9.35 Spring 2024
Problem Set 4

The goals of this problem set are to:
a) gain more experience generating filters and applying them to images, and thinking about their relationship to perception
b) think through the steps of processing involved in the Retinex algorithm, and the extent to which it can account for the perception of reflectance (lightness)
c) review the idea of metamers and think through how they relate to cone transduction in the retina
d) practice thinking about motion in velocity space, and the role of constraint lines
e) review amplitude modulation and think about how its perception relates to cochlear filtering

Please show your work for all problems.

1. The image below was generated by overlaying concentric squares of gradually increasing luminance. When we view this image we see a glowing X along the diagonals, even though the corners of the squares are no higher in luminance than the sides (each square "ring" is of constant luminance).

In this problem you will test whether the perception of this illusion can be explained with simple receptive field structures.

a) First, write a Matlab/Python function to generate a "center-surround" receptive field by taking the difference of 2D Gaussian functions. Choose widths for the two Gaussians to give something that looks like the center-surround RFs we've looked at in class. Convolve center-surround filters of different scales with the image and discuss whether their responses could account for the image appearance. Note whether you used on-center or off-center receptive fields, or both, and why any variation in their response is interpreted the way it is by the brain.

b) Use the Matlab/Python functions you wrote to implement Gabor filters for the last problem set to measure the energy (i.e., the sum of the squared responses of two Gabor filters differing in phase by 90 degrees) at horizontal, vertical, left-diagonal, and right-diagonal orientations, at a couple spatial scales. Can the orientation energy explain the illusion?

2. The two stimuli below (**a** and **b**) contain identical luminance profiles in the region between the two arrows (apart from the ovals at the top of **a** and the parallelograms at the top of **b**), shown in the graph at the bottom.

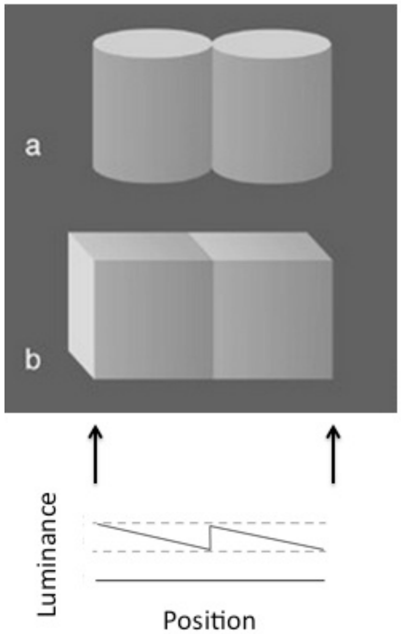

What would the Retinex algorithm do when applied to these two images (considering only a horizontal strip in the middle of each stimulus)?

What could account for the fact that humans perceive a lightness difference in **b** but not in **a**?

How does this illusion indicate the Retinex theory of lightness perception is incomplete?

3. Consider the color matching experiment depicted below, in which observers adjust the intensities of three "primary" lights (symbolized with R, G, and B) to match the appearance of a test light T. The spectra of the primaries are set by the experimenter; the observer adjusts only their intensities.

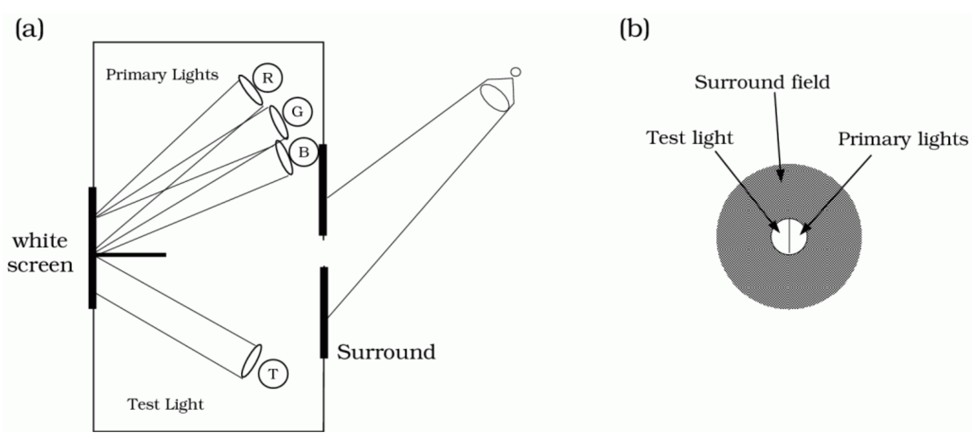

a) What choice of primaries would make the color matching task impossible, in the sense that there would exist lights that could not be matched by adjusting the intensities of the primaries?

b) Show that the intensities of the primaries needed to produce a match with a test light T are related to the spectrum of T by a linear transform (i.e., matrix multiplication). Assume the spectrum is represented by a vector and that the cone responses depend linearly on the spectrum.

c) Suppose a participant in the experiment establishes a metameric match. They then put on a pair of sunglasses. Will the match be preserved? Why or why not?

