Midterm solutions

CS265 / CME309 Autumn 2026

100 points + 5 bonus

Question	Points
1. Quick checks	28
2. Short arguments	32
3. Checking a very large matrix power	20
4. A large sum-free subset	20 + 5 bonus

1. Quick checks **28 points**

(a) (4 pt.) Suppose a Las Vegas algorithm always returns the correct answer and has expected running time at most 8 seconds on a fixed input. If we stop it after 80 seconds and return FAIL if it has not finished, what is the strongest upper bound on the probability of returning FAIL that holds for an such algorithm?

(A) 0 (B) 1/10 (C) 1/2 (D) 1

Solution. (B). If T is the running time, Markov's inequality gives

$$\Pr[T > 80] \leq \frac{\mathbb{E}T}{80} \leq \frac{1}{10}.$$

(b) (4 pt.) Suppose we divide a class of n into two rooms by independently flipping a fair coin for each student. Let X, Y denote the number of students assigned to each of the two rooms. Which of the following is true (circle all true statements):

(A) X and Y are independent.

(B) X and Y are dependent.

(C) The marginal distributions of X and Y are each distributed like a binomial random variable.

Solution. (B) and (C). Both X and Y have distribution $\mathrm{Bin}(n, 1/2)$. They are dependent: $Y = n - X$, so knowing one count determines the other.

(c) (4 pt.) Suppose that the number of students who show up, S, is distributed like a Poisson random variable of expectation n. And we divide the S students who show up into two rooms by independently flipping a fair coin for each student. Let X, Y denote the number of students assigned to each of the two rooms. Which of the following is true (circle all true statements):

(A) X and Y are independent.

(B) X and Y are dependent.

(C) The marginal distributions of X and Y are each distributed like Poisson random variables.

(D) The distribution of X conditioned on S is distributed like a binomial random variable.

Solution. (A), (C), and (D). By Poissonization, X and Y are independent, each with distribution $\mathrm{Poi}(n/2)$. Given $S = s$, each of the s students independently chooses the first room with probability $1/2$, so X has conditional distribution $\mathrm{Bin}(s, 1/2)$.

1. Quick checks (continued)

(d) (4 pt.) Contract any non-loop edge of a connected undirected multigraph with at least three vertices, deleting self-loops and keeping parallel edges. Which statement about the minimum-cut value is always true?

(A) It cannot decrease.

(B) It strictly decreases.

(C) It cannot increase.

(D) It is unchanged.

Solution. (A). Every cut after contraction corresponds to a cut before contraction with the same number of crossing edges. Thus contraction cannot decrease the minimum-cut value.

(e) (4 pt.) Suppose C, D are strictly positive random variables over some finite number of possible outcomes (over some common probability space, in the sense that one can take about events like $Pr[C = 17 \text{ and } D = 3]$). Suppose $\mathbb{E}[C] = 2, \quad \mathbb{E}[D] = 1$. Which statements (if any) must hold?

(A) $\mathbb{E}[C/D] \leq 2$.

(B) $Pr[C/D \leq 2] > 0$.

(C) $C/D \leq 2$ for every outcome.

Solution. (B) only. If $C > 2D$ at every outcome of positive probability, then $\mathbb{E}C > 2\mathbb{E}D$, contradicting the given expectations. To rule out (A) and (C), take $C = 2$ always and let D equal $1/2$ or $3/2$, each with probability $1/2$: then $\mathbb{E}D = 1$, but C/D equals 4 or $4/3$, and $\mathbb{E}[C/D] = 8/3 > 2$.

(f) (4 pt.) Let $x \in \mathbb{R}^d$ be nonzero and let $y = 3x$. Choose a $k \times d$ matrix A with independent $N(0, 1)$ entries, where $k \geq 1$, and define the length k vector $h(z) = \text{sign}(Az)$ coordinatewise (ie the ith coordinate is $sign(A_i z)$ where A_i is the ith row of A). What is $\Pr[h(x) = h(y)]$?

