

Midterm Exam

- Don't Panic.
- The midterm contains six problems (and one just for fun). You have 100 minutes to earn 100 points.
- The midterm contains 20 pages, including this one and 4 pages of scratch paper.
- The midterm is closed book. You may bring one double-sided sheet of A4 paper to the midterm. (You may not bring any magnification equipment!) You may not use a calculator, your mobile phone, or any other electronic device.
- Write your solutions in the space provided. If you need more space, please use the scratch paper at the end of the midterm. Do not put part of the answer to one problem on a page for another problem.
- Read through the problems before starting. Do not spend too much time on any one problem.
- Show your work. Partial credit will be given. You will be graded not only on the correctness of your answer, but also on the clarity with which you express it. Be neat.
- Draw pictures and give examples.
- Good luck!

Problem #	Name	Possible Points	Achieved Points
1	A Few Random Variables	10	
2	A Simple Game of Chance	10	
3	In the Dungeon	10	
4	Average Case Insertion Sort	20	
5	Scheduling Exams	15	
6	Random Graphs	35	
Total:		100	

Student Number: _____

Problem 1. A few random variables [10 points]

Let x_1, x_2, \dots, x_n be a collection of n independent indicator random variables where each $x_j = 1$ with probability p and $x_j = 0$ with probability $(1 - p)$. Let $X = \sum_{i=1}^n x_i$.

Let S be a subset of k of the random variables, and let E_S be the event that:

$$\forall x_j \in S, x_j = 1 .$$

For example, if $S = \{x_1, x_2\}$, then $k = 2$ and E_S is the event that the first two random variables are 1.

For each part below, give an exact answer (not an approximation) as a function of n , p , and k .

Problem 1.a. What is $E[X]$?

Problem 1.b. What is variance $\text{Var}[X]$?

Problem 1.c. What is $\Pr[X = 0]$?

Problem 1.d. For an arbitrary $k > 0$, what is $\Pr[X = k]$?

Problem 1.e. Given some S and k , what is $E[X|E_S]$?

Problem 2. A simple game of chance [10 points]

Alice has a 5-sided die. (Weird, I know.) When you roll the die, each number from 1 to 5 comes up with equal probability. Alice also has a coin. On one side of the coin is the number 2 and on the other side of the coin is the number 4.



5-sided die

coin with 2 and 4 on each side

Problem 2.a. (*Expected Value.*)

What is the expected value of rolling the die?

What is the expected value of flipping the coin?

Problem 2.b. (*Variance.*)

What is the variance of rolling the die?

What is the variance of flipping the coin?

Problem 2.c. Alice proposes the following fun game:

- You roll the die and get a number from 1 to 5.
- You flip the coin and get either a 2 or a 4.
- You multiply the results together.

For example, if the die rolls a 5 and the coin flips a 4, then the result is 20.

Alice reasons that we can find the expected value of the game using your answers from the previous part:

If d is the expected value of the die and c is the expected value of the coin, then $d \cdot c$ is the expected value of the game.

Is Alice's reasoning right or wrong? (Circle your answer.)

RIGHT **WRONG**

Explain your answer: if Alice is correct, explain why; if Alice is wrong, give the correct expected value of the game.

Problem 2.d. Alice also reasons that we can find the variance of the game using your answers from the previous part:

If v_1 is the variance of the die and v_2 is the variance of the coin, then $v_1 \cdot v_2$ is the variance of the game.

Is Alice's reasoning right or wrong? (Circle your answer.)

RIGHT **WRONG**

Explain your answer: if Alice is correct, explain; if Alice is wrong, give the correct variance.

Problem 3. In the dungeon [10 points]

Redmore the Red, an adventurer of yore, is arrested by the king and thrown in the dungeon. Also in the dungeon are two other adventurers, Alice and Bob.¹ The prison guard informs them that one of them has been selected, uniformly at random, for execution.

