# Colour Tutorial: Sep 20, '04

**Colour appearance is an inference** — a perceptual representation of the surface reflectance of an object.

References:

- colourTutorial.m
- Wandell B - Foundations of Vision. Ch 4 and 9.
- Foley/van Dam - Computer Graphics 2nd Edition pg 574
- Munsell Book of Colour - Matte Finish Collection (Munsell Colour, Baltimore, Md., 1976).
- http://munsell.com

# Matlab set-up

```
> more ~jepson/pub/matlab/README

> cd ~jepson/pub/matlab
> matlab -nojvm

>> clear startup
>> which startup
>>  startup

% Then you can run the tutorials by opening, say,
% utvisToolbox/tutorials/colourTutorial.m in an
% editor and pasting the commands into the matlab
% console.

% A simple way to check if this is going to work
% is to type:
>>  which colourTutorial
% It should refer to the one in utvisToolbox
>>  colourTutorial  % Execute the whole script.
```

# Visible Spectrum

Wavelengths from 400 to 700 nm roughly. This tutorial uses a slightly larger range [370, 730], sampled every 1 nm.

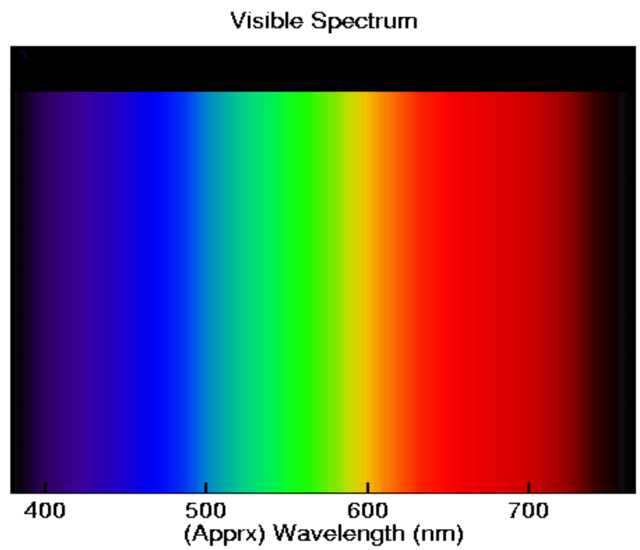


Figure 1: The visible light spectrum (A.K.A. The Colour Spectrum).

Neither the image above nor the one displayed in Matlab are calibrated, but they offer a reasonable idea of which colours correspond to what wavelengths.

# Colour Image Perception

The colour perceived by an observer is the result of the combination of four components: The illuminant, the reflectance function of any objects in the observer's field of view, the geometry of the scene (i.e. the angles between the illuminant and the objects' surfaces, and between those surfaces and the observer's eye), and the spectral sensitivity of the photo-receptors in the observer's eye.

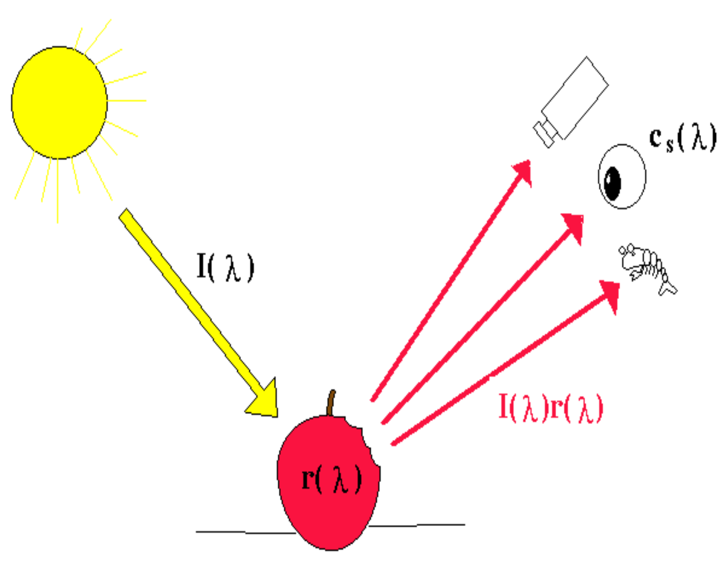


Figure 2: Stimulus formation and perception.

For simplicity, we will assume in this tutorial that the object of interest is simply a flat surface, and that the term $|\vec{N} \cdot \vec{L}|$ is constant and can be neglected.

Any light source has a particular **Spectral Power Distribution** or **Colour Temperature** $I(\lambda)$, when light from the light source hits an object, part of it is absorbed and part of it reflected. The proportion of reflected light at each wavelength is given by the object's spectral reflectance distribution $r(\lambda)$.

