Final Exam Study Guide 2018

For all plots label your axes and whenever possible include a numerical scale with units.

Useful formulas: *eigenvalues and eigenvectors of 2 by 2 symmetric matrices.*

A symmetric 2 by 2 matrix has this general form:

$$A = \begin{bmatrix} a & b \\ b & c \end{bmatrix}$$

The eigenvalues of A are obtained by solving:

$$\det(A - \lambda I) = 0$$

Which leads to these formulas:

$$\lambda_1 = \frac{(a+c)}{2} + \frac{\sqrt{(a-c)^2 + 4b^2}}{2}$$

$$\lambda_2 = \frac{(a+c)}{2} - \frac{\sqrt{(a-c)^2 + 4b^2}}{2}$$

The eigenvectors span the following directions:

$$f_1 = \begin{bmatrix} b \\ \lambda_1 - a \end{bmatrix}, \quad f_2 = \begin{bmatrix} \lambda_2 - c \\ b \end{bmatrix}$$

Note that these are **not unit vectors!**
These formulas will be given to you and are only valid for 2 by 2 **symmetric matrices**.

1. Perceptrons

You will be presented with a data set plotted in the input space u_1-u_2, where a perceptron can correctly classify the examples.

We expect you to be able to:

A. Sketch a possible decision boundary for the classification problem presented.
B. Compute the synaptic weight vector corresponding to the decision boundary you sketched in part A for a specified neural threshold.
C. Sketch the perceptron circuit with synaptic weights labeled.

Some classification problems cannot be solved by a single layer perceptron.
For this type of problem we expect you to:

D. Be able to explain why this is the case.

E. Be able to construct a network that combines two single layer perceptrons to solve this problem.

Exercise:

Design a two-layer perceptron that performs the XOR operation specified by the following truth table:

u_1	u_2	output
0	0	0
1	0	1
0	1	1
1	1	0

Assume that the neurons have a threshold of 0.5.

2. Feed-forward linear neural networks

A. Be able to write down a matrix that perform the following operations for data in a 2-dimensional space:

i. Rotations through the origin for any arbitrary specified angle.
ii. Stretch or compression along the *x* or *y* axes for any arbitrary stretch or compression factor.

B. Be able to combine the above matrices to obtain stretches and/or compressions along any direction.
C. Be able to construct a feed-forward neural network that carries out the above matrix operations.

Exercise:

1) Design a feed-forward linear neural network to implement a 60-degree clockwise rotation in two dimensions. Write down the weights connectivity matrix W and sketch the network indicating input neurons, output neurons, synaptic connections and weights.
2) A cloud of points in the *x-y* plane is linearly transformed via matrix multiplication. Construct a matrix M in two dimensions that stretches the cloud by a factor of 5 along a line 60º from horizontal. The figure below shows the untransformed set of points and the stretch direction.

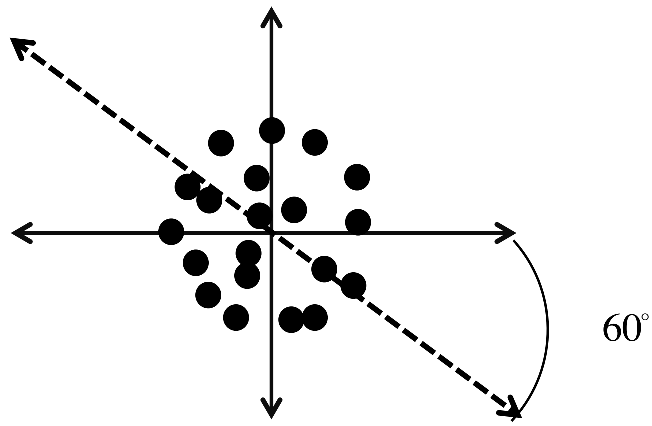

3. Principal Components Analysis

A fundamental problem in neurobiology is how the defining aspects of a stimulus, such as its quality, quantity, and temporal structure, are encoded by the activity of sensory receptors. The problem is particularly intriguing in the case of olfactory stimuli, which are not related according to a single, continuous function such as wavelength or frequency. Odorants, rather, have discrete molecular structures and differ with respect to a variety of physicochemical properties. Moreover, the natural odors that an olfactory system must encode are generally complex mixtures of varying concentration and duration.

Odorant receptors are transmembrane proteins encoded by large gene families. The fruit fly *Drosophila* has a highly diverse family of 60 odorant receptor genes which are primarily expressed at the olfactory receptor neurons (ORNs). ORNs that express the same odorant receptor project to the same glomerulus, or functional processing unit, in the antennal lobes (one of the two primary olfactory organs in flies).

