EECS120 - Fall 2003
Homework No. 4 Solutions
You are strongly urged to check your copy of the homework against the solutions before the exam.
Questions can be asked in ucb.class.ee120 or in office hours

Problem 4.1 *Book Problems from Lee and Varaiya, chapter 10.*
Problems: 11, 13, 14, 15
See Attached Images

Problem 4.2 *Projections and Least Squares.*
This problem is designed to help you review your linear algebra by exploring the special properties of the projection operation and see how it relates to the Discrete Fourier Transform and Fourier Series.

a. *One dimensional projections. Given an* n*-dimensional column vector* $\vec{v}$*, derive a matrix* L *that projects column vectors onto* $\vec{v}$*. (i.e. If* $\vec{y} = L\vec{x}$*, then* $\vec{y}$ *is the projection of* $\vec{x}$ *in the direction of* $\vec{v}$*.)*

   Simple trigonometry shows that the projection of $\vec{x}$ into the direction of $\vec{v}$ has length $\|\vec{v}\|\cos\theta$ along the unit vector in the direction $\vec{v}$ where θ is the angle between $\vec{v}$ and $\vec{x}$.

   We know that $\|\vec{v}\|\|\vec{x}\|\cos\theta = <\vec{v},\vec{x}> = \vec{v}^*\vec{x}$, and also that unit vector in the direction $\vec{v}$ is just $\frac{\vec{v}}{\|\vec{v}\|}$. As such $\frac{<\vec{v},\vec{x}>}{\|\vec{v}\|}$ gives the length of the desired projection. So the answer as a vector is:
$$(\frac{\vec{v}}{\|\vec{v}\|})\frac{<\vec{v},\vec{x}>}{\|\vec{v}\|} = (\frac{\vec{v}}{\|\vec{v}\|^2})\vec{v}^*\vec{x} = (\frac{\vec{v}\vec{v}^*}{\vec{v}^*\vec{v}})\vec{x}$$

   So the matrix $L = \frac{\vec{v}\vec{v}^*}{\vec{v}^*\vec{v}}$.

b. *You perform an experiment that yields a column vector of* n *measurements denoted by* $\vec{d}$*. You believe this data to come from a physical process that should give measurements of the form* $\alpha\vec{a}$ *where* α *is some unknown scalar gain parameter. In order to estimate* α*, you decide to minimize the standard Euclidean norm of the error vector* $\vec{e} = \vec{d} - \alpha\vec{a}$ *where* α *is the parameter you get to set.*

   Show that the optimal α **gives rise to** $\alpha\vec{a}$ **being the projection of** $\vec{d}$ **onto the direction given by** $\vec{a}$**.**

   We follow the hint: Take the derivative of the squared norm of the error vector and set it to zero. Then show that this is the global minimum.

   The norm of the error vector is $\|\vec{e}\|$. Because the norm is positive, and squaring is strictly monotonic over a positive domain, any local extrema of $\|\vec{e}\|$ must also correspond to a local extrema of $\|\vec{e}\|^2 = \vec{e}^*\vec{e} = (\vec{d} - \alpha\vec{a})^*(\vec{d} - \alpha\vec{a}) = \vec{d}^*\vec{d} + \alpha^*\alpha\vec{a}^*\vec{a} - \alpha\vec{d}^*\vec{a} - \alpha^*\vec{a}^*\vec{d}$.

   The complex vector case might be unsettling to some of you, so we will first work out the solution in the real case. In that case, the conjugate transpose above is the same as the simple transpose and the derivative with respect to the single real parameter α is $2\alpha\vec{a}^T\vec{a} - \vec{d}^T\vec{a} - \vec{a}^T\vec{d}$. Setting this to zero and solving for α gives the unique local extremum of

$$\alpha = \frac{\vec{d}^T\vec{a} + \vec{a}^T\vec{d}}{2\vec{a}^T\vec{a}} = \frac{\vec{a}^T\vec{d} + \vec{a}^T\vec{d}}{2\vec{a}^T\vec{a}} = \frac{\vec{a}^T\vec{d}}{\vec{a}^T\vec{a}}$$

and thus

$$\alpha\vec{a} = \vec{a}\alpha = \frac{\vec{a}\vec{a}^T}{\vec{a}^T\vec{a}}\vec{d}$$

which is the same as the projection derived earlier. To verify that this is indeed a minimum and not a maximum, we can just take the second derivative and see that it is always $\vec{a}^T\vec{a} > 0$.