The multiplication of these two yields the stimulus $s(\lambda) = I(\lambda)r(\lambda)$ received by the observer's eye.

The response of a photo-receptor inside the observer's eye (or in a camera's CCD) is given by $r_s = \int(c_s(\lambda)s(\lambda))d\lambda$, where $c_s(\lambda)$ is the photo-receptor's spectral response.

# Questions

**colourTutorial.m** is roughly organized to answer the following sequence of questions and more!

- What are the typical values for $I(\lambda)$ (daylight,sunny day,sunrise/sunset) and $r(\lambda)$ (Munsell/paint chips)?
- Generate stimulus $s(\lambda)$ by computing: $s(\lambda) = r(\lambda) \times I(\lambda)$
- What are spectral sensitivities for an average human? Three types of cones (sensors): short(S), medium(M), long(L)
- What is the cone response?
- Why only 3 cones?
- What are metamers? What is colour constancy?
- What is CIE XYZ and xy-colour coordinates?

# Spectral Power Distributions

As mentioned before, each light source has a typical spectral power distribution that specifies how much energy is given at each wavelength. The figure below shows typical SPDs for daylight.

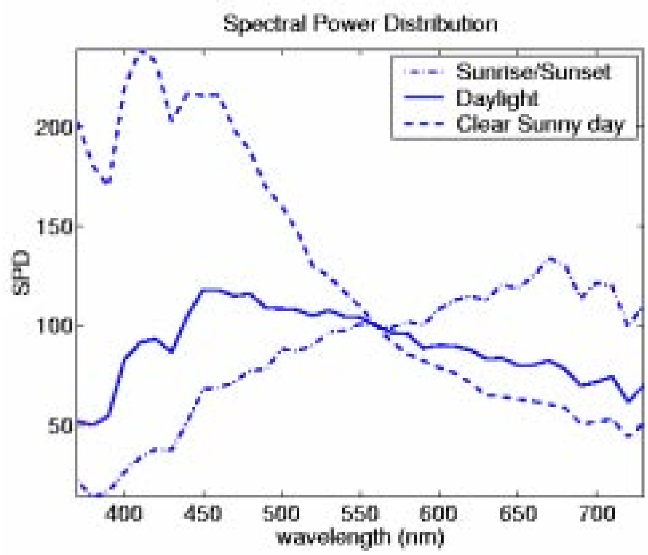


Figure 3: Spectral power distributions for daylight.

The spectral power distribution is affected by the medium through which the emitted light is propagating, in the case of daylight, the atmosphere has a strong influence in the resulting SPD due to factors such as the amount of pollution and dust in the air, the temperature and humidity, and the time of day.

# Munsell Chips

The Munsell set is a large collection of calibrated colour chips (as in paint 'chips'), the colours are systematically chosen to span a wide range of possibilities. There are 1269 chips in the database.

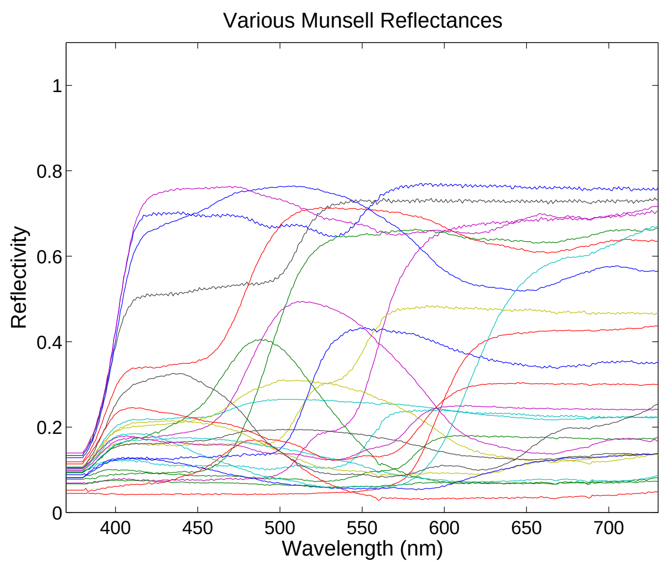


Figure 4: Typical spectral reflectance distributions for a few Munsell chips.

In the tutorial, each column of the matrix "munsell" is the spectral reflectance distribution of one chip for a range of wavelengths.

Munsell chips named for their perceived colour, the colour is specified by three parameters: Hue, lightness, and saturation.

hue: Specifies the name of the perceived colour.

lght : Indicates the brightness of the colour

sat : saturation of colour (how different from white?)

