EE263 Final Exam

Summer 2026

Name: ____________________________ **Student ID number:** __________________

The Stanford Honor Code

The Honor Code is an undertaking of the Stanford academic community, individually and collectively. Its purpose is to uphold a culture of academic honesty.

Students will support this culture of academic honesty by neither giving nor accepting unpermitted academic aid in any work that serves as a component of grading or evaluation, including assignments, examinations, and research.

Instructors will support this culture of academic honesty by providing clear guidance, both in their course syllabi and in response to student questions, on what constitutes permitted and unpermitted aid. Instructors will also not take unusual or unreasonable precautions to prevent academic dishonesty.

Students and instructors will also cultivate an environment conducive to academic integrity. While instructors alone set academic requirements, the Honor Code is a community undertaking that requires students and instructors to work together to ensure conditions that support academic integrity.

I acknowledge and accept the Stanford Honor Code. I affirm that I have neither given nor received unpermitted aid on this examination.

Signature: _________________________________ **Date:** __________________

Instructions

- This is a closed-book, closed-notes exam. You may use two double-sided sheets of notes. No calculators, phones, or other electronic devices are allowed.
- You have 150 minutes. The exam has 5 problems worth 100 points in total, followed by a sixth problem that is entirely extra credit. The point value of each problem and of each part is stated. Budget your time for problems 1 through 5 and turn to problem 6 only if you have time left.
- There are 10 extra-credit points in all: 2 at the end of problem 2, part (d), and the 8 of problem 6. Anything you earn there is added to your score on the 100 points above, and skipping them costs you nothing.
- Problem 1 is a set of multiple-choice items and needs no explanation. Problems 2 through 6 ask you to derive, compute, and explain. For those problems, an answer with no reasoning receives little credit, and a clearly reasoned answer with a small arithmetic slip receives most of the credit.
- Write your answers on the separate answer sheet, not in this booklet. Anything written in this booklet will not be graded. The answer sheet has a page for each part, with room sized for the answer we have in mind. If a part runs over, use the extra pages at the end of the answer sheet and write *continued on page ...* in that part's space.
- Wherever we ask for a specific answer, such as a number, a matrix, a line, or a formula in terms of a parameter, the answer sheet gives you a box for it, labelled with what belongs in it. Put that answer in the box and keep the reasoning in the space around it.
- Every part of every problem has a short solution, and every number is chosen so that the arithmetic can be done by hand. If you find yourself in a calculation that looks prohibitively long, or that seems to need a calculator, there is a better way, and you are missing the structure that makes the part easy.
- Keep your answers short. We will deduct points from answers that are technically correct but far more complicated than they need to be.
- We have made every effort to state the problems clearly. If you believe a problem is unclear or ambiguous, you may make a reasonable assumption that resolves the ambiguity and solve the problem under that assumption. State your assumption. If your interpretation is reasonable, it may receive credit even if it differs from the one we intended.
- Good luck!

1. Multiple choice *(20 points)*. For each item, write the item number and the letter of your answer in your blue book. Each item has exactly one correct answer. No explanation is required, and none will be read. *(2 points each)*

1. Let $u_1, \cdots, u_4 \in \mathbb{R}^n$ be unit-norm pairwise orthogonal vectors, and define $A = \begin{bmatrix} u_1 & u_2 \end{bmatrix}$, $B = \begin{bmatrix} u_3 & u_4 \end{bmatrix}$. Which of the following statements is correct about the matrix $AB^\top$?

   (a) It has exactly two non-zero eigenvalues
   (b) It can have up to two non-zero eigenvalues
   (c) All its eigenvalues are zero, but the matrix itself isn't necessarily the zero matrix
   (d) It is possible for it to have no non-zero eigenvalues, in which case it will be the zero matrix

2. Let $X \in \mathbb{R}^n$ be a vector of random variables with the covariance matrix $\Sigma \in \mathbb{R}^{n \times n}$. If $\operatorname{rank}(\Sigma) = n - 2$, which of the following statements is the most accurate about these random variables?

