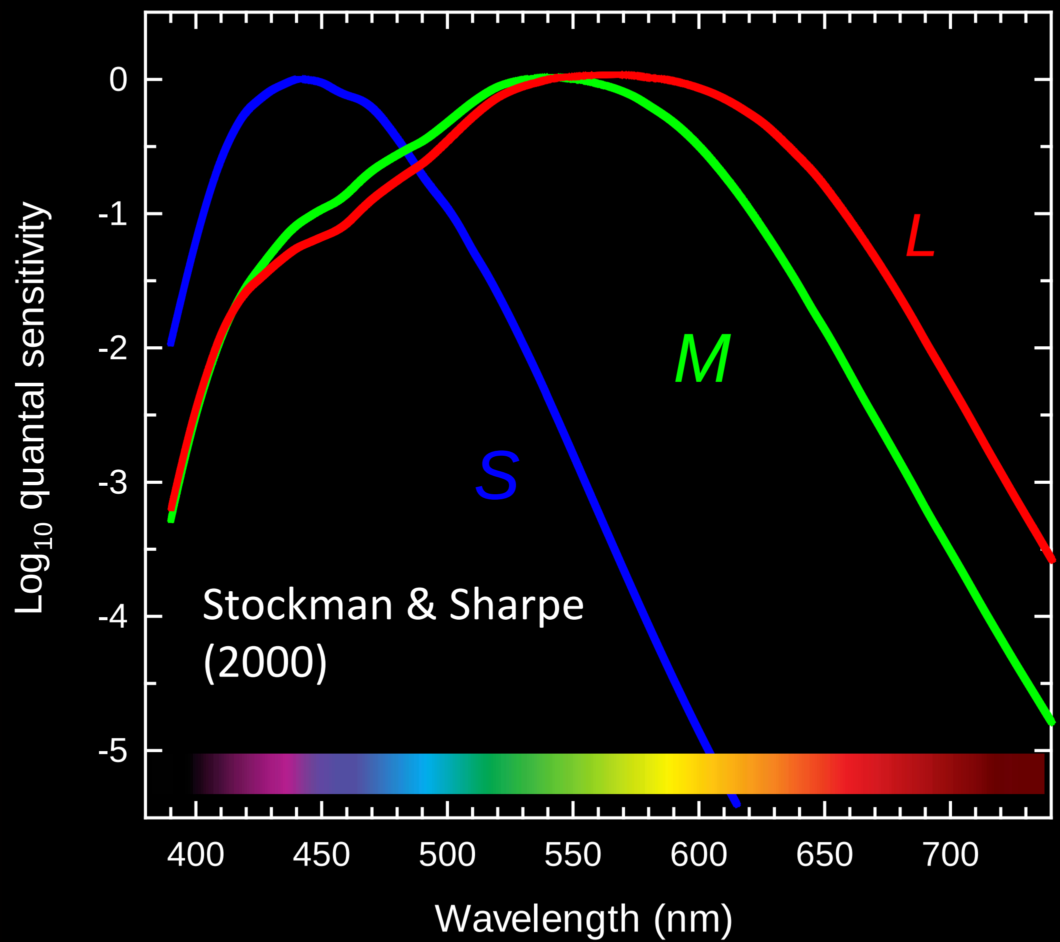

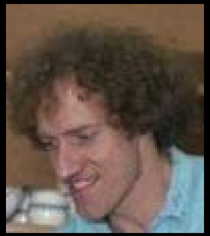

are the recent physiologically relevant LMS CIE (2006) color matching functions.

If we know these three spectral sensitivities, and thus the effects that lights have on the three cones, we can completely specify those lights.

Defining colours in a 3-dimensional cone sensor colour space

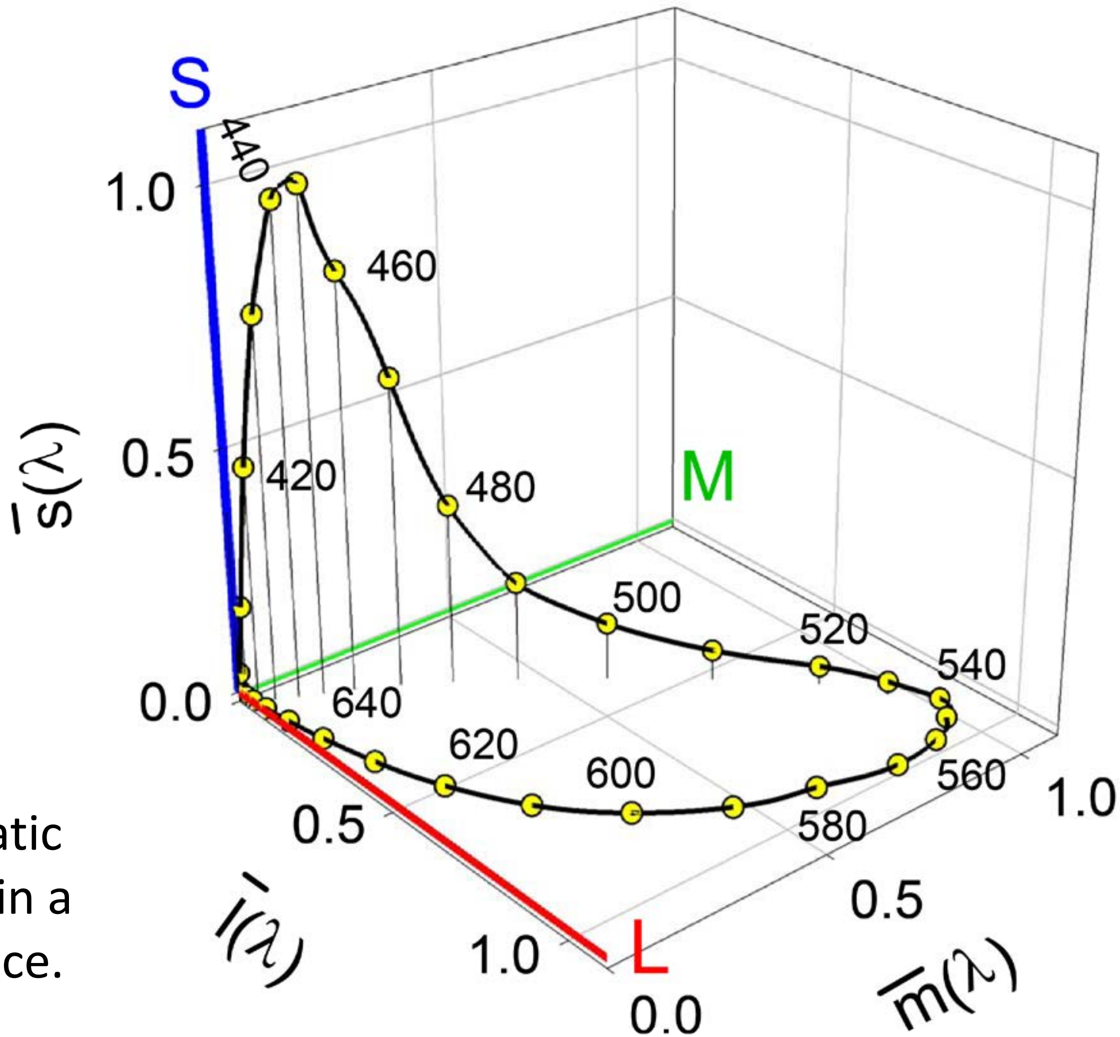

SPECTRUM LOCUS:
Plot of monochromatic "spectral" lights (as in a rainbow) in LMS space.

Defining colours in a 2-dimensional cone sensor colour space (a plane of 3-D space)