To determine the stimuli corresponding to each Munsell chip, we illuminate the chips with the SPD for standard daylight. This gives a spectral distribution function $s(\lambda) = r(\lambda) \times I(\lambda)$ for each chip as shown in the figure below.

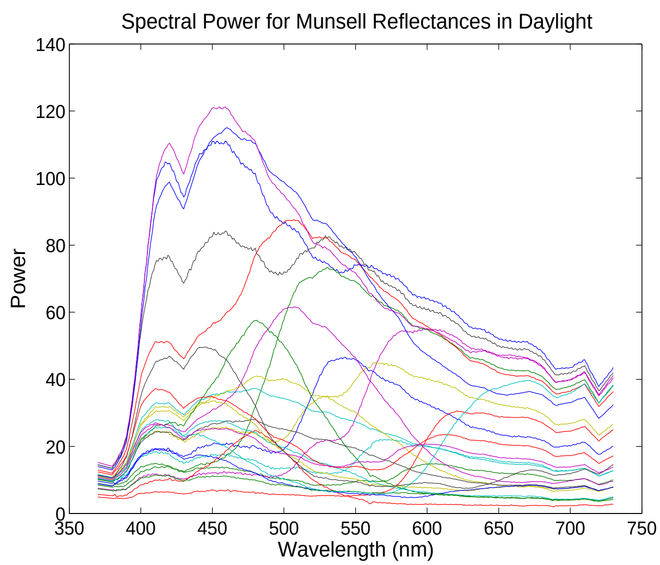


Figure 5: Stimuli resulting from multiplying the spectral power distribution for standard daylight with the spectral reflectance distributions for various Munsell chips.

# Sensor Spectral Sensitivities

The spectral sensitivity of a sensor characterizes the response from the sensor to each wavelength, for the 3 types of human colour receptors (long wavelength (L), medium wavelength (M), and short wavelength (S), usually referred to as the red, green, and blue receptors respectively) we have

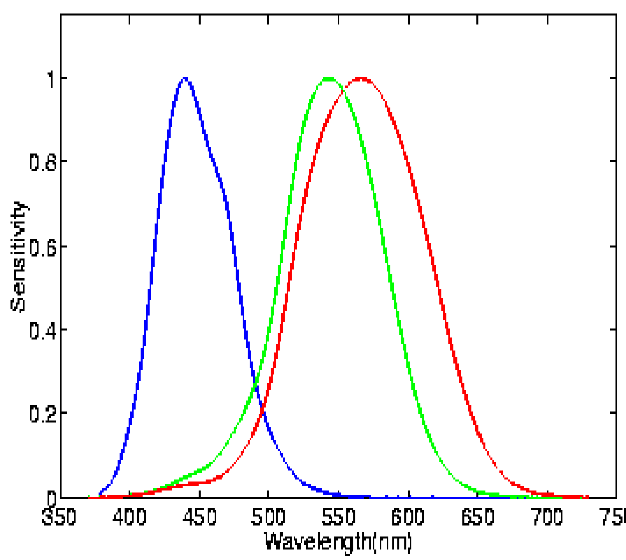


Figure 6: Spectral sensitivities for the L, M, and S photo-receptors in the human eye.

# Cone Responses

The sensors integrate the received stimulus according to their spectral sensitivity. The output of the sensor is proportional to the energy received at each wavelength times the sensitivity of the sensor at that particular wavelength. The proportionality factor corresponds to the exposure interval. In short $r_s = \int (c_s(\lambda) \times s(\lambda))d\lambda$.

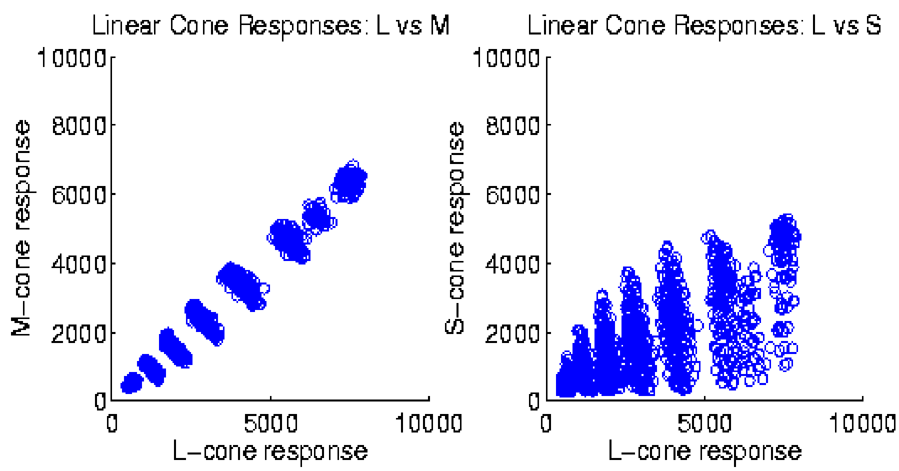


Figure 7: Scatter plots of cone responses. a) L vs. M, b) L vs. S

Notice that the responses of the L and M cones are more highly correlated than the responses of the L and S cones (which makes sense, consider the overlap in sensitivity between the L and M cones!).