Redmore is naturally quite afraid, and he asks the guard which of them is to be executed. The guard responds that he is not allowed to tell Redmore his fate, but that he would provide the following information:

- If Redmore *was not* selected for execution, then the guard will tell Redmore accurately whether Alice or Bob is to be executed.
- If Redmore *was* selected for execution, then the guard will randomly flip a coin and with probability $1/2$ indicate Alice, and with probability $1/2$ indicate Bob.

Redmore agrees to this plan. (He did not have much choice, did he?) The guard then tells Redmore that Alice is going to be executed.

Given the guard's information:

- What is the probability that Alice is executed?

- What is the probability that Bob is executed?

- What is the probability that Redmore is executed?

Explain your answer:

¹Alice was arrested for cheating in a weird game involving a five-sided die.

Problem 4. Average-Case Analysis of InsertionSort [20 points]

One of the first methods that we teach for sorting is *InsertionSort*. The algorithm proceeds through the array, sorting each element into place, one at a time, by comparing to all the preceding items in the list. The details are unimportant for this problem, but the pseudocode is as follows:

```
InsertionSort(Array A, integer n)
  for i=1 to (n-1) do
    item = A[i];
    int slot = i;
    while (slot > 0) and (A[slot] > item) do
      A[slot] = A[slot-1];
      slot = slot-1;
    A[slot] = item;
}
```

The key property that is important is that the running time of InsertionSort depends on the number of *inversions* in the permutation being sorted. Given a sequence of integers $\{a_1, a_2, \dots, a_n\}$, an **inversion** is a mis-ordered pair (a_j, a_k) where $j < k$ but $a_j > a_k$. For example, in the sequence:

$$S_1 = \{1 \ 2 \ 5 \ 3 \ 10\}$$

there is one inversion: the pair $(5, 3)$. As another example, consider the sequence:

$$S_2 = \{1 \ 2 \ 10 \ 5 \ 3\}$$

This sequence contains three inversions: $(10, 5)$, $(10, 3)$, and $(5, 3)$. The following theorem relates the running time of InsertionSort to the number of inversions in the input:

Theorem 1 *If a permutation S of length n contains I inversions, then $\text{InsertionSort}(S, n)$ runs in time $\Theta(n + I)$.*

Continued on the next page.

Analyze the *average-case* performance of InsertionSort. Let S be a permutation of the integers $[1, \dots, n]$ chosen uniformly at random from the set of all permutations. Show that the expected running time of InsertionSort on S is $\Theta(n^2)$.

Solution to Problem 5:

Problem 5. Scheduling exams [15 points]

Imagine a university with n students and m modules. Each module has a final exam. And each student takes at least k different modules. (Some may take more than k modules!) You may assume that $k \geq c \log n$ for some constant c . (Do not worry if c is a large integer.)

The exam period consists of two weeks: Week A and Week B. Each exam is randomly assigned to Week A with probability $1/2$ or Week B with probability $1/2$.

Your goal is to prove that, with high probability, no student has too many exams in either week: if a student is taking x exams, then no more than $3/4$ of her exams should be in Week A and no more than $3/4$ of her exams should be in Week B.

Example: Perhaps Alice is taking {CS1001, CS2001, CS3001, and CS4001}. Perhaps Bob is taking {CS2001, CS2002, CS2003, CS2004}. Perhaps the random assignment assigns to Week A the following exams: {CS1001, CS2001, CS2002, CS2003}; it assigns to Week B the following exams: {CS2004, CS3001, CS4001}.

In this case, both students are happy since they have $\leq 3/4$ of their exams in each week: Alice has two exams in Week A and two exams in Week B; Bob has three exams in Week A and one exam in Week B.

Problem 5.a. Give a one sentence overview of your proof strategy.

Problem 5.b. Define the random variables that you will use.

Continued on the next page.

Problem 5.c. What probabilistic tools / tail bounds / etc. will you use to prove this? Give the complete definition of the bound (not just its name).