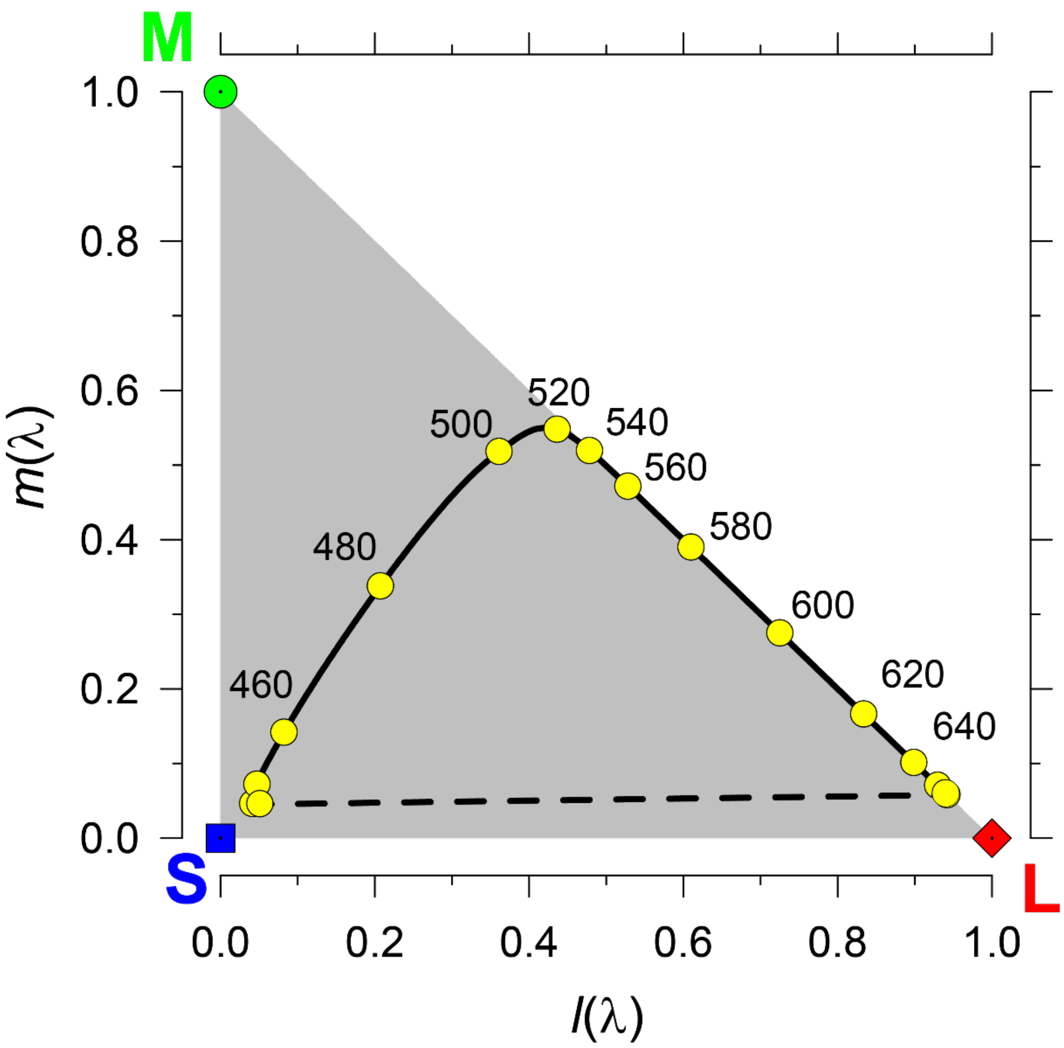

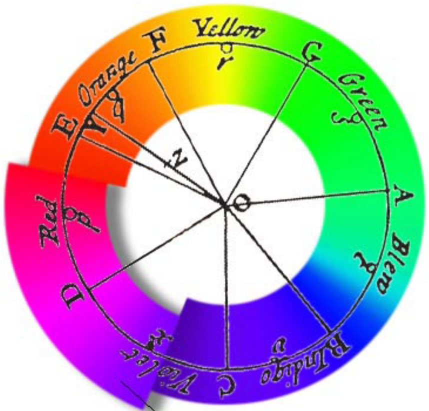

extraspectral hues mixed by overlapping opposite spectrum ends

2-dimensional CIE XYZ colour space

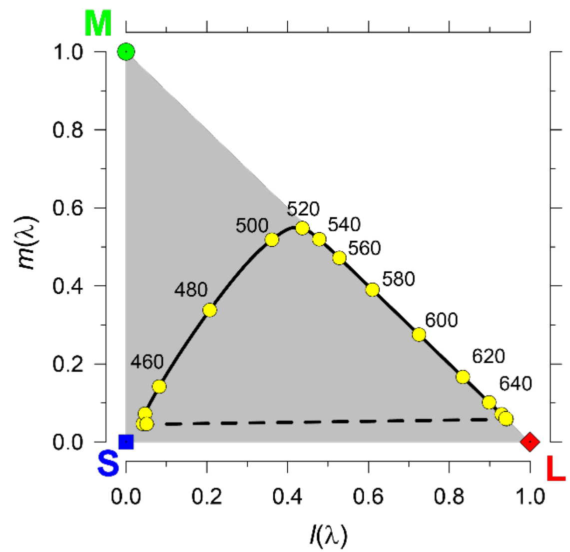

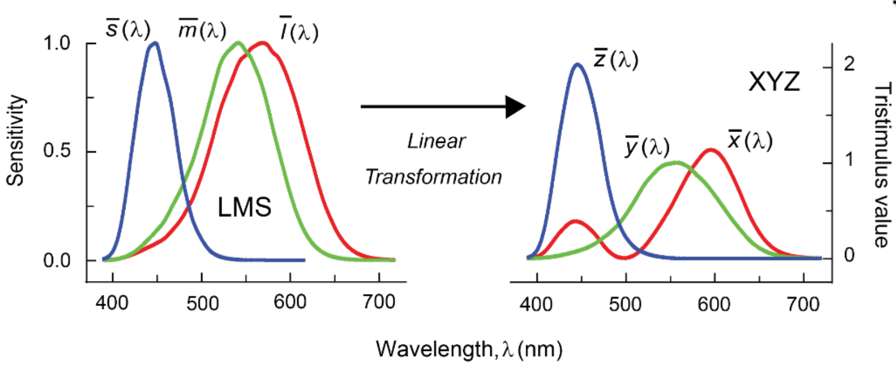

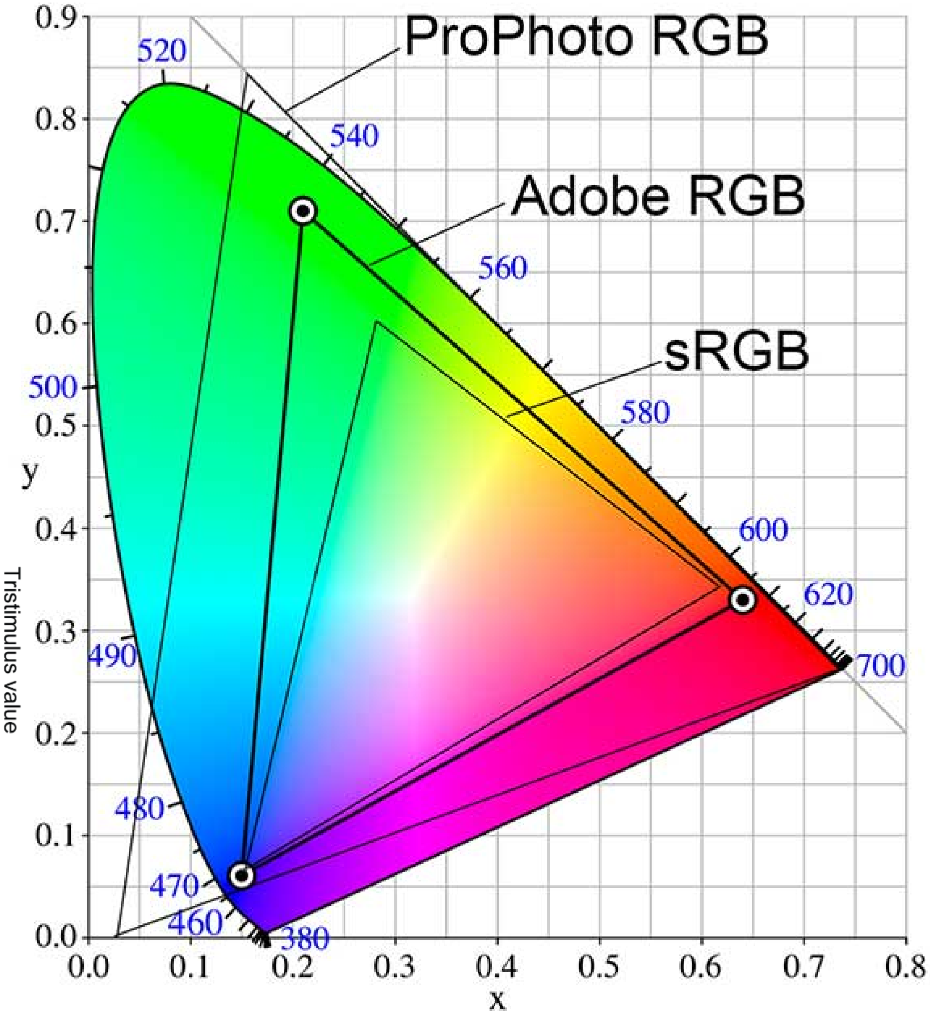

Simple linear transformation from LMS sensors to XYZ sensors

Photo Review

But the xy diagram is only a plane of a 3-dimensional XYZ sensor colour space.

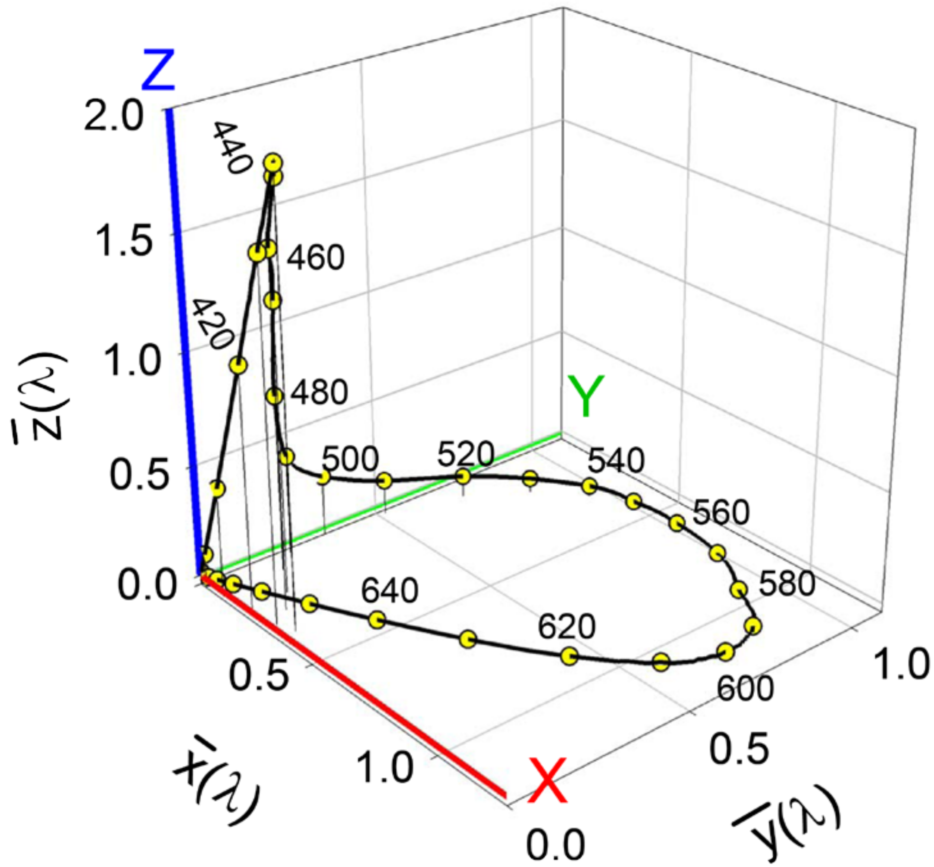

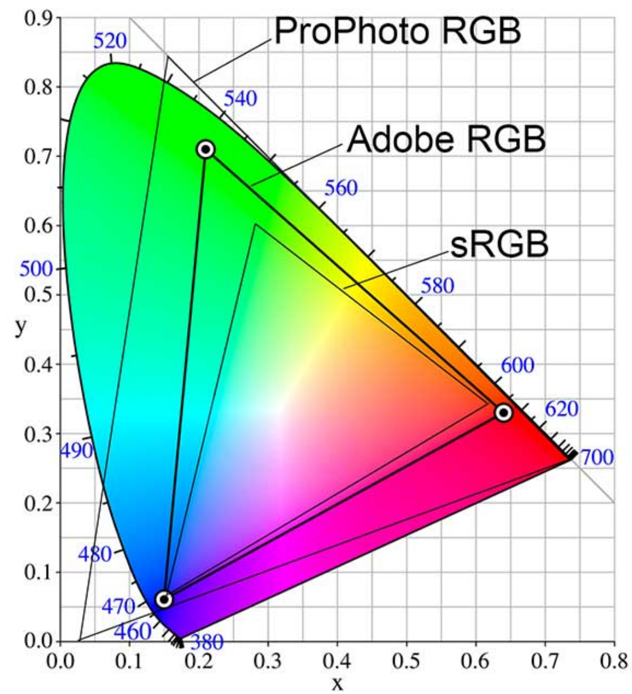

3. COLOUR AFTER THE SENSORS

How is colour encoded after the cones?

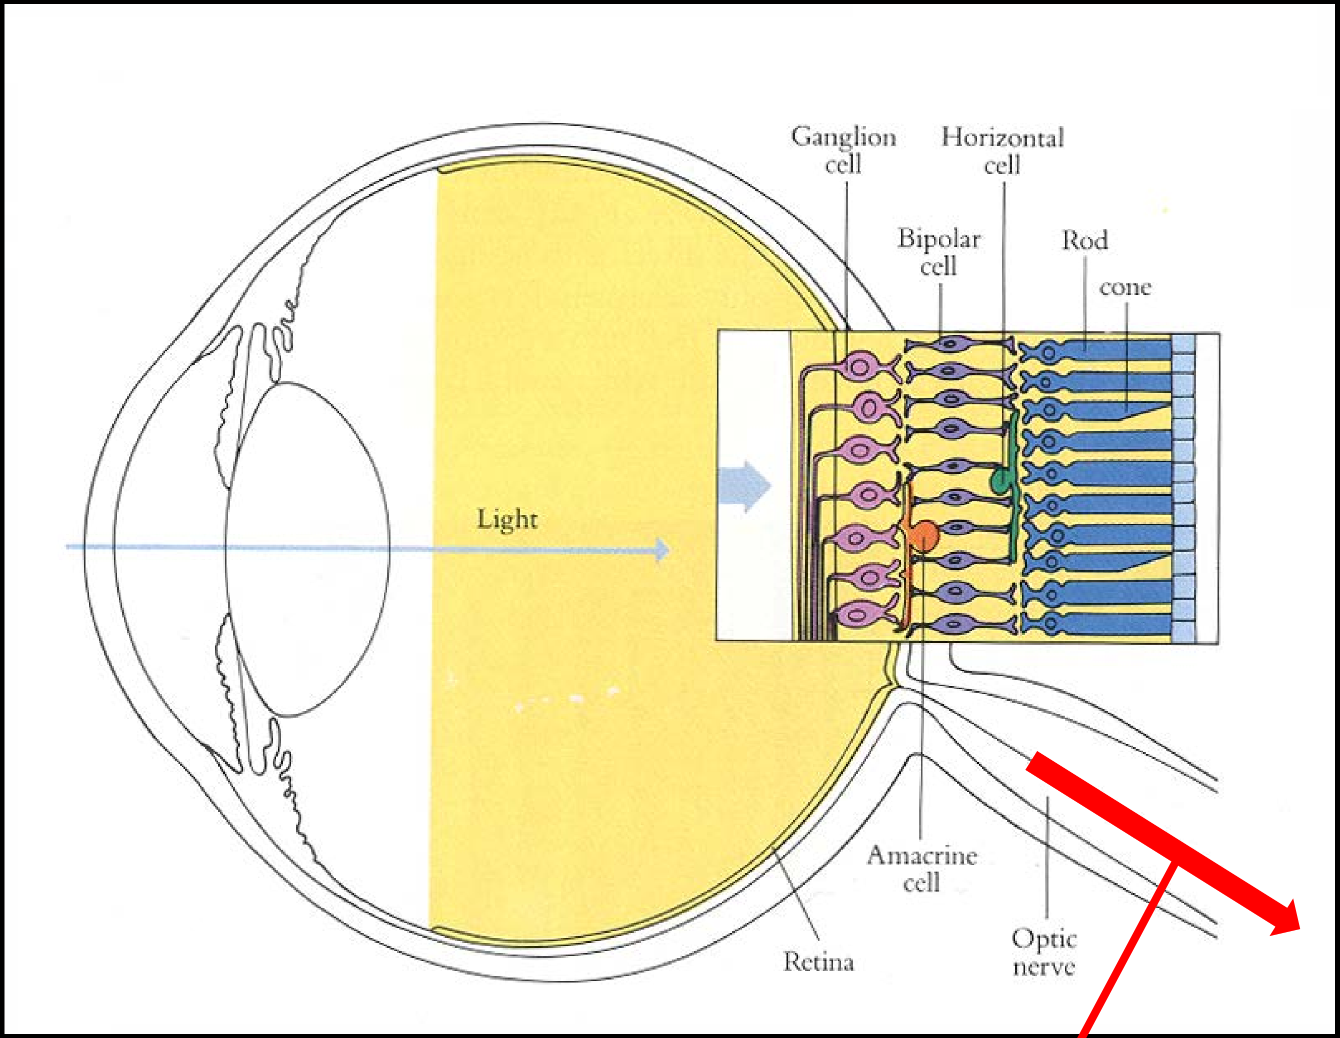

Before the signals are transmitted to the brain

Eye, Brain & Vision

Colour phenomenology

Can provide clues about how colours are processed after the photoreceptors…

Imagine a single patch of colour inside a dark surround

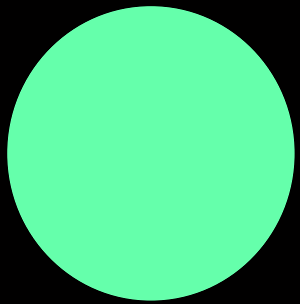

- Which pairs of colours coexist in a single, uniform patch of colour?
- Which pairs never coexist?