With the real case in hand, we can move on to the complex case. Here, we can not simply use the "complex derivative" since the Cauchy-Riemann equations are not satisfied and so the "complex derivative" does not exist.[1] Instead, we will just consider the complex parameter α as a pair of real parameters α_r, α_i. Since we are just looking for an extremal value, all that is necessary is for us to take the partial derivatives with respect to the real and imaginary parts of α.

Writing out $\|\vec{e}\|^2$ in terms of the real and imaginary parts gives us:

$$\begin{aligned}
\|\vec{e}\|^2 &= \vec{d}^*\vec{d} + \alpha^*\alpha\vec{a}^*\vec{a} - \alpha\vec{d}^*\vec{a} - \alpha^*\vec{a}^*\vec{d} \\
&= \vec{d}^*\vec{d} + (\alpha_r^2 + \alpha_i^2)\vec{a}^*\vec{a} - \alpha_r\mathrm{Re}(\vec{d}^*\vec{a}) + \alpha_i\mathrm{Im}(\vec{d}^*\vec{a}) - \alpha_r\mathrm{Re}(\vec{a}^*\vec{d}) - \alpha_i\mathrm{Im}(\vec{a}^*\vec{d}) \\
&= \vec{d}^*\vec{d} + (\alpha_r^2 + \alpha_i^2)\vec{a}^*\vec{a} - 2\alpha_r\mathrm{Re}(\vec{a}^*\vec{d}) - 2\alpha_i\mathrm{Im}(\vec{a}^*\vec{d}) \\
&= \vec{d}^*\vec{d} + (\alpha_r^2\vec{a}^*\vec{a} - 2\alpha_r\mathrm{Re}(\vec{a}^*\vec{d})) + (\alpha_i^2\vec{a}^*\vec{a} - 2\alpha_i\mathrm{Im}(\vec{a}^*\vec{d}))
\end{aligned}$$

Setting the partial derivatives equal to zero gives us the pair of equations:

$$\begin{aligned}
2\alpha_r\vec{a}^*\vec{a} - 2\mathrm{Re}(\vec{a}^*\vec{d}) &= 0 \\
2\alpha_i\vec{a}^*\vec{a} - 2\mathrm{Im}(\vec{a}^*\vec{d}) &= 0
\end{aligned}$$

Solving the above pair immediately gives us the unique complex

$$\begin{aligned}
\alpha_r &= \frac{\mathrm{Re}(\vec{a}^*\vec{d})}{\vec{a}^*\vec{a}} \\
\alpha_i &= \frac{\mathrm{Im}(\vec{a}^*\vec{d})}{\vec{a}^*\vec{a}}
\end{aligned}$$

which can be expressed in more compact complex notation as

$$\alpha = \frac{\vec{a}^*\vec{d}}{\vec{a}^*\vec{a}}$$

[1] If you do not remember what these conditions are, do not worry. They just assure that the limit of $\epsilon \to 0$ in the definition of the derivative exists in the sense that it does not depend on which direction in the complex plane is used to have the complex $\epsilon \to 0$.

and so

$$\alpha\vec{a} = \vec{a}\alpha = \frac{\vec{a}\vec{a}^*}{\vec{a}^*\vec{a}}\vec{d}$$

which is the same as the projection we wanted.

c. *Show that the error vector $\vec{e}$ is orthogonal to $\vec{a}$.*

To see this, we just calculate the inner product of $\vec{e}$ with $\vec{a}$.

$$\begin{aligned}
< \vec{a}, \vec{e} > &= \vec{a}^*\vec{e} \\
&= \vec{a}^*\vec{d} - \frac{\vec{a}^*\vec{a}\vec{a}^*}{\vec{a}^*\vec{a}}\vec{d} \\
&= \vec{a}^*\vec{d} - \vec{a}^*\vec{d} \\
&= 0
\end{aligned}$$

d. *Repeat part (b) above, but now your model for the measurements has two unknowns (α, β) and is of the form $\alpha\vec{a} + \beta\vec{b}$ where $\vec{a}$ and $\vec{b}$ are linearly independent. What is the choice of α and β that minimizes the norm of the error vector?*

