MIT CSAIL

6.869 Advances in Computer Vision

Fall 2017

Problem Set 3

Posted: Thursday, September 21, 2017 **Due:** Thursday, September 28, 2017

Submission Instructions: Please submit **three separate files:** 1) a report named ⟨your_kerberos⟩.pdf, including your responses to all required questions with images and/or plots showing your results, 2) a video file named ⟨your_kerberos⟩.avi containing your output for Problem 1d, and 3) a file named ⟨your_kerberos⟩.zip, containing relevant source code. **Submissions that do not adhere to these instructions are subject to an additional penalty.**

Late Submission Policy: We do not accept late submissions. The submission deadline has a 50-minute soft cut-off; after midnight Thursday, submissions are penalized 2% per minute late.

Collaborators: You are free to discuss problems with other students but all writing must be done individually. Please list all collaborators at the top of your report.

Readings: Lectures 3, 4, 5, and the corresponding class notes chapters.

Problem 1: Motion Magnification In this problem we will investigate motion magnification in videos. Recall that position shifts in image space correspond to phase shifts in the frequency domain of the Fourier transform. This means that for two images, we can compare the Fourier transform of the two images to find the phase shift between the images. Amplifying the phase shift by a fixed factor in the Fourier transform frequency domain will amplify the position shift by the same factor in the image domain after we perform the inverse Fourier transform. We will use this idea to exaggerate the motions in videos.

(a) For a purely horizontal offset of an impulse signal, magnifying the phase shift will result in a magnified horizontal offset after the inverse transfom. Please fill in lines 8 and 11 in `magnifyChange.m`. You should find the phase shift between the two input images and magnify it by the specified `magnificationFactor`. When complete, the function `magnifyChange` should return an image showing what image 2 would look like with the magnified offset. Please run `part_a.m` and submit the generated plot.

(b) If there is motion in more than one direction between two images, we will see that naively magnifying the phase shift of the whole images will not work. In `part_b.m`, we have set up a vertical offset of an impulse signal as well as the horizontal one from part a. Please run `part_b.m` and submit the generated plot, then explain why the two offsets were not properly magnified.

(c) One strategy we can use if there are multiple motions between two images is to do a localized Fourier transform by independently magnifying the offsets on small windows of the images and aggregating the results across the windows. When we restrict our window of consideration, it is more likely for everything in the window to be moving the same way. We will use Gaussian filters to mask small windows of the image and perform magnification on each window independently. In `part_c.m`, please fill in the Gaussian filter in line 30 and the appropriately windowed input images in line 32. Since we are working with images, we will use the discrete Gaussian filter rather than the continuous one. Run `part_c.m` to confirm that the two motions were properly magnified and submit the generated plot.

(d) We are now ready to apply motion magnification to videos. We will use the same approach as in part c of magnifying Gaussian windowed regions of the video frames. Rather than directly finding the phase shifts between consecutive video frames, we will keep a moving average of the Fourier transform phases and compare each new frame's DFT phase with the current moving average of phase. The moving average is an IIR low-pass filter, averaging 0.5 times the previous average with 0.5 times the current phase. For simplicity, each of the RGB channels are processed independently and identically. In `video_momag.m`, you will need to fill in the Gaussian filter in line 49, the DFT phase of the magnified window in line 69, and the DFT of the magnified window in line 72. Please run `video_momag.m` and submit the generated video. Note that the code may take some time to run - you can temporarily modify `sigma` to decrease the number of windowed regions to process.

Problem 2: Color

For this problem you should use Matlab for computations and plotting. Please include all plots as part of your report. Recall that two colors are a perceptual match if they look like the same color to our eyes. We say that a set of primaries, P, is associated with a set of color matching functions, C, if: $Ct_1 = Ct_2$ for two light spectra t_1 and t_2 that are perceptual matches, and if $CP = eye(3)$. In words, we project an input spectrum onto the color matching function associated with each primary to determine the amount of that primary needed to give a perceptual match to the input spectrum. In all data provided in this example, the data points are sampled at wavelengths $[360 : 5 : 730]nm$.

