

Problem Set 5

This problem set is all about induction and its sheer breadth of applications. By the time you're done with this problem set, you will have a much deeper understanding of how to think inductively. Plus, you'll learn a bit about data structures, recurrence relations, puzzles, games, and tournament structures!

We recommend that you read the Handout #24, "Guide to Induction," before starting this problem set. It contains a lot of useful advice about how to approach problems inductively, how to structure inductive proofs, and how to not fall into common inductive traps.

As a note on this problem set – normally, you're welcome to use any proof technique you'd like to prove results in this course. On this problem set, we've specifically requested on some problems that you prove a result inductively. For those problems, you should prove those results using induction or complete induction, even if there is another way to prove the result. (If you'd like to use induction in conjunction with other techniques like proof by contradiction or proof by contrapositive, that's perfectly fine.)

As always, please feel free to drop by office hours, visit Piazza, or send us emails if you have any questions. We'd be happy to help out.

Good luck, and have fun!

**Due Friday, February 12 at the start of class.
There is no checkpoint problem.**

Problem One: Recurrence Relations (4 Points)

A *recurrence relation* is a recursive definition of the terms in a sequence. Typically, a recurrence relation specifies the value of the first few terms in a sequence, then defines the remaining terms from the previous terms. For example, the *Fibonacci sequence* can be defined by the following recurrence relation:

$$\begin{aligned} F_0 &= 0 \\ F_1 &= 1 \\ F_{n+2} &= F_n + F_{n+1} \end{aligned}$$

The first terms of this sequence are $F_0 = 0$, $F_1 = 1$, $F_2 = 1$, $F_3 = 2$, $F_4 = 3$, $F_5 = 5$, $F_6 = 8$, etc.

Some recurrence relations define well-known sequences. For example, consider the following recurrence relation:

$$\begin{aligned} a_0 &= 1 \\ a_{n+1} &= 2a_n \end{aligned}$$

The first few terms of this sequence are 1, 2, 4, 8, 16, 32, ..., which happen to be powers of two. It turns out that this isn't a coincidence – this recurrence relation perfectly describes the powers of two.

- i. Prove by induction that for any $n \in \mathbb{N}$, we have $a_n = 2^n$.

Minor changes to the recursive step in a recurrence relation can lead to enormous changes in what numbers are generated. Consider the following two recurrence relations, which are similar to the a_n sequence defined above but with slight changes to the recursive step:

$$\begin{array}{ll} b_0 = 1 & c_0 = 1 \\ b_{n+1} = 2b_n - 1 & c_{n+1} = 2c_n + 1 \end{array}$$

- ii. Find non-recursive definitions for b_n and c_n , then prove by induction that your definitions are correct.

Finding non-recursive definitions for recurrences (often called “solving” the recurrence) is useful in the design and analysis of algorithms. Commonly, when trying to analyze the runtime of an algorithm, you will arrive at a recurrence relation describing the runtime on an input of size n in terms of the runtime on inputs of smaller sizes. Solving the recurrence then lets you precisely determine the runtime. To learn more, take CS161, Math 108, or consider reading through *Concrete Mathematics* by Graham, Knuth, and Patashnik.

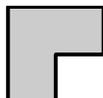
Problem Two: A Lot of Rocks? (3 Points)

Let's say that *a lot of rocks* is a number of rocks that isn't zero (surely, zero rocks is not a lot of rocks) and where if you remove one of the rocks, you're left with a lot of rocks (if there really are a lot of rocks, removing one rock shouldn't make a difference).

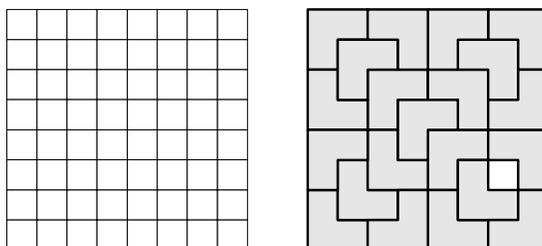
Prove by induction that no natural number of rocks is a lot of rocks. This hopefully teaches you to be careful with new definitions – just because we can define something doesn't mean it actually exists!