   (a) Out of the n random variables, 2 of them have zero variances
   (b) X_n and X_{n-1} are deterministic linear combinations of $X_1, \cdots, X_{n-2}$
   (c) Some i, j exist, such that if we remove X_i, X_j from X, no information will be lost
   (d) With probability one, X lies in an affine subspace of $\mathbb{R}^n$ of dimension 2

3. Let $A \in \mathbb{R}^{n \times n}$, and suppose $A = P \Sigma P^{-1}$ for some invertible P and some diagonal Σ. Which of the following is true?

   (a) If A is symmetric, then $P^{-1} = P^\top$
   (b) If $P^{-1} = P^\top$, then the eigenvalues of A are distinct
   (c) If the eigenvalues of A are distinct, then $P^{-1} = P^\top$
   (d) None of the above

4. Let $A \in \mathbb{R}^{m \times n}$, $B \in \mathbb{R}^{p \times n}$, $\mu > 0$, and $C = A^\top A + B^\top B + \mu I$. Which of the following statements is correct about the matrix C?

   (a) $\lambda_{\max}(C) = \max\left(\sigma^2_{\max}(A), \sigma^2_{\max}(B)\right) + \mu$
   (b) $\lambda_{\max}(C) \leq \max\left(\sigma^2_{\max}(A), \sigma^2_{\max}(B)\right) + \mu$
   (c) $\lambda_{\max}(C) = \sigma^2_{\max}(A) + \sigma^2_{\max}(B) + \mu$
   (d) $\lambda_{\max}(C) \leq \sigma^2_{\max}(A) + \sigma^2_{\max}(B) + \mu$

5. A unit mass sits at rest on a frictionless surface. A force is applied to it over ten unit-time intervals and held constant on each one, with value x_i for $i-1 < t \le i$. If y_1 and y_2 are the position and the velocity at $t = 10$, then $y = Ax$ with $A \in \mathbb{R}^{2\times 10}$, where $y_2 = \sum_i x_i$ and $y_1 = \sum_i \left(\frac{21}{2} - i\right) x_i$. To move the mass a unit distance and leave it at rest, we solved

$$\text{minimize} \quad \|x\| \qquad \text{subject to} \qquad Ax = \begin{bmatrix} 1 \\ 0 \end{bmatrix},$$

and the solution came out as $x = \begin{bmatrix} x_0 & 0 & \cdots & 0 & -x_0 \end{bmatrix}^\mathsf{T}$: one push in the first interval, an equal pull in the last, and nothing in between. Which norm was minimized?

(a) $\|x\|_1$
(b) $\|x\|_2$
(c) $\|x\|_p$ for some $p > 2$
(d) $\|x\|_\infty$

6. For an analytic function $f(a) = \beta_0 + \beta_1 a + \beta_2 a^2 + \cdots$ and a square matrix A, the matrix $f(A)$ is defined by the same series, $f(A) = \beta_0 I + \beta_1 A + \beta_2 A^2 + \cdots$. Take

$$\cos a = 1 - \frac{a^2}{2!} + \frac{a^4}{4!} - \frac{a^6}{6!} + \cdots,$$

and let $A \in \mathbb{R}^{3\times 3}$ be symmetric with eigenvalues $\pi/3$, $2\pi/3$ and π. Which of the following is true of $\cos A$?

(a) It is positive semidefinite
(b) It is negative semidefinite
(c) It is indefinite
(d) It cannot be determined from the given information

7. Let $A = \begin{bmatrix} 1 & 1/2 \\ 1/2 & 1 \end{bmatrix}$. In each figure below the solid curve is the set of $x \in \mathbb{R}^2$ with $x^\mathsf{T} A x = 1$, and the dashed curve is the set of $x \in \mathbb{R}^2$ with $x^\mathsf{T} A^{-1} x = 1$. The tick marks are at ± 1. Which figure is correct?

