

☞ Midterm 1 Study Questions ☞

This is a “core dump” of potential questions for Midterm 1. This should give you a good idea of the *types* of questions that we will ask on the exam—in particular, there *will* be a series of True/False questions—but the actual exam questions may or may not appear in this handout. This list intentionally includes a few questions that are too long or difficult for exam conditions; most of these are indicated with a *star.

Questions from Jeff’s past exams are labeled with the semester they were used: **⟨⟨S14⟩⟩**, **⟨⟨F14⟩⟩**, or **⟨⟨F16⟩⟩**. Questions from this semester’s homework are labeled **⟨⟨HW⟩⟩**. Questions from this semester’s labs are labeled **⟨⟨Lab⟩⟩**. Some unflagged questions may have been used in exams by other instructors.

☞ How to Use These Problems ☞

Solving every problem in this handout is **not** the best way to study for the exam. Memorizing the solutions to every problem in this handout is the **absolute worst** way to study for the exam.

What we recommend instead is to work on a *sample* of the problems. Choose one or two problems at random from each section and try to solve them from scratch under exam conditions—by yourself, in a quiet room, with a 30-minute timer, *without* your notes, *without* the internet, and if possible, even without your cheat sheet. If you’re comfortable solving a few problems in a particular section, you’re probably ready for that type of problem on the exam. Move on to the next section.

Discussing problems with other people (in your study groups, in the review sessions, in office hours, or on Piazza) and/or looking up old solutions can be *extremely* helpful, but **only after** you have (1) made a good-faith effort to solve the problem on your own, and (2) you have either a candidate solution or some idea about where you’re getting stuck.

If you find yourself getting stuck on a particular type of problem, try to figure out *why* you’re stuck. Do you understand the problem statement? Are you stuck on choosing the right high-level approach, are you stuck on the technical details, or are you struggling to express your ideas clearly?

Similarly, if feedback suggests that your solutions to a particular type of problem are incorrect or incomplete, try to figure out what you missed. For induction proofs: Are you sure you have the right induction hypothesis? Are your cases obviously exhaustive? For regular expressions, DFAs, NFAs, and context-free grammars: Is your solution both exclusive and exhaustive? Did you try a few positive examples *and* a few negative examples? For fooling sets: Are you imposing enough structure? Are x and y really *arbitrary* strings from F ? For language transformations: Are you transforming in the right direction? Are you using non-determinism correctly? Do you understand the formal notation for DFAs and NFAs?

Remember that your goal is *not* merely to “understand” the solution to any particular problem, but to become more comfortable with solving a certain *type* of problem on your own. **“Understanding” is a trap; aim for mastery.** If you can identify specific steps that you find problematic, read more *about those steps*, focus your practice *on those steps*, and try to find helpful information *about those steps* to write on your cheat sheet. Then work on the next problem!

Recursion and Dynamic Programming

Elementary Recursion/Divide and Conquer

1. **⟨⟨Lab⟩⟩**

- (a) Suppose $A[1..n]$ is an array of n distinct integers, sorted so that $A[1] < A[2] < \dots < A[n]$. Each integer $A[i]$ could be positive, negative, or zero. Describe a fast algorithm that either computes an index i such that $A[i] = i$ or correctly reports that no such index exists..
- (b) Now suppose $A[1..n]$ is a sorted array of n distinct **positive** integers. Describe an even faster algorithm that either computes an index i such that $A[i] = i$ or correctly reports that no such index exists. [*Hint: This is really easy.*]

2. **⟨⟨Lab⟩⟩** Suppose we are given an array $A[1..n]$ such that $A[1] \geq A[2]$ and $A[n-1] \leq A[n]$. We say that an element $A[x]$ is a **local minimum** if both $A[x-1] \geq A[x]$ and $A[x] \leq A[x+1]$. For example, there are exactly six local minima in the following array:

