

EE263 Final Exam, August 2023

- This is a 12-hour take-home exam. Please turn it in on Gradescope.
Be aware that you *must* turn it in within 12 hours of downloading it. After that, Gradescope will not let you turn it in.
- You may use any books, notes, or computer programs, including searching online, *except for ChatGPT or similar tools*. You may not discuss the exam or course material with others, or work in a group.
- Do not discuss the exam until 8/23, after everyone has taken it.
- If you have a question, please email the staff. We have done our best to make the exam unambiguous. So unless there is a mistake, we are unlikely to say much.
- Please check your email during the exam, in case we need to send a clarification or announcement.
- We expect your solutions to be legible, neat, and clear. Do not hand in your rough notes. Please try to simplify your solutions as much as you can. We will deduct points from solutions that are technically correct, but much more complicated than they need to be.
- Start each solution on a new page. Individual parts can be on same page.
- When a problem involves some computation, we expect a clear discussion and justification of what you did as well as the final numerical result.
- Because this is an exam, **you must turn in your code**. Include the code in your pdf submission. We will deduct points for missing code.
- Good luck!

6.2320. The Sinewave Scaling. The Sinewave Scaling is a rare mathematically-inclined sea serpent which can only fly in patterns of the form

$$x(t) = \sum_{k=1}^N a_k^{(1)} \sin\left(\frac{k}{30}\pi t\right) + a_k^{(2)} \cos\left(\frac{k}{30}\pi t\right)$$

$$y(t) = \sum_{k=1}^N b_k^{(1)} \sin\left(\frac{k}{30}\pi t\right) + b_k^{(2)} \cos\left(\frac{k}{30}\pi t\right)$$

where $a^{(1)}, a^{(2)}, b^{(1)}, b^{(2)} \in \mathbb{R}^N$. For simplicity, define $a = \begin{bmatrix} a^{(1)} \\ a^{(2)} \end{bmatrix}$ and $b = \begin{bmatrix} b^{(1)} \\ b^{(2)} \end{bmatrix}$.

The scaling is a solitary creature and regularly surveys its territory for intruders. To do so, it plans to start at its home base $(0, 0)$, fly over some islands with (x, y) coordinates $(s_1^{(x)}, s_1^{(y)}), \dots, (s_5^{(x)}, s_5^{(y)})$ and return home (a total of $T = 7$ waypoints). It plans to depart at time t_0 , reach island i at time t_i and return at time t_6 . You can assume $0 \leq t_i < t_{i+1}$ for all $i = 0, \dots, T - 1$. You may use $T = 7$, $N = 6$ and the following waypoints and times:

$$s^{(x)} = [0 \ 17 \ 25 \ 21 \ 26 \ 8 \ 0]^\top$$

$$s^{(y)} = [0 \ -5 \ -3 \ 5 \ 11 \ 7 \ 0]^\top$$

$$t = [0 \ 4 \ 7 \ 19 \ 25 \ 30 \ 34]^\top.$$

a) Define two matrices $U \in \mathbb{R}^{T \times N}$ and $V \in \mathbb{R}^{T \times N}$ such that

$$\begin{bmatrix} x(t_0) \\ \vdots \\ x(t_{T-1}) \end{bmatrix} = [U \ V] a, \quad \begin{bmatrix} y(t_0) \\ \vdots \\ y(t_{T-1}) \end{bmatrix} = [U \ V] b.$$

b) Propose a method to plan a trajectory that reaches each island at the correct time while minimizing the loss function

$$\|a\|_2^2 + \|b\|_2^2 \tag{1}$$

which represents the energy the scaling expends while flying. Clearly describe any matrices you construct. You are encouraged to use block matrix notation and the definitions of U and V from part (a).

c) Use your method to compute a trajectory for the scaling. Report the minimum energy defined in (??) the scaling must expend to visit all 5 islands. Plot the islands and $(x(\hat{t}), y(\hat{t}))$ using a smaller timestep; for example, `τ_hat = τ[1]:0.5:τ[end]`. You should generate a 2D plot showing the trajectory of the scaling and the locations of the islands.

