

A practice exam starts on the next page. We've done our best to draw this exam independently from the same distribution¹ as the real exam. Inspired by this, here is a warm-up question about how best to use this practice exam.

(Note: this is not a serious question :))

0. **(0 pt.)** Suppose that a class covers n topics. Each student i in the class has studied up on a set $S(i) \subseteq \{1, \dots, n\}$ of topics. To create the exam, the instructors choose k topics to test, independently and uniformly at random (with replacement) among the n topics, and give $100/k$ points to each. Suppose that n and k are sufficiently large², and that $k = o(n)$.
- (a) What is the expected score of student i , in terms of $|S(i)|$? (Assuming that student i acs any problem on a topic they have studied).
 - (b) Suppose that a student wants to use the practice exam to assess if they have studied enough. That is, the student studies first, and then scores s on the practice exam taken under the same conditions as the actual final. Bound the probability that they score worse than $s - 1$ (out of 100) on the actual final, asymptotically in terms of k and/or n .
 - (c) Suppose on the other hand that a student wants to use the practice exam to help guide their study. That is, they student studies the topics that appear on the practice exam as they take it. Say the student scores s on the practice exam taken in this open-book way. What can you say about how well the student will score on the actual final?
 - (d) How many topics should a student study up on to maximize their expected score on the final exam?

¹Okay, to be honest, this practice exam is not as vetted as a real exam would be, and it might involve a bit more reading/writing (i.e., the problems take longer to state and answer) than we'd ideally put on a timed exam. So it's not quite the same distribution. But we are shooting for the same distribution of, say, difficulty.

²It seems reasonable to us that there would be upwards of a million topics on the exam...

Instructions

- **DO NOT OPEN THE EXAM UNTIL YOU ARE INSTRUCTED TO.** (*Note: since this is a practice exam, go ahead when you feel ready!*)
- Answer all of the questions as well as you can. You have three hours to complete this exam.
- The exam is **non-collaborative**; you must complete it on your own. If you have any clarification questions, please ask the course staff (we are outside the exam room). We cannot provide any hints or help.
- This exam is **closed-book**, except for:
 - **Up to three double-sided sheets of paper** that you have prepared ahead of time. You can have anything you want written on these sheets of paper.
 - We have also provided a “cheat-sheet” with some helpful theorems and inequalities. **You can find this as the last page of this exam. Feel free to rip it off of the exam.**
- **Please DO NOT separate pages of your exam** (except for the cheat sheet at the back). The course staff is not responsible for finding lost pages, and you may not get credit for a problem if it goes missing.
- There are a few pages of extra paper at the back of the exam in case you run out of room on any problem. If you use them, please clearly indicate on the relevant problem page that you have used them, and please clearly label any work on the extra pages.

General Advice

- If you get stuck on a question or a part, move on and come back to it later. The questions on this exam have a wide range of difficulty, and you can do well on the exam even if you don't get a few questions.
- Pay attention to the point values. Don't spend too much time on questions that are not worth a lot of points.
- There are 100 points total on this exam.

Name (please print clearly):

Honor Code

The following is a statement of the Stanford University Honor Code:

1. *The Honor Code is an undertaking of the students, individually and collectively:*
 - (1) *that they will not give or receive aid in examinations; that they will not give or receive unpermitted aid in class work, in the preparation of reports, or in any other work that is to be used by the instructor as the basis of grading;*
 - (2) *that they will do their share and take an active part in seeing to it that others as well as themselves uphold the spirit and letter of the Honor Code.*
2. *The faculty on its part manifests its confidence in the honor of its students by refraining from proctoring examinations and from taking unusual and unreasonable precautions to prevent the forms of dishonesty mentioned above. The faculty will also avoid, as far as practicable, academic procedures that create temptations to violate the Honor Code.*
3. *While the faculty alone has the right and obligation to set academic requirements, the students and faculty will work together to establish optimal conditions for honorable academic work.*

For my part, I believe that we have upheld our end of the agreement in Item 2. I don't think we are taking unusual or unreasonable precautions that would indicate a lack of confidence in the honor of students, and I believe that the in-person setting avoids temptations to violate the honor code to the extent practicable.