9	7	7	2	1	3	7	5	4	7	3	3	4	8	6	9
	▲			▲				▲		▲	▲			▲	

Describe and analyze a fast algorithm that returns the index of one local minimum. For example, given the array above, your algorithm could return the integer 5, because $A[5]$ is a local minimum. [*Hint: With the given boundary conditions, any array must contain at least one local minimum. Why?*]

3. **⟨⟨Lab⟩⟩** Suppose you are given two sorted arrays $A[1..n]$ and $B[1..n]$ containing distinct integers. Describe a fast algorithm to find the median (meaning the n th smallest element) of the union $A \cup B$. For example, given the input

$$A[1..8] = [0, 1, 6, 9, 12, 13, 18, 20] \quad B[1..8] = [2, 4, 5, 8, 17, 19, 21, 23]$$

your algorithm should return the integer 9. [*Hint: What can you learn by comparing one element of A with one element of B ?*]

4. **⟨⟨F14, S14⟩⟩** An array $A[0..n-1]$ of n distinct numbers is **bitonic** if there are unique indices i and j such that $A[(i-1) \bmod n] < A[i] > A[(i+1) \bmod n]$ and $A[(j-1) \bmod n] > A[j] < A[(j+1) \bmod n]$. In other words, a bitonic sequence either consists of an increasing sequence followed by a decreasing sequence, or can be circularly shifted to become so. For example,

4	6	9	8	7	5	1	2	3	is bitonic, but
3	6	9	8	7	5	1	2	4	is not bitonic.

Describe and analyze an algorithm to find the index of the *smallest* element in a given bitonic array $A[0..n-1]$ in $O(\log n)$ time. You may assume that the numbers in the input array are distinct. For example, given the first array above, your algorithm should return 6, because $A[6] = 1$ is the smallest element in that array.

5. **⟨⟨F16⟩⟩** Suppose you are given a sorted array $A[1..n]$ of distinct numbers that has been rotated k steps, for some **unknown** integer k between 1 and $n - 1$. That is, the prefix $A[1..k]$ is sorted in increasing order, the suffix $A[k + 1..n]$ is sorted in increasing order, and $A[n] < A[1]$. For example, you might be given the following 16-element array (where $k = 10$):

9	13	16	18	19	23	28	31	37	42	-4	0	2	5	7	8
---	----	----	----	----	----	----	----	----	----	----	---	---	---	---	---

Describe and analyze an efficient algorithm to determine if the given array contains a given number x . The input to your algorithm is the array $A[1..n]$ and the number x ; your algorithm is **not** given the integer k .

6. **⟨⟨F16⟩⟩** Suppose you are given two unsorted arrays $A[1..n]$ and $B[1..n]$ containing $2n$ distinct integers, such that $A[1] < B[1]$ and $A[n] > B[n]$. Describe and analyze an efficient algorithm to compute an index i such that $A[i] < B[i]$ and $A[i + 1] > B[i + 1]$. [Hint: Why does such an index i always exist?]
7. Suppose you are given a stack of n pancakes of different sizes. You want to sort the pancakes so that smaller pancakes are on top of larger pancakes. The only operation you can perform is a *flip*—insert a spatula under the top k pancakes, for some integer k between 1 and n , and flip them all over.




Figure 1. Flipping the top four pancakes.


- (a) Describe an algorithm to sort an arbitrary stack of n pancakes using as few flips as possible. *Exactly* how many flips does your algorithm perform in the worst case?
- (b) Now suppose one side of each pancake is burned. Describe an algorithm to sort an arbitrary stack of n pancakes, so that the burned side of every pancake is facing down, using as few flips as possible. *Exactly* how many flips does your algorithm perform in the worst case?

[Hint: This problem has **nothing** to do with the Tower of Hanoi!]

8. (a) Describe an algorithm to determine in $O(n)$ time whether an arbitrary array $A[1..n]$ contains more than $n/4$ copies of any value.
- (b) Describe and analyze an algorithm to determine, given an arbitrary array $A[1..n]$ and an integer k , whether A contains more than k copies of any value. Express the running time of your algorithm as a function of both n and k .