d) Due to its unusual flight pattern, the scaling has difficulty controlling its altitude $z(t)$, which we can model as

$$z(t) = \frac{1}{10}x(t) + \frac{1}{10}y(t).$$

Suppose the scaling wants to make a more precise flyover reaching specific (x, y, z) coordinates at each waypoint. Given the coordinates

$$s^{(z)} = [0 \ 5 \ 3 \ 6 \ 2 \ 1 \ 0]^\top,$$

are the three-dimensional waypoints reachable? Why or why not? Provide a short explanation. Your answer may use facts about the matrices you constructed earlier.

e) Propose a method to minimize the sum of errors

$$J = \sum_{i=0}^{T-1} \left(x(t_i) - s_i^{(x)} \right)^2 + \left(y(t_i) - s_i^{(y)} \right)^2 + \left(z(t_i) - s_i^{(z)} \right)^2.$$

Provide a short justification for your proposed method (not a formal proof).

f) Use your method to compute a three-dimensional trajectory for the scaling. Display your trajectory, along with the locations of the islands, on a three-dimensional plot using the same smaller timestep as in part (c).

Submit your plot, the value of J , and the energy required for this trajectory.

Hint: In Julia you can import `Plots.jl`, then use `plot3d(x, y, z)` and `scatter3d(x,y,z)`.

18.2310. Particle in magnetic chamber. You have just been employed by Magneto Inc. to work on their Small Particle Collider (SPC). The SPC generates magnetic fields in order to apply forces to particles inside of it. These forces can be summed to produce a resultant force in any direction in 3D space. Your assigned task is to develop a sequence of input forces $u(0), \dots, u(t)$, generated by the SPC walls, that will move particles in specific ways. We will use the vector $x(t)$ to represent the position and velocity of a particle in the x , y , and z directions. Namely,

$$x(t) = \begin{bmatrix} p(t) \\ \dot{p}(t) \end{bmatrix} \in \mathbb{R}^6,$$

where $p(t)$ and $\dot{p}(t)$ represent the particle's position and velocity, respectively, at time t . Note that $u(t) \in \mathbb{R}^3$.

a) We know that the motion dynamics for a particle with mass m can be expressed as $u = m\ddot{p}$. Use the forward Euler approximation, given by $x(t+1) = x(t) + h\dot{x}(t)$, to express this system as a discrete-time linear dynamical system. Explicitly define your matrices A , B , C , and D in terms of m and h . For our vector of outputs, $y(t)$, we are interested in just the particle's position for now, *i.e.*, $y(t) = p(t)$.

b) Write a general expression for

$$\tilde{y} = \begin{bmatrix} y(0) \\ y(1) \\ \vdots \\ y(t) \end{bmatrix} \in \mathbb{R}^{3(t+1)}$$

given the initial position and velocity $x(0)$. Your answer may involve $x(0)$, A , B , C , D , and the inputs $u(0), \dots, u(t)$.

c) We would like to apply forces to the particle such that it returns back to its initial position after some time T has passed. We also want the particle to pass by a detector at time $T/2$ (assuming that T is even).

We would like to find

$$\tilde{u} = \begin{bmatrix} u(0) \\ u(1) \\ \vdots \\ u(T) \end{bmatrix} \in \mathbb{R}^{3(t+1)}$$

such that $y(T) = p_0$ and $y(T/2) = p_{T/2}$, where $p_0, p_{T/2} \in \mathbb{R}^3$ are the initial position and the position of the detector respectively. Also, assume that the initial velocity $\dot{p}_0 \in \mathbb{R}^3$ is known. Describe a method to find the forces that satisfy these constraints while minimizing $\|\tilde{u}\|^2$.

d) Apply your method to the following numerical values and plot the x , y , and z coordinates of the particle from time $t = 0$ to $t = T$. Make sure to plot the $T/2$ way-point as well. You should have one figure with 3 plots, one for each of the axes directions.

$$p_0 = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}, \quad \dot{p}_0 = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \quad p_{T/2} = \begin{bmatrix} 2 \\ -1.5 \\ -1 \end{bmatrix}, \quad T = 200, \quad m = 5, \quad h = 1$$

e) The lead engineer has a new goal; they would like the particle to not only pass through the desired positions as explained in the previous parts, but also to do so with precise velocities denoted by \dot{p}_0 and $\dot{p}_{T/2}$. Adjust your definitions from part (a) for x , y , A , B , C , or D to suit

this new task if necessary.