A group of researchers examined the antennal odorant receptors of *Drosophila* using a mutant antennal neuron that lacks endogenous odorant receptors, an in vivo expression system called ''the empty neuron''. Individual receptors were expressed one at a time in the empty neuron, and odorant responses were assayed electrophysiologically by measuring the number of spikes fired during odorant presentation. In this study, 24 odorant receptors were included, and were each tested with 150 different odorants. The odorants encompassed 5 different chemical families (alcohols, aldehydes, terpenes, amines, esters) that were chosen to represent a broad range of ecologically relevant odors.

After the experiment, the researchers performed PCA to test the idea that a few components would allow them to cluster the responses by odorant chemical family. To do this, they constructed a matrix **X**, suitable for doing PCA, with the response of each odorant to each odor.

A. Make a sketch of matrix X, and label what the rows and columns of this matrix represent. Also, indicate the dimensions of matrix X.
B. Write a recipe (i.e. a step by step prescription/algorithm) to perform PCA on this data set. For each step on your recipe:

    a. Write one sentence explaining what the step is doing.
    b. Write either a matrix operation or a line of MATLAB code that implements the step.

After performing PCA, it turns out that the largest 2 PC's explain roughly 80% of the total variance in the data. The other components are not informative.

C. How do you determine the percentage of variance explained by each principal component? How do you plot the eigenvalue 'spectrum'?
D. Write down the necessary matrix operation to rotate the data set into the space spanned by the first 2 largest principal components.
E. How would you interpret the following plot of the data in the space of PC1 vs. PC2?

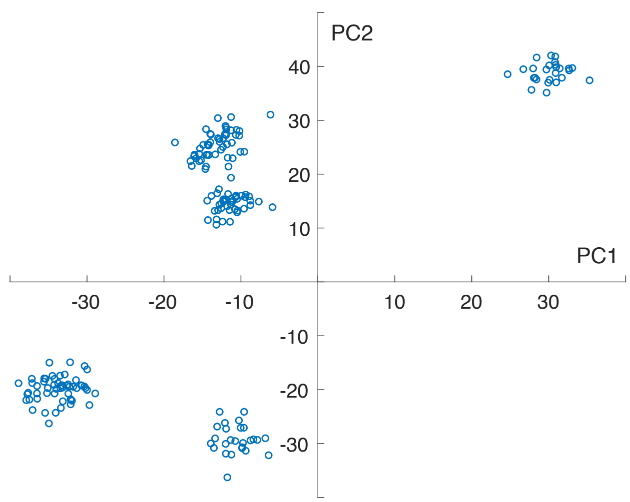

Additional exercise:

You will use Principal Components Analysis to find the vector along which the following data points exhibit maximum variance: {(0,0), (0.5,1.5), (1.5,0.5), (2,2)}. We can write the points in matrix notation as:

$$X = \begin{bmatrix} 0 & 0.5 & 1.5 & 2 \\ 0 & 1.5 & 0.5 & 2 \end{bmatrix}$$

A. Find the mean vector for this data set, and write out the mean-subtracted data matrix *Z*.
B. Sketch in the *x-y* plane the mean-subtracted data points.
C. Write down the covariance matrix as an expression involving the matrix *Z*. Compute the elements of this matrix numerically using the data points above.
D. Calculate (or guess) the eigenvectors of this matrix. Write them as unit vectors. Compute the eigenvalues of the covariance matrix. Show that the eigenvector equations are satisfied. Write the eigenvectors as a change of basis matrix $\mathbf{\Phi}$.
E. What do the eigenvalues of the covariance matrix represent?
F. Write the matrix expression to project the mean subtracted data points (*Z*) onto the basis set of the eigenvectors.
G. Sketch the data points in the new basis set.

4. Linear recurrent neural networks

Consider the symmetric recurrent neural network diagrammed below. The two neurons are linear and the intrinsic time-constant of each neuron is τ_n =100 msec. The autapses are inhibitory with weight -0.1. The recurrent excitatory connections have weight 1.1. The firing rates of the neurons are given by v_1 and v_2.

