Engr108: Introduction to Matrix Methods J. Duchi
Fall 2024

Midterm Examination

This is a 90 minute in-class midterm examination.

You may use one piece of paper (double-sided), but otherwise this examination is closed book. While you may ask us questions if you find a question confusing, we've tried pretty hard to make the exam unambiguous and clear, so we're unlikely to say much.

Please respect the honor code.

All problems have equal weight. Some are (quite) straightforward. Others, not so much.

Some problems involve applications. But you do not need to know *anything* about the problem area to solve the problem; the problem statement contains everything you need.

The problems do *not* appear in order of increasing difficulty.

Name: ____

Stanford ID #: ____

M.1 Let x and y be n-vectors, where $n \geq 2$, and α, β be scalars. Under each of the following conditions, write whether x and y are linearly dependent (meaning that they must be linearly dependent based on the given statement); linearly independent (meaning that they must be linearly dependent based on the given statement); or that it cannot be determined based on the available information.

(a) $\alpha x + \beta y = \mathbf{0}$

(b) $\alpha x + \beta y = \mathbf{0}$ and $\alpha > 0$

(c) $\alpha x + \beta y \neq \mathbf{0}$ and $\alpha > 0, \beta > 0$

(d) $\alpha x + \beta y = y$ and $\alpha > 0, \beta > 0$

M.2 Consider the *Householder matrix*

$$H_u := I - 2uu^T,$$

where $u \in \mathbf{R}^n$ is a unit vector, that is, $\|u\| = 1$. (Note that this has flipped a sign from the transformations in the exercises!) Let x be an n-vector.

(a) Compute $\|H_u x\|^2$. Simplify as much as possible.

(b) Draw a picture of the linear transformation $H_u x$ for $x, u \in \mathbf{R}^2$. In your picture, you should include at least (i) the vector x, (ii) the vector u, (iii) the projection $\mathbf{proj}(x)$ of x onto the span $\{tu \mid t \in \mathbf{R}\}$ of u, (iv) the plane perpendicular to u, and (v) the transformed point $H_u x$.

M.3 Let $A \in \{-1, 0, 1\}^{n \times m}$ be the incidence matrix for a network with n nodes, $i = 1, \ldots, n$, and m edges $j = 1, \ldots, m$, so

$$A_{ij} = \begin{cases} 1 & \text{if edge } j \text{ points into node } i \\ -1 & \text{if edge } j \text{ points away from node } i \\ 0 & \text{otherwise.} \end{cases}$$

Let $x \in \mathbf{R}^m$ be a vector representing a flow along the edges of the network, so that x_j is the flow in the direction of edge j ($x_j < 0$ means the flow is against the direction of the edge).

(a) If $(Ax)_i = 0$, then what does this say about the flow across node i?

(b) We wish to have node 1 be a *source* node generating 2.3 units of flow, meaning that the total flow out of node 1 is 2.3. We wish node 2 to be a *sink* node accepting this flow, meaning that the total flow into node 2 is 2.3. All other nodes conserve flow: they neither generate nor absorb additional flow. Give a vector $b \in \mathbf{R}^m$ so that solving

$$Ax = b$$

for x gives a flow x satisfying these constraints.

(c) Suppose $b \in \mathbf{R}^n$ is a vector for which b_i represents flow into node i. Assume there is a vector x solving $Ax = b$. Give the most concise description of $\mathbf{1}^T b$ you can. In particular, what does $\mathbf{1}^T b$ equal?

M.4 Consider a signal x_t taken at times $t = 1, 2, \ldots, n$. On each plot below, (approximately) draw the convolution $x * h$ with the given vector h, where x is the given line. (Recall that for vectors $a \in \mathbf{R}^m, b \in \mathbf{R}^n$, $c = a * b \in \mathbf{R}^{m+n-1}$ has entries $c_k = \sum_{i=1}^{\min\{n,k\}} a_i b_{k+1-i}$.)

(a) $h = [1]$, the vector of a single 1.