[signed, Mary Wootters]

Please acknowledge that you have held up your end of the agreement in Item 1:

I have abided by the Honor Code, and in particular the policies listed above, both in letter and in spirit, while taking this exam.

signed, _____

Good Luck!

1. (26 pt.)

- (a) (7 pt.) Let $\varepsilon \in (0, 1/2)$. Suppose you have an algorithm A that outputs answers in $\{0, 1\}$ and is correct with probability $\frac{1}{2} + \varepsilon$. You decide to make a more robust algorithm, \tilde{A} , that just runs A independently T times and returns the most frequent answer. Show that \tilde{A} is correct with probability 0.99 for some value T that is $O(1/\varepsilon^2)$.

- (b) (7 pt.) Now suppose that A can output answers in $\{0, 1, \dots, n\}$, instead of just $\{0, 1\}$. Suppose that A is correct with probability at least $p \geq \frac{C \log n}{n \log \log n}$, where C is some constant that you get to choose. Further suppose that for any *incorrect* answer $i \in \{0, 1, \dots, n\}$, the probability that A outputs i is at most $1/n$. **Sketch a proof** that, for sufficiently large n , \tilde{A} (which still returns the most frequent answer out of the T trials) is correct with probability at least 0.99 when $T = n$.

Note: you don't need to give a super formal proof, but explain the steps you would go through to give a formal proof. (That is, what theorems/inequalities would you use, on what random variables, and how would they fit together to prove this result?)

[More parts on next page]

[Continued from previous page]

- (c) **(7 pt.) Sketch a proof** that shows that, in the previous part, you cannot take p to be substantially smaller. That is, explain why there are some constants $C', n_0 > 0$ so that, for any $n \geq n_0$, if $p \leq \frac{\log n}{C'n \log \log n}$, then when $T = n$, \tilde{A} could³ be incorrect with probability at least $1/2$.

Note: you don't need to give a super formal proof, but explain the steps you would go through to give a formal proof. (That is, what theorems/inequalities would you use, on what random variables, and how would they fit together to prove this result?)

- (d) **[May be more difficult] (5 pt.)** As above, say that A outputs answers in $\{0, 1, \dots, n\}$. Now suppose that A is correct with probability $1/4$, and can output any particular incorrect answer i with probability at most $1/8$. How small can you take T to still allow the guarantee that \tilde{A} is correct with probability at least 0.99 ?

³Here, the “could” is with respect to the probabilities of returning each answer. You should show that there exists a way of setting up these probabilities, consistent with the requirements in the previous part that the probability of each wrong answer be at most $1/n$, so that \tilde{A} is likely incorrect.

2. (24 pt.) For each of the following tasks, **briefly** sketch a randomized algorithm that does it and **briefly** explain why it works. You do not need to give a formal proof that it works. You can use any algorithm we have seen in class as a black box (unless otherwise noted), and your answer should be no more than a few sentences and possibly some very short pseudocode for each part.

(a) (6 pt.) Given a connected, undirected, unweighted graph G on n vertices and m edges, find a cut (S, \bar{S}) so that the number of edges crossing the cut is minimized, with probability at least 0.9. The algorithm should run in time $\text{poly}(n)$.

(b) (6 pt.) Given a data set $X \subseteq \mathbb{R}^N$ of size N with $\|x\|_2 = 1$ for all $x \in X$, give a randomized algorithm that returns estimates of $\|x - y\|_2$ for *all* pairs $x, y \in X$. With probability at least 0.99, your estimates should all be accurate up to a multiplicative factor of (1 ± 0.01) . Your algorithm should run in time $O(N^2 \log N)$.