Do not use hashing, or radix sort, or any other method that depends on the precise input values, as opposed to their order.

9. For this problem, a *subtree* of a binary tree means any connected subgraph. A binary tree is *complete* if every internal node has two children, and every leaf has exactly the same depth. Describe and analyze a recursive algorithm to compute the *largest complete subtree* of a given binary tree. Your algorithm should return both the root and the depth of this subtree.



The largest complete subtree of this binary tree has depth 2.

Dynamic Programming

1. **⟨⟨Lab⟩⟩** Describe and analyze efficient algorithms for the following problems.
 - (a) Given an array $A[1..n]$ of integers, compute the length of a longest **increasing** subsequence of A . A sequence $B[1..l]$ is *increasing* if $B[i] > B[i - 1]$ for every index $i \geq 2$.
 - (b) Given an array $A[1..n]$ of integers, compute the length of a longest **decreasing** subsequence of A . A sequence $B[1..l]$ is *decreasing* if $B[i] < B[i - 1]$ for every index $i \geq 2$.
 - (c) Given an array $A[1..n]$ of integers, compute the length of a longest **alternating** subsequence of A . A sequence $B[1..l]$ is *alternating* if $B[i] < B[i - 1]$ for every even index $i \geq 2$, and $B[i] > B[i - 1]$ for every odd index $i \geq 3$.
 - (d) Given an array $A[1..n]$ of integers, compute the length of a longest **convex** subsequence of A . A sequence $B[1..l]$ is *convex* if $B[i] - B[i - 1] > B[i - 1] - B[i - 2]$ for every index $i \geq 3$.
 - (e) Given an array $A[1..n]$, compute the length of a longest **palindrome** subsequence of A . Recall that a sequence $B[1..l]$ is a *palindrome* if $B[i] = B[l - i + 1]$ for every index i .

2. **⟨⟨F16, HW⟩⟩** It's almost time to show off your flippin' sweet dancing skills! Tomorrow is the big dance contest you've been training for your entire life, except for that summer you spent with your uncle in Alaska hunting wolverines. You've obtained an advance copy of the the list of n songs that the judges will play during the contest, in chronological order.

You know all the songs, all the judges, and your own dancing ability extremely well. For each integer k , you know that if you dance to the k th song on the schedule, you will be awarded exactly $Score[k]$ points, but then you will be physically unable to dance for the next $Wait[k]$ songs (that is, you cannot dance to songs $k + 1$ through $k + Wait[k]$). The dancer with the highest total score at the end of the night wins the contest, so you want your total score to be as high as possible.

Describe and analyze an efficient algorithm to compute the maximum total score you can achieve. The input to your sweet algorithm is the pair of arrays $Score[1..n]$ and $Wait[1..n]$.

3. **⟨⟨S16⟩⟩** After the Revolutionary War, Alexander Hamilton's biggest rival as a lawyer was Aaron Burr. (Sir!) In fact, the two worked next door to each other. Unlike Hamilton, Burr cannot work non-stop; every case he tries exhausts him. The bigger the case, the longer he must rest before he is well enough to take the next case. (Of course, he is willing to wait for it.) If a case arrives while Burr is resting, Hamilton snatches it up instead.

Burr has been asked to consider a sequence of n upcoming cases. He quickly computes two arrays $profit[1..n]$ and $skip[1..n]$, where for each index i ,

 - $profit[i]$ is the amount of money Burr would make by taking the i th case, and
 - $skip[i]$ is the number of consecutive cases Burr must skip if he accepts the i th case. That is, if Burr accepts the i th case, he cannot accept cases $i + 1$ through $i + skip[i]$.

Design and analyze an algorithm that determines the maximum total profit Burr can secure from these n cases, using his two arrays as input.