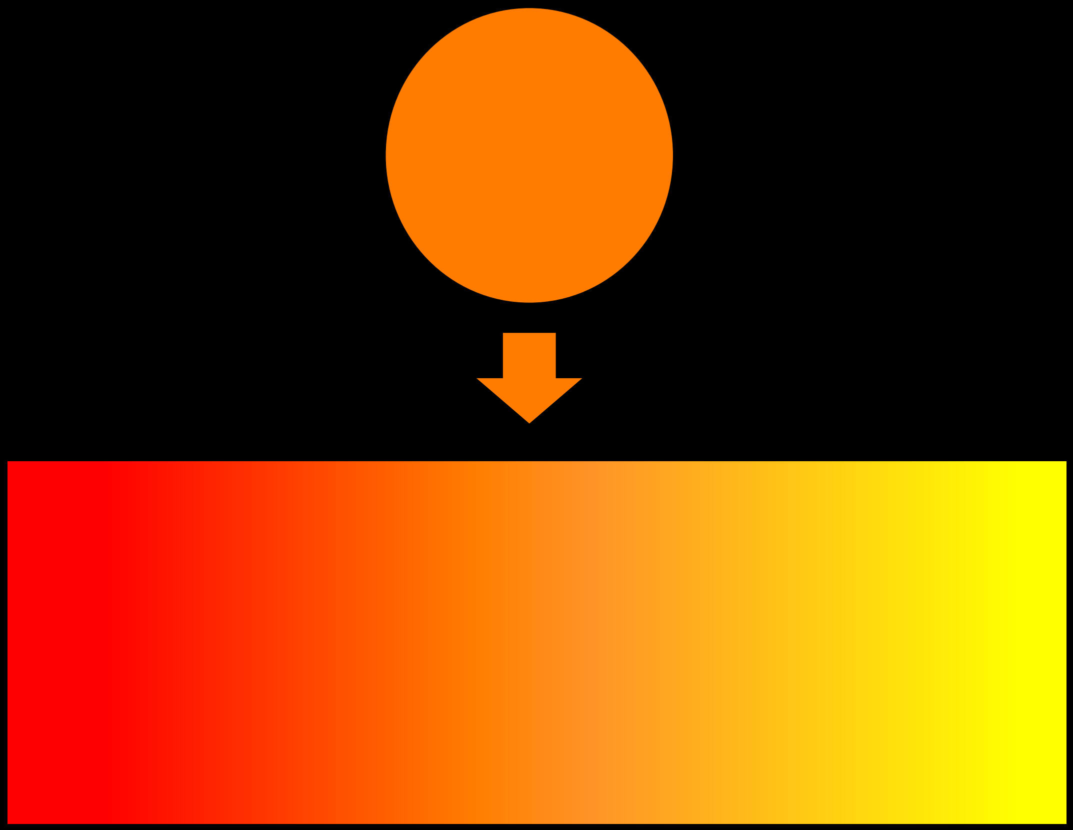

Can a patch be reddish-yellow? ✓

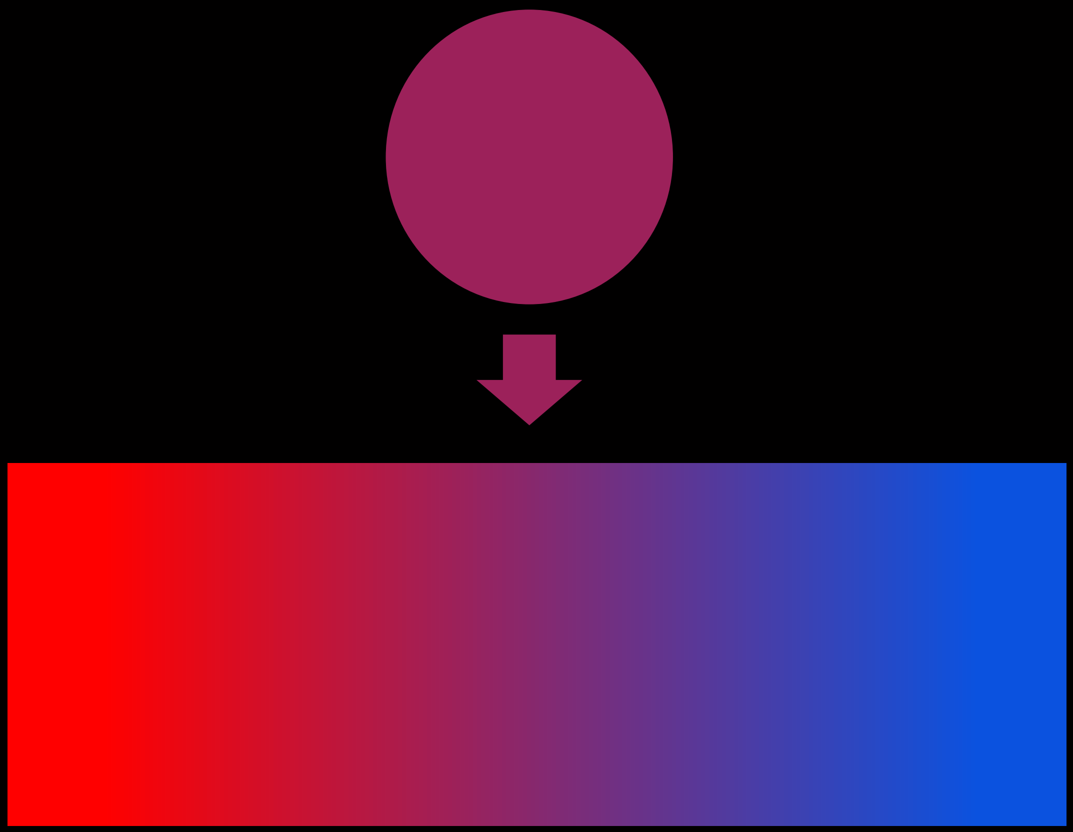

Can it be reddish-blue?

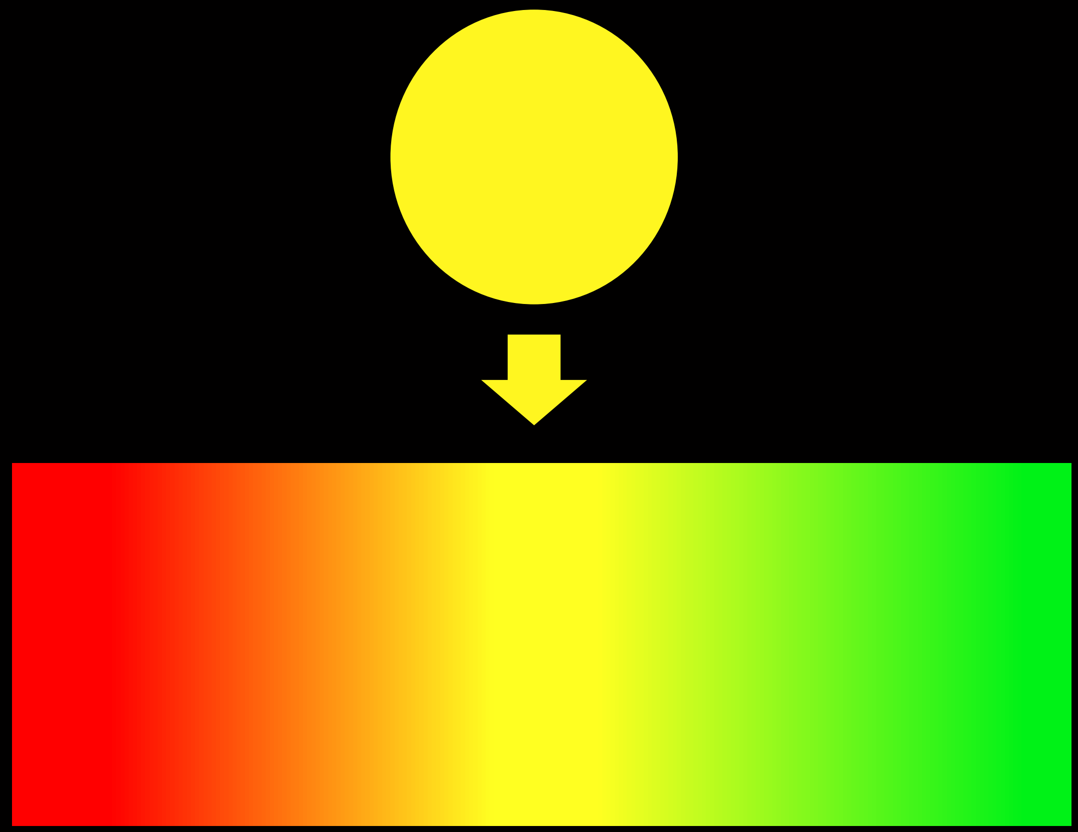

Can it be reddish-green?

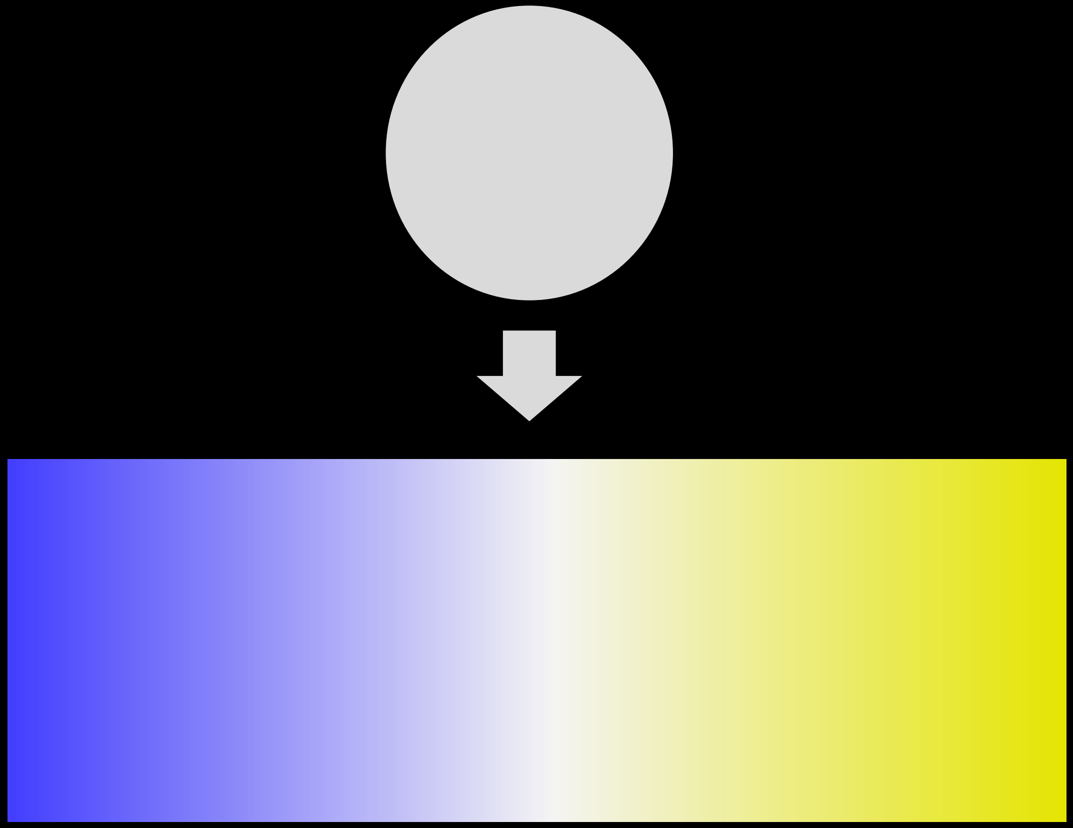

Can it be bluish-yellow? ✗

The colour opponent theory of Hering

is opposed to R-G

is opposed to Y-B

And indeed cells in the early visual pathway oppose the signals from different cone sensors and can be loosely classified as “red-green” or “blue-yellow” opponent:

R-G AND B-Y

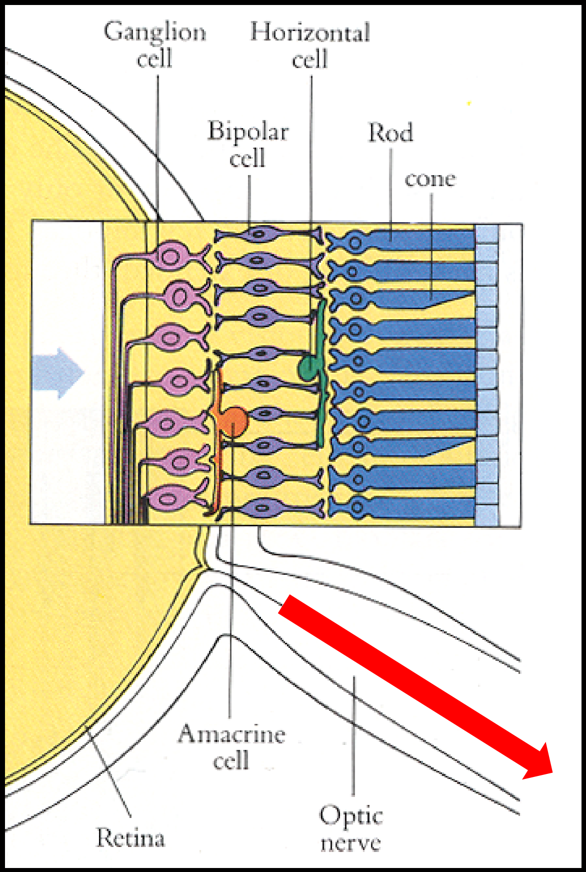

But why is this processing going on in the retina?

Because there are only about 800,000 to 1.7 million fibres in the optic nerve, but 6-7 million cone sensors and 120 million rod sensors (plus the encoding changes from analogue at the sensor outputs to digital in the optic nerve).

Thus, there is a BANDWIDTH PROBLEM. Differencing solves some of this by “decorrelating” the sensor signals.