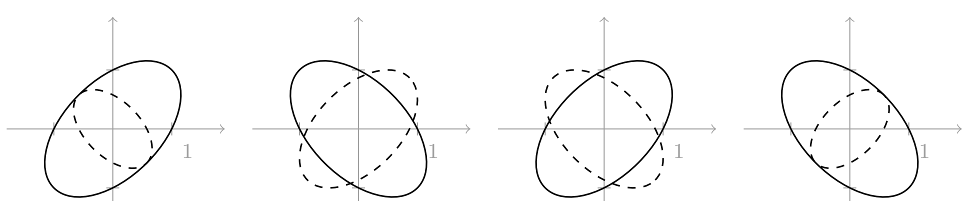

(a) (b) (c) (d)

8. Let $A \in \mathbb{R}^{10\times 10}$ with $\operatorname{rank}(A) = 7$, and consider

$$\text{minimize} \quad \|A - X\| \qquad \text{subject to} \qquad \operatorname{rank}(X) = k,$$

over $X \in \mathbb{R}^{10\times 10}$, where the constraint asks for rank exactly k, not at most k. Let B be a solution for $k = 6$ and C a solution for $k = 8$. Which of the following is true?

   (a) Both B and C exist
   (b) B exists, and C does not exist
   (c) B exists, and C may or may not exist, depending on A
   (d) Each of B and C may or may not exist, depending on A

9. A centered data set of $N = 10^6$ points in $\mathbb{R}^{10000}$ is reduced to 20 dimensions in two ways.

   - *One shot:* keep the top 20 principal components of the data.
   - *Two stage:* keep the top 500 principal components, then run PCA again on the resulting 500-dimensional data and keep its top 20 principal components.

   How much variance do the two pipelines capture?

   (a) The two-stage pipeline captures less, because the first stage discards directions that the second one can no longer use
   (b) Both capture exactly the same variance
   (c) Both capture the same variance only if the data has rank at most 500
   (d) The two-stage pipeline captures more: dropping the weakest directions first clears away clutter, and the second PCA then works in the 500-dimensional space, where it can pick up structure in the data that was not visible in the original one

10. Let G be an undirected graph on $n \geq 4$ nodes with positive edge weights, let $L = D - W$ be its Laplacian, and let $0 = \lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n$ be the eigenvalues of L. Which of the following statements about λ_2 is correct?

   (a) Adding an edge to G cannot decrease λ_2
   (b) $\lambda_2 = 0$ if and only if some node of G has degree zero
   (c) $\lambda_2 \geq \min_i d_i$, where d_i is the weighted degree of node i
   (d) Multiplying every edge weight by 2 leaves λ_2 unchanged

2. Numerical computations *(20 points)***.** What the four parts of this problem have in common is not a topic but a task: each one asks you to carry a calculation through to actual numbers, and the four are drawn from four different corners of the course. Beyond that they are independent. Nothing in one part is needed in another, so you may do them in any order.

a) *(4 points)* **A graph Laplacian.** The undirected, unweighted graph below has 9 nodes. Let W be its adjacency matrix, D the diagonal matrix of degrees, and $L = D - W$ its Laplacian.

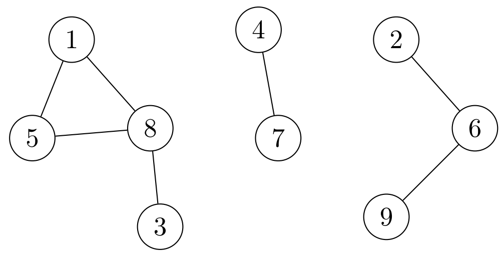

Find a matrix $Z \in \mathbb{R}^{9\times k}$, with k as large as possible, such that

- $LZ = 0$,
- every entry of Z is 0 or 1,
- the columns of Z are nonzero and mutually orthogonal.