4. **⟨⟨S14⟩⟩** Recall that a *palindrome* is any string that is the same as its reversal. For example, I, DAD, HANNAH, AIBOHPHOBIA (fear of palindromes), and the empty string are all palindromes.
- (a) Describe and analyze an algorithm to find the length of the longest substring (not *subsequence*!) of a given input string that is a palindrome. For example, **BASEESAB** is the longest palindrome substring of **BUBBASEESABANANA** (“Bubba sees a banana.”). Thus, given the input string **BUBBASEESABANANA**, your algorithm should return the integer 8.
- (b) **⟨⟨Lab, F16⟩⟩** Describe and analyze an algorithm to find the length of the longest subsequence (not *substring*!) of a given input string that is a palindrome. For example, the longest palindrome subsequence of **MAHDYNAMICPROGRAMZLETMESHOWYOUTHEM** is **MHYMRORMYHM**, so given that string as input, your algorithm should output the number 11.
- (c) **⟨⟨HW⟩⟩** Any string can be decomposed into a sequence of palindrome substrings. For example, the string **BUBBASEESABANANA** can be broken into palindromes in the following ways (and many others):

BUB + BASEESAB + ANANA
B + U + BB + A + SEES + ABA + NAN + A
B + U + BB + A + SEES + A + B + ANANA
B + U + B + B + A + S + E + E + S + A + B + A + N + A + N + A

Describe and analyze an algorithm to find the smallest number of palindromes that make up a given input string. For example, if your input is the string **BUBBASEESA-BANANA**, your algorithm should return the integer 3.

5. **⟨⟨F16⟩⟩** A *shuffle* of two strings X and Y is formed by interspersing the characters into a new string, keeping the characters of X and Y in the same order. For example, the string **BANANAANANAS** is a shuffle of the strings **BANANA** and **ANANAS** in several different ways.

BANANA_{ANANAS} **BAN**_{ANA}**ANA**_{NAS} **B**_{AN}**AN**_A**A**_{NA}**NA**_S

Similarly, the strings **PROGYRNAMAMMIINCG** and **DYPRONGARMAMMICING** are both shuffles of **DYNAMIC** and **PROGRAMMING**:

PRO_D**G**_Y**R**_{NAM}**AMMI**_I**N**_C**G** **DY**_{PRO}**N**_G**A**_R**M**_{AMM}**I**_C**ING**

Describe and analyze an efficient algorithm to determine, given three strings $A[1..m]$, $B[1..n]$, and $C[1..m+n]$, whether C is a shuffle of A and B .

6. Suppose we are given an n -digit integer X . Repeatedly remove one digit from either end of X (your choice) until no digits are left. The *square-depth* of X is the maximum number

of perfect squares that you can see during this process. For example, the number 32492 has square-depth 3, by the following sequence of removals:

$$32492 \rightarrow \underline{3249}2 \rightarrow \underline{324}9 \rightarrow 324 \rightarrow \underline{24} \rightarrow 4.$$

Describe and analyze an algorithm to compute the square-depth of a given integer X , represented as an array $X[1..n]$ of n decimal digits. Assume you have access to a subroutine `IS SQUARE` that determines whether a given k -digit number (represented by an array of digits) is a perfect square **in $O(k^2)$ time**.

7. Suppose you are given a sequence of non-negative integers separated by $+$ and \times signs; for example:

$$2 \times 3 + 0 \times 6 \times 1 + 4 \times 2$$

You can change the value of this expression by adding parentheses in different places. For example:


$$\begin{aligned} 2 \times (3 + (0 \times (6 \times (1 + (4 \times 2)))))) &= 6 \\ (((((2 \times 3) + 0) \times 6) \times 1) + 4) \times 2 &= 80 \\ ((2 \times 3) + (0 \times 6)) \times (1 + (4 \times 2)) &= 108 \\ (((2 \times 3) + 0) \times 6) \times ((1 + 4) \times 2) &= 360 \end{aligned}$$

Describe and analyze an algorithm to compute, given a list of integers separated by $+$ and \times signs, the smallest possible value we can obtain by inserting parentheses.

