University of California
College of Engineering
Department of Electrical Engineering
and Computer Sciences

E. Alon, B. Ayazifar, C. Tomlin
G. Ranade

Thurs., Sep. 24, 2015
11:00-12:30pm

EECS 16B: FALL 2015—MIDTERM 1

Important notes: Please read every question carefully and completely – the setup may or may not be the same as what you have seen before. Also, be sure to show your work since that is the only way we can potentially give you partial credit.

NAME	Last Solutions First

SID	

Login	

Problem 1: _____/ 13

Problem 2: _____/ 18

Problem 3: _____/ 23

Problem 4: _____/ 10

Total: _____/ 64

PROBLEM 1. Signal Analysis with DFT (13 points)

In this problem we will look at a few examples of signal analysis with DFT and how changes in the characteristics of the signals and/or the DFT window impact the resulting frequency-domain representation.

Throughout this problem, you must explain any answers you give in order to receive any credit – i.e., simply guessing an answer will result in zero points.

a) **(3 pts)** Shown below are two time domain signals (x[n], y[n]) and two magnitude plots of DFT coefficients (A[k], B[k]) – which DFT coefficients correspond with which time domain signal? (I.e., does A[k] correspond to x[n], or to y[n]?) Be sure to explain your answer.

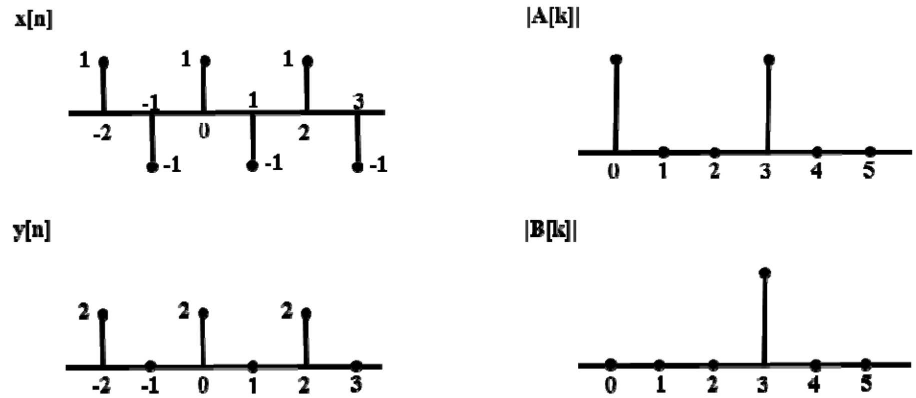

$y[n] \longleftrightarrow B[k]$ Because $y[n]$ has non-zero average value and hence must have a non-zero coefficient for $k=0$

$x[n] \longleftrightarrow A[k]$ since $x[n]$ is a single tone and has no DC content.

b) **(5 pts)** For the x[n]/y[n] and A[k]/B[k] shown below, now which DFT coefficients correspond to which time domain signal? Hint: you should pay particular attention to the frequencies (k's) where the coefficients have zero value.

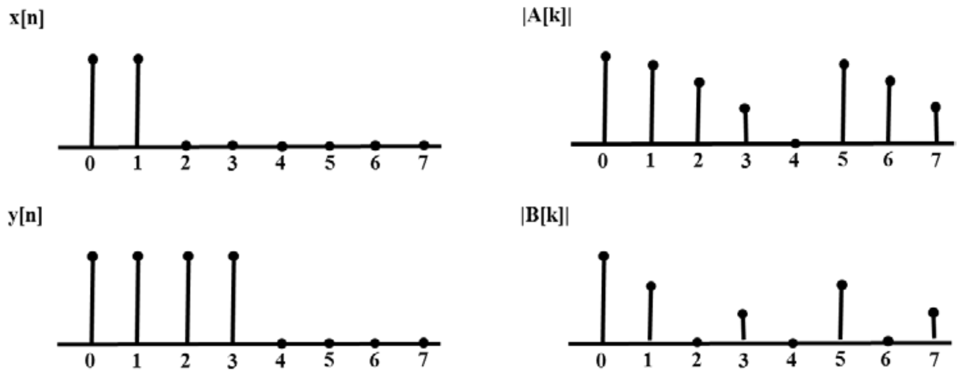

$x[n] \longleftrightarrow A[k]$ because for $x[n]$ only $e^{i\frac{2\pi}{8}\cdot 4n} = e^{i\pi n} = (-1)^n$ should have a value of zero.