Describe a method that achieves the constraints

$$\begin{bmatrix} p(T/2) \\ \dot{p}(T/2) \end{bmatrix} = \begin{bmatrix} p_{T/2} \\ \dot{p}_{T/2} \end{bmatrix}, \quad \begin{bmatrix} p(T) \\ \dot{p}(T) \end{bmatrix} = \begin{bmatrix} p_0 \\ \dot{p}_0 \end{bmatrix}$$

while minimizing $\|\tilde{u}\|^2$. Similar to part (d), run your method on the given data and plot the x , y , and z coordinates of the particle from time $t = 0$ to $t = T$. Use the same numerical values as in part (d), along with

$$\dot{p}_{T/2} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}.$$

- f) Due to the limitations of the machine, we want to define an objective that represents the rate of change of the force produced. Define

$$J_2 = \sum_{k=1}^{n-1} \|u(k+1) - u(k-1)\|^2$$

as the central difference approximation for the force rate of change. Also define $J_1 = \|\tilde{u}\|^2$.

Devise a method that minimizes $J_1 + \gamma J_2$ while adhering to the position and velocity constraints defined in part (e).

- g) Plot the new x , y , and z coordinates of the particle over time for $\gamma = 10^p$ for $p \in \{-2, 0, 2, 4, 6\}$. What do you see as γ is increased?

19.2330. Multi-Armed bandits. A multi-armed bandit is a concept often used in probability theory and decision-making. Imagine you're in a casino facing multiple slot machines (the "arms" of the bandit). Each machine has a different unknown probability of giving you a reward when you pull its lever. You have a limited number of pulls to make and you want to maximize your total reward.

The challenge is to decide which machines to pull and in what order, to strike a balance between exploring (trying out different machines to learn their probabilities) and exploiting (focusing on the machine that seems to give the best rewards based on your current knowledge).

In real-world scenarios beyond casinos, multi-armed bandits can be used to model situations where you have limited resources (like time, money, or experiments) and need to make decisions to optimize outcomes while dealing with uncertainty. This concept finds applications in various fields, such as online advertising, clinical trials, and recommendation systems.

An extension of this setup, often referred to as a stochastic linear bandit, can be formulated as follows: Each potential decision, denoted as an arm of the bandit, is represented by a vector $x \in \mathbb{R}^d$. It is assumed that there exists an unknown parameter vector $\theta \in \mathbb{R}^d$ that influences the observed rewards. To put it more formally, at each time step t , the player is presented with a decision set $\mathcal{D}_t \subseteq \mathbb{R}^d$ from which they must select an action x_t . Consequently, they observe a reward $y_t = x_t^\top \theta + v_t$, where v_t is a zero-mean random noise.

- a) Suppose the value of θ is known and $\mathcal{D}_t = \{x \in \mathbb{R}^d : \|x\| \leq 1\}$. Propose an optimal choice for x_t to maximize the reward and briefly explain why it is a good choice.
- b) In practice, θ is unknown and we try to estimate it. Suppose we have T pairs of observations (x_i, y_i) for $i = 1, \dots, T$ from the past decisions and rewards. We aim to find an estimate of θ , namely $\hat{\theta}$, to minimize the cost

$$J = \|X_T \hat{\theta} - y\|^2,$$

where $X_T = \begin{bmatrix} x_1^\top \\ \vdots \\ x_T^\top \end{bmatrix} \in \mathbb{R}^{T \times d}$ and $y = \begin{bmatrix} y_1 \\ \vdots \\ y_T \end{bmatrix} \in \mathbb{R}^T$.

Which of the following statements most accurately describes the problem of minimizing J when $T < d$? Briefly justify your answer.

