Note: Your TA probably will not cover all the problems. This is totally fine, the discussion worksheets are not designed to be finished in an hour. They are deliberately made long so they can serve as a resource you can use to practice, reinforce, and build upon concepts discussed in lecture, readings, and the homework.

1 Dynamic Programming Introduction: Fibonacci Numbers

The Fibonacci sequence is defined by the following recurrence relation:

$$F_n = F_{n-1} + F_{n-2},$$

with base cases $F_0 = 0$ and $F_1 = 1$. Back in CS 61A, we learned how to write a program to find the nth fibonacci number, which would look something like:

```python
def fibo(n):
    if n <= 1:
        return n
    return fibo(n-1) + fibo(n-2)
```

However, this program is actually super slow! In the box below, show that calling `fibo(n)` takes $2^{\Theta(n)}$ time.

Challenge: show that the runtime is $\Theta\left(\left(\frac{1+\sqrt{5}}{2}\right)^n\right)$.

If you didn't above, in the box below draw out the recurrence tree produced when calling `fibo(5)`. Do you notice any repeated computations (i.e. nodes)?

In the recurrence tree, we notice that we end up recomputing some of same values multiple. For instance, we end up computing F_1 5 times! To reduce the number of recomputing we have to do, we can **store each fibonacci number in an array after computing it**. This way, we can simply index into that array when we need that value, rather than recomputing it every time we recurse. To implement this, fill out the blank lines in the code below:

```python
def optimized_fibo(n):
    stored_fibos = [-1 for _ in range(n+1)]

    def fibo(n):
        # base case

        if _____:

            return _____

        # if we've already computed fibo(n) before, we can reuse it via
            stored_fibos!

        if _____:

            return _____

        # if we haven't already computed fibo(n), then we need to recurse as before:
        # make sure to store it in stored_fibos so that we can use it in the future!

        _____

        return _____
```

What is the runtime of this new algorithm?

Congratulations, you've just implemented your first **dynamic programming (DP)** algorithm! This is essentially all that DP is: recursion plus storing stuff (memoization), so that we don't have to fully solve any subproblems more than once.

Now, there are actually two ways to implement DP algorithms. The implementation that you've completed above uses a **top-down** approach, i.e. you start from the largest subproblem (top) and repeatedly recurse on smaller subproblems (going down). The other implementation method uses a **bottom-up** approach, which starts from the smallest subproblems (i.e. the base cases), and builds up larger subproblems in an iterative manner.

Referencing your previous top-down approach, fill in the blank lines in the code below to complete the bottom-up implementation of the Fibonacci DP algorithm:

```python
def fibo_dp_bottom_up(n):
    stored_fibos = [-1 for _ in range(n+1)]

    # define base cases at the "bottom"
    stored_fibos[0] = 0
    stored_fibos[1] = 1

    # build your subproblems "up" from your smaller subproblems
    for i in range(2, n+1):

        _____

    # what element in stored_fibos represents the nth fibonacci number?

    answer = _____
    return answer
```

Yay! You've now learned how to implement both types of DP algorithms.

(Challenge) What is the space complexity of your algorithm? Can you modify it to only use O(1) extra space?

```python
def fibo_dp_bottom_up(n):
    # define base cases here

    if _____:

        _____

    _____

    _____

    # build your subproblems "up" from your smaller subproblems

    for i in range(2, n+1):

        _____

    # where are you storing the nth fibonacci number?

    answer = _____
    return answer
```

2   Planting Trees

This problem will guide you through the process of writing a dynamic programming algorithm.

You have a garden and want to plant some apple trees in your garden, so that they produce as many apples as possible. There are n adjacent spots numbered 1 to n in your garden where you can place a tree. Based on the quality of the soil in each spot, you know that if you plant a tree in the ith spot, it will produce exactly x_i apples. However, each tree needs space to grow, so if you place a tree in the ith spot, you can't place a tree in spots $i - 1$ or $i + 1$. What is the maximum number of apples you can produce in your garden?

(a) Give an example of an input for which:

- Starting from either the first or second spot and then picking every other spot (e.g. either planting the trees in spots $1, 3, 5 \ldots$ or in spots $2, 4, 6 \ldots$) does not produce an optimal solution.

- The following algorithm does not produce an optimal solution: While it is possible to plant another tree, plant a tree in the spot where we are allowed to plant a tree with the largest x_i value.

(b) To solve this problem, we'll think about solving the following, more general problem: "What is the maximum number of apples that can be produced using only spots 1 to i?". Let $f(i)$ denote the answer to this question for any i. Define $f(0) = 0$, as when we have no spots, we can't plant any trees. What is $f(1)$? What is $f(2)$?

(c) Suppose you know that the best way to plant trees using only spots 1 to i does not place a tree in spot i. In this case, express $f(i)$ in terms of x_i and $f(j)$ for $j < i$.

Hint: What spots are we left with? What is the best way to plant trees in these spots?

(d) Suppose you know that the best way to plant trees using only spots 1 to i places a tree in spot i. In this case, express $f(i)$ in terms of x_i and $f(j)$ for $j < i$.

(e) Describe a linear-time algorithm to compute the maximum number of apples you can produce.

Hint: Compute $f(i)$ for every i. You should be able to combine your results from the previous two parts to perform each computation in $O(1)$ time.

3 Longest Common Subsequence

In lecture, we covered the longest increasing subsequence problem (LIS). Now, let us consider the longest common subsequence problem (LCS), which is a bit more involved. Given two arrays A and B of integers, you want to determine the length of their longest common subsequence. If they do not share any common elements, return 0.

For example, given $A = [1, 2, 3, 4, 5]$ and $B = [1, 3, 5, 7]$, their longest common subsequence is $[1, 3, 5]$ with length 3.

We will design an algorithm that solves this problem in $O(nm)$ time, where n is the length of A and m is the length of B.

(a) Define your subproblem.

Hint: looking at the Edit Distance subproblem may be helpful.

(b) Write your recurrence relation.

(c) In what order do we solve the subproblems?

(d) What is the runtime of this dynamic programming algorithm?

(e) What is the space complexity of your DP algorithm? Is it possible to optimize it?

4 Change making

You are given an unlimited supply of coins of denominations $v_1, \ldots, v_n \in N$ and a value $W \in N$. Your goal is to make change for W using the minimum number of coins, that is, find a smallest set of coins whose total value is W.

(a) Design a dynamic programming algorithm for solving the change making problem. What is its running time?

(b) You now have the additional constraint that there is only one coin per denomination. Does your previous algorithm still work? If not, design a new one.