(A) 0 (B) 2^{-k} (C) 1/2 (D) 1

Solution. (D). Since $Ay = 3Ax$, corresponding coordinates have the same sign. Hence the probability is 1.

(g) (4 pt.) Let $A : \mathbb{R}^d \to \mathbb{R}^k$ be a random Gaussian Johnson–Lindenstrauss map, with $1 \leq k < d$. What is the probability that **there exists** a unit vector x, possibly chosen after seeing A, for which $Ax = 0$?

(A) 0 (B) 0.01 (C) 0.99 (D) 1

Solution. (D). Since A has more columns than rows, its columns are linearly dependent, so some nonzero vector v satisfies $Av = 0$. Then $x = v/\|v\|_2$ is a unit vector with $Ax = 0$, so the probability is 1.

2. Short arguments **32 points**

For each part, justify your answer in at most three sentences.

(a) (4 pt.) Let U be a finite nonempty set, and let $S_1, \dots, S_n \subseteq U$ be fixed subsets, each containing at least half of U. Draw $t = \log_2(100n)$ independent uniform elements of U, with replacement. Show that, with probability at least 0.99, every S_i contains at least one sampled element. The subsets can overlap arbitrarily.

Solution. A fixed set S_i is missed on all draws with probability at most $2^{-t} \le 1/(100n)$. A union bound over the n sets gives probability at most $1/100$ of missing any set, proving the claim.

(b) (3 pt.) Let $Z_1, \dots, Z_n$ be independent, each uniformly distributed on $\{-1, +1\}$, and let $S = \sum_i Z_i$. Prove that $\Pr[S \ge 2\sqrt{n}] \le 1/4$..

Solution. Each Z_i has mean 0 and variance 1, so independence gives $\mathbb{E}S = 0$ and $\mathrm{Var}(S) = n$. Chebyshev's inequality then gives

$$\Pr[S \ge 2\sqrt{n}] \le \Pr[|S| \ge 2\sqrt{n}] \le \frac{n}{(2\sqrt{n})^2} = \frac{1}{4}.$$

(c) (3 pt.) For S as defined above, prove that $\Pr[S \ge 2\sqrt{n}] \le 1/8$..

Solution. Flipping every sign shows that S and $-S$ have the same distribution. Therefore

$$\Pr[S \ge 2\sqrt{n}] = \frac{1}{2}\Pr[|S| \ge 2\sqrt{n}] \le \frac{1}{2} \cdot \frac{1}{4} = \frac{1}{8},$$

where the inequality uses the two-sided Chebyshev bound from part (b).

2. Short arguments (continued)

(d) (6 pt.) For S as defined above, prove that $\Pr[S \geq n/2] \leq e^{-n/8}$. You must prove this 'from scratch' without appealing to Chernoff bounds. Feel free to assume Markov's inequality, and you may also assume that

$$\mathbb{E}[e^{tZ_i}] \leq e^{t^2/2} \qquad \text{for every real } t.$$

If you leverage the above fact, do state the value of t you use.

Solution. For $t > 0$, apply Markov's inequality to e^{tS}, then use independence and the supplied bound:

$$\begin{aligned}\Pr[S \geq n/2] &\leq e^{-tn/2}\mathbb{E}[e^{tS}] \\ &= e^{-tn/2}\prod_i \mathbb{E}[e^{tZ_i}] \\ &\leq \exp\left(-\frac{tn}{2} + \frac{nt^2}{2}\right).\end{aligned}$$

Taking $t = 1/2$ gives $e^{-n/8}$.

(e) (3 pt.) A map f embeds a finite metric space (X, d) into Euclidean space with distortion at most D. A second map g satisfies, for all $x, y \in X$,

$$\|f(x) - f(y)\|_2 \leq \|g(f(x)) - g(f(y))\|_2 \leq 1.1\,\|f(x) - f(y)\|_2.$$

What upper bound on the distortion of $g \circ f$ follows?