[More parts on next page]

[Continued from previous page]

- (c) **(6 pt.)** Say you are given a 2-CNF formula φ (that is, φ is of the form $(x_1 \vee x_2) \wedge (\bar{x}_1 \vee x_3) \wedge \dots$, with n variables and m clauses that each contain two distinct literals). Describe a randomized algorithm that runs in time $\text{poly}(n, m)$ and, if there is a satisfying assignment to φ , returns it with probability at least 0.99. If there is no satisfying assignment, your algorithm should return “NOPE” with probability 1.

Note: Please describe your algorithm, don’t use an algorithm from class as a black box.

- (d) **(6 pt.)** Let (X, d) be an arbitrary finite metric space with $|X| = n$. We say that k points $x_1, \dots, x_k \in X$ form an r -cluster if $d(x_i, x_j) \leq r$ for all $i, j \in \{1, \dots, k\}$. For a general metric space (X, d) , the problem of finding an r -cluster of size k , given k and r , seems pretty hard; but fortunately you have access to a magic genie who can do it for (\mathbb{R}^d, ℓ_1) in time polynomial in n and in $2^{\sqrt{d}}$. (If there is no r -cluster of size k , the genie outputs “Sorry, no such cluster.”)

Give a randomized algorithm that runs in time $\text{poly}(n)$, takes as input k and r , and satisfies the following guarantee with probability at least 0.99, for some $D = O(\log n)$: If there is r/D -cluster in X of size k , the algorithm must output an r -cluster in X of size k . (Otherwise it can do whatever it wants).

3. **(15 pt.)** Let G be a cycle with $16n$ vertices. (That is, the vertices are labeled $0, 1, \dots, 16n-1$, and each vertex i is connected only to $i \pm 1 \pmod{16n}$). The vertices are each colored one of n colors, with 16 occurrences of each color. Show that it is always possible to find n vertices in G so that all are distinct colors, and no two are connected by an edge.

4. **(15 pt.)** Consider the following procedure for shuffling a deck of n cards: Choose two indices $i, j \in \{1, \dots, n\}$ uniformly and independently at random, and switch the card at position i and the card at position j . (Note that it is possible that $i = j$).

In this problem, we will use a coupling argument to bound the mixing time τ_{mix} of this procedure. Let X_t denote the state of the deck after we have swapped t cards.

- (a) **(5 pt.)** Your friend suggests the following coupling (which we also encountered on a quiz). Let X_t be the walk described above, and define Y_t to be a walk that makes the *same* choice of i and j at each step. Unfortunately, this isn't a great idea for bounding τ_{mix} . In at most a few sentences, explain why not.

- (b) **(10 pt.)** Here's another coupling to consider. We will view our shuffling procedure slightly differently: Instead of choosing i, j at random, choose a *card* c (like "the ace of spades") uniformly at random, and choose an index $i \in \{1, \dots, n\}$ uniformly at random. Then switch the card c with whatever card at index i . Note that this is an alternative way of defining the same Markov chain $\{X_t\}$. Now define a coupling (X_t, Y_t) by choosing the *same* choice of c and i in both chains.

Use this coupling to show that $\tau_{mix} = O(n^2)$.

[HINT: Keep track of a variable D_t which is defined to be the number of positions in which the decks X_t and Y_t differ. Show that $\Pr[D_{t+1} < D_t] \geq (D_t/n)^2$.]

[More space on next page]

[Continued from previous page; more space for part (b)]

- (c) **(0 pt.) BONUS** [We wouldn't put this on a real exam, but it might be fun to think about :)]. Show that $\tau_{mix} = O(n \log n)$.

5. **(20 pt.)** Let $s_1, s_2 \in \{0, 1\}^n$ denote two independent and uniformly random length n Boolean strings.

A *subsequence* of a string $s \in \{0, 1\}^n$ is any sequence of the form $s[i_1]s[i_2] \cdots s[i_\ell]$ for $i_1 < i_2 < \cdots < i_\ell$. For example, 000 is a subsequence of 010101. A *common subsequence* between strings s_1 and s_2 is a subsequence that's common to both. For example, $s_1 = 010101$ and $s_2 = 001100$ have a common subsequence 000. They also have a longer common subsequence, 0011. The *longest common subsequence* is a common subsequence with the most bits in it: in this example, one happens to be 0011.