But in addition to the difference signals:

R-G AND B-Y

You also need an intensity or “luminance” signal, which is achieved by adding together the L- and M-cone (red and green sensor) signals:

L+M

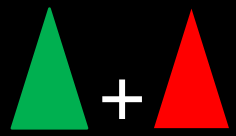

4. ACHROMATIC AND CHROMATIC CONE VISION

LUMINANCE AND COLOUR

Luminance is encoded by summing the L- and M-cone signals:

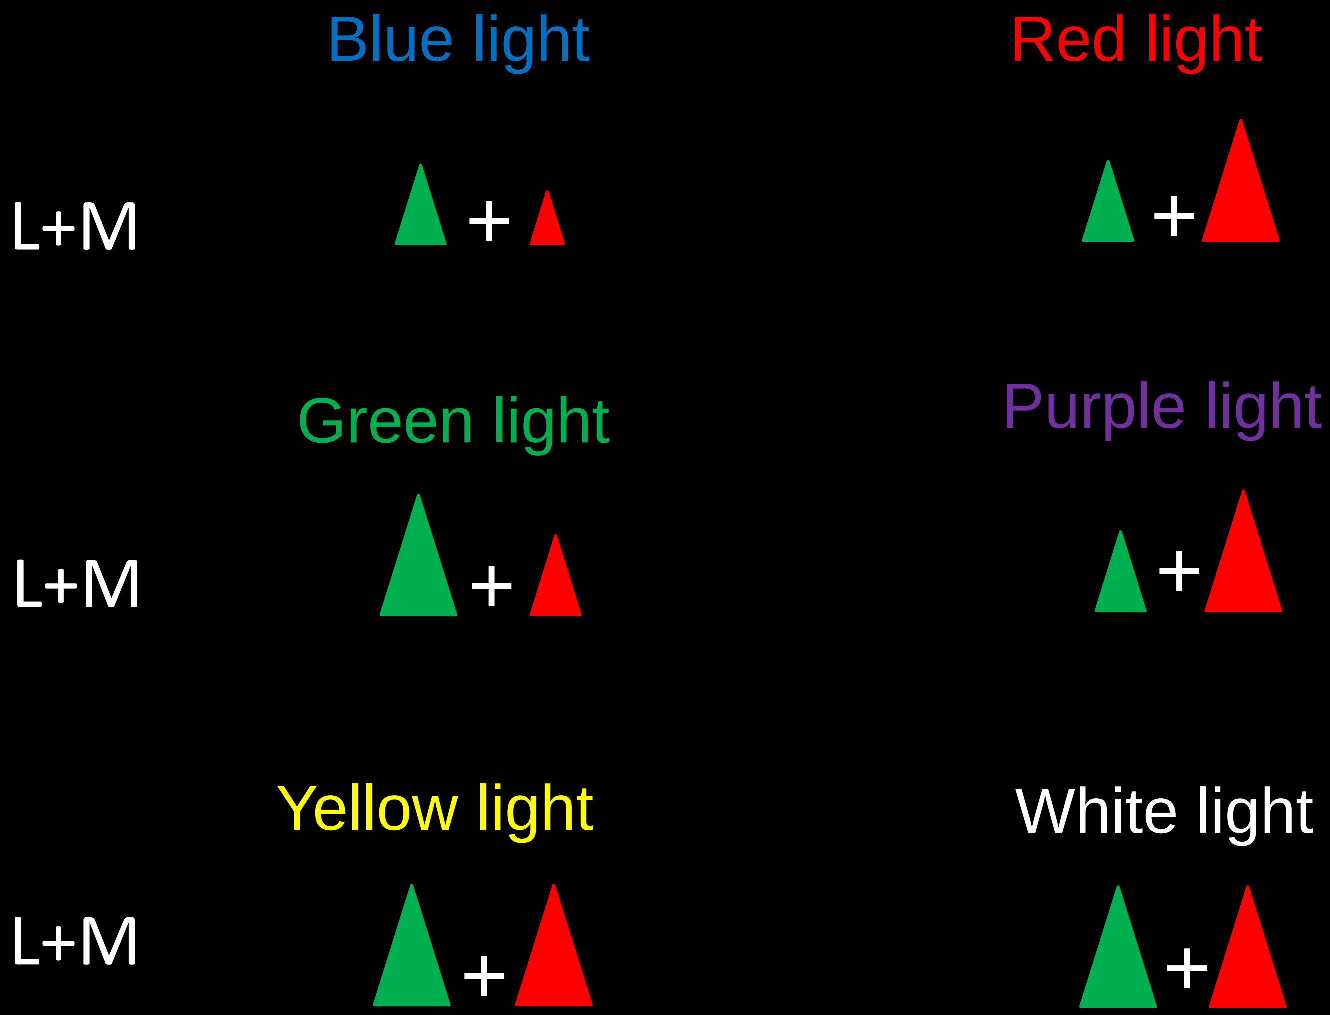

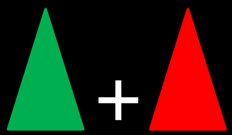

The luminance spectral sensitivity function is called "V(λ)", and is achieved simply by adding together the red and green sensor outputs (at least in principle).

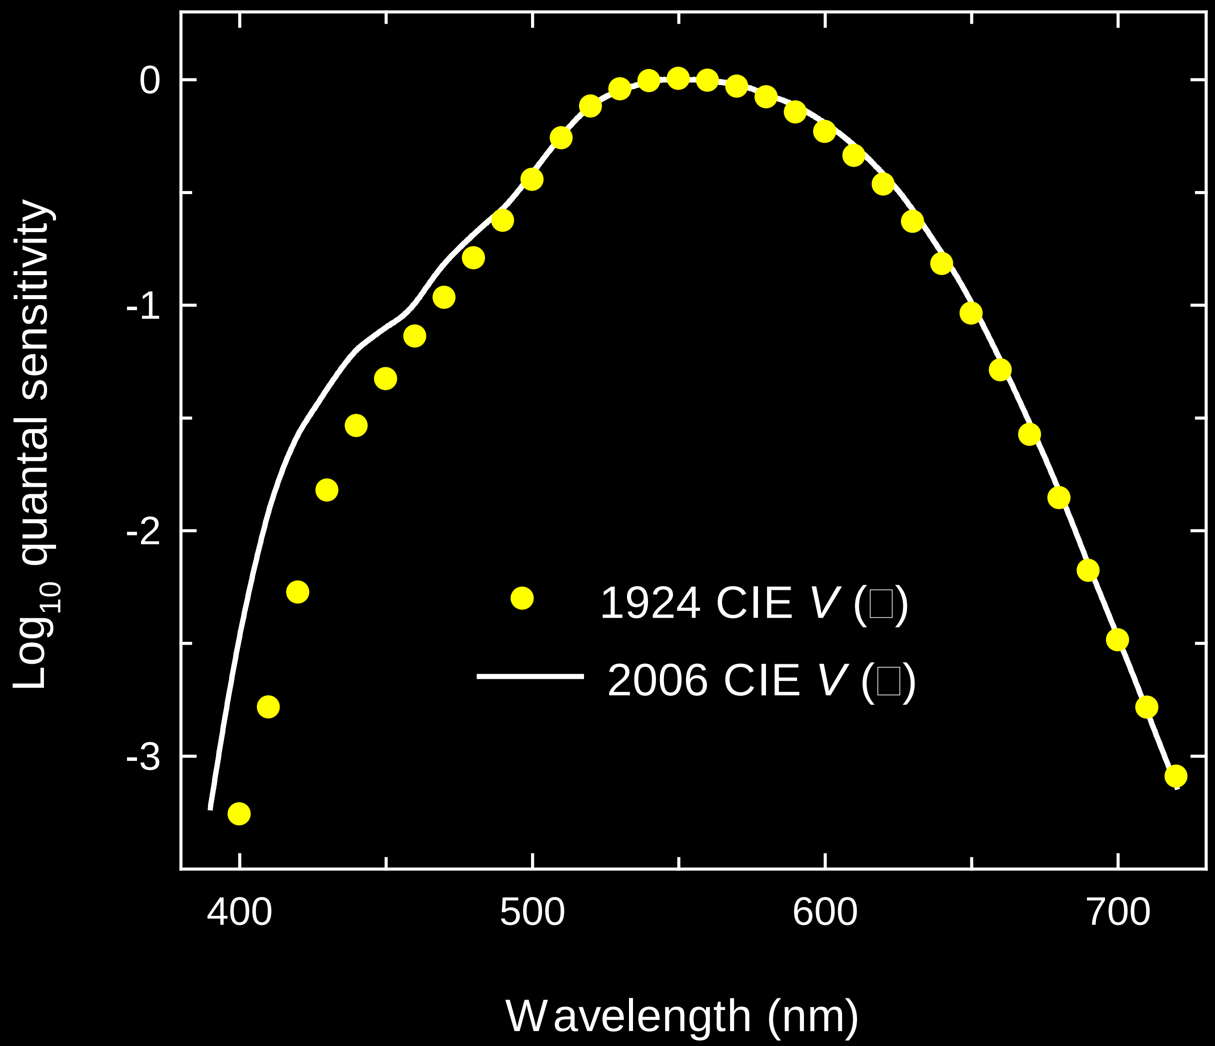

Defines photopic luminosity:
for example, lumens, lux, cd/m^2.

Colour is in many ways secondary to luminance

ACHROMATIC COMPONENTS

CHROMATIC COMPONENTS

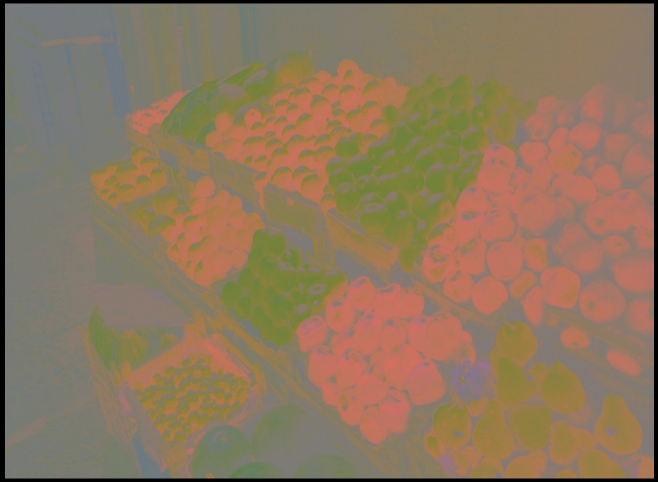

Demonstration:

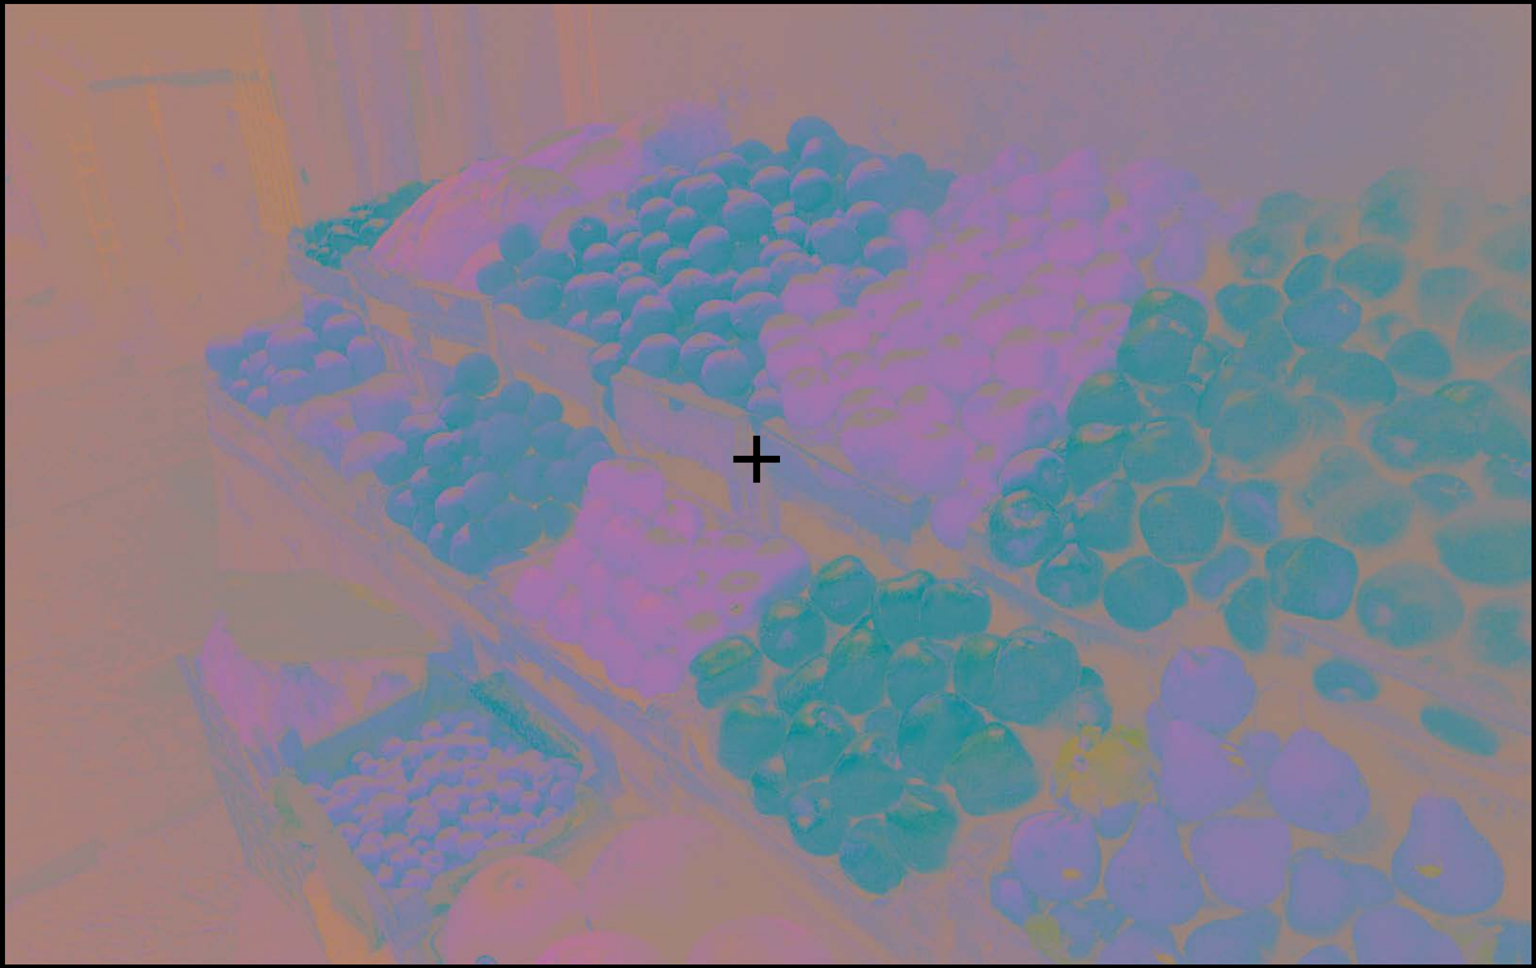


Souvlaki

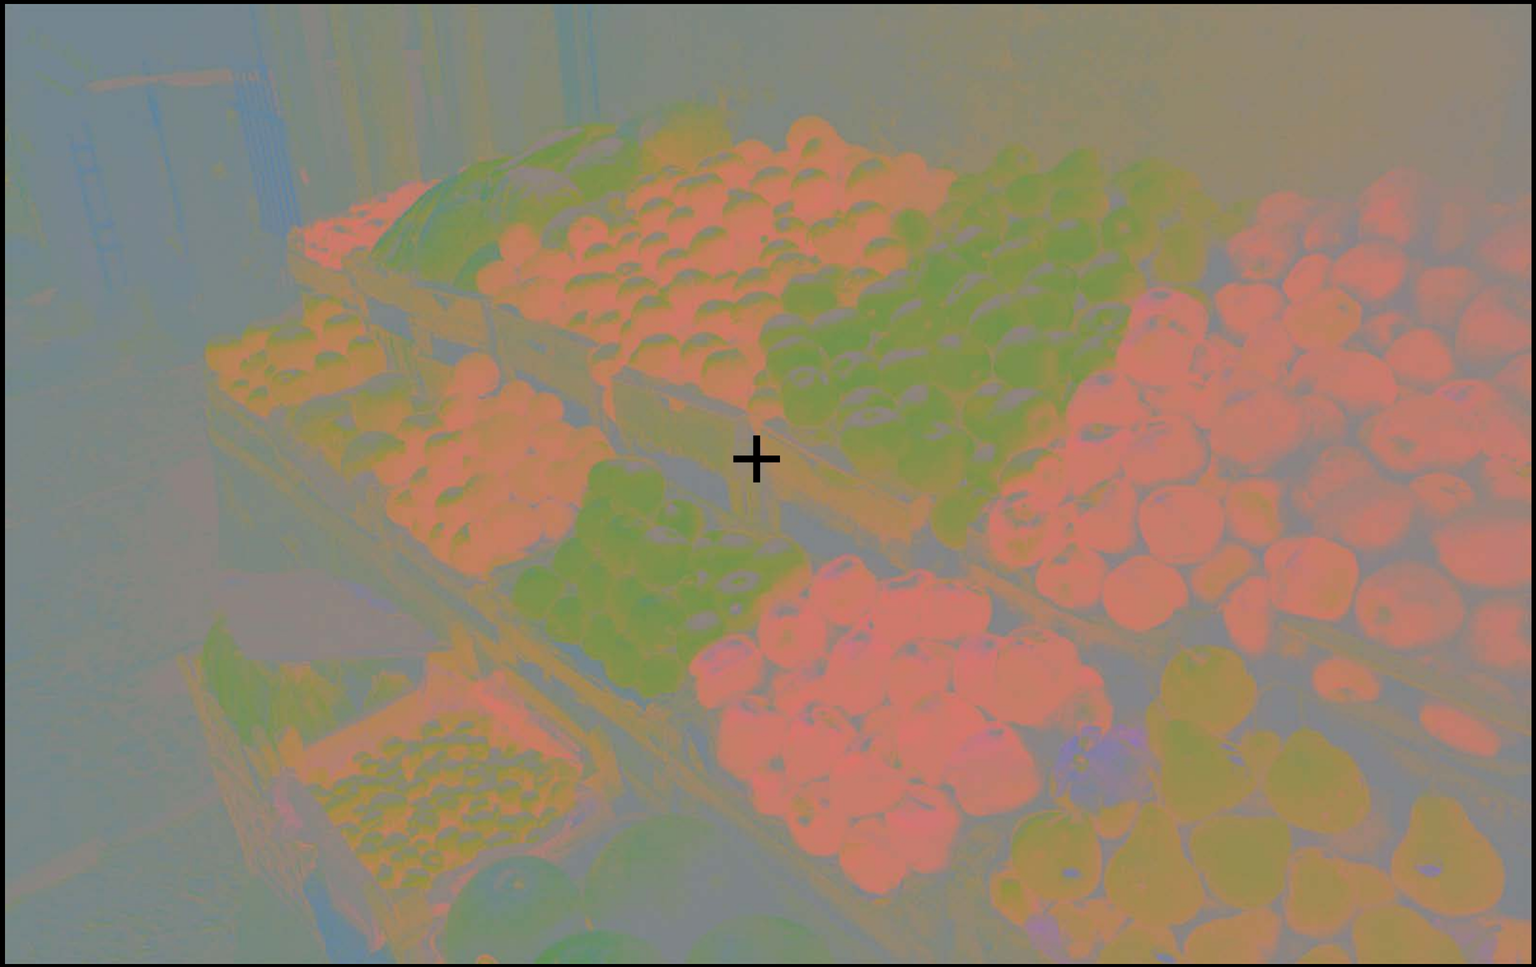


Souvlaki


Souvlaki

Demonstration 2:

LUMINANCE

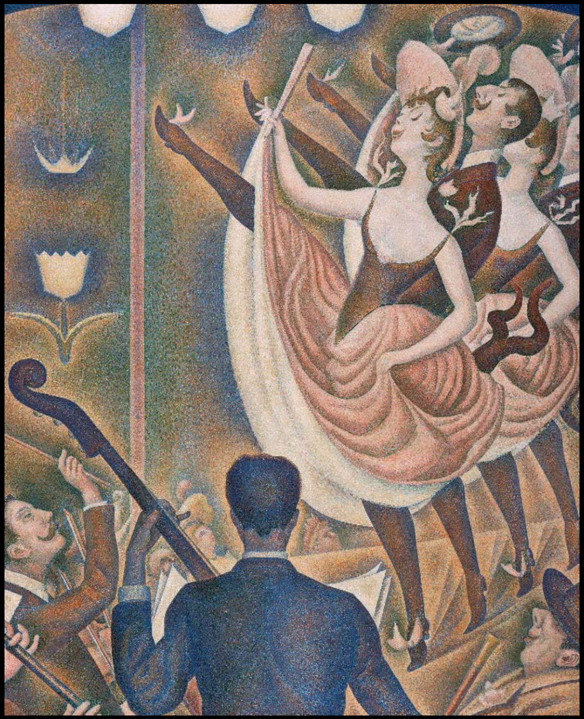

G. Seurat (*c.* 1889-90)
Le Chahut
[The High-Kick, Can-Can]

COLOUR

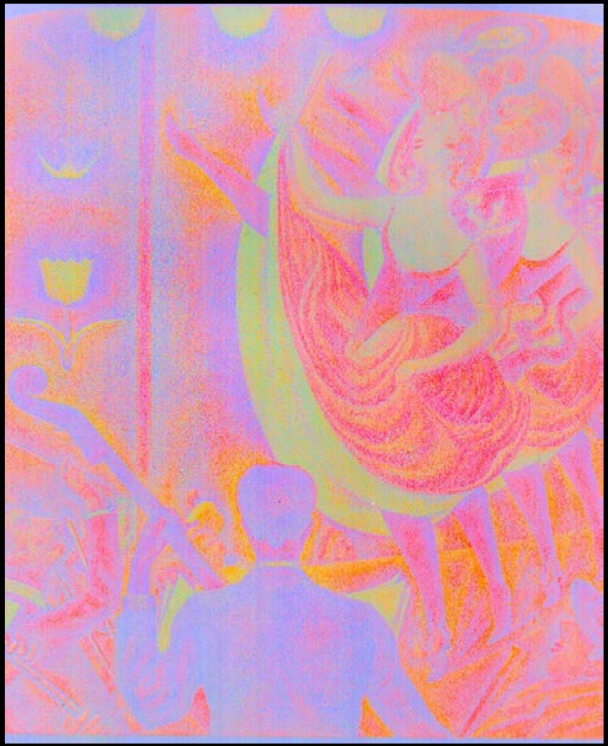

Courtesy: Jack Werner

LUMINANCE

COLOUR

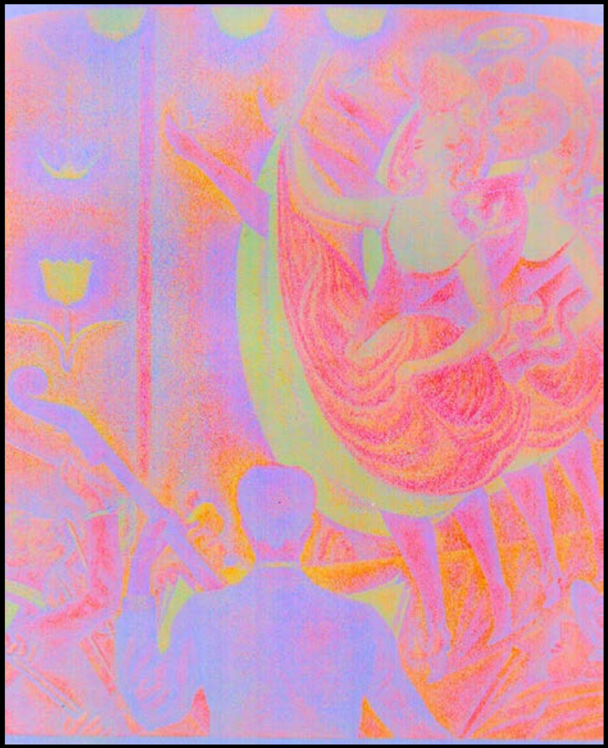

G. Seurat (*c.* 1889-90)
Le Chahut
[The High-Kick, Can-Can]

Courtesy: Jack Werner

Watercolour effect

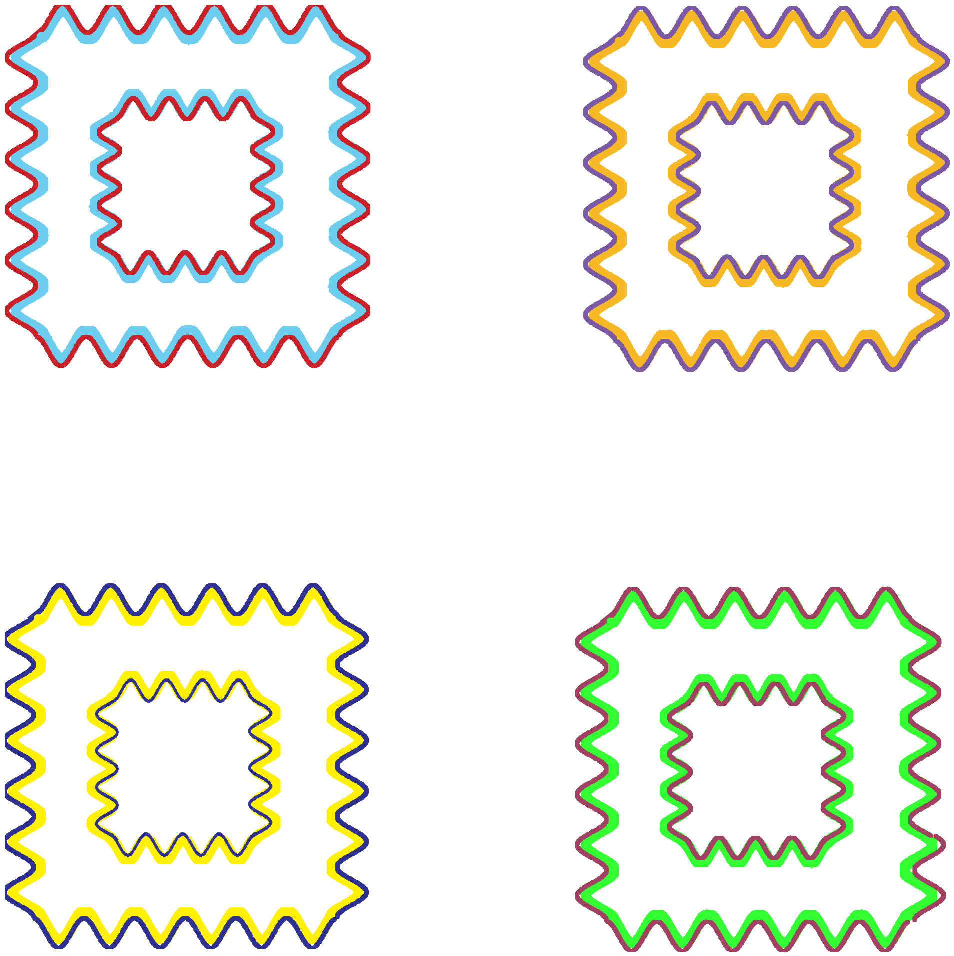

Wassily Kandinsky
Untitled, also known as 'Bagatellen' 1916
Private Collection, London

5. COLOUR IN TIME AND SPACE

So far, we've mainly been talking about the colours of isolated patches of light. But the colour of a patch depends also upon:

(i) What precedes it (in time)

COLOUR AFTER-EFFECTS

(ii) What surrounds it (in space)

COLOUR CONTRAST

COLOUR ASSIMILATION

COLOUR AFTER-EFFECTS

(what precedes the patch)