$y[n] \longleftrightarrow B[k]$ because for $y[n]$ $e^{i\frac{2\pi}{8}2n}$, $e^{i\frac{2\pi}{8}4n}$, and $e^{i\frac{2\pi}{8}6n}$ $\left(= e^{-i\frac{2\pi}{8}2n}\right)$ should have a value of zero

c) **(5 pts)** If x[n] = $e^{i(4\pi/\mathrm{M})n} + e^{i(6\pi/\mathrm{M})n}$ where M is some positive integer, what is the minimum number of points you must use in your DFT in order for X[k] to contain exactly two non-zero coefficients?

For general M, the only way to guarantee no spectral leakage (integer number of periods) is to make $\omega_0 = \frac{2\pi}{M}$ (s.t. first component corresponds to $k=2$ and second corresponds to $k=3$).

So, need to use an M-point DFT.

PROBLEM 2. Spectral Leakage and Windowing (18 pts)

As we saw in lecture and homework, if the signal we are taking the DFT of is periodic, but an exact integer period of the signal does not fit within the window we are using, we will end up with what is known as *spectral leakage*. In practice this situation is actually extremely common; e.g., consider a BMI system where the physical mechanisms that set the frequency of an oscillation in say an EEG signal have nothing to do with the sampling rate (and hence time window) of our electronics. In this problem we will examine a technique known as "windowing" that tries to mitigate spectral leakage by multiplying the original signal $x[n]$ with a "window function" $w[n]$. Conceptually, the window function attempts to taper the signal near the boundaries of the DFT interval in order to make the signal "fit better".

Let's assume that our signal of interest is $x[n] = e^{i\pi n/2}$, and that we will be taking a DFT of length 6 over the interval n = 0, 1, … , 5. The magnitudes of the DFT coefficients for this signal are $|X[0]| = |X[3]| = 1.41$, $|X[1]| = |X[4]| = 3.86$, and $|X[2]| = |X[5]| = 1.04$.

In this problem we'll examine the simplest type of window known as a "boxcar". Specifically, we will be multiplying $x[n]$ by a window $w[n] = 1$ for $0 \le n < 4$ and zero otherwise.

a) **(2 pts)** Sketch the real part of the original signal $x[n]$ and the real part of the new windowed signal $\hat{x}[n]$

$$\hat{x}[n] = x[n] \cdot w[n] = e^{i\pi n/2} \text{ for } 0 \le n \le 3, \text{ 0 otherwise}$$

over the same interval we used to take the DFT as the original signal.

$x[0] = 1$, $x[1] = e^{i\frac{\pi}{2}}$, $x[2] = e^{i\pi}$, $x[3] = e^{i\frac{3\pi}{2}}$, $x[4] = 1$, $x[5] = e^{i\frac{\pi}{2}}$

real part: 1 0 -1 0 1 0

$Re\{x[n]\}$:

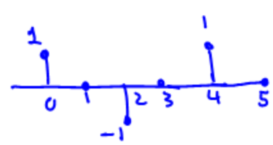

$Re\{\hat{x}[n]\}$:

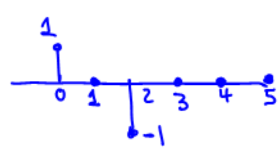

b) **(6 pts)** If we use the same length-6 interval as we did for the original signal to compute the DFT coefficients $\widehat{X}[k]$ of the new signal, for which two values of k will $\widehat{X}[k]$ be equal to zero? You must show your work to receive any credit for this problem.

* Know by looking at the plot from (a) that $\hat{X}[0] = 0$ since the signal has an average value of 0.

* Also know that $e^{i\frac{2\pi}{6}3n} = (-1)^n$ when projected on to $\hat{x}[n]$ will result in $\hat{X}[3] = 0$ because:

$$\underbrace{\begin{bmatrix} 1 & -1 & 1 & -1 & 1 & -1 \end{bmatrix}}_{(-1)^n} \underbrace{\begin{bmatrix} 1 \\ 0 \\ -1 \\ 0 \\ 0 \\ 0 \end{bmatrix}}_{\hat{x}[n]} = 1 \cdot 1 - 1 \cdot 1 = 0$$

c) **(5 pts)** Assuming that $\left|\widehat{X}[0]\right| = \left|\widehat{X}[3]\right|$=0, $\left|\widehat{X}[1]\right| = \left|\widehat{X}[4]\right|$=3.35, $\left|\widehat{X}[2]\right| = \left|\widehat{X}[5]\right|$=0.9, explain how applying the window *w*[n] captures the nature of the original input signal *x*[n] better than not windowing.