Your input is an array $A[0..2n]$ where each $A[i]$ is an integer if i is even and $+$ or \times if i is odd. Assume any arithmetic operation in your algorithm takes $O(1)$ time.

8. Suppose you are given an array $A[1..n]$ of numbers, which may be positive, negative, or zero, and which are **not** necessarily integers.
- Describe and analyze an algorithm that finds the largest sum of elements in a contiguous subarray $A[i..j]$.
 - Describe and analyze an algorithm that finds the largest *product* of elements in a contiguous subarray $A[i..j]$.

For example, given the array $[-6, 12, -7, 0, 14, -7, 5]$ as input, your first algorithm should return the integer 19, and your second algorithm should return the integer 504.



For the sake of analysis, assume that comparing, adding, or multiplying any pair of numbers takes $O(1)$ time.

[Hint: Problem (a) has been a standard computer science interview question since at least the mid-1980s. You can find many correct solutions on the web; the problem even has its own [Wikipedia page](#)! But at least in 2016, a significant fraction of the solutions I found on the web for problem (b) were either significantly slower than necessary or actually incorrect. Remember that the product of two negative numbers is positive.]

9. Suppose you are given three strings $A[1..n]$, $B[1..n]$, and $C[1..n]$.
- (a) Describe and analyze an algorithm to find the length of the longest common subsequence of all three strings. For example, given the input strings

$$A = \text{AxxBxxCDxEF}, \quad B = \text{yyABCDyEyFy}, \quad C = \text{zAzzBCDzEFz},$$

your algorithm should output the number 6, which is the length of the longest common subsequence **ABCDEF**.

- (b) Describe and analyze an algorithm to find the length of the shortest common supersequence of all three strings. For example, given the input strings

$$A = \text{AxxBxxCDxEF}, \quad B = \text{yyABCDyEyFy}, \quad C = \text{zAzzBCDzEFz},$$

your algorithm should output the number 21, which is the length of the shortest common supersequence **zyAxzzxBxxCDxyzEFzy**.

10. (a) Suppose we are given a set L of n line segments in the plane, where each segment has one endpoint on the line $y = 0$ and one endpoint on the line $y = 1$, and all $2n$ endpoints are distinct. Describe and analyze an algorithm to compute the largest subset of L in which no pair of segments intersects.
- (b) Suppose we are given a set L of n line segments in the plane, where each segment has one endpoint on the line $y = 0$ and one endpoint on the line $y = 1$, and all $2n$ endpoints are distinct. Describe and analyze an algorithm to compute the largest subset of L in which *every* pair of segments intersects.
11. Suppose you are given an $m \times n$ bitmap, represented by an array $M[1..n, 1..m]$ of 0s and 1s. A *solid square block* in M is a subarray of the form $M[i..i+w, j..j+w]$ containing only 1-bits. Describe and analyze an algorithm to find the largest solid square block in M .
12. You and your six-year-old nephew Elmo decide to play a simple card game. At the beginning of the game, the cards are dealt face up in a long row. Each card is worth a different number of points. After all the cards are dealt, you and Elmo take turns removing either the leftmost or rightmost card from the row, until all the cards are gone. At each turn, you can decide which of the two cards to take. The winner of the game is the player that has collected the most points when the game ends.

Having never taken an algorithms class, Elmo follows the obvious greedy strategy—when it's his turn, Elmo *always* takes the card with the higher point value. Your task is to find a strategy that will beat Elmo whenever possible. (It might seem mean to beat up on a little kid like this, but Elmo absolutely *hates* it when grown-ups let him win.)

- (a) Prove that you should not also use the greedy strategy. That is, show that there is a game that you can win, but only if you do *not* follow the same greedy strategy as Elmo.
- (b) Describe and analyze an algorithm to determine, given the initial sequence of cards, the maximum number of points that you can collect playing against Elmo.

(c) Five years later, thirteen-year-old Elmo has become a *much* stronger player. Describe and analyze an algorithm to determine, given the initial sequence of cards, the maximum number of points that you can collect playing against a *perfect* opponent.