Solution. For some common scale $\beta > 0$, the bounds for f and g give

$$\begin{aligned}\|g(f(x)) - g(f(y))\|_2 &\geq \|f(x) - f(y)\|_2 \geq \beta d(x, y), \\ \|g(f(x)) - g(f(y))\|_2 &\leq 1.1\|f(x) - f(y)\|_2 \leq 1.1D\beta d(x, y).\end{aligned}$$

Thus the distortion is at most $1.1D$.

(f) (3 pt.) In the setting of the previous part, If f is Bourgain's embedding into ℓ_2 and g is a Johnson–Lindenstrauss map, what are the properties of $g \circ f$ in terms of distortion?

Solution. Bourgain's embedding gives $D = O(\log |X|)$, so part (e) gives distortion $O(\log |X|)$ for $g \circ f$. JL reduces the dimension while changing the distortion bound by only a constant factor.

2. Short arguments (continued)

(g) (3 pt.) Describe a practical/efficient algorithm for sampling a uniformly random 1024-bit prime number (e.g. to use as a secret cryptographic key). [Feel free to use a randomized primality testing algorithm as a black-box, no need to describe how that works.]

Solution. Draw a fresh uniform integer from $\{2^{1023}, \ldots, 2^{1024} - 1\}$, apply an always-correct primality test with polynomial expected running time, and repeat until a prime is found. Every prime in this range has the same probability of being chosen and is always accepted, so the returned prime is uniform.

(h) (3 pt.) Describe the runtime of your algorithm above, in big-Oh notation, assuming you'll adapt it to sample a uniformly random n-bit prime number [ie give your answer in terms of n.]. Justify your answer in a few sentences.

Solution. The fraction of n-bit integers that are prime is $\Theta(1/n)$, so the expected number of candidates tested is $\Theta(n)$. If the chosen primality test takes expected time $O(n^c)$ for some constant c, the total expected time is $O(n^{c+1})$.

For the algorithm we saw in class, $c = 3$, though any solution that either just claimed a polynomial runtime or that attempted an analysis of the runtime of the algorithm we saw in class and concluded a polynomial runtime got full credit.

(i) (4 pt.) Consider the four-point metric space $X = \{a, b, c, d\}$. The distance from a to each of b, c, d is 1, and the distance between any two distinct points in $\{b, c, d\}$ is 2.

Give an embedding of this metric into the Euclidean plane $\mathbb{R}^2$ with distortion at most 2. You may give coordinates or a clearly labeled drawing. Briefly justify the distortion bound.

Solution. Take

$$f(a) = (0, 0), \qquad f(b) = (2, 0), \qquad f(c) = (-2, 0), \qquad f(d) = (0, 2).$$

The three distances from $f(a)$ are 2, while the other three distances are $4, 2\sqrt{2}, 2\sqrt{2}$. Each new distance lies between its original value and twice its original value, so the distortion is at most 2.

3. Checking a very large matrix power **20 points**

All matrices in this problem are real $n \times n$ matrices, where $n \geq 2$. Use exact arithmetic and count arithmetic operations. A matrix-vector product takes $O(n^2)$ operations. All matrices are fixed before any test randomness is chosen and stay fixed throughout the test.

(a) (6 pt.) Let $D \neq 0$ be a fixed matrix. Choose $r \in \{0,1\}^n$ with independent uniform coordinates. Prove that

$$\Pr[Dr = 0] \leq \frac{1}{2}.$$

Hint: Focus on a nonzero entry in one row, and condition on all but one coordinate of r.

Solution. Choose an entry $D_{ij} \neq 0$, and fix every coordinate of r except r_j. The ith coordinate of Dr is then $c + D_{ij} r_j$ for a fixed c, and at most one of the two equally likely values of r_j makes it zero. This holds for every choice of the other coordinates, so

$$\Pr[Dr = 0] \leq \Pr[(Dr)_i = 0] \leq \frac{1}{2}.$$

(b) (6 pt.) Given A, B, C, design an $O(n^2)$-time randomized test for the claim $AB = C$. It must always accept a correct claim and accept an incorrect claim with probability at most $1/2$. Explain the running time and both guarantees.