Give k and Z, explain why no larger k is possible, and say in one sentence what the columns of Z represent.

b) *(6 points)* **A noisy sensor.** A single sensor measures the sum of the two components of an unknown $x \in \mathbb{R}^2$:

$$y = Ax + w, \qquad A = \begin{bmatrix} 1 & 1 \end{bmatrix},$$

where $x \sim \mathcal{N}(0, \Sigma_x)$ with $\Sigma_x = 9I$, and the noise $w \sim \mathcal{N}(0, 9/4)$ is independent of x. You observe $y = 9$. Find the MMSE estimate $\hat{x}$ of x and the posterior covariance P. Then sketch the two ellipsoids

$$\{z : z^\mathsf{T}\Sigma_x^{-1}z \le 1\} \quad \text{and} \quad \{z : z^\mathsf{T}P^{-1}z \le 1\},$$

and give the direction in which the second has shrunk the most relative to the first and the direction in which it has shrunk the least, with the factor in each case. Show that the least shrinkage is in fact no shrinkage at all, and explain why the measurement left that direction untouched.

c) *(6 points)* **Three ways to fit a line.** You are given the four data points

$$(2,2), \quad (4,0), \quad (4,4), \quad (6,6),$$

plotted below. Fit a line $y = mx + b$ to them three times, determining both the slope m and the intercept b in each case. The first fit minimizes the sum of squared *vertical* distances from the four points to the line, the second minimizes the sum of squared *horizontal* distances, and the third minimizes the sum of squared *perpendicular* distances. Give the three lines, and say in one sentence why they are not all the same.

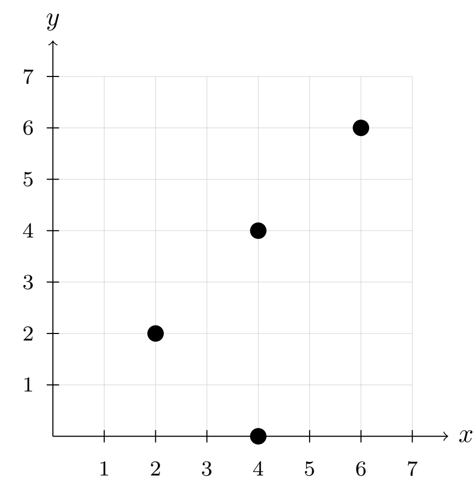

d) *(4 points)* **A trade-off between two objectives.** With $x \in \mathbb{R}^2$, let

$$A = \begin{bmatrix} 1 & 1 \end{bmatrix}, \quad b = 0, \qquad C = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad d = \begin{bmatrix} 1 \\ 2 \end{bmatrix},$$

and consider the two objectives

$$J_1 = \|Ax - b\|^2, \qquad J_2 = \|Cx - d\|^2.$$

For $\mu > 0$, let $x^\star(\mu)$ be the minimizer of the composite objective $J_1 + \mu J_2$. Find $x^\star(\mu)$ explicitly as a function of μ. Then sketch the optimal trade-off curve in the J_1–J_2 plane, and mark the two points on it that correspond to $\mu \to 0$ and $\mu \to \infty$, giving their coordinates.

Extra credit (2 points): find the equation of the trade-off curve, in the form $J_1 = f(J_2)$.

3. Bringing back a population ***(20 points).*** A discrete-time linear dynamical system has a *state* $x(t) \in \mathbb{R}^n$ at times $t = 0, 1, 2, \ldots$ and an *input* $u(t) \in \mathbb{R}^m$ that we choose at each time. The state evolves according to

$$x(t+1) = Ax(t) + Bu(t), \qquad A \in \mathbb{R}^{n\times n}, \quad B \in \mathbb{R}^{n\times m},$$

so once $x(0)$ and the inputs are given, the whole trajectory is determined.

A nature reserve holds a threatened bird species, counted once a year just after the young have fledged. Let $j(t)$ be the number of juveniles, meaning birds in their first year, and let $a(t)$ be the number of adults at the count in year t. In our model, from one count to the next each adult leaves one young that survives to the next count, half of the juveniles survive and become adults, and half of the adults survive.