Colour after-effects

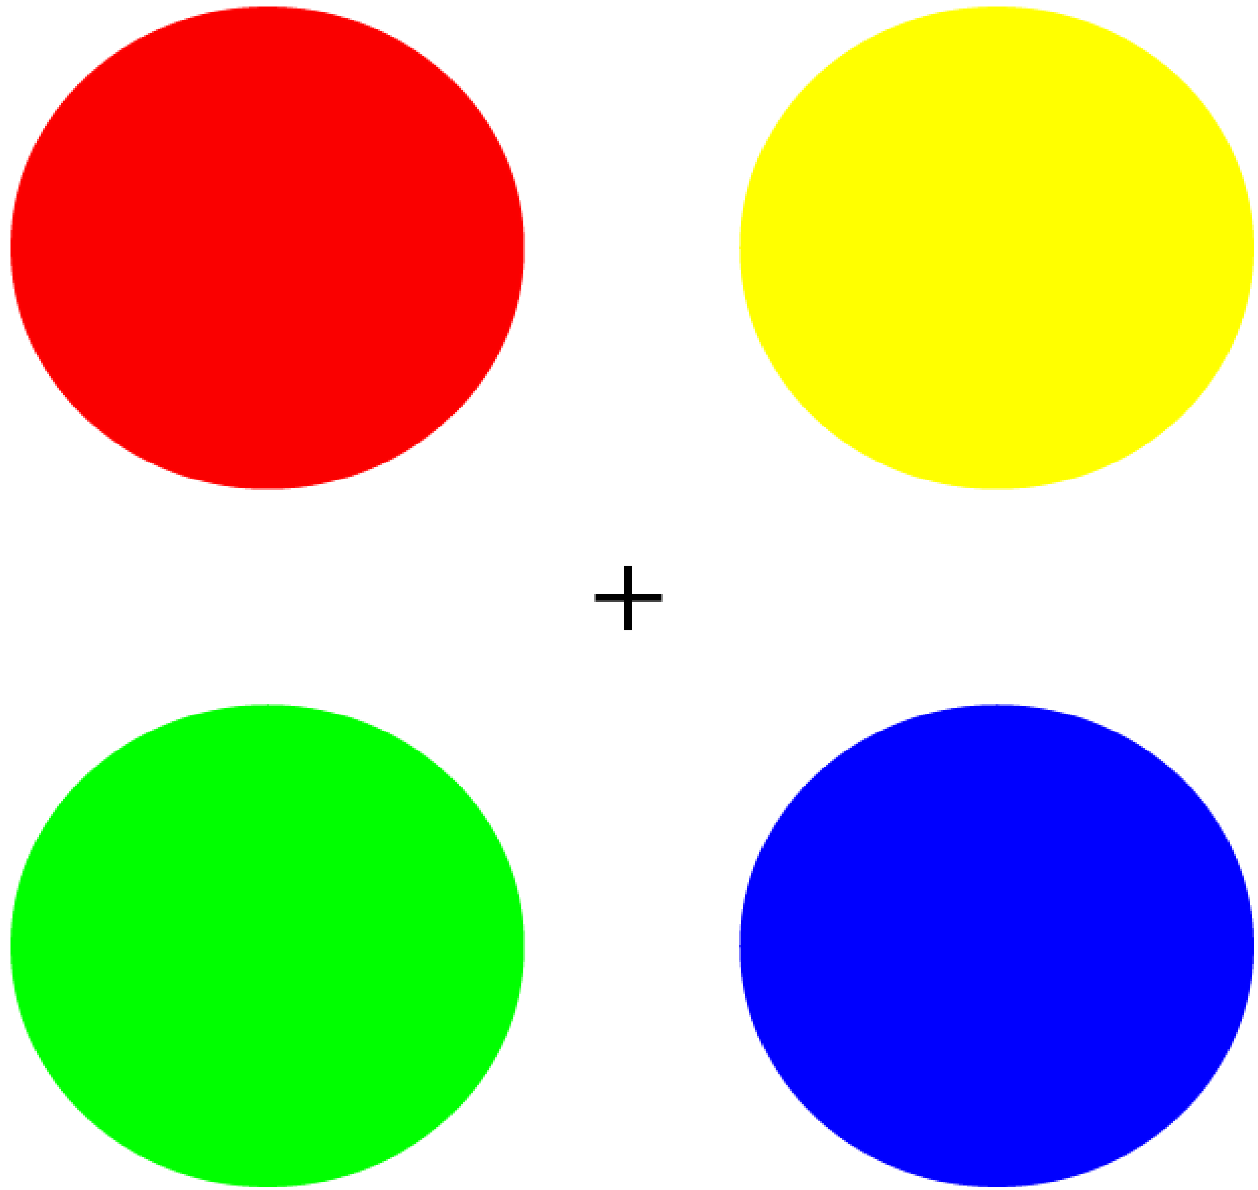

+

Mediafire

Lilac chaser or Pac-Man illusion

The "phi" phenomenon

The basis of cinema!

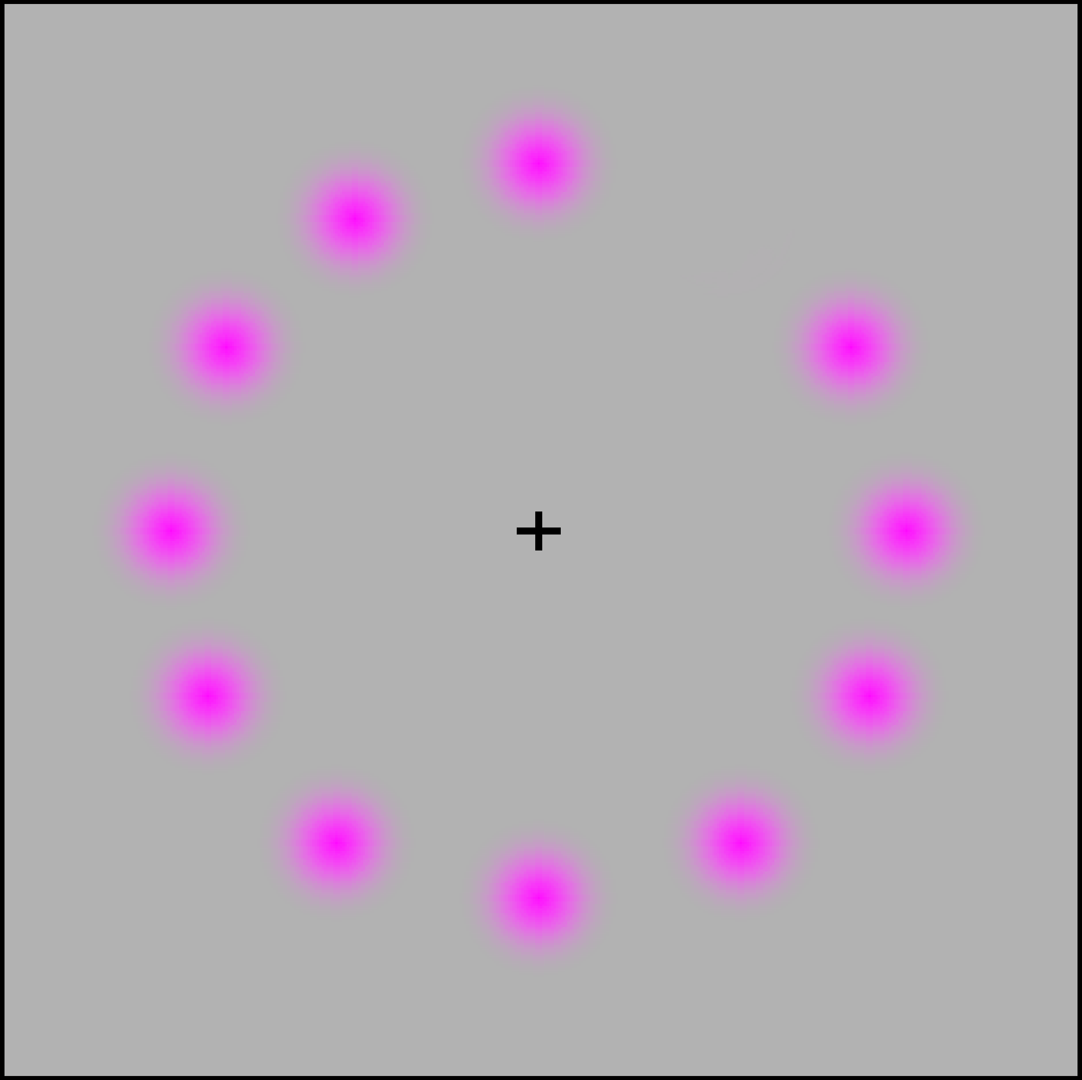

Jeremy Hinton

Lilac chaser or Pac-Man illusion

http://michaelbach.de/ot/index.html

Michael Bach and Jeremy Hinton

COLOUR CONTRAST

(what surrounds the patch)