- i. There are infinitely many solutions that achieve $J = 0$.
 - ii. There are two possible cases: one unique solution, or infinitely many solutions that achieve $J = 0$.
 - iii. There are two possible cases: no solution, or infinitely many solutions that achieve $J = 0$. In the case of no solution with $J = 0$, the optimal $\hat{\theta}$ with $J > 0$ can be unique.
 - iv. There are two possible cases: no solution, or infinitely many solutions that achieve $J = 0$. In the case of no solution with $J = 0$, the optimal $\hat{\theta}$ with $J > 0$ is not unique.
 - v. There are three possible cases: No solutions, one unique solution, or infinitely many solutions that achieve $J = 0$.
- c) It is mostly common to solve the regularized version of the problem, *i.e.*, to minimize the objective

$$J_\mu = \|X_T \hat{\theta} - y\|^2 + \mu \|\hat{\theta}\|^2,$$

where $\mu > 0$ is the regularization parameter.

Give a closed form solution for the optimal $\hat{\theta}$ that minimizes J_μ . Under what circumstances is this solution unique?

- d) In some applications, we might have some prior knowledge about θ . A simple case is to have a prior guess, namely θ_0 , and we would like our estimated $\hat{\theta}$ to be close to θ_0 . In this case, we can minimize the cost

$$J_{\mu,\lambda} = \|X_T \hat{\theta} - y\|^2 + \mu \|\hat{\theta}\|^2 + \lambda \|\hat{\theta} - \theta_0\|^2,$$

where $\mu, \lambda > 0$ are the regularization parameters.

Give a closed form solution for the optimal $\hat{\theta}$ that minimizes $J_{\mu,\lambda}$. Under what circumstances is this solution unique?

- e) A practical approach for selecting an action x_t at time t involves estimating θ using $\hat{\theta}_{t-1}$ based on all past observations (x_τ, y_τ) for $\tau = 1, \dots, t-1$. This estimation is performed by minimizing J_μ . Subsequently, x_t is chosen to maximize the reward under the assumption that $\theta = \hat{\theta}_{t-1}$. Afterward, we observe the reward y_t and incorporate the new pair (x_t, y_t) into the set of observations. This procedure can be repeated for as many steps as needed.

It has been demonstrated that if we continue this process for T steps, there exists a probability of at least 0.99 that the following inequality holds:

$$\|\theta - \hat{\theta}_T\|_{V_T} \leq R,$$

where R is a linear function of \sqrt{d} , $\sqrt{\log T}$, and $\sqrt{\mu}$. This means that while we do not know the exact value of θ , there exists a region within which we can assert, with high probability, that θ resides. Here, $\|\theta - \hat{\theta}_T\|_{V_T}$ represents the norm of the difference between θ and $\hat{\theta}_T$ with respect to the matrix V_T , defined as $V_T = X_T^\top X_T + \mu I$. The norm of a vector x with respect to a matrix A is commonly denoted as $\|x\|_A = \sqrt{x^\top A x}$.

What conditions are required for a matrix A to be valid for defining a norm? In simpler terms, what condition must matrix A fulfill to ensure that $\|x\|_A$ is defined for every $x \in \mathbb{R}^d$ and is zero only if $x = 0$? Does V_T meet this condition?

- f) According to the previous part, we know that $\|\theta - \hat{\theta}_T\|_{V_T} \leq R$ holds with a high probability. In the worst case scenario, the inequality holds with equality, *i.e.*, $\|\theta - \hat{\theta}_T\|_{V_T} = R$. In such a situation, what is the smallest and largest possible values of $\|\theta - \hat{\theta}_T\|$? Here, the norm denotes the standard Euclidean norm. Your answers should be in terms of R , μ , and possibly the eigenvalues of $X_T^\top X_T$.

21.2300. True or false. For each of the statements below, state whether it is true or false. If true, provide a concise explanation (two or three sentences at most) as to why it holds. If false, give a counterexample. Correct true/false answers with incorrect/missing explanations will receive half credit.