The reserve can also bring birds in from a captive breeding centre. Let $u(t)$ be the number of captive-bred young released just before the count in year $t+1$, so the effect of $u(t)$ shows up in $j(t+1)$, one count later, and from then on those birds face the same odds as the wild ones.

Take the state to be $x(t) = \begin{bmatrix} j(t) \\ a(t) \end{bmatrix}$. The halves in the model will sometimes give counts that are not whole numbers, which is nothing to worry about: rates like half of the adults surviving are averages over a population, not statements about one bird. Treat $j(t)$, $a(t)$ and $u(t)$ as real numbers and do not round anything.

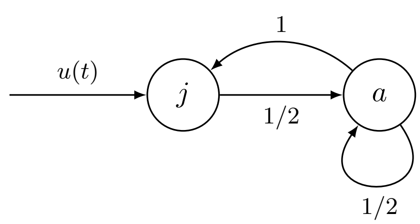

The reserve was founded with 24 adult birds, released just before the year-0 count, so $x(0) = \begin{bmatrix} 0 \\ 24 \end{bmatrix}$. The recovery plan calls for 28 juveniles and 28 adults at the year-3 count, 56 birds in all; part (b) explains why the plan asks for that particular split. Captive-bred birds are the expensive part of the program: a large release in a single year costs more per bird than a small one, since the birds have to come from farther afield and more of them are lost to dispersal and competition on arrival. We model the cost of a release of u birds in one year as u^2. The model ignores crowding, which is reasonable at these numbers.

a) *(4 points)* Write the model in the form $x(t+1) = Ax(t) + Bu(t)$, and give A and B. Then suppose the reserve releases no birds at all, so $u(t) = 0$. Compute $x(1)$, $x(2)$ and $x(3)$, and say whether the reserve reaches the plan on its own.

b) *(5 points)* Suppose the program has ended, so $u(t) = 0$ from then on. Find the eigenvalues and eigenvectors of A, say whether A is diagonalizable, and give a closed form for the matrix power A^t. What do these say about the population: does it grow, and what does it look like in the long run? What is special about the state $\begin{bmatrix} 28 & 28 \end{bmatrix}^{\mathsf{T}}$ the plan aims at?

c) *(5 points)* Starting from $x(0) = \begin{bmatrix} 0 \\ 24 \end{bmatrix}$, you want to meet the plan exactly at the year-3 count:

$$x(3) = \begin{bmatrix} 28 \\ 28 \end{bmatrix}.$$

Show that $x(3) = A^3x(0) + \mathcal{C}u$, where $u = \begin{bmatrix} u(0) & u(1) & u(2) \end{bmatrix}^\mathsf{T}$, and give the matrix $\mathcal{C} \in \mathbb{R}^{2\times 3}$ explicitly. Explain why such a release plan exists, and why there is more than one. Then find the cheapest plan, the one that minimizes $u(0)^2 + u(1)^2 + u(2)^2$, and describe it in words.

d) *(6 points)* The breeding centre has already published what it will send: 4 birds the first year, 20 the second, 12 the third. That schedule lands at $x(3) = \begin{bmatrix} 32 & 26 \end{bmatrix}^\mathsf{T}$ rather than on the plan. You have to meet the plan exactly, and you want to depart from the published schedule as little as possible. The first year's birds are already committed, so a change there counts double:

$$\begin{aligned} &\text{minimize} \quad && 4\big(u(0) - 4\big)^2 + \big(u(1) - 20\big)^2 + \big(u(2) - 12\big)^2 \\ &\text{subject to} \quad && x(3) = \begin{bmatrix} 28 \\ 28 \end{bmatrix}. \end{aligned}$$

Find the optimal release plan.

One route is to write the optimality conditions for this problem as a single linear equation $Wz = y$, which you could then solve for z. Giving W and y explicitly, saying what the unknown z is, and explaining how solving $Wz = y$ would give you the optimal plan is worth 3 of the 6 points on its own: you do not have to solve that system, which may well be larger than is reasonable to do on paper. Full credit is for the plan itself, by any method you like. If you can find it some other way, do that and skip $Wz = y$ entirely.