# Colour Spaces and Metamers

The **perceived** stimulus is equal to the power absorbed by the L, M, and S cones, different spectral power distributions that yield identical absorbed power result in identical perceived stimuli, such perceptually identical stimuli are known as **metamers**.

## CIE XYZ-Colour Coordinates

- CIE has carefully calibrated 2 and 3D colour spaces, to specify perceptually identical chromatic stimuli.

- XYZ colour matching functions related to spectral sensitivities of human cones

- For the average human observer, stimuli with the same XYZ coordinates are perceptually identical.

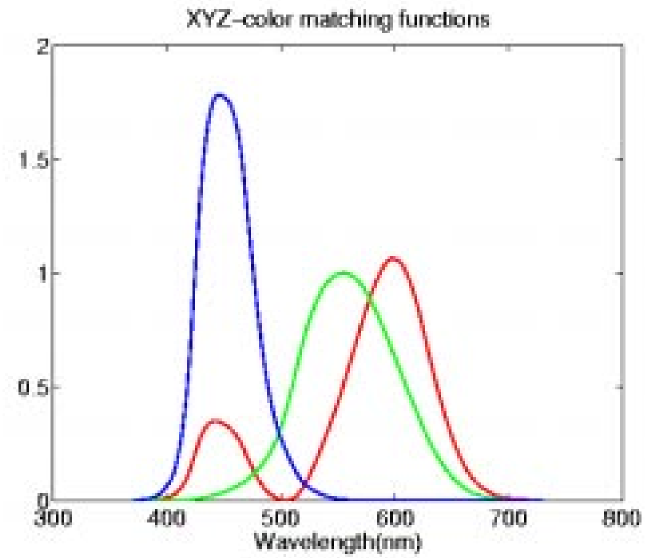


Figure 8: CIE colour matching functions that yield CIE colour coordinates X, Y, and Z.

## CIE xy-Colour Coordinates

Divide X and Y coords by the sum of the X,Y, and Z coords, this has the effect of normalizing the colour coordinates by the total brightness.

# Exploring CIE xy space

The $xy$ colour coordinates define a 2D colour space as shown below.

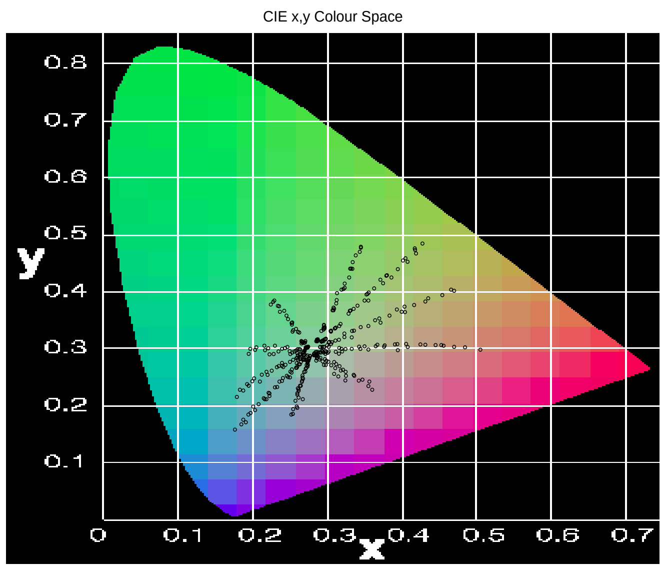


Figure 9: The CIE x-y colour space and points corresponding to some Munsell chips under standard daylight.

Points along the curved boundary correspond to monochromatic stimuli, points within the diagram correspond to linear mixtures of monochromatic stimuli. This colour space does not indicate overall brightness, only colour and saturation are represented.