- a) The continuous-time linear dynamical system $\dot{x} = x + Bu$ is never controllable if $u \in \mathbb{R}$, but it can be controllable if $u \in \mathbb{R}^2$.
- b) A symmetric matrix $A \in \mathbb{R}^{n \times n}$ with eigenvectors v_1, v_2 corresponding to distinct eigenvalues λ_1, λ_2 exists such that the matrix $P = v_1 v_2^\top$ has exactly one non-zero eigenvalue.
- c) If $A \in \mathbb{R}^{m \times n}$, then $\sigma_i(A^\top A + \mu I) \geq \mu$, where $\sigma_i(\cdot)$ is the i th singular value.
- d) If $A \in \mathbb{R}^{n \times n}$ and $\kappa(A) = 1$, where $\kappa(\cdot)$ is the condition number of a matrix, then, $c \in \mathbb{R}$ exists such that $\|Ax\| = c\|x\|$ for any $x \in \mathbb{R}^n$.
- e) Let A^\dagger be the pseudo-inverse of $A \in \mathbb{R}^{m \times n}$. Then, $A^\dagger Ax = x$ if $x \notin \text{null}(A)$.
- f) Let A^\dagger be the pseudo-inverse of $A \in \mathbb{R}^{m \times n}$. Then, $AA^\dagger x = x$ if $x \in \text{range}(A)$.
- g) Suppose $h(t)$ is the impulse response of a continuous-time linear dynamical system with scalar inputs and outputs. If $h(t) = 0$ for $t > 5$, then the output at time t does not depend on any input from more than 5 seconds ago.
- h) Suppose $A = [a_1 \ a_2 \ \cdots \ a_n] \in \mathbb{R}^{m \times n}$ is given with the QR factorization $A = QR$. If the top left 3×4 block of R is given by $\begin{bmatrix} 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}$, then $a_1 = 0$, $\{a_2, a_4\}$ is orthonormal, and $\text{rank}(A) \leq n - 2$.
- i) Suppose $A \in \mathbb{R}^{n \times n}$ and S is a single-point invariant set of the autonomous linear dynamical system $\dot{x} = Ax$. Then, $S = \{0\}$.
- j) If $A \in \mathbb{R}^{4 \times 4}$, then $\beta_0, \beta_1, \beta_2, \beta_3$ exist such that $e^A = \beta_0 I + \beta_1 A + \beta_2 A^2 + \beta_3 A^3$.

Please read the following problem setup and then answer the questions (k) through (o) based on it.

Consider the following discrete-time time dynamical system with nonlinear output:

$$\begin{aligned} x(t+1) &= Ax(t) + u(t) \\ y(t) &= x(t)^\top Cx(t), \end{aligned}$$

where $A, C \in \mathbb{R}^{n \times n}$ are symmetric and $x(0) = 0$.

- k) The problem of maximizing the output at time $t = 1$ with control inputs $u(t)$ such that $\|u(t)\| \leq 1$ is equivalent to maximizing $u(0)^\top Cu(0)$ subject to the constraint $\|u(0)\| \leq 1$.
- l) Even if C is not positive semi-definite, the maximum value of $y(1)$ is attained when $u(0) = \lambda_1 q_1$, where λ_1 is the largest eigenvalue of C with the corresponding unit-length eigenvector q_1 .
- m) Suppose $C > 0$. As a measure of the output robustness, we are interested in finding the maximum-norm input $u(0)$ that keeps $y(1) \leq 1$. This maximum-norm input is the longest semi-axis of the ellipsoid $\mathcal{E} = \{x : x^\top Cx \leq 1\}$.
- n) The maximum-norm input described in the previous part is given by $u(0) = \lambda_1^{-1/2} q_1$, where λ_1 is the largest eigenvalue of C with the corresponding unit-length eigenvector q_1 .
- o) Suppose $C > 0$ and the cost of applying the control input $u(t)$ is $u(t)^\top Du(t)$, where D is symmetric and positive semi-definite. The problem of finding the most costly control input that maintains $y(1) \leq 1$ is equivalent to maximizing $u(0)^\top Cu(0)$ subject to the constraint $u(0)^\top Du(0) \leq 1$.