- (a) **(10 pt.)** Letting L denote the length of the longest common subsequence of the two strings, prove that $\Pr[|L - \mathbb{E}[L]| \geq \lambda] \leq 2e^{-\frac{\lambda^2}{2n}}$.

- (b) **(5 pt.)** Prove that for sufficiently large n , with probability tending to 1 as $n \rightarrow \infty$, the length of the longest common subsequence of the two strings is at least $0.49n$.

[Another part on next page]

[Continued from previous page]

- (c) **(5 pt.) [May be more difficult]** Find a constant $c > 1/2$ and prove $\mathbb{E}[L] \geq cn$.

This is the end!

This is the end of the exam! You can use this page for extra work on any problem. **Keep this page attached** to the exam packet, and if you want work on it graded, clearly label which question your extra work is for.

This page is for extra work on any problem. **Keep this page attached** to the exam packet, and if you want work on it graded, clearly label which question your extra work is for.

This page is for extra work on any problem. **Keep this page attached** to the exam packet, and if you want work on it graded, clearly label which question your extra work is for.

Some useful inequalities, definitions and theorem statements

Note: We have not always stated full theorems here, just the quantitative punchlines. You are responsible for knowing when each theorem applies.

Inequalities and Series

- $1 - x \leq e^{-x}$ for any x .
- $(n/k)^k \leq \binom{n}{k} \leq (en/k)^k$ for all $k \leq n$.
- $\binom{n}{k} \leq \frac{n^k}{k!}$ for all $k \leq n$.
- $\sum_{i=1}^n 1/i = \Theta(\log n)$
- $\sum_{i=1}^n 1/i^c = O(1)$ for all $c > 1$.

Definitions

- $f(n) = O(g(n))$ means that there are some constants $c, n_0 > 0$ so that for all $n \geq n_0$, $f(n) \leq cg(n)$.
- $f(n) = \Omega(g(n))$ means that there are some constants $c, n_0 > 0$ so that for all $n \geq n_0$, $f(n) \geq cg(n)$.
- $f(n) = o(g(n))$ means that $\frac{f(n)}{g(n)} \rightarrow 0$ as $n \rightarrow \infty$.
- $f(n) = \omega(g(n))$ means that $\frac{f(n)}{g(n)} \rightarrow \infty$ as $n \rightarrow \infty$.
- If $X \sim \text{Poi}(\lambda)$, then $\Pr[X = k] = \frac{e^{-\lambda} \lambda^k}{k!}$.
- If $X \sim N(\mu, \sigma^2)$, then $\Pr[X = x] = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{1}{2} \left(\frac{x-\mu}{\sigma}\right)^2\right)$
- If $X \sim \text{Ber}(p)$, then $X \in \{0, 1\}$ and $\Pr[X = 1] = p$.

Concentration Inequalities

- Markov's inequality: For a non-negative random variable X , $\Pr[X > t] \leq \frac{\mathbb{E}X}{t}$.
- Chebyshev's inequality: For any random variable X , $\Pr[|X - \mathbb{E}X| > t] \leq \frac{\text{Var}(X)}{t^2}$.
- A few Chernoff bounds: For independent $X_i \in \{0, 1\}$, if $X = \sum_{i=1}^n X_i$, then:
 - For $\delta > 0$, $\Pr[X \geq (1 + \delta)\mathbb{E}[X]] \leq \left(\frac{e^\delta}{(1+\delta)^{1+\delta}}\right)^{\mathbb{E}[X]}$. If $\delta \in (0, 1]$ this is $\leq \exp(-\delta^2\mathbb{E}[X]/3)$.
 - For $\delta \in (0, 1]$, $\Pr[X \leq (1 - \delta)\mathbb{E}[X]] \leq \left(\frac{e^{-\delta}}{(1-\delta)^{1-\delta}}\right)^{\mathbb{E}[X]}$. If $\delta \in (0, 1]$, this is $\leq \exp(-\delta^2\mathbb{E}[X]/2)$.
 - For $c \geq 6$, $\Pr[X \geq c\mu] \leq 2^{-c\mu}$.