Problem Three: Tiling with Triominoes (4 Points)*

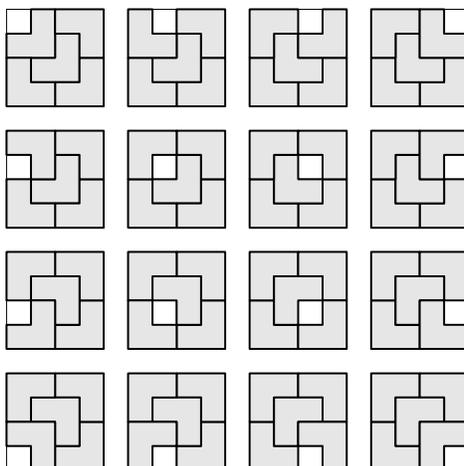
Recall from Problem Set One that a *right triomino* is an L-shaped tile that looks like this:



Suppose that you are also given a square grid of size $2^n \times 2^n$ and want to *tile* it with right triominoes by covering the grid with triominoes such that all triominoes are completely on the grid and no triominoes overlap. Here's an attempt to cover an 8×8 grid with triominoes, which fails because not all squares in the grid are covered:



Amazingly, it turns out that it is always possible to tile any $2^n \times 2^n$ grid that's missing exactly one square with right triominoes. It doesn't matter what n is or which square is removed; there is always a solution to the problem. For example, here are all the ways to tile a 4×4 grid that has a square missing:

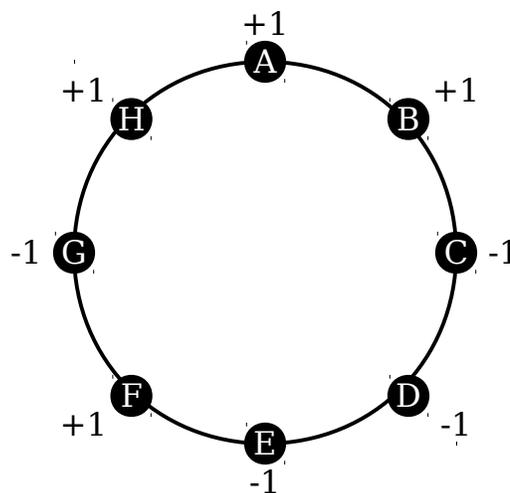


Prove by induction that any $2^n \times 2^n$ grid with any one square removed can be tiled by right triominoes.

* I originally heard this problem from David Gries of Cornell University.

Problem Four: The Circle Game (4 Points)

Here's a game you can play. Suppose that you have a circle with $2n$ arbitrarily-chosen points on its circumference. n of these points are labeled $+1$, and the remaining n are labeled -1 . One sample circle with eight points, of which four are labeled $+1$ and four are labeled -1 , is shown to the right.



Here's the rules of the game. First, choose one of the $2n$ points as your starting point. Then, start moving clockwise around the circle. As you go, you'll pass through some number of $+1$ points and some number of -1 points. You lose the game if at any point on your journey you pass through more -1 points than $+1$ points. You win the game if you get all the way back around to your starting point without losing.

For example, if you started at point A, the game would go like this:

Start at A: $+1$.
 Pass through B: $+2$.
 Pass through C: $+1$.
 Pass through D: 0 .
 Pass through E: -1 . (*You lose.*)

If you started at point G, the game would go like this:

Start at G: -1 (*You lose.*)

However, if you started at point F, the game would go like this:

Start at F: $+1$.
 Pass through G: 0 .
 Pass through H: $+1$.
 Pass through A: $+2$.
 Pass through B: $+3$.
 Pass through C: $+2$.
 Pass through D: $+1$.
 Pass through E: $+0$.
 Return to F. (*You win!*)

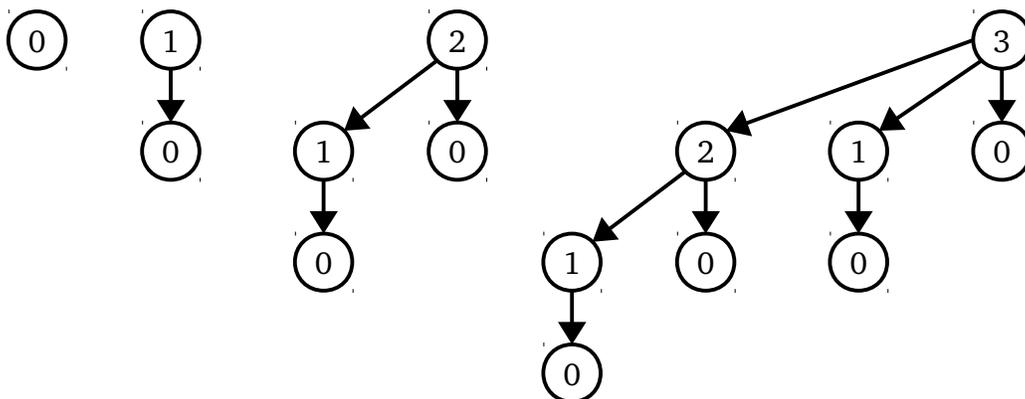
Interestingly, it turns out that no matter which n points are labeled $+1$ and which n points are labeled -1 , there is always at least one point you can start at to win the game.