4. The gain of a block matrix *(20 points)*. Let $A, B, C, D \in \mathbb{R}^{n\times n}$, and assemble them into the block matrix

$$M = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \in \mathbb{R}^{2n\times 2n}.$$

Throughout, $\|\cdot\|$ is the matrix norm, that is, the largest singular value, which we read as a gain:

$$\|A\| = \sigma_{\max}(A) = \max_{x\neq 0} \frac{\|Ax\|}{\|x\|}.$$

We know the gain of each of the four blocks, and the question is what those four numbers tell us about M itself.

a) *(8 points)* Show that

$$\max\{\|A\|,\ \|B\|,\ \|C\|,\ \|D\|\} \ \le\ \|M\|,$$

and that

$$\|M\| \ \le\ \sqrt{\|A\|^2 + \|B\|^2 + \|C\|^2 + \|D\|^2}.$$

Each of the two bounds is worth 4 points.

b) *(6 points)* Show that

$$\|M\| \ \le\ \left\| \begin{bmatrix} \|A\| & \|B\| \\ \|C\| & \|D\| \end{bmatrix} \right\|,$$

where the matrix on the right-hand side is 2×2.

c) *(6 points)* Parts (a) and (b) give two upper bounds on $\|M\|$. Decide which of the two is smaller, prove your answer, and say exactly when the two agree. This is worth 3 points.

Then, for the remaining 3 points, give one numerical example: choose n and give explicit numbers for A, B, C, D for which the lower bound of part (a) holds with equality, and exactly one of the two upper bounds holds with equality.

5. Isolating one effect *(20 points).* A retailer wants to know what a dollar of advertising is worth. Over n weeks you record the sales y_i, the advertising spend d_i, and the conditions of the week that are known to move sales on their own: the season, holidays, the price charged, and so on. Stack the first two into $y, d \in \mathbb{R}^n$, and put week i's conditions in the ith row of a feature matrix $\Phi \in \mathbb{R}^{n \times p}$; assume $p < n$ and Φ full rank. The records satisfy

$$y = \theta d + \Phi\alpha + u, \qquad d = \Phi\gamma + v, \qquad \Phi^\mathsf{T} u = 0, \qquad \Phi^\mathsf{T} v = 0, \qquad v \neq 0.$$

The scalar θ is what you want: the sales brought in by one more dollar of advertising. The vectors $\alpha, \gamma \in \mathbb{R}^p$ are of no interest in themselves, and may be large. The first says how the season and the price move sales by themselves; the second says how the marketing team sets its budget from that same information, since it advertises hardest in the weeks it expects to be strong. That second equation is the difficulty. It makes d and the columns of Φ move together, so a plain fit of y on d credits advertising with the sales the holidays would have brought anyway. The two leftovers carry nothing the features could have explained, which is what the orthogonality conditions say: v is the spending the team's own rule does not account for, and u is the part of sales that neither advertising nor the conditions explain. Write

$$P = \Phi(\Phi^\mathsf{T}\Phi)^{-1}\Phi^\mathsf{T}, \qquad M = I - P.$$

These should look familiar: for any vector $z \in \mathbb{R}^n$, Pz is the least-squares approximation of z by the columns of Φ, and $Mz = z - Pz$ is the residual, the part of z that no combination of features can account for. That is the one idea the whole problem runs on. If we can subtract off everything the features explain, what is left should be linked only by advertising. Part (a) works out what M does to the data; parts (b) and (c) ask what estimate of θ you get when you do that subtraction on one of y and d, and then on both.

a) *(8 points) The tool.* Before using M we should know exactly what it does, so this part is eight short questions about it, one point each.

   i. Show that $M^\mathsf{T} = M$ and $M^2 = M$, so that M is an orthogonal projector.

   ii. Give range(M), as a subspace built from Φ.

   iii. Give null(M), as a subspace built from Φ.

   iv. Give $\operatorname{rank} M$, as a number in terms of n and p.

   v. Let $\Phi = QR$ be a QR factorization. Express P in terms of Q, with no matrix inverse left in the formula.

   vi. Give the eigenvalues of M, and the multiplicity of each in terms of n and p.

   vii. Show that $Md = v$.

   viii. Show that $My = \theta v + u$.