- Tail bound for Poisson random variables: If $X \sim \text{Poi}(\lambda)$, then for any $c > 0$, $\Pr[|X - \lambda| \geq c] \leq 2 \exp\left(\frac{-c^2}{2(c+\lambda)}\right)$.
- Azuma-Hoeffding Inequality: Let $\{Z_t\}$ be a martingale with respect to $\{X_t\}$, and suppose $|Z_i - Z_{i-1}| \leq c_i$ for all $i \leq n$. For any $\lambda > 0$, $\Pr[|Z_n - Z_0| \geq \lambda] \leq 2 \exp\left(\frac{-\lambda^2}{2 \sum_{i=1}^n c_i^2}\right)$.

Dimension Reduction

- Bourgain's Embedding: for any finite metric space (X, d) with $|X| = n$, there is an embedding of (X, d) into \mathbb{R}^k under the ℓ_1 metric with distortion $O(\log n)$, where $k = O((\log n)^2)$.
- Johnson-Lindenstrauss Lemma: for any $\varepsilon \in (0, 1)$, for any $X \subseteq \mathbb{R}^d$ with $|X| = n$, there is a linear map $f : \mathbb{R}^d \rightarrow \mathbb{R}^m$ with $m = O(\varepsilon^{-2} \log n)$ that embeds (X, ℓ_2) into (\mathbb{R}^m, ℓ_2) with distortion at most $(1 + \varepsilon)$.

Probabilistic Method

- Second moment method: for real-valued X , $\Pr[X = 0] \leq \frac{\text{Var}[X]}{(\mathbb{E}[X])^2}$.
- LLL: Let A_1, \dots, A_n be events so that $\Pr[A_i] \leq p$ for all i , and where each A_i is mutually independent of all but d other events. Then:
 - If $pd \leq 1/4$, then $\Pr[\bigcap_i \overline{A_i}] > 0$
 - If $p(d+1) \leq 1/e$, then $\Pr[\bigcap_i \overline{A_i}] > 0$.

Markov Chain / Martingale Theorems

- Fundamental theorem of Markov chains: Let $\{X_t\}$ be an irreducible aperiodic Markov chain over a finite state space with transition matrix P . Then there is a unique stationary distribution π so that $\Pr[X_t = i | X_0 = j] \rightarrow \pi_i$ for all states i, j . Further, π_i is the expected return time of state i , and $\pi P = \pi$.
- Let $\{X_t\}$ be a finite irreducible aperiodic Markov chain with a coupling $\{(X_t, Y_t)\}$. Then $\Delta(t) \leq \max_{s, s'} \Pr[X_t \neq Y_t | X_0 = s, Y_0 = s']$.
- Let $\{X_t\}$ be a finite irreducible aperiodic Markov chain and let T be a strong stationary stopping time. Then $\Delta(t) \leq \Pr[T > t]$.
- The *Doob Martingale* for a quantity A is $Z_t = \mathbb{E}[A | X_0, \dots, X_t]$. Theorem: it is a martingale.
- Martingale stopping theorem: Let $\{Z_t\}$ be a martingale with respect to $\{X_t\}$. Let T be a stopping time for $\{X_t\}$. Then $\mathbb{E}[Z_T] = \mathbb{E}[Z_0]$ if at least one of the following holds:
 1. There is a constant c s.t. $|Z_i| \leq c$ for all i .
 2. There is a constant c s.t. $T < c$ with probability 1.
 3. $\mathbb{E}[T] < \infty$ and there is a constant c s.t. for all i , $\mathbb{E}[|Z_{i+1} - Z_i| | X_0, \dots, X_i] < c$.