(a) We can transform the color coordinate in CIE XYZ space to RGB space using a transformation matrix T:

$$\begin{pmatrix} R \\ G \\ B \end{pmatrix} = T \times \begin{pmatrix} X \\ Y \\ Z \end{pmatrix}, \text{where } T = \begin{bmatrix} 3.24 & -1.54 & -0.50 \\ -0.97 & 1.88 & 0.04 \\ 0.06 & -0.207 & 1.06 \end{bmatrix}.$$

(i) Using the transformation matrix T (*CIE2RGB.mat*) and the color matching functions for the CIE XYZ color space (*CIEMatch.mat*), compute the color matching functions associated with those specified by RGB primaries. Plot the RGB color matching functions as a function of wavelength and comment on their positivity.

(ii) Find a valid set of primary light spectra associated with the RGB color space. Plot them as a function of wavelength. Comment on the positivity of the power spectra. *Hint:* Applying the color matching functions to a spectrum tells how much of each primary is needed to perceptually match that spectrum. If the spectrum happens to be exactly that of one of the primaries, then one times that primary, and zero times the others, will match.

(iii) Plot the CIE color matching functions (*CIEMatch.mat*) as a function of wavelength. Comment on their positivity.

(iv) Find a valid set of primary light spectral associated with the CIE color space. Plot them as a function of wavelength. Comment on the positivity of the power spectra.

(b) Figure 1 shows the spectral response curves for eye photoreceptors (we have also provided the response curves in *LMSResponse.mat*). Find a set of primary lights that correspond to the spectral sensitivity curves of the eye. Comment on the positivity of the power spectra.

(c) Show that if the spectral response curves of the eye (assumed to be nonnegative) were orthogonal to each other (with a zero dot product), there would exist a corresponding set of primaries with power spectra that were nonnegative. *Hint:* Try to construct nonnegative primaries from the color matching functions. It may help to review parts (a) and (b) of this problem.

(d) *6.869 only:* For a set of primaries, P, and an associated set of color matching functions, C: If the primaries P are specified, are the color matching functions C uniquely determined? If the color matching functions are specified, are the corresponding primaries P uniquely determined?

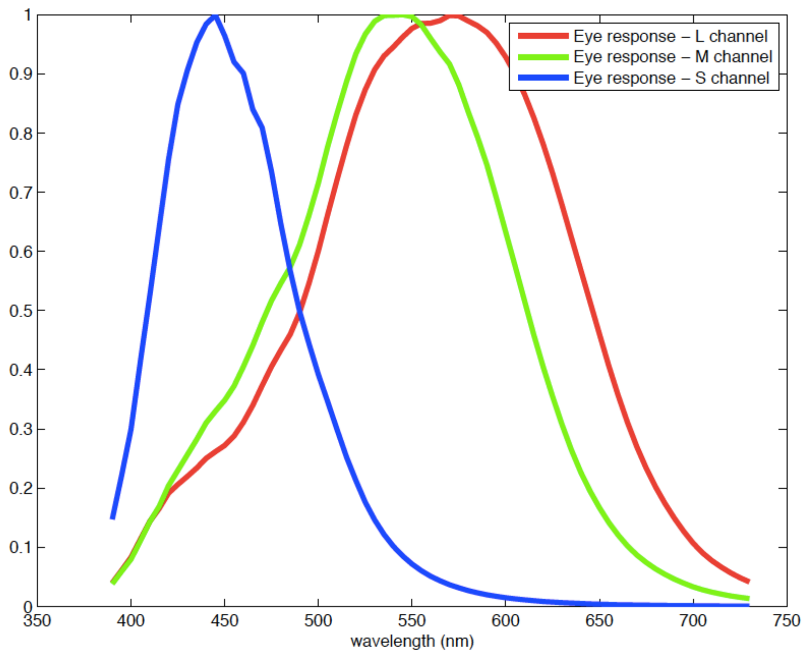

Figure 1: The spectral response curves for eye photoreceptors.