## Changing the Illuminant

- sunset/sunrise, $\Longrightarrow$ shifts xy coords upward and to the right (more reddish)
- bright blue sky, $\Longrightarrow$ shifts xy coords downward and to the left (more blueish)

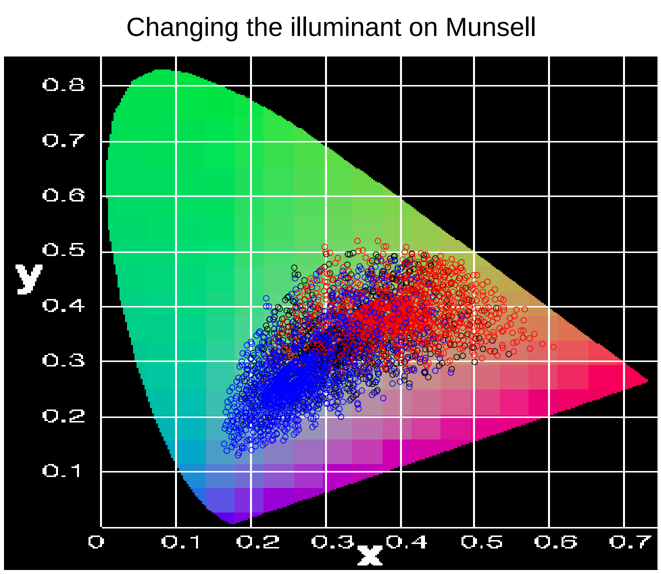


Figure 10: Points corresponding to the same Munsell chips under different illuminants, standard daylight (black), bright blue day (blue), and sunrise/sunset (red).

# Colour Constancy

This refers to the ability to neglect changes in the illuminant when estimating the colours of objects (i.e. a beach ball at different times of day). Alternatively, we can think of colour constancy as the ability to discern whether a change in the appearance of an object or surface is the result of a change in illumination or a change in the material of which the object is made.

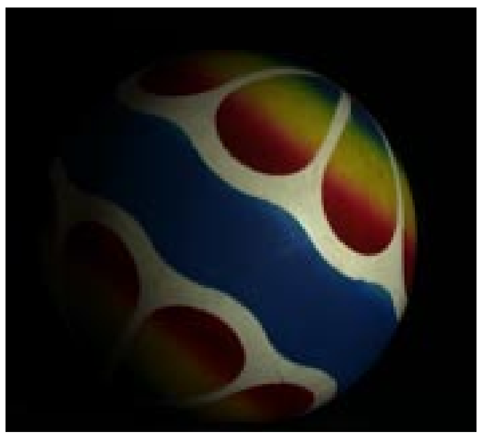

Figure 11: What are the colours of this ball? (*)

How can we do this at all?

(*) Image from: Kobus Barnard, Lindsay Martin, Brian Funt, and Adam Coath, A Data Set for Colour Research, Color Research and Application, Volume 27, Number 3, pp. 147-151, 2002.

# Why only three cone types?

Natural reflectances are well modelled by a three dimensional space.

```
>> [U, S, V] = svd(munsell, 0);
>> sv = diag(S);
```

Compute variance of each component, total variance of the data set, and cumulative variance of the first k components

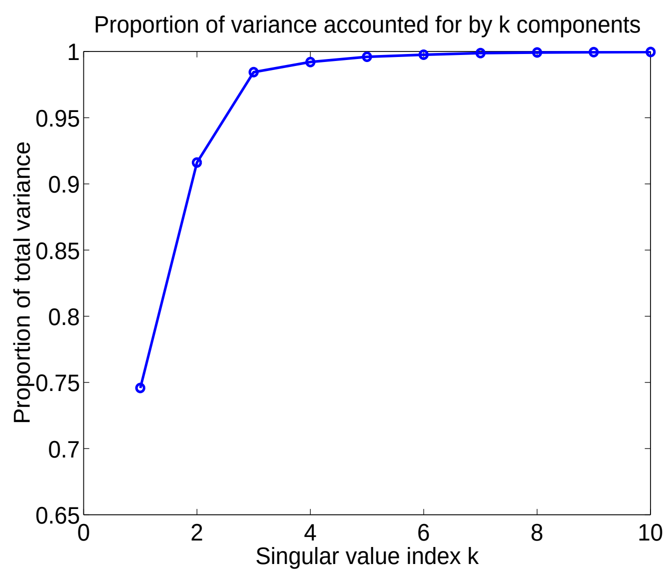


Figure 12: Percent of variance accounted for by succeeding principal components.

Just 3 linear components accounts for more than 95% of the total variance of the training data.

The first four principal vectors encode (approx):

1. mean brightness,
2. yellow vs blue,
3. green vs purple
4. more rapid fluctuations of reflectance