Originally had large values for X[0] and X[3], neither of which are even "close" in frequency to the actual signal. These have now both been zero'ed out by the windowing, more accurately reflecting that the signal is actually a single tone. We still of course have 4 non-zero coefficients instead of just one, but it is an improvement over 6 non-zero coefficients (especially since the two we knocked out with windowing were relatively large).

d) **(5 pts)** Defining the energy of a signal y[n] as $E_y = \sum_{n=0}^{5} |y[n]|^2$, compare the energy of the original signal x[n] against the energy of the windowed signal $\hat{x}[n]$. Based on this comparison, comment on how well windowing x[n] preserves the original signal.

Directly in time: $E_x = 6$

(don't forget imaginary part of the signal) $E_{\hat{x}} = 4$

Using Parseval: $E_x = \dfrac{2 \cdot 1.41^2 + 2 \cdot 3.86^2 + 2 \cdot 1.04^2}{6} = 6$

$E_{\hat{x}} = \dfrac{2 \cdot 3.35^2 + 2 \cdot 0.9^2}{6} = 4$

$\hat{x}$ has approximately 2/3 of the energy that x had – this makes sense since the length-4 window is 2/3 of the total DFT window (6). So "most" of the signal energy is preserved, although obviously it isn't perfect.

PROBLEM 3. Three Letter Acronyms (23 points)

In this problem we will explore how using either DFT or SVD to analyze and perhaps even compress (by keeping track of only non-zero coefficients/components) a signal of interest can give us two different views of the same underlying characteristics of a signal. Throughout this problem, we will be considering the signal $x[n] = \cos(\omega_0 n) + \cos(2\omega_0 n)$, where $\omega_0 = (2\pi/10)$.

a) **(1 pts)** If every sample of the signal requires 4 bytes of memory and we keep 1000 total samples of the signal, how many bytes will it take to store the signal?

4 bytes / sample · 1000 samples = 4000 bytes

b) **(8 pts)** Compute the DFT coefficients $X[k]$ corresponding to the signal $x[n]$ for a DFT interval of n=0, 1, … , 9. How many non-zero coefficients are there? If each coefficient and its corresponding frequency index (i.e., $X[k]$ and k) require a total of 12 bytes of memory to store, how many bytes does it take to completely represent the signal?

Should be 4 non-zero coefficients; easiest to find with inverse Euler:

$$x[N] = \frac{1}{2}e^{i\frac{2\pi}{10}N} + \frac{1}{2}e^{-i\frac{2\pi}{10}N} + \frac{1}{2}e^{i\frac{2\pi}{5}N} + \frac{1}{2}e^{-i\frac{2\pi}{5}N}$$

$X[1] = 5$ $\qquad$ $X[2] = 5$

$$e^{\frac{2\pi}{10}\cdot 9N} = e^{i\frac{2\pi}{10}\cdot 10N} \cdot e^{-i\frac{2\pi}{10}\cdot N} \qquad e^{\frac{2\pi}{10}\cdot 8N} = e^{i\frac{2\pi}{5}\cdot 5N} \cdot e^{-i\frac{2\pi}{5}N}$$

$X[9] = 5$ $\qquad$ $X[8] = 5$

Need 48 bytes total

c) **(5 pts)** Now let's use a procedure similar to what we did with the neural data in the lab to construct a matrix based off of this signal that we can then analyze with SVD. Specifically, imagine that we construct a matrix A by taking 10-sample long intervals of the signal x[n] and using each one of those to populate the rows of the matrix – i.e.,:

$$A = \begin{bmatrix} x[0] & x[1] & \ldots & x[9] \\ x[10] & x[11] & \ldots & x[19] \\ \vdots & \vdots & \vdots & \vdots \end{bmatrix}$$

If we were to take the SVD of the matrix A, how many non-zero singular values will the matrix have in this case?

Key is to realize that $x[n]$ is periodic with period 10, so every row of A will actually be identical.

Hence, the matrix (no matter how many rows it has) will be rank 1, and will thus have only 1 non-zero singular value.

d) **(3 pts)** Given your answer to part c), if every singular value or entry in a singular vector requires 4 bytes of memory to store, how many total bytes does it take to completely represent the signal?