Problem 5.d. Recall that we assumed each student is taking at least $k \geq c \log n$ modules, for some constant c . What value of c will you require in your proof? (You can choose any constant. Do not worry if it is unrealistically large.)


Choose $c =$

Problem 5.e. Give the complete proof that with probability at least $(1 - 2/n)$, every student has no more than $3/4$ of their exams in either week.

Problem 6. Random Graphs [30 points]

Imagine we have a random graph $G(n, p)$ that we construct as follows:

- Begin with n nodes and no edges. Let V be the set of n nodes. Let edges $E = \emptyset$.
- For each pair of nodes $u, v \in V$, add (undirected) edge (u, v) to E with probability p .




Problem 6.a. Assume $p = \frac{18 \log n}{n}$. Prove that with probability at least $1 - 1/n$, every node in graph $G(n, p)$ has degree at most $O(\log n)$.

Problem 6.b. Assume $p = \frac{1}{3n}$. Prove that with probability at least $1/2$, graph $G(n, p)$ is *not connected*. (You may assume that n is sufficiently large, e.g., $n > 10$.)

Hint: What is the expected number of edges?

Problem 6.c. How many triangles are there in $G(n, p)$?

A triangle is a set of nodes $u, v, w \in V$ where all three nodes are connected. That is, we say that $T = (u, v, w)$ is a triangle in graph $G = (V, E)$ if: $(u, v) \in E, (v, w) \in E, (u, w) \in E$.



Consider the random graph $G(n, p)$ where $p = \frac{1}{2n}$. Show that with probability at least $1/2$, graph $G(n, p)$ has no triangles.

Problem 6.d. Now we want to show the opposite claim: that if $p = \frac{\alpha}{n}$, for some constant $\alpha > 1$, then graph $G(n, p)$ has at least one triangle with constant probability. We will do this in two steps: first, we will bound the variance on the number of triangles, and then we will do the actual probability calculation. Assume here that α is some constant (e.g., $\alpha \geq 10$), and that n is sufficiently large (e.g., $n > 100\alpha^{10}$).

Let C be the random variable representing the number of triangles in $G(n, p)$. Assume that we have already proven that the variance $\text{Var}[C] \leq 2\alpha^3$. Use Chebychev's Inequality to show that with probability at least $1/2$, graph $G(n, p)$ has at least one triangle.

Problem 6.e. Prove that variance $\text{Var}[C] \leq 2\alpha^3$.

Hint 1: For a triangle T , let x_T be an indicator random variable where $X_T = 1$ if triangle T is in the graph $G(n, p)$ and $x_T = 0$ otherwise. Let $C = \sum_T x_T$. The variance of C is:

$$\begin{aligned} \text{Var}[C] &= \sum_{T_1} \sum_{T_2} \text{CoVar}[T_1, T_2] \\ \text{CoVar}[T_1, T_2] &= \text{E}[x_{T_1} x_{T_2}] - \text{E}[x_{T_1}] \text{E}[x_{T_2}] \end{aligned}$$

Hint 2: You may want to divide this summation into three components: the set of triangle pairs (T_1, T_2) that share: (i) zero edges, (ii) one edge, and (iii) three edges.

Hint 3: In some cases, you also may find it useful to observe that $\text{CoVar}[T_1, T_2] \leq \text{E}[x_{T_1} x_{T_2}]$.

Note: You will observe that we have shown something interesting here! The probability $1/n$ is a critical threshold for the random graph $G(n, p)$. If $p > 1/n$, then we are likely to have one triangle; if $p < 1/n$ we are likely to have no triangles. And a triangle is really just a cycle of length 3. What if we considered a cycle of length k ? Does graph $G(n, p)$ contain a cycle of length k ? Would it be surprising if $1/n$ were a critical threshold for this question as well?

Scratch Paper

Scratch Paper

Scratch Paper

Scratch Paper