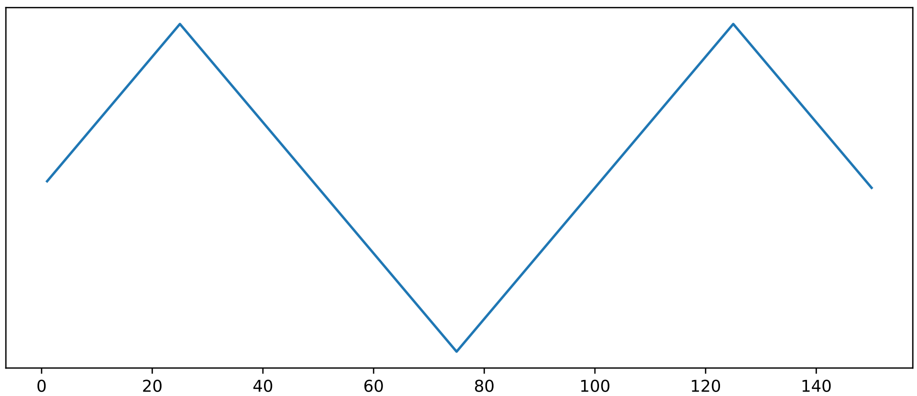

(b) $h = \mathbf{1}_7/7$, the 7-vector with all entries 1/7.

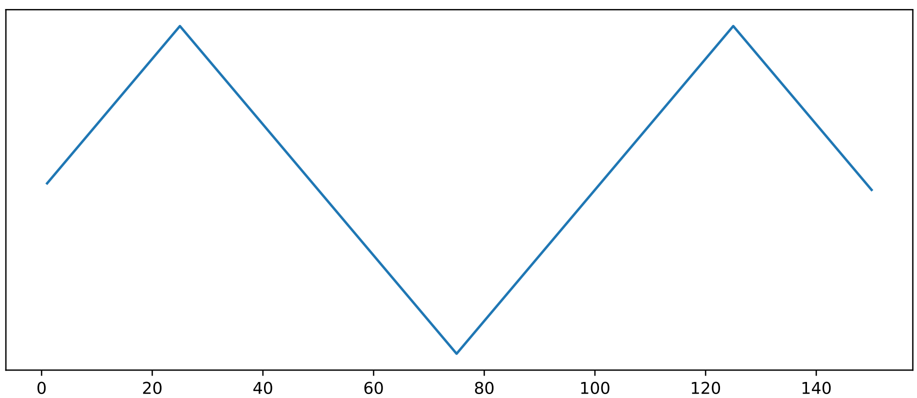

(c) $h = \mathbf{1}_n/n$, the n-vector with all entries $1/n$.

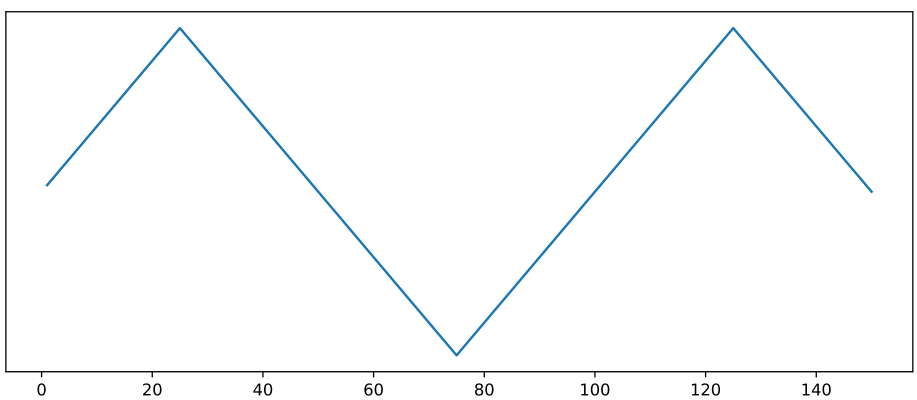

M.5 An intern at a real estate firm collects a dataset of 50 houses and their sale prices, recording 100 features of each house, including the square footage of the house, the lot width and depth, the number of bathrooms, etc. The ith sale price is $y_i \in \mathbf{R}$ and $x^{(i)} \in \mathbf{R}^{100}$ represents the ith home. The intern excitedly tells their boss that "I can predict *perfectly* the sale price of each house based on these 100 features! In fact, I have a vector w such that $y_i = (x^{(i)})^T w$ exactly for each house. Let's use w to give us an edge in the market." Should the boss use the intern's vector? Why or why not?