Color contrast

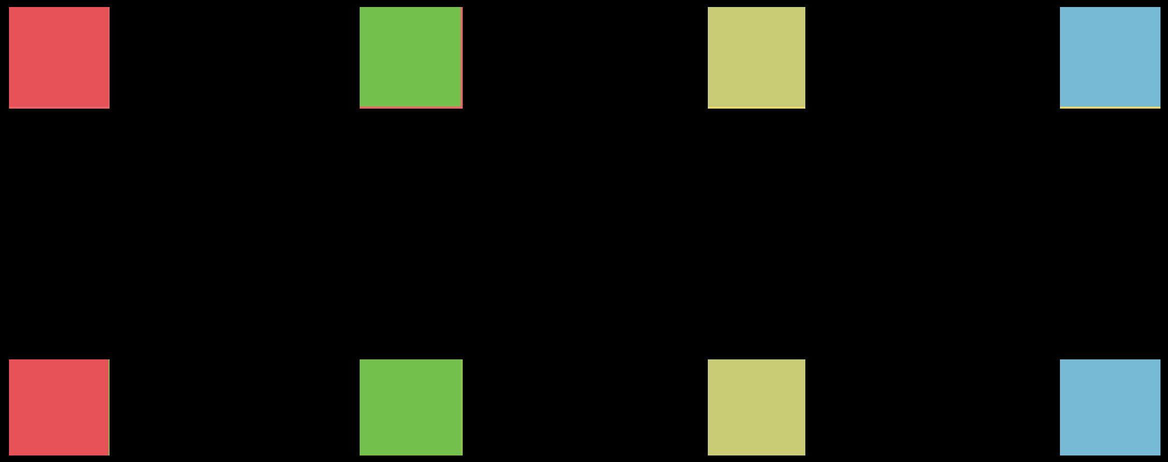

Color contrast

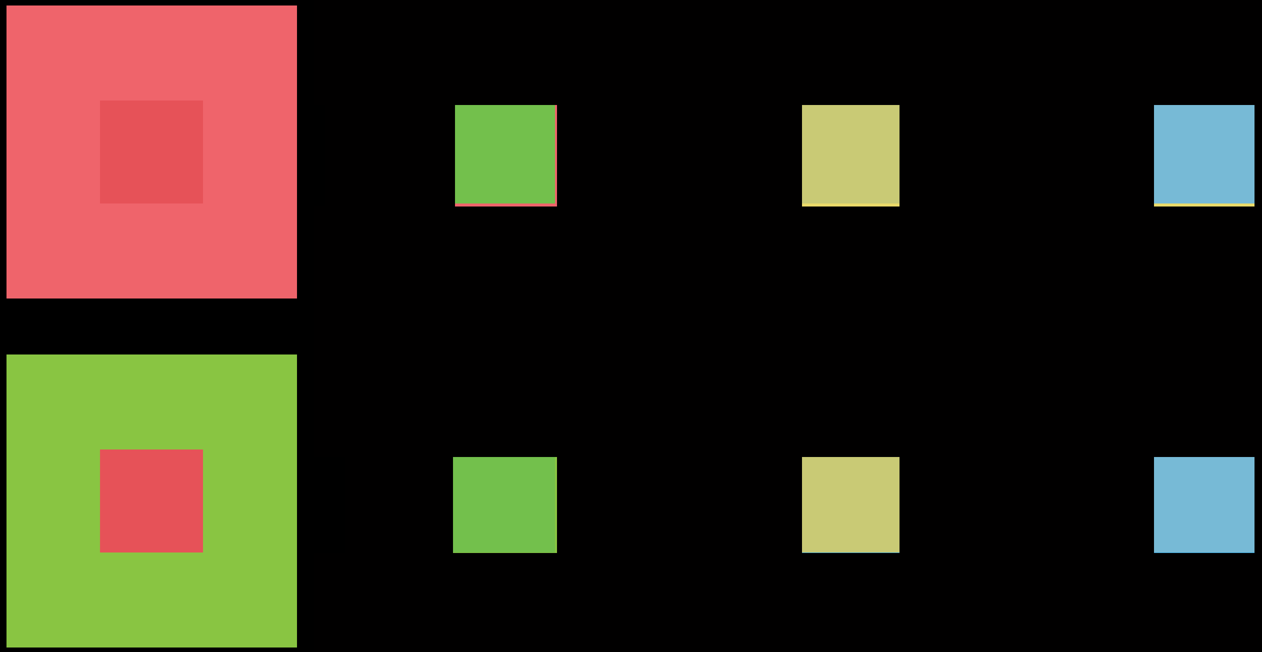

Color contrast

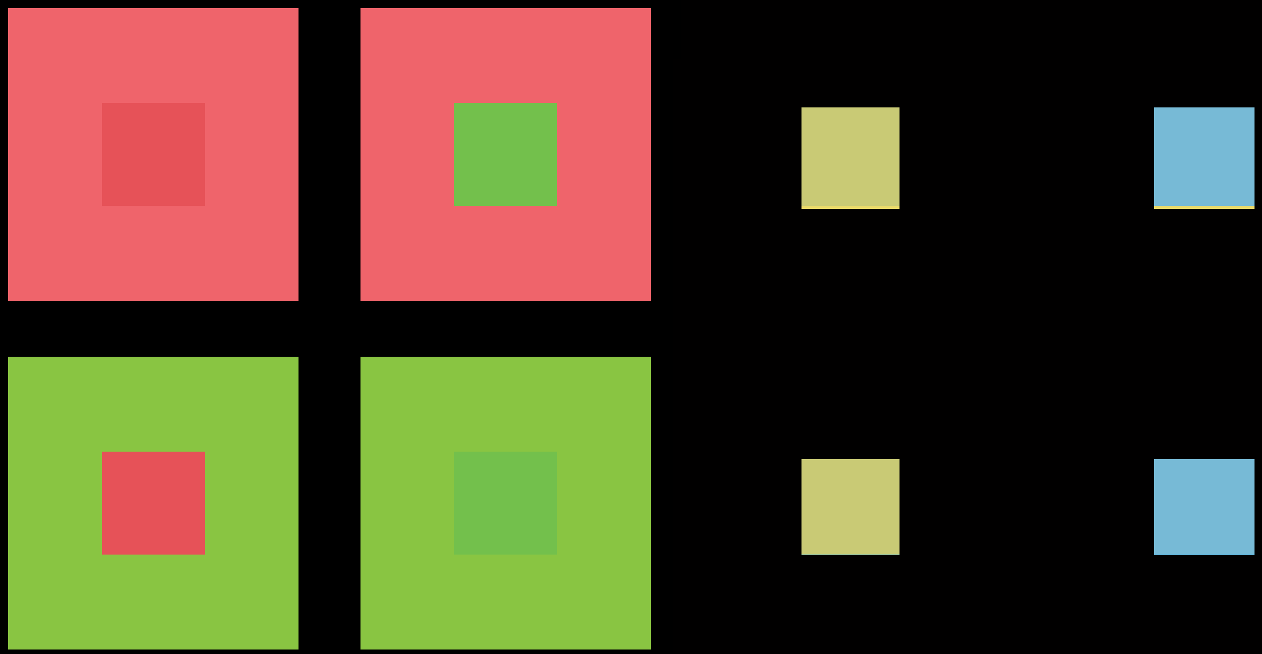

Color contrast

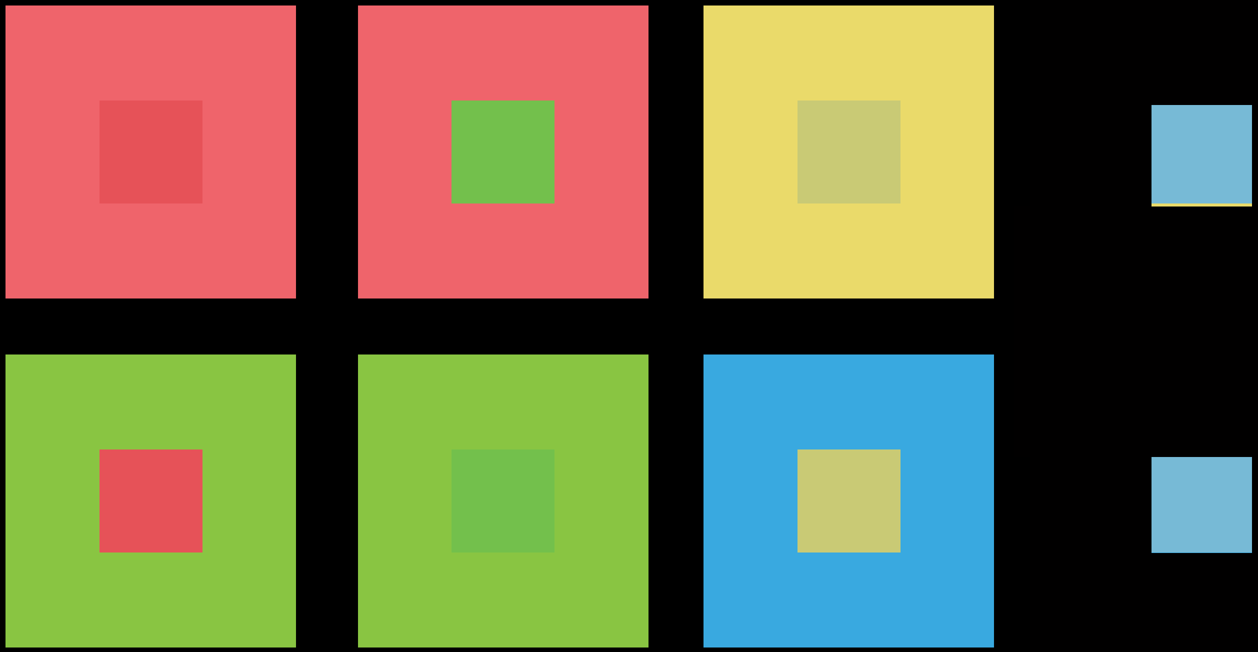

Color contrast

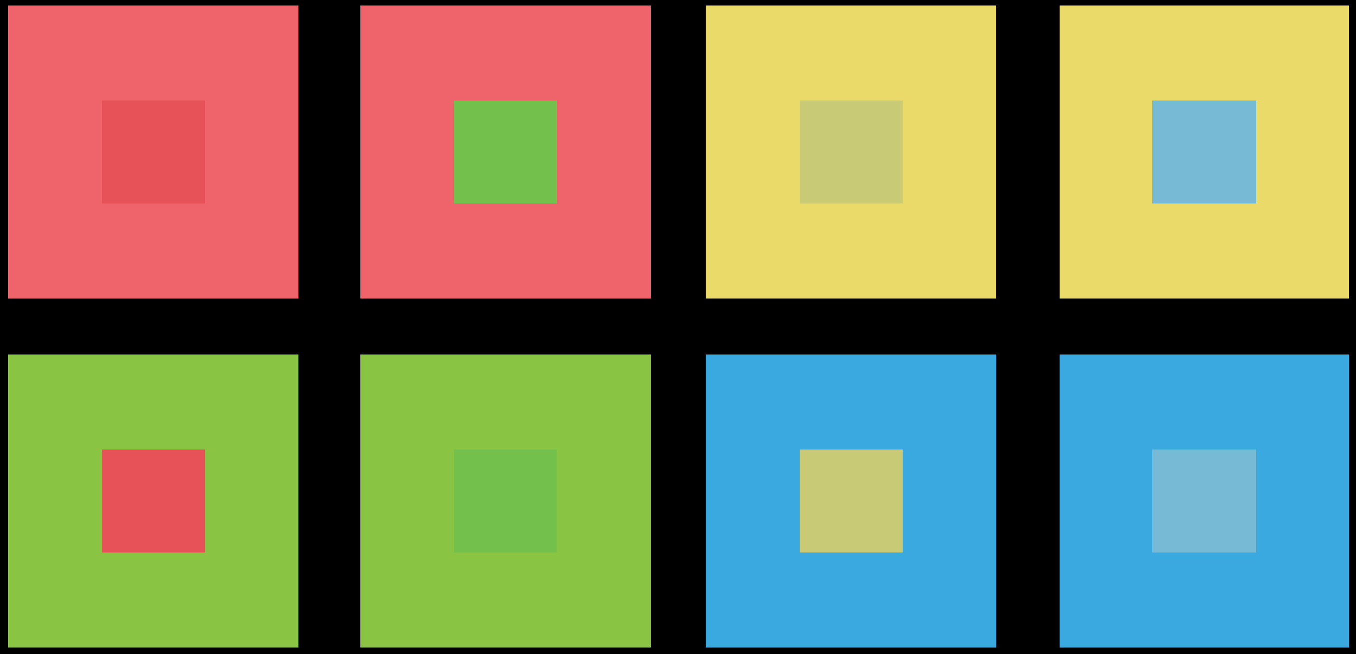

The Night Cafe in Arles, *by Vincent Van Gogh* Watercolour, 1888.

Claude Monet, The Regatta at Argenteuil, c. 1872
Musee d'Orsay, Paris

COLOUR ASSIMILATION

Colour assimilation

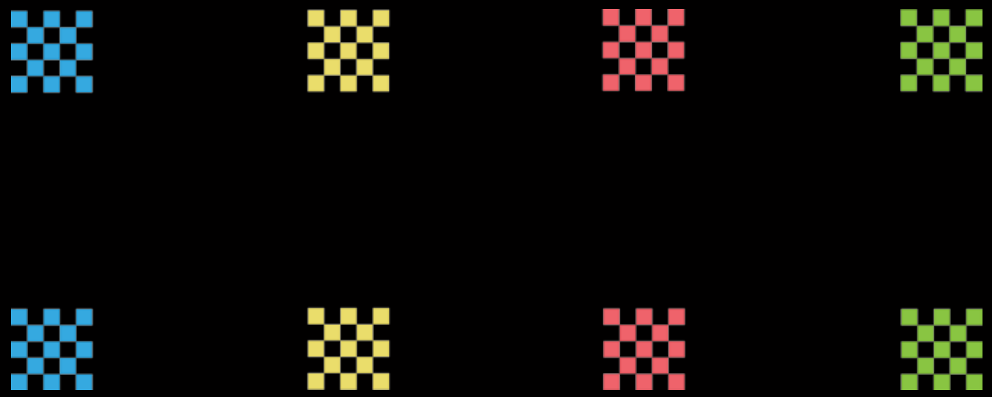

Colour assimilation

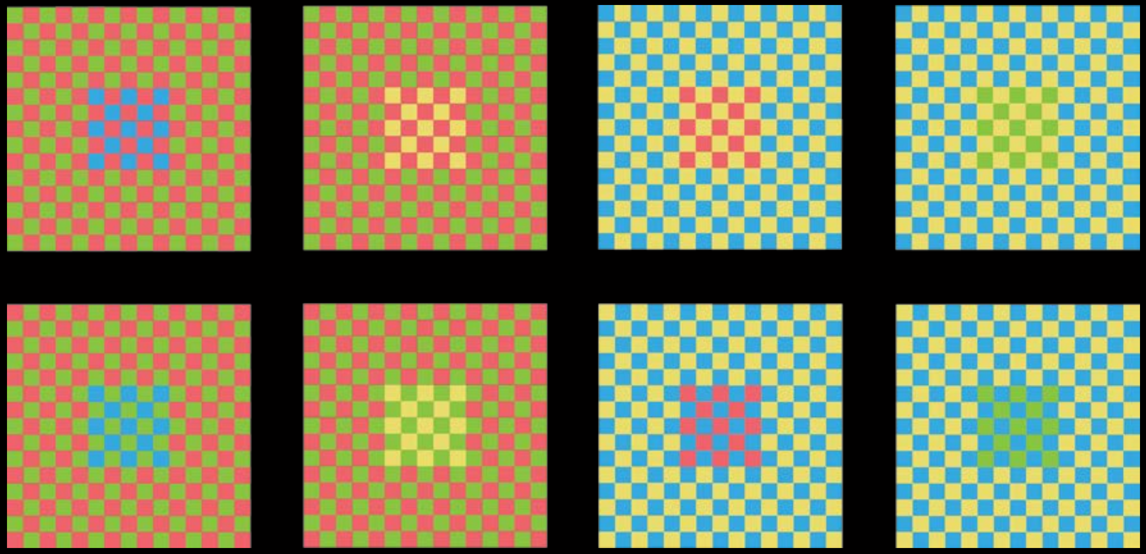

6. COLOUR IN THE MIND

Colour and the illuminant

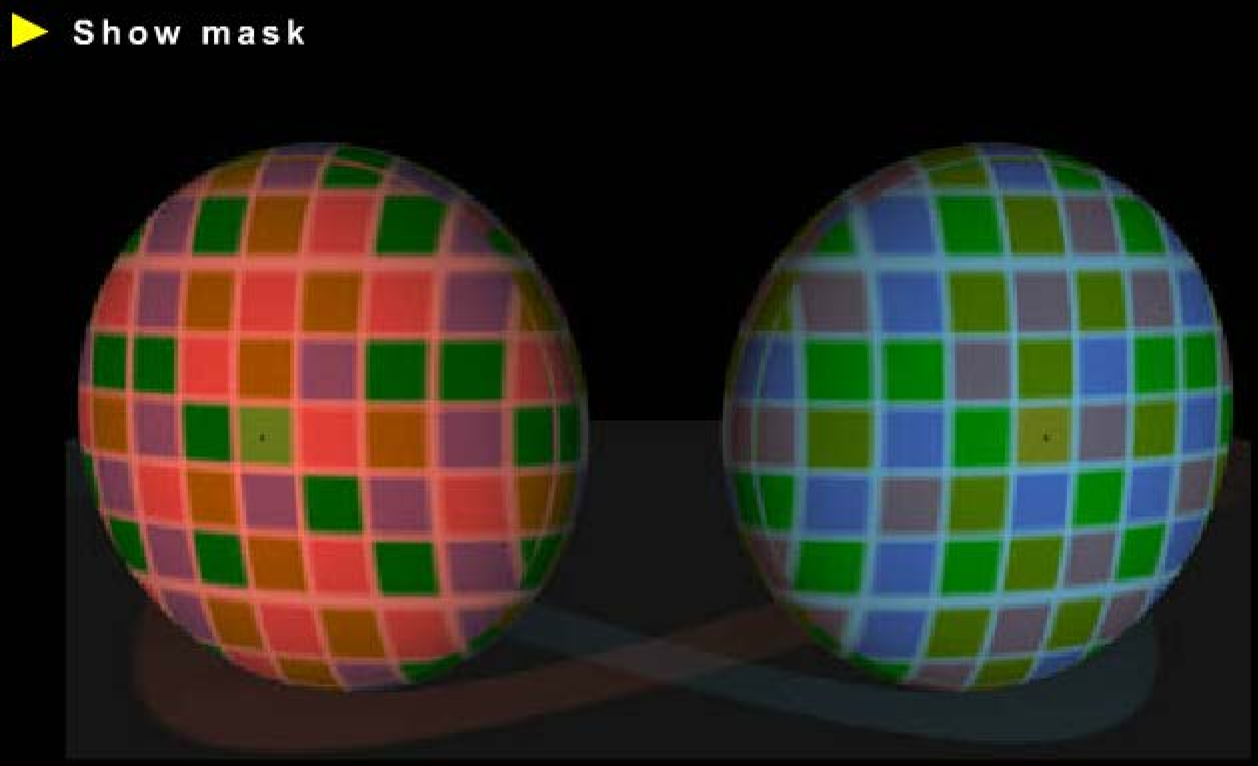

Beau Lotto

The dress…

www.wired.com

Colour and brightness

THE EFFECT OF COLOR ON BRIGHTNESS PERCEPTION

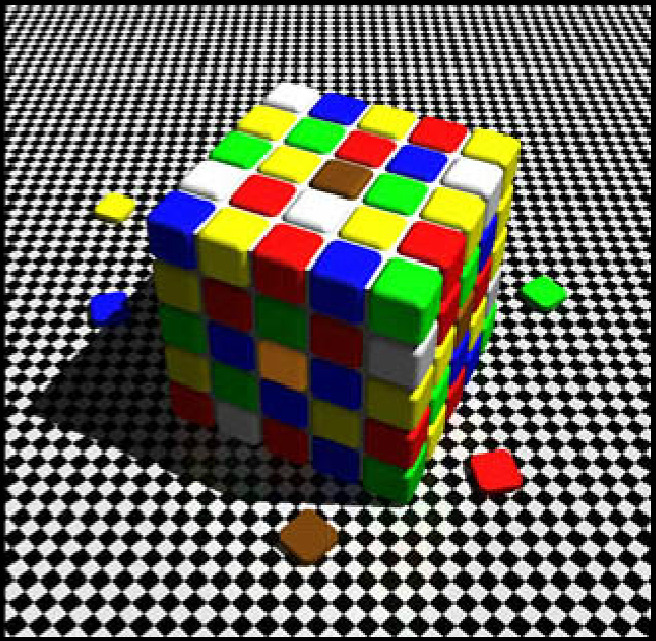

The color of the "brown" Chiclet-like square in the middle of the upper face of the cube is identical to the "orange" square in the middle of the shaded face. To prove this, click on the "Play" button (top) to view an animation in which all but the center two squares are covered by a mask, or click on the "Move mask" button (bottom) to manually position the mask over the center squares.