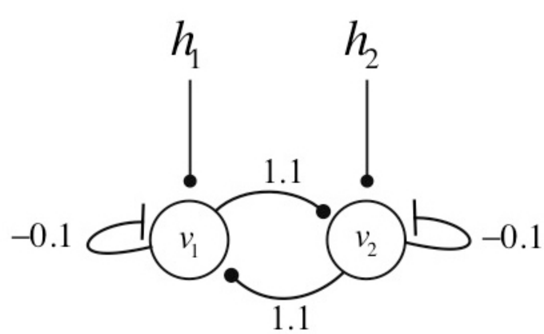

A. Write down the weight matrix **M**.
B. The eigenvectors of the recurrent connection matrix **M** can be written as:

$$\hat{f}_1 = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \hat{f}_2 = \frac{1}{\sqrt{2}}\begin{pmatrix} -1 \\ 1 \end{pmatrix}$$

Find the eigenvalue of M associated with each eigenvector and show that the these eigenvalues solve the eigenvalue equation.

C. For each mode (c_1 and c_2) write down a differential equation describing the response of mode activity to an input vector $\vec{h}$.
D. What is the time constant of each mode if the intrinsic time constant is 100 msec?
E. We present to this network the following input:

$$\vec{h}(t) = \begin{cases} \vec{0}\text{ Hz} & t < 0\text{ sec} \\ (-15,15)\text{ Hz} & 0\text{ sec} \le t < 5\text{ sec} \\ \vec{0}\text{ Hz} & t \ge 5\text{sec} \end{cases}$$

i) Plot the vector $\vec{h}(t = 1\text{ sec})$ in the h_1-h_2 plane. On the same axes, indicate the directions of the eigenvectors $\widehat{f_1}$ and $\widehat{f_2}$.
ii) What are the firing rates of each mode to this input 5 seconds after the stimulus is turned off (t = 10 sec)? Assume that the initial firing rates are zero.
iii) What is the steady state of this network (i.e. the firing rates, v_1 and v_2, of the output neurons). 5 seconds after the stimulus is turned off (t = 10 sec)? Assume that the initial firing rates are zero.

F. Now consider the following input:

$$\vec{h}(t) = \begin{cases} \vec{0}\text{ Hz} & t < 0\text{ sec} \\ (15,15)\text{ Hz} & 0\text{ sec} \le t < 5\text{ sec} \\ \vec{0}\text{ Hz} & t \ge 5\text{ sec} \end{cases}$$

i) Plot the vector $\vec{h}(t = 1\text{ sec})$ in the h_1-h_2 plane. On the same axes, indicate the directions of the eigenvectors $\widehat{f_1}$ and $\widehat{f_2}$.
ii) Sketch qualitatively the activities of mode 1 and as a function of time during an interval from t = -5 to t = 10 seconds. Assume that the initial firing rates are 0. Label your axes.
iii) What are the activities of mode 1 and mode 2 (in Hz) at t=10 seconds, i.e. 5 seconds after the input is turned off? Assume that the initial firing rates are zero. Please provide a numerical answer to the nearest 0.1Hz.
iv) What are the firing rates (v_1, and v_2) of the output neurons at this time (t = 10sec)? Remember that we defined the mode activities (c_1 and c_2) as a function of the neural firing rate vector (v_1, v_2) as follows:

$c_1 = \vec{v} \cdot \widehat{f_1}$, $c_2 = \vec{v} \cdot \widehat{f_2}$

5. Topics from previous midterm exams.

As outlined on the syllabus there will be two questions on the final exam covering topics from the previous midterm exams. Each of these questions is worth 10% of the final exam score.

The first question will be based on the Hodgkin-Huxley model for action potential generation.

The second question will cover: Fourier transforms and the convolution theorem.

Please review the corresponding lecture materials and the study guides to be well prepared for these questions.

Problem 4, numerical answers:

A. $M = \begin{bmatrix} -0.1 & 1.1 \\ 1.1 & -0.1 \end{bmatrix}$

B. $\lambda_1 = 1, \lambda_2 = -1.2$

D. $\tau_{eff,2} = 45.4$ msec . Mode 1 does not have a time constant (integrator).

E. ii) $\vec{c}\ (t = 10 \text{ sec}) = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$ Hz

iii) $\vec{v}(t = 10 \text{ sec}) = \vec{0}$ Hz

F. iii) $\vec{c}(t = 10 \text{ sec}) = \begin{bmatrix} 1060.7 \\ 0 \end{bmatrix}$ Hz

iv) $\vec{v}(t = 10 \text{ sec}) = \begin{bmatrix} 750 \\ 750 \end{bmatrix}$ Hz

MIT OpenCourseWare
https://ocw.mit.edu/

9.40 Introduction to Neural Computation
Spring 2018

For information about citing these materials or our Terms of Use, visit: https://ocw.mit.edu/terms.