Now, we have two unknowns and

$$\begin{aligned}
\|\vec{e}\|^2 &= (\vec{d} - \alpha\vec{a} - \beta\vec{b})^*(\vec{d} - \alpha\vec{a} - \beta\vec{b}) \\
&= \vec{d}^*\vec{d} + \alpha^*\alpha\vec{a}^*\vec{a} - \alpha\vec{d}^*\vec{a} - \alpha^*\vec{a}^*\vec{d} + \beta^*\beta\vec{b}^*\vec{b} - \beta\vec{d}^*\vec{b} - \beta^*\vec{b}^*\vec{d} + \alpha^*\beta\vec{a}^*\vec{b} + \beta^*\alpha\vec{b}^*\vec{a} \\
&= \vec{d}^*\vec{d} + (\alpha_r^2\vec{a}^*\vec{a} - 2\alpha_r\mathrm{Re}(\vec{a}^*\vec{d})) + (\alpha_i^2\vec{a}^*\vec{a} - 2\alpha_i\mathrm{Im}(\vec{a}^*\vec{d})) \\
&\quad +(\beta_r^2\vec{b}^*\vec{b} - 2\beta_r\mathrm{Re}(\vec{b}^*\vec{d})) + (\beta_i^2\vec{b}^*\vec{b} - 2\beta_i\mathrm{Im}(\vec{b}^*\vec{d})) \\
&\quad +2(\alpha_r\beta_r + \alpha_i\beta_i)\mathrm{Re}(\vec{a}^*\vec{b}) + 2(\beta_i\alpha_r - \alpha_i\beta_r)\mathrm{Im}(\vec{a}^*\vec{b})
\end{aligned}$$

Taking partial derivatives and setting to zero again gives rise to four linear equations:

$$\begin{aligned}
\alpha_r\vec{a}^*\vec{a} + \beta_r\mathrm{Re}(\vec{a}^*\vec{b}) + \beta_i\mathrm{Im}(\vec{a}^*\vec{b}) &= \mathrm{Re}(\vec{a}^*\vec{d}) \\
\beta_r\vec{b}^*\vec{b} + \alpha_r\mathrm{Re}(\vec{a}^*\vec{b}) + \alpha_i\mathrm{Im}(\vec{b}^*\vec{a}) &= \mathrm{Re}(\vec{b}^*\vec{d}) \\
\alpha_i\vec{a}^*\vec{a} + \beta_i\mathrm{Re}(\vec{a}^*\vec{b}) + \beta_r\mathrm{Im}(\vec{b}^*\vec{a}) &= \mathrm{Im}(\vec{a}^*\vec{d}) \\
\beta_i\vec{b}^*\vec{b} + \alpha_i\mathrm{Re}(\vec{a}^*\vec{b}) + \alpha_r\mathrm{Im}(\vec{a}^*\vec{b}) &= \mathrm{Im}(\vec{b}^*\vec{d})
\end{aligned}$$

written in more compact complex notation the above are just:

$$\begin{aligned}
\alpha\vec{a}^*\vec{a} + \beta\vec{b}^*\vec{a} &= \vec{a}^*\vec{d} \\
\alpha\vec{a}^*\vec{b} + \beta\vec{b}^*\vec{b} &= \vec{b}^*\vec{d}
\end{aligned}$$

which we can rewrite in matrix form to give even more insight:

$$\begin{bmatrix} \vec{a}^*\vec{a} & \vec{b}^*\vec{a} \\ \vec{a}^*\vec{b} & \vec{b}^*\vec{b} \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \begin{bmatrix} \vec{a}^* \\ \vec{b}^* \end{bmatrix} \vec{d}$$

The above can be rewritten again by letting the complex matrix $V = [\vec{a}, \vec{b}]$ and then we have:

$$V^*V \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = V^*\vec{d}$$

which can immediately be solved to give us

$$\begin{bmatrix} \alpha \\ \beta \end{bmatrix} = (V^*V)^{-1}V^*\vec{d}$$

To see that this is the global minimum rather than the maximum, we can just notice that it is the only local extreme point and the matrix of second partial derivatives is clearly positive definite since $\vec{a}$ and $\vec{b}$ are linearly independent.