Solution. Choose r as in part (a), and accept exactly when $A(Br) = Cr$. Three matrix-vector products take $O(n^2)$ operations. If $AB = C$, the test always accepts; otherwise acceptance means $(AB - C)r = 0$, which has probability at most $1/2$ by part (a).

3. Checking a very large matrix power (continued)

(c) (8 pt.) Someone claims to have computed A^{2^t} for an integer $t \geq 1$. They supply a sequence

$$B_0 = A,\ B_1, \ldots, B_t,$$

claiming $B_i = B_{i-1}^2$ for every $i = 1, \ldots, t$.

Design an $O(tn^2)$-time test using **one random vector r shared across all t checks**. Your test must:

- always accept when every claimed equality is correct; and
- accept with probability at most $1/2$ whenever $B_t \neq A^{2^t}$.

Explain carefully why sharing r is sufficient.

Solution. Choose one vector r as in part (a), and accept only if

$$B_{i-1}(B_{i-1}r) = B_i r \qquad \text{for every } i = 1, \ldots, t.$$

This costs $O(tn^2)$, and every correct sequence passes.

If $B_t \neq A^{2^t}$, some equality $B_i = B_{i-1}^2$ must be false, since otherwise repeated squaring gives $B_i = A^{2^i}$ for every i. Fix one such index j using only the supplied matrices, before choosing r. Acceptance requires

$$(B_{j-1}^2 - B_j)r = 0.$$

The matrix in parentheses is fixed and nonzero, so part (a) bounds this probability by $1/2$. The checks need not be independent: accepting all of them requires passing this particular false equality.

4. A large sum-free subset **20 + 5 bonus points**

Let A be a set of $n \geq 1$ positive integers. A subset B is **sum-free** if there are no $x, y, z \in B$ satisfying $x + y = z$.

We will prove that A contains a sum-free subset with at least $n/3$ elements, and strengthen this bound in the bonus part.

Choose a real number α uniformly at random from $[0, 1)$. For each $x \in A$, define

$$r_\alpha(x) = \alpha x - \lfloor \alpha x \rfloor.$$

Thus, $r_\alpha(x)$ is the fractional part of αx, a number in $[0, 1)$. For example, if $x = 10$ and $\alpha = 0.123$, then $r_\alpha(x) = 0.23$ which is the fractional part of $10 \cdot 0.123 = 1.23$.

For a given α, define the set $B_\alpha \subset A$ to consist of those $x \in A$ for which

$$\frac{1}{3} < r_\alpha(x) < \frac{2}{3}.$$

(a) (3 pt.) For a fixed $x \in A$, show that $r_\alpha(x)$ is uniformly distributed on $[0, 1)$.

Solution. Partition $[0, 1)$ into the x equal intervals $[j/x, (j + 1)/x)$, for $j = 0, \ldots, x - 1$. Within any one interval, $r_\alpha(x) = x\alpha - j$ is uniform on $[0, 1)$. Since this is true whichever interval contains α, $r_\alpha(x)$ is uniform on $[0, 1)$.

(b) (5 pt.) Using the previous part (even if you didnt prove it), compute the expected number of elements in B_α.

Solution. By part (a), each $x \in A$ is selected with probability $2/3 - 1/3 = 1/3$. Linearity of expectation gives

$$\mathbb{E}|B_\alpha| = \sum_{x \in A} \Pr[x \in B_\alpha] = \frac{n}{3}.$$

The selections need not be independent.

4. A large sum-free subset (continued)

(c) (5 pt.) Prove that B_α is a "sum-free" set for every choice of α.

Solution. Suppose $x, y, z \in B_\alpha$ satisfy $x + y = z$, and put $s = r_\alpha(x) + r_\alpha(y)$, so $2/3 < s < 4/3$. The fractional part $r_\alpha(z)$ is s if $s < 1$, and $s - 1$ otherwise; in either case it lies outside $(1/3, 2/3)$, contradicting $z \in B_\alpha$. This also covers $x = y$.

(d) (7 pt.) Prove that A contains a sum-free subset of size at least $n/3$.