[From Lotto, R. B. & Purves, D. The Effects of Color on Brightness. *Nature Neuroscience* 2, 1010-1014 (1999)]

Why are S-cones sparse?

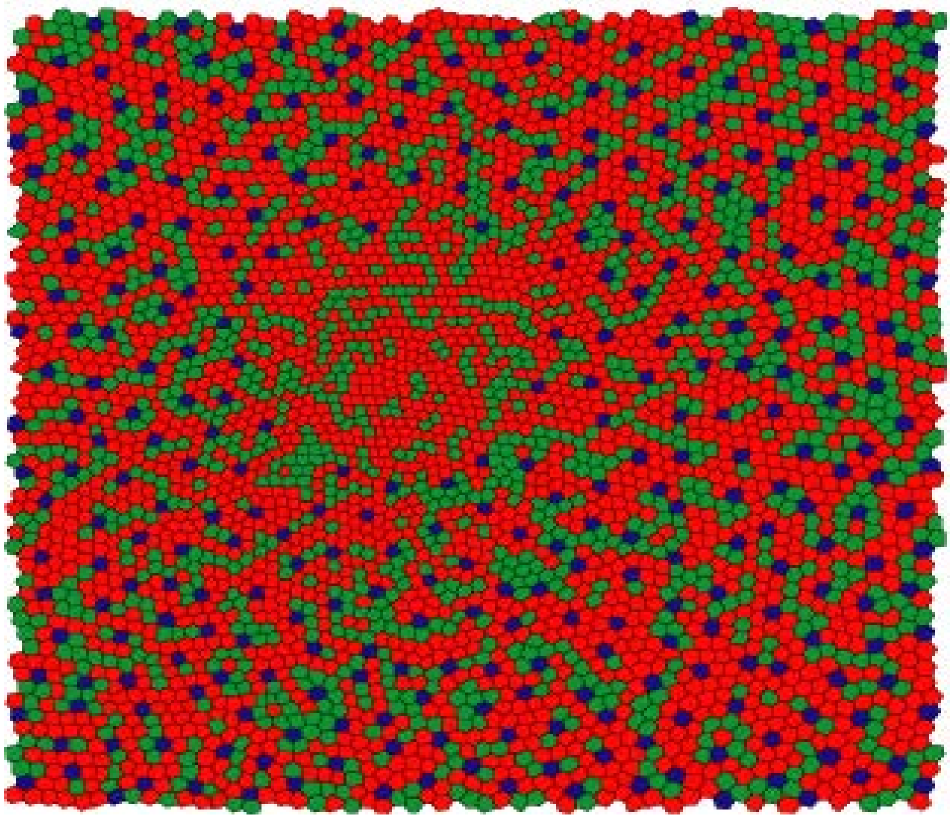

Chromatic aberration

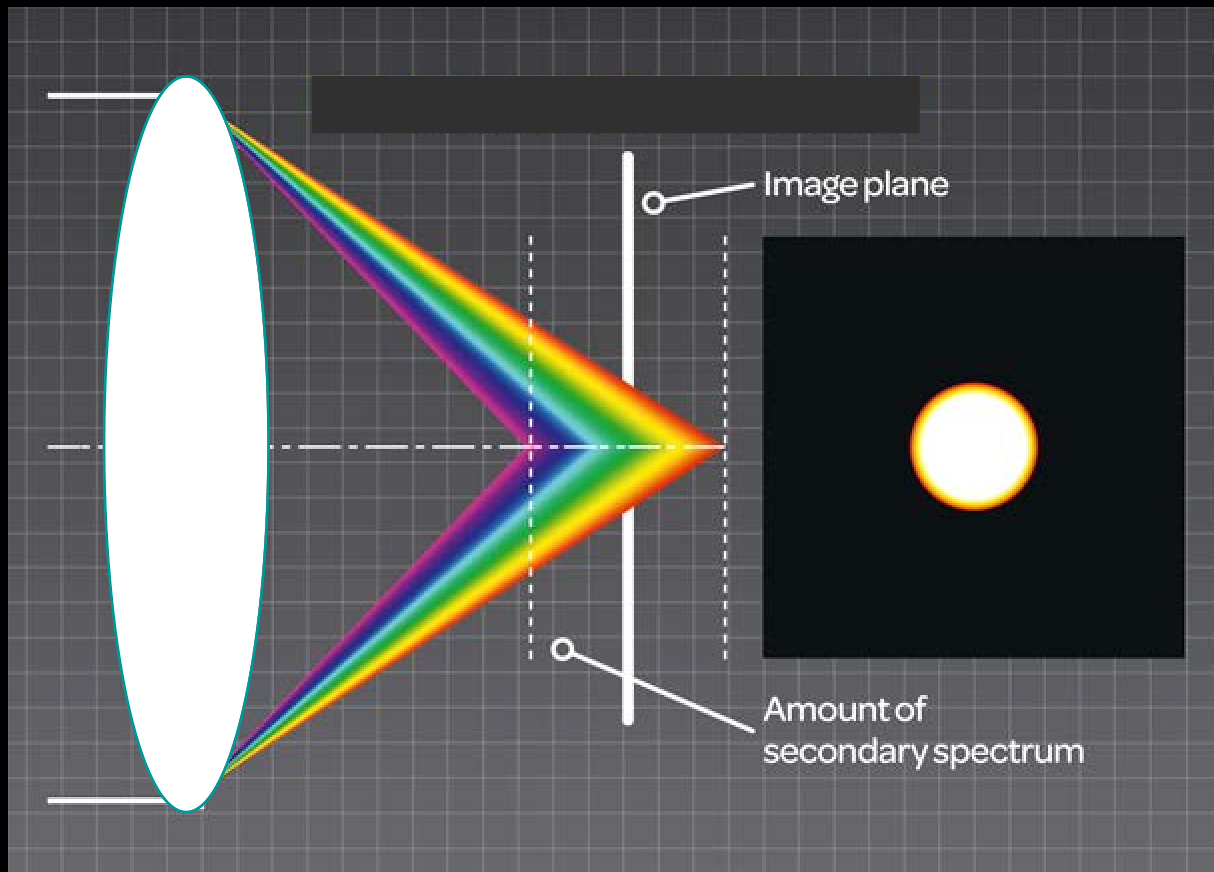

Base picture: Digital camera world

Effect of chromatic blur on eye chart

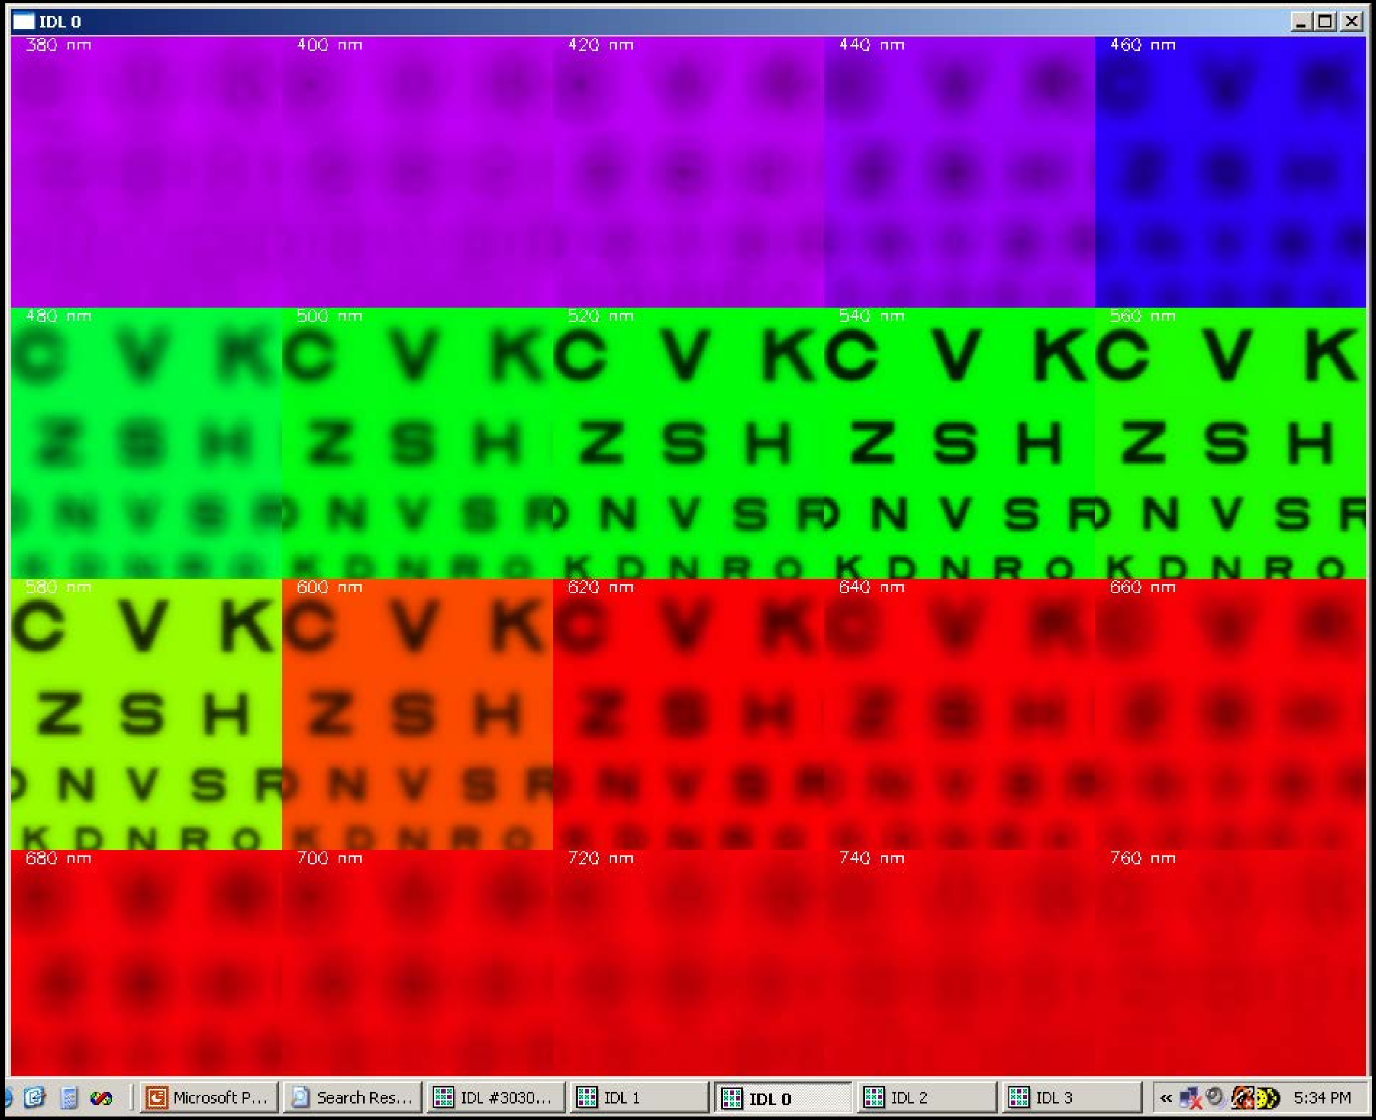

Jim Schwiegerling, U. Arizona

Chromostereopsis

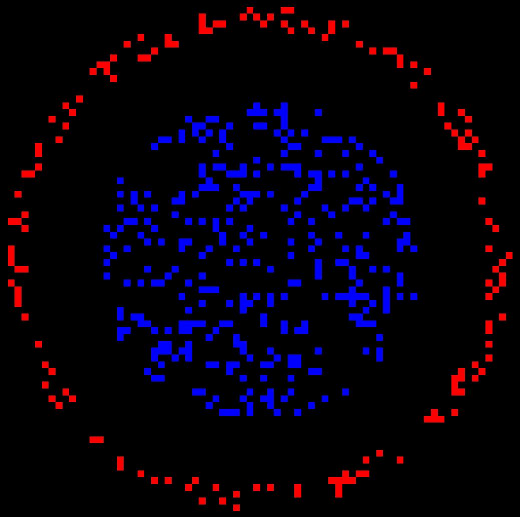

Different colours are perceived
at different depths…

Akitaoka Kitaoka

CONCLUSIONS

The human visual system works well (mostly, but there are some wonderful visual illusions or mistakes). However, it is far from perfect.

It is an evolutionary compromise, many aspects of which have been preserved since the first eyes evolved.

Many properties of cameras now exceed those of the visual system for which they are generating input!

But not always…

There is a need to understand how the visual system works (or is limited).