Prove, by induction, that the above fact is true for any $n \geq 1$.

Problem Five: Binomial Trees (3 Points)

In computer science, a *directed tree* is a structure used to represent hierarchical information. A directed tree consists of a set of *nodes*. Each node can have zero or more *child nodes*. Every node has exactly one parent except for a special node called the *root node*, which has no parent.

Binomial trees are a specific family of directed trees defined as follows: a binomial tree of order n is a single node with n children, which are binomial trees of order $0, 1, 2, \dots, n - 1$. For example, here are pictures of binomial trees of orders 0, 1, 2, and 3:



Prove by induction that a binomial tree of order n has exactly 2^n nodes. Your proof should use to the formal definition of binomial trees given above – if you want to use any other facts about binomial trees in your proof of the main result, you will need to prove those other facts first.

Binomial trees are used as a building block in the *binomial heap* data structure, which can be used to efficiently determine the smallest element out of a group of values. Binomial heaps form the basis for a huge number of different data structures, such as the soft heap (used in a fast algorithm for the minimum spanning tree problem), the Fibonacci heap (used to implement fast versions of Dijkstra's shortest paths algorithm), and several others. If you're interested in learning more about the binomial heap and its applications, consider picking up a copy of *Introduction to Algorithms, Second Edition* by Cormen, Leiserson, Rivest, and Stein or taking CS166.

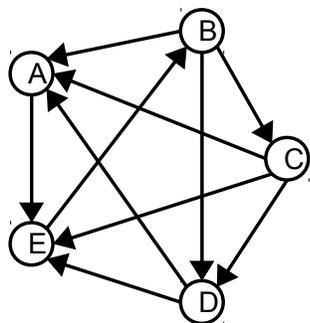
Problem Six: Nim (4 Points)

Nim is a family of games played by two players. The game begins with several piles of stones, each of which has zero or more stones in it, that are shared between the two players. Players alternate taking turns removing any nonzero number of stones from any single pile of their choice. If at the start of a player's turn all the piles are empty, then that player loses the game.

Prove, by induction, that if the game is played with just two piles of stones, each of which begins with exactly the same number of stones, then the second player can always win the game if she plays correctly.

Before trying to write up your answer to this question, we recommend playing this game with a partner until you can find a winning strategy.

Problem Seven: Tournament Graphs (5 Points)



In Problem Set Two and Problem Set Five, you explored *tournaments*, contests among $n \geq 0$ players where each player plays exactly one game against each other player and no game ends in a tie.

From a graph-theoretic perspective, we can treat tournaments as graphs. Each player is a node, and there's exactly one edge between each pair of nodes. The edge is directed such that the edge goes from the winner of the game to the loser of the game. For example, in the graph to the left, player D beat players A and E, but lost to players B and C. These sorts of graphs are called *tournament graphs*.

Let's introduce a new definition. If T is a tournament graph, a *subtournament* of T is a tournament graph T' formed by picking a subset of the players in T and considering just the games those players played against each other. (Note that, just like any set is a subset of itself, any tournament is a subtournament of itself).

Now, let's introduce one final definition. A tournament graph T is called *transitive* if it contains no cycles. Prove by induction that for any natural number n , if T is a tournament graph with exactly 2^n nodes, then T contains a transitive subtournament T^* with at least $n + 1$ nodes. As a hint, you may want to use the generalized pigeonhole principle in your inductive step.

Extra Credit Problem: Pie Fights! (1 Point Extra Credit)

Here's how a pie fight works: $n \geq 2$ people, each of whom has a pie in their hands, stands in a big open field. Everyone then throws their pie at the person closest to them. (To make sure that "the person closest to them" is well-defined, let's assume that no two pairs of people are standing at the same distance from one another). In the end, lots of people will end up covered in pie.

Here's a surprising result: in any pie fight with an odd number of people, at least one person won't have a pie thrown at them. Prove this result by induction. (As a note: it is *not* the case that if person A throws a pie at person B, then person B will throw a pie at person A. Make sure you see why this is!)