Lecture 1 Monday 1/24

A set is a collection of objects. We can describe one in two ways: either by listing the elements, such as $\{A, B, C, A\}$, or by saying which elements we want from a bigger set, such as "even numbers". We will always have a "universal set" in mind, which will contain all the objects we are talking about, e.g. all letters, or all numbers.

The *union* of two sets A and B contains all the elements of both, and is denoted $A \cup B$. The *intersection* contains only those elements that are present in both sets