Need 1 singular value and 10 entries for the one right singular vector, so a total of 44 bytes.

e) **(6 pts)** If we were to construct a new matrix B using the same approach as we had taken to construct A in part b), but using intervals of length 5 (instead of length 10), how many non-zero singular values will the new matrix B have?

Now the rows will be:

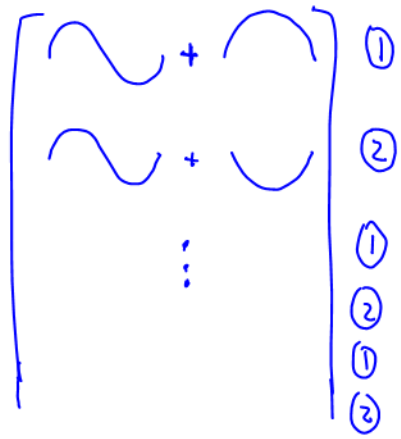

or, conceptually

$$\begin{bmatrix} x^T+y^T \\ x^T-y^T \\ x^T+y^T \\ x^T-y^T \\ \vdots \end{bmatrix}$$

So the matrix will be rank 2, and we will have 2 non-zero singular values.

PROBLEM 4. Fortune Telling (10 pts + BONUS 5 pts)

Using your newfound knowledge of PCA and the SVD, you decide to play a little trick on a few (or even a few hundred) of your closest friends and gather some data that will allow you to make an educated prediction about whether or not they will successfully buy a house in the ultra-competitive Bay Area real estate market within the next 5 years. Knowing that each of these pieces of information in some way are likely to impact this, you collect the following data from each of your friends:

- Age (in years)
- How many years they have been with their current partner (0 if they are single)
- Age of their partner (0 if they are single)
- How many siblings they have
- Approximate salary (in \$/year)
- Approximate salary of their partner (in \$/year, 0 if single)
- Their height (in m)
- The height of their partner (in m, 0 if single)
- The number of times per week they go clubbing (0 if not single)
- Their overall satisfaction with life (rated on a scale from 0-10, with 10 being satisfied to the point of annoying everyone around them)

For the rest of this problem, let's assume that you successfully gathered this information from 100 of your friends.

a) **(3 pts)** Describe how you would construct a matrix *A* out of the data collected above that we might later be able to analyze in order to make a prediction about whether your friend will end up buying a house. You should arrange *A* such that the information from each friend is arranged in a row of the matrix. Be sure to indicate what the dimensions of the matrix would be.

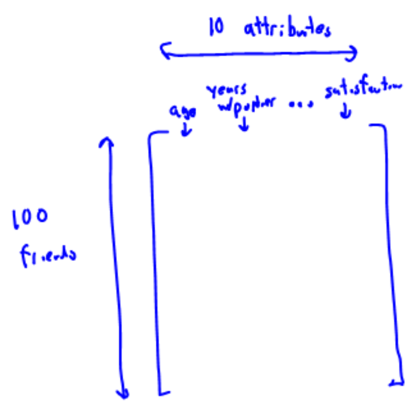

b) **(7 pts)** Assuming that the A matrix has one very dominant singular value and that you find that the data is actually indicative by running a few test cases where you know whether your friend bought a house or not, for any one particular friend i whose data you have collected, describe how you would use the matrix A and their individual data vector a_i^T to predict whether or not they will buy a house in the next 5 years.

* Decompose A in to $U \Sigma V^T$

* Know that have only one principal component since have one dominant singular value, and know that this component is indicative of the question we want answered (by statement in problem).

* So, if we take the top right singular vector V_1 (which is a 10x1 vector) and project a_i onto it by taking a dot product - i.e., $a_i^T \cdot V_1$, we can make a prediction based on whether the result is above or below a certain threshold.
(This threshold would be set by looking for "clusters" in the data and/or by looking at the value of the dot product for friends whose outcome we already know.)

c) **(BONUS: 5 pts)** Using the same method as part b), suggest some other outcome (besides buying a house) that you might be able to make a prediction about, what data you might need, and under what conditions it is likely for the predicted outcome to be accurate. Note that most of the bonus credit will be assigned to addressing the final issue (i.e., under what conditions your prediction about the outcome is likely to be accurate).

The point of this problem was to show that SVD/PCA can be used to predict just about anything as long as you have the appropriate data. "Appropriate" in this case means that not only are there a (typically small) set of dominant principal components, but that when we project onto those components we can identify clusters of observations. If we then have a few known test cases to tell us which clusters correspond to which outcome, we have a good chance that actual predictions (where we don't know the outcome in advance) will be accurate as well.