## PSYCH-GA.2211/NEURL-GA.2201 – Fall 2021 Mathematical Tools for Neural and Cognitive Science

## Homework 2

Due: 12 Oct 2021 (late homeworks penalized 10% per day)

See the course web site for submission details. Reminder: rather than using the functions pinv() and norm(), use the linear algebra tools we learned in class. Please: don't wait until the day before the due date... start *now*!

1. Trichromacy. Load the file colMatch.mat in your MATLAB environment (or use scipy.io.loadmat for Python). This file contains matrices and vectors related to the color matching experiment presented in class. In particular, the variable P is an  $N \times 3$  matrix containing wavelength spectra for three "primary" lights, that could be used in a color-matching experiment. For these problems N = 31, corresponding to samples of the visible wavelength spectrum from 400nm to 700nm in increments of 10nm.

The function humanColorMatcher.p simulates a normal human observer in a color matching experiment. For Python, download the file and use from human\_color\_matcher import human\_color\_matcher. The input variable light should contain the wavelength spectrum of a test light (a 31-dimensional column vector). The input variable primaries should contain the wavelength spectra of a set of primary lights (typically, a  $31 \times 3$  matrix, as for matrix P described above). The function returns a 3-vector containing the observer's "knob settings" - the intensities of each of the primaries that, when mixed together, appear identical to the test light. The function can also be called with more than one test light (by passing a matrix whose columns contain 31-dimensional test lights), in which case it returns a matrix whose columns are the knob settings corresponding to each test light.

- (a) Create a test light with an arbitrary wavelength spectrum, by generating a random column vector with 31 positive components (use rand in MATLAB or np.random.rand in Python). Use humanColorMatcher to "run an experiment", asking the "human" to set the intensities of the three primaries in P to match the appearance of the test light. Compute the 31-dimensional wavelength spectrum of this combination of primaries, plot it together with the original light spectrum, and explain why the two spectra are so different, even though they appear the same to the human.
- (b) Your colleague down the hall, Dr. Evoprimary, proudly tells you about a new colormatching configuration using primaries derived from pigments that were most prevalent in the environment of our evolutionary ancestors. You respond that this is a beautiful concept, but seems unlikely to offer new insights into human trichromacy, since you can precisely predict the color matches that will be obtained with these new primaries. Dr. E allows you to measure the wavelength spectra of the new primaries (stored in variable eP). Derive (write math, and explain logic) an expression for the color-matching matrix (31x3, maps a light to 3 knob settings) that predicts matches that would be obtained in Dr E's lab. Compute this matrix, making use of humanColorMatcher.p with your own primaries, P. Check, for a random test light, that the predicted mixture of eP primaries

matches (produces the same knob settings when tested with *your* primaries, P). If it does not, explain why.

- (c) The variable **Cones** contains (in the rows) approximate spectral sensitivities of the three color photoreceptors (cones) in the human eye: Cones(1,:) is for the L (longwavelength, or red) cones, Cones(2,:) the M (green) cones, and Cones(3,:) the S (blue) cones (for Python users, the indexing starts from 0). Applying the matrix Cones to any light l yields a 3-vector containing the average number of photons absorbed by that cone (try plot(Cones') to visualize them!). Verify that the cones provide a physiological explanation for the matching experiment, in that the cone absorptions are equal for any pair of lights that are perceptually matched. First, do this informally, by checking that randomly generated lights and their corresponding 3-primary matching lights produce equal cone absorptions. Then, provide a few lines of matlab code that provide a more mathematical demonstration, along with an extended comment explaining your reasoning using concepts from linear algebra. [Hints for two possible approaches: (1) write math/code that computes cone responses for any test light and then computes the weighted combination of primaries that would produce the same cone responses - show that this is numerically the same as the color-matching matrix; (2) convince yourself, and explain why, it is sufficient to show that M and Cones have the same nullspace. Then use SVD to demonstrate that this is true!]
- Polynomial regression. Load the file regress1.mat into your MATLAB or Jupyter note-book environment. Plot variable Y as a function of X. Find a least-squares fit of the data with polynomials of order 0 (a constant), 1 (a line, parameterized by intercept and and slope), 2, 3, 4, and 5. [Compute this using svd and basic linear algebra manipulations that you've learned in class!] On a separate graph, plot the squared error as a function of the order of the polynomial. Which fit do you think is "best"? Explain.
- 3. Constrained Least Squares Optimization. Load the file constrainedLS.mat into MAT-LABor Jupyter notebook. This contains an  $N \times 2$  data matrix, data, whose columns correspond to horizontal and vertical coordinates of a set of 2D data points,  $\vec{d_n}$ . It also contains a 2-vector w. Consider a constrained optimization problem:

$$\min_{\vec{\beta}} \sum_{n} \left( \vec{\beta}^T \vec{d}_n \right)^2, \quad \text{s.t.} \quad \vec{\beta}^T \vec{w} = 1.$$

There is a family of possible vectors  $\vec{\beta}$  that satisfy the *constraint*  $\vec{\beta}^T \vec{w} = 1$ . Geometrically, any  $\vec{\beta}$  whose arrow-tip lies on a specific line perpendicular to  $\vec{w}$  will satisfy the constraint. The perpendicular distance of this constraint line from the origin will be  $1/||\vec{w}||$  from the origin (think about the dot product, draw the vector  $\vec{w}$  and the constraint line to prove this to yourself).

- (a) Rewrite the optimization problem in matrix form. Then rewrite the problem in terms of a new optimization variable,  $\tilde{\beta}$  (i.e. 'beta tilde', a linear transformation of  $\vec{\beta}$ ), such that the quantity to be minimized is now  $||\tilde{\beta}||^2$ . Note: you must also rewrite the constraint in terms of  $\tilde{\beta}$ .
- (b) The transformed problem is one that you should be able to solve. In particular, you must find the shortest vector  $\tilde{\beta}$  that lies on the constraint line. Compute the solution for  $\tilde{\beta}$ , and plot the solution vector, the constraint line and the transformed data points.

- (c) Transform the solution back into the original space (i.e., solve for  $\vec{\beta}$ ). Plot  $\vec{\beta}$ , the original constraint line, and the original data points. Is the optimal vector  $\vec{\beta}$  perpendicular to the constraint line? On the same graph, plot the total least squares solution (i.e., the vector that minimizes the same objective function, but that is constrained to be a unit vector). Are the two solutions the same?
- 4. Dimensionality reduction with PCA. Professors Hugh Bell and Wi Zell were recording extracellular action potentials (i.e. spikes) from cat primary visual cortex late one evening when their computer malfunctioned. It had already isolated a set of 400 time windows in which voltages had crossed a threshold, indicating the presence of spike. But these traces likely arose from multiple cells, with each cell producing a characteristic waveform, and the computer failed before sorting the voltage traces to determine how many cells were present, and which spikes arose from each cell. The professors come to you (the only math-tools-enabled graduate student still in the building at that hour), asking for help. They provide you with a file windowedSpikes.mat containing a 400 × 150 matrix, data, whose rows contain the electrode measurements (voltages recorded for each 150 msec window, at 1msec intervals). Your task is to determine how many neurons produced these 400 spikes.
  - (a) Plot the 400 waveforms superimposed and describe what you see. Be sure to label your axes! Using these spike waveform plots, can you devise a way to deduce how many neurons produced these spikes? Feel free to include an additional plot containing just a subset of the waveforms in order to aid in your explanation.
  - (b) Perform principal components analysis (PCA) on your data, and plot the eigenvalues in descending order (alternatively, compute the SVD of data). It might help to display the eigenvalues on a log-scale. Interpret what you see.
  - (c) Measure the length of the projection of each of the 400 spike waveforms onto the top two principal components of the dataset, and plot the resulting values as points in 2 dimensions. Describe what you see. Can you deduce how many distinct neurons produced the 400 voltage traces?
  - (d) Now project each waveform onto the top three principal axes, and plot in 3 dimensions (you may want to spin it around, using rotate3d in matlab). Are there any significant changes you see? Using the 3D plot, can you inform Drs. Bell and Zell how many neurons they likely recorded from?