e. *Show that the error vector from part (d) is orthogonal to both $\vec{a}$ and $\vec{b}$ and hence to the entire subspace spanned by the two of them.*

Use the notation we had at the end of (d). $\vec{e} = \vec{d} - V \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \vec{d} - V(V^*V)^{-1}V^*\vec{d} = (I - V(V^*V)^{-1}V^*)\vec{d}$. Now, any vector in the subspace spanned by $\vec{a}$ and $\vec{b}$ can be written as $V\vec{u}$ where $\vec{u}$ is 2-d complex column vector. To see orthogonality, we calculate:

$$\begin{aligned} < V\vec{u}, \vec{e} > &= \vec{u}^*V^*(I - V(V^*V)^{-1}V^*)\vec{d} \\ &= \vec{u}^*(V^* - V^*V(V^*V)^{-1}V^*)\vec{d} \\ &= \vec{u}^*(V^* - V^*)\vec{d} \\ &= \vec{u}^*0\vec{d} \\ &= 0 \end{aligned}$$

f. *Repeat part (d) for the case of a model with m linearly independent vectors $\vec{a}_i$ and m corresponding unknowns α_i.*

We follow the hint: Put the m column vectors $\vec{a}_i$ into a matrix A and then collect all the α_i unknown parameters into a column vector $\vec{x}$. Now the linear combination is given by the matrix multiplication $A\vec{x}$ and the error is $\vec{d} - A\vec{x}$.

The norm squared is $(\vec{d}^* - \vec{x}^*A^*)(\vec{d} - A\vec{x}) = \vec{d}^*\vec{d} + \vec{x}^*A^*A\vec{x} - \vec{x}^*A^*\vec{d} - \vec{d}^*A\vec{x}$. Taking the partial derivatives as we did earlier and setting it to zero will give rise to the complex equation[2]:

$$2A^*A\vec{x} - 2A^*\vec{d} = 0$$

which has the unique solution

$$\vec{x} = (A^*A)^{-1}A^*\vec{d}$$

This gives us the values for the α_i.

[2] Notice that this is the only way the dimensionality of the matrices will match up in a way that makes sense.

To see that it is a local minimum, notice that the matrix of second partial derivatives is just A^*A which has to be positive definite since $\vec{w}^*A^*A\vec{w} =< A\vec{w}, A\vec{w} >= \|A\vec{w}\|^2 > 0$ since A consists of linearly independent columns.

g. *Now, suppose in part (f) above that the vectors $\vec{a}_i$ are all orthogonal to each other. Simplify your answer to part (f) for this case. What is special about orthogonality?*

If the vectors are orthogonal, then the matrix (A^*A) is diagonal. As such, its inverse $(A^*A)^{-1}$ is also diagonal where the ith diagonal entry is just $\frac{1}{\vec{a}_i^*\vec{a}_i}$. As such, the equations for the α_i become particularly simple:

$$\alpha_i = \frac{\vec{a}_i^*\vec{d}}{\vec{a}_i^*\vec{a}_i}$$

The orthogonality lets us calculate the α_i individually without having to calculate any of the others.

h. *BONUS: Show that the error vector in part (f) is orthogonal to the entire subspace spanned by the $\vec{a}_i$.*

One answer is to just repeat the calculation in part (e) above with A playing the role of V above.

The other, more intuitive solution, is to follow the hint. We use part (g) and feed it a set of orthogonal vectors generated from the original $\vec{a}_i$'s by using Gramm-Schmidt on them. We don't actually have to do it, just realize that we can and that it would give you m orthogonal vectors $\vec{o}_i$ spanning the same subspace as the $\vec{a}_i$. So we know that the answers from part (g) and part (f) must describe the same vector and hence:

$$A(A^*A)^{-1}A^*\vec{d} = \sum_{i=1}^{m} \vec{o}_i \frac{\vec{o}_i^*\vec{d}}{\vec{o}_i^*\vec{o}_i}$$