   The last two identities are what the rest of the problem lives on, so they are worth reading: after the subtraction, all that survives of the spend is the team's own v, and all that survives of the sales is θ times that same v, plus u.

b) *(4 points) A first idea.* Sales move with the season and the price as well as with advertising, so clean that out of the sales: use My in place of y, and fit θ against the spend by least squares. This is an ordinary least-squares problem with one unknown and the single-column matrix d:

$$\hat{\theta}_{\text{naive}} = \underset{\theta}{\operatorname{argmin}} \ \|My - \theta d\|^2.$$

Give a closed form for $\hat{\theta}_{\text{naive}}$. Now test the idea: set $u = 0$, so the records follow the model exactly and any error we see is the fault of the method rather than of noise. What we would like is $\hat{\theta}_{\text{naive}} = \theta$. Show that instead the estimate comes out proportional to θ,

$$\hat{\theta}_{\text{naive}} = \kappa \, \theta,$$

and find the factor κ. Say what κ is geometrically, and state the one condition under which the idea does return θ.

c) *(8 points) A better idea.* Part (b) cleaned the features out of the sales but left them sitting inside the spend, which is what the factor κ is measuring. So clean both, and fit what is left of the sales against what is left of the spend. This is again a least-squares problem with one unknown, now with the single-column matrix Md:

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} \ \|My - \theta \, Md\|^2.$$

Give a closed form for $\hat{\theta}$, and test it the same way as before: show that when $u = 0$ it recovers θ exactly, for every α and every γ. Then compare it with the approach you would probably have reached for first, which is to fit θ and α together in one least-squares problem and read off the first entry. Show that these give the same answer, that is, that $\hat{\theta}$ is the first component of

$$\underset{\theta,\alpha}{\operatorname{argmin}} \ \left\| y - \begin{bmatrix} d & \Phi \end{bmatrix} \begin{bmatrix} \theta \\ \alpha \end{bmatrix} \right\|^2.$$

This problem is **entirely extra credit**. Attempt it only if you have time left over.

6. Extra credit: a huge matrix you can only multiply by *(8 points)*. A symmetric matrix $A \in \mathbb{R}^{n\times n}$ is given, with n in the millions. Its eigenvalues satisfy

$$|\lambda_1| > |\lambda_2| > \cdots > |\lambda_n|,$$

and you want λ_1 together with an eigenvector belonging to it.

You never get to see the entries of A. What you have is a module that takes any vector $z \in \mathbb{R}^n$ and returns Az, and you may call it many times. Beyond that you can do the cheap things with vectors: add them, scale them, take inner products and norms.

Two limits come with the size, and both matter. Your machine has room for a few vectors of length n and nothing like room for n^2 numbers, so you cannot keep a copy of A, however you might come by it: forming it or factoring it is out of the question. And the module is not free. Every call to it, one matrix-vector product, costs you a dollar. Aim to get the job done cheaply: a good answer for a couple of hundred dollars would be a success.

Give a procedure that produces both λ_1 and an eigenvector for it, and justify it. Your answer should say what you compute and in what order, how the eigenvalue and the eigenvector are read off from what you have computed, why what you get belongs to λ_1 and not to one of the other eigenvalues, which of the assumptions stated above your justification uses and where, and what you have to assume about anything you are free to choose yourself. Since this has to run on a real machine, say something about the sizes of the numbers your procedure carries along, and about the work and the storage each stage costs.

You are not being asked for an exact answer after a fixed amount of arithmetic. A procedure whose two outputs can be brought as close as you like to λ_1 and to an eigenvector for it is a full answer.