13. **⟨⟨S16⟩⟩** Your nephew Elmo is visiting you for Christmas, and he's brought a different card game. Like your previous game with Elmo, this game is played with a row of n cards, each labeled with an integer (which could be positive, negative, or zero). Both players can see all n card values. Otherwise, the game is almost completely different.

On each turn, the current player must take the leftmost card. The player can either keep the card or give it to their opponent. If they keep the card, their turn ends and their opponent takes the next card; however, if they give the card to their opponent, the current player's turn continues with the next card. In short, the player that does *not* get the i th card decides who gets the $(i + 1)$ th card. The game ends when all cards have been played. Each player adds up their card values, and whoever has the higher total wins.

For example, suppose the initial cards are $[3, -1, 4, 1, 5, 9]$, and Elmo plays first. Then the game might proceed as follows:

- Elmo keeps the 3, ending his turn.
- You give Elmo the -1 .
- You keep the 4, ending your turn.
- Elmo gives you the 1.
- Elmo gives you the 5.
- Elmo keeps the 9, ending his turn. All cards are gone, so the game is over.
- Your score is $1 + 4 + 5 = 10$ and Elmo's score is $3 - 1 + 9 = 11$, so Elmo wins.

Describe an algorithm to compute the highest possible score you can earn from a given row of cards, assuming Elmo plays first and plays perfectly. Your input is the array $C[1..n]$ of card values. For example, if the input is $[3, -1, 4, 1, 5, 9]$, your algorithm should return the integer 10.

14. **⟨⟨F14⟩⟩** The new swap-puzzle game *Candy Swap Saga XIII* involves n cute animals numbered 1 through n . Each animal holds one of three types of candy: circus peanuts, Heath bars, and Ciocolateria Gardini chocolate truffles. You also have a candy in your hand; at the start of the game, you have a circus peanut.

To earn points, you visit each of the animals in order from 1 to n . For each animal, you can either keep the candy in your hand or exchange it with the candy the animal is holding.

- If you swap your candy for another candy of the *same* type, you earn one point.
- If you swap your candy for a candy of a *different* type, you lose one point. (Yes, your score can be negative.)
- If you visit an animal and decide not to swap candy, your score does not change.

You *must* visit the animals in order, and once you visit an animal, you can never visit it again.

Describe and analyze an efficient algorithm to compute your maximum possible score. Your input is an array $C[1..n]$, where $C[i]$ is the type of candy that the i th animal is holding.

15. **⟨⟨F14⟩⟩** Farmers Boggis, Bunce, and Bean have set up an obstacle course for Mr. Fox. The course consists of a row of n booths, each with an integer painted on the front with bright red paint, which could be positive, negative, or zero. Let $A[i]$ denote the number painted on the front of the i th booth. Everyone has agreed to the following rules:
- At each booth, Mr. Fox *must* say either “Ring!” or “Ding!”.
 - If Mr. Fox says “Ring!” at the i th booth, he earns a reward of $A[i]$ chickens. (If $A[i] < 0$, Mr. Fox pays a penalty of $-A[i]$ chickens.)
 - If Mr. Fox says “Ding!” at the i th booth, he pays a penalty of $A[i]$ chickens. (If $A[i] < 0$, Mr. Fox earns a reward of $-A[i]$ chickens.)
 - Mr. Fox is forbidden to say the same word more than three times in a row. For example, if he says “Ring!” at booths 6, 7, and 8, then he *must* say “Ding!” at booth 9.
 - All accounts will be settled at the end; Mr. Fox does not actually have to carry chickens through the obstacle course.
 - If Mr. Fox violates any of the rules, or if he ends the obstacle course owing the farmers chickens, the farmers will shoot him.

Describe and analyze an algorithm to compute the largest number of chickens that Mr. Fox can earn by running the obstacle course, given the array $A[1..n]$ of booth numbers as input.