Now, we can generate an o_{m+1} by throwing in $\vec{d}$ to the the list for Gramm-Schmidt orthogonalization. At this point, we know by construction that $\vec{d}$ is in the span of the $m+1$ orthogonal vectors $\vec{o}_i$. So we have:

$$\vec{d} = \sum_{i=1}^{m+1} \vec{o}_i \frac{\vec{o}_i^*\vec{d}}{\vec{o}_i^*\vec{o}_i} = \vec{o}_{m+1}\frac{\vec{o}_{m+1}^*\vec{d}}{\vec{o}_{m+1}^*\vec{o}_{m+1}} + \sum_{i=1}^{m} \vec{o}_i \frac{\vec{o}_i^*\vec{d}}{\vec{o}_i^*\vec{o}_i}$$

Subtracting the two equations gives us:

$$\vec{e} = \vec{o}_{m+1}\frac{\vec{o}_{m+1}^*\vec{d}}{\vec{o}_{m+1}^*\vec{o}_{m+1}}$$

which is clearly orthogonal to all the $\vec{o}_i$ for $0 < i < m+1$ by construction. So the error vector is perpendicular to the entire subspace spanned by the original $\vec{a}_i$.

At this point, you should have a firm understanding of the two major characterizations of the projection: the geometric, as the component of a vector inside a given subspace — so that the residual error is orthogonal to that subspace — and as the solution to a least squares problem. In later courses and in algorithmic use of signal processing ideas, you will have occasion to use both characterizations.

i. *Interpret (g) and extend it to the case of trying to model a finite duration continuous time complex signal $s(t)$ (defined for $t \in [0, T)$ you can also interpret this as a periodic signal with period T and defined for all real t) as a weighted sum of the first $2m + 1$ of the T-periodic complex exponentials $e^{j\frac{2\pi i}{T}t}$, for $i = -m, -m + 1, \ldots, -1, 0, 1, \ldots, +m$. How would you pick the α_i? Does your choice minimize the energy in the error?*

We will follow the hint and think about the results of (g) in terms of inner products and norms as we did in class. We can use the inner product notation to write

$$\alpha_i = \frac{< \vec{a}_i, \vec{d} >}{< \vec{a}_i, \vec{a}_i >}$$

So, we can use the inner product on complex functions defined on $[0, T)$ given by

$$< f, g >= \int_0^T f^*(t)g(t)dt$$

for which we know that the complex exponentials are orthogonal since

$$\begin{aligned} < e^{j\frac{2\pi i}{T}t}, e^{j\frac{2\pi k}{T}t} > &= \int_0^T e^{-j\frac{2\pi i}{T}t} e^{j\frac{2\pi k}{T}t} dt \\ &= \int_0^T e^{j\frac{2\pi(k-i)}{T}t} dt \\ &= \begin{cases} 0 & \text{if } i \neq k \\ T & \text{if } i = k \end{cases} \end{aligned}$$

By the arguments given in the previous section, our resulting choice of

$$\alpha_i = \frac{\int_0^T s(t)e^{-j\frac{2\pi i}{T}t}dt}{T}$$

minimizes the energy left in the error signal.

j. *BONUS: Suppose that you wished to represent a finite duration discrete time signal $s(t)$ with $t \in \{0, 1, \ldots, T - 1\}$ but instead of using the first $2m + 1$ of the T-periodic complex exponentials ($2m + 1 < T$), you decided to use the $2m + 1$ of them that gave the smallest energy in the error. Give an algorithm and justification for why it works. (Variations of this are actually used in doing lossy image and audio compression by combining this idea with another norm that accounts for perceptual effects.)*

This is remarkably straightforward with the results we already have. We can represent finite duration discrete-time signals as T-dimensional complex vectors. To verify that the complex exponentials are an orthogonal basis we just check:

$$\begin{aligned}
< e^{j\frac{2\pi i}{T}t}, e^{j\frac{2\pi k}{T}t} > &= \sum_{0}^{T-1} e^{-j\frac{2\pi i}{T}t} e^{j\frac{2\pi k}{T}t} \\
&= \sum_{0}^{T-1} e^{j\frac{2\pi(k-i)}{T}t} dt \\
&= \begin{cases} 0 & \text{if } i \neq k \\ T & \text{if } i = k \end{cases}
\end{aligned}$$

and then we know that we can represent our signal as $s(t) = \sum_{i=0}^{T-1} \alpha_i e^{j\frac{2\pi i}{T}t}$ where the $\alpha_i = \frac{1}{T}\sum_{t=0}^{T-1} s(t)e^{-j\frac{2\pi i}{T}t}$. If we wanted to represent it by a smaller subset of $i \in \Omega$ where $|\Omega| = 2m+1$, then you know that the least squares representation would be