4. Consider a superposition of three drifting sine wave gratings, one vertical and drifting to the left, one oriented at -120 degrees and drifting downward, and one oriented at +120 degree and drifting upwards.

a) Draw a velocity space diagram of the resulting constraint lines.

b) Make a prediction for what you should perceive when you view such a stimulus, assuming the visual system will use the intersection of constraints algorithm to determine the potential velocities of objects in the world.

c) Should anything change if the gratings are all low contrast, assuming a Bayesian model of motion estimation using a prior favoring slow speeds? Why or why not?

5. A sinusoidally amplitude-modulated (SAM) pure tone is obtained by multiplying a “carrier” tone of frequency f_c by a “modulator” of frequency f_m:

$$\sin(2\pi f_c t) * [1 + \sin(2\pi f_m t)]/2.$$

The power spectrum (i.e. the constituent frequencies, and their amplitudes) of a SAM tone has three components (at $f_c - f_m$, f_c, and $f_c + f_m$), meaning that the SAM tone can also be generated by adding two sideband tones to the carrier. The sideband tones are half the amplitude of the carrier tone. Using trigonometric identities, derive the frequencies that are present in the SAM tone, and their relative amplitudes.

6. a) Using the empirical relationship between center frequency and the equivalent rectangular bandwidth (ERB) of the cochlear filter at that frequency:

$ERB(f) = 24.7 + f/9.265$, where f is in Hz,

make a plot of the fastest SAM modulation that should be audible to a human listener as a function for the carrier frequency f_c. For simplicity assume that the cochlear filters have rectangular frequency responses (i.e., that the response is zero to frequencies that fall outside the pass-band).

b) What is the fastest SAM modulation of a 1000 Hz carrier that should be audible as such to a human listener? Write a Matlab/Python function to generate a SAM tone with an rms level of 0.1. What is the fastest modulation that you can hear for a carrier frequency of 1000 Hz? Does this correspond to the theoretical prediction?

8. EXTRA CREDIT: Consider the task of detecting a tone. Assume that internal noise is normally distributed with a variance that is independent of whether the tone is present or not:

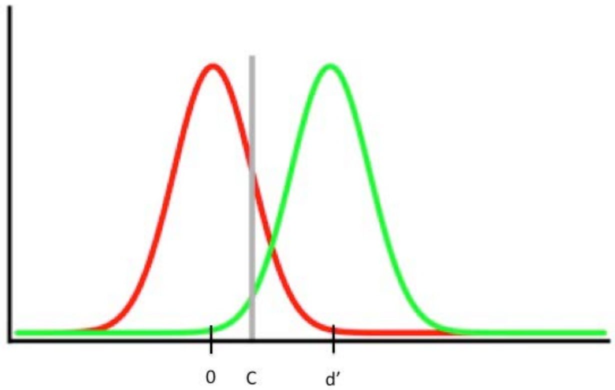

where the graph plots the probability of a particular internal response level when the tone is present (green) and absent (red), and the gray line represents the observer's criterion.

Write Matlab/Python functions to convert the data you would collect in a psychophysical experiment (hit-rate, miss-rate, false-alarm-rate, and correct-rejection-rate) into:

a) d', the standard measure of sensitivity (the distance between the mean of the signal and noise distributions, in z-scores)

b) the observer's criterion (in units of standard deviation of the internal noise distribution, relative to the point where signal and noise are equiprobable given the internal response)

c) the area under the receiver-operating-characteristic (ROC) curve (which plots the hit rate vs. the false alarm rate as they vary for different settings of the criterion).

Hints:
i) make use of *normcdf* (look into scipy.stats.norm), the cumulative distribution function for a standard Normal distribution (mean 0, standard deviation 1):

$$normcdf(c) = \int_{-\infty}^{c} p(x)dx = P(X \leq c)$$

where $p(x)$ is the standard Normal distribution
as well as its inverse, *norminv (look into scipy.stats.ppf)*.

ii) Assume the coordinate system defined in the figure, where the noise distribution is centered at zero and the internal response is specified in units of z-scores.