Solution. By part (b), the average size of B_α is $n/3$. Some choice of α therefore gives $|B_\alpha| \geq n/3$. By part (c), that subset is sum-free.

(e) (5 bonus points) Strengthen the conclusion of part (d): prove that A contains a sum-free subset with **strictly more than** $n/3$ **elements**.

Hint: What happens to B_α when α is sufficiently small?

Solution. Let $M = \max A$. If $0 < \alpha < 1/(3M)$, every $x \in A$ has $0 < \alpha x < 1/3$, so B_α is empty; this event has probability $1/(3M) > 0$. If every B_α had size at most $n/3$, then

$$\mathbb{E}|B_\alpha| \leq \left(1 - \frac{1}{3M}\right)\frac{n}{3} < \frac{n}{3},$$

contradicting part (b). Hence some B_α has strictly more than $n/3$ elements, and part (c) says it is sum-free.

End of the questions.

Reference sheet

These statements may be used as black boxes. Check their hypotheses when applying them.

Expectation and variance.

$$\mathbb{E}\sum_i a_i X_i = \sum_i a_i \mathbb{E}X_i, \qquad \mathrm{Var}(X) = \mathbb{E}[X^2] - (\mathbb{E}X)^2.$$

$$\mathrm{Cov}(X,Y) = \mathbb{E}[XY] - \mathbb{E}X\mathbb{E}Y, \qquad \mathrm{Var}\left(\sum_i X_i\right) = \sum_i \mathrm{Var}(X_i) + 2\sum_{i<j}\mathrm{Cov}(X_i, X_j).$$

For $X \sim \mathrm{Ber}(p)$, $\mathbb{E}X = p$ and $\mathrm{Var}(X) = p(1-p)$.

Markov and Chebyshev. For nonnegative X and $a > 0$, $\Pr[X \geq a] \leq \mathbb{E}X/a$. For a real random variable of finite variance,

$$\Pr[|X - \mathbb{E}X| \geq a] \leq \frac{\mathrm{Var}(X)}{a^2}.$$

Exponential moments. If X_i are independent, then

$$\mathbb{E}e^{t\sum_i X_i} = \prod_i \mathbb{E}e^{tX_i}.$$

For $t > 0$, $\Pr[X \geq a] \leq e^{-ta}\mathbb{E}e^{tX}$ whenever the expectation is finite.

Chernoff. If X is a sum of independent 0/1 variables and $\mu = \mathbb{E}X$, then for $0 < \delta \leq 1$,

$$\Pr[X \geq (1+\delta)\mu] \leq e^{-\mu\delta^2/3}, \qquad \Pr[X \leq (1-\delta)\mu] \leq e^{-\mu\delta^2/2}.$$

Poissonization. If $T \sim \mathrm{Poi}(\lambda)$ and each of T objects is independently placed uniformly into one of b bins, the bin counts are independent $\mathrm{Poi}(\lambda/b)$ variables. For $Y \sim \mathrm{Poi}(a)$,

$$\Pr[Y = j] = e^{-a}a^j/j!, \qquad \mathbb{E}Y = \mathrm{Var}(Y) = a.$$

Metric distortion. An embedding f has distortion at most D if there is a common scale $\beta > 0$ such that

$$\beta d(x,y) \leq \|f(x) - f(y)\| \leq D\beta d(x,y) \quad \text{for all } x, y.$$

Bourgain's construction embeds an n-point metric into ℓ_2 (also into ℓ_1) in $O(\log^2 n)$ dimensions with distortion $O(\log n)$, with high probability. Gaussian JL maps preserve all distances in a fixed n-point Euclidean set within factors $1 \pm \varepsilon$ with probability at least 0.99, using $O(\varepsilon^{-2}\log n)$ dimensions.

Basic facts. $\Pr[\bigcup_i E_i] \leq \sum_i \Pr[E_i]$, with no independence needed; $1 - x \leq e^{-x}$ for all real x; a finite collection of real numbers contains one at least as large as its average.