$$\hat{s}(t) = \sum_{i\in\Omega} \alpha_i e^{j\frac{2\pi i}{T}t}$$

where the α_i are as before. Moreover, we know that the error signal would be:

$$e(t) = \sum_{i\in\bar{\Omega}} \alpha_i e^{j\frac{2\pi i}{T}t}$$

The total energy in the error is just

$$\begin{aligned}
< e, e > &= \sum_{t=0}^{T-1} e^*(t)e(t) \\
&= \sum_{t=0}^{T-1} (\sum_{i\in\bar{\Omega}} \alpha_i^* e^{-j\frac{2\pi i}{T}t})(\sum_{k\in\bar{\Omega}} \alpha_k e^{j\frac{2\pi k}{T}t}) \\
&= \sum_{i\in\bar{\Omega}} \sum_{k\in\bar{\Omega}} \alpha_i^* \alpha_k \sum_{t=0}^{T-1} e^{-j\frac{2\pi i}{T}t} e^{j\frac{2\pi k}{T}t} \\
&= \sum_{i\in\bar{\Omega}} \sum_{k\in\bar{\Omega}} \alpha_i^* \alpha_k \sum_{t=0}^{T-1} e^{j\frac{2\pi(k-i)}{T}t} \\
&= T \sum_{i\in\bar{\Omega}} \alpha_i^* \alpha_i \\
&= T \sum_{i\in\bar{\Omega}} |\alpha_i|^2
\end{aligned}$$

And so, the choice of Ω that minimizes the energy in the error signal is the one that has the smallest sum of square magnitudes for $\sum_{i\in\bar{\Omega}} |\alpha_i|^2$. But squaring is monotonic and the α_i do not depend on the choice of Ω.

This is now like asking if you want to maximize the weight of candy by choosing $2m+1$ different candy pieces, which candy pieces should you choose. The answer is obviously[3] to choose the $2m+1$ biggest pieces of candy. How? Order all the candy and pick the first $2m+1$.[4]

So, the algorithm is to sort the frequencies by $|\alpha_i|$ in descending order and to choose the $2m+1$ frequencies that have the most energy in them.

k. *Show that if $\vec{x} = \sum_{i=0}^{n-1} \alpha_i \vec{v}_i$ and the $\vec{v}_i$ from an orthonormal set, then using the regular Euclidean norm, $\|\vec{x}\|^2 = \sum_{i=0}^{n-1} |\alpha_i|^2$.*

   We already did this in the last part except we had a T coming from the norm of the basis vectors. Here that is just 1.

l. *BONUS: Extend (d) to the continuous-time finite duration case. Suppose that you have a continuous real-valued measurement $d(t)$ defined over the interval $t \in [0, T]$ and you want to model it as the weighted sum of m real-valued signals $a_i(t)$ all defined over the same interval. Give and justify an algorithm involving integrals and finite-sized matrix operations that will give the choice of α_i coefficients that minimizes the energy in the error signal $e(t) = d(t) - \sum_{i=0}^{m-1} \alpha_i a_i(t)$.*

   Here we will use the inner product view and reason by analogy. While earlier we had equations of the form:

$$A^* A\vec{x} - A^* \vec{d} = 0$$

   these can be reinterpreted on a row by row basis as:

$$\vec{a}_i^*(\vec{d} - A\vec{x}) = 0$$

   or in inner product form as:

$$< \vec{a}_i, \vec{d} - A\vec{x} >= 0$$

   This says that the error must be orthogonal to all the vectors $\vec{a}_i$. We can use the inner product on continuous time signals to immediately get a matrix M where

$$m_{i,j} =< a_i, a_j >= \int_0^T a_i^*(t) a_j(t) dt$$

   and similarly a column vector $\vec{c}$ where

$$c_i =< a_i, d >= \int_0^T a_i^*(t) d(t) dt$$

   Then, our desired coefficients can be found by solving $M\vec{x} = \vec{c}$ or

$$\vec{x} = M^{-1}\vec{c}$$

[3] and you can use induction to prove it if you are so inclined

[4] You can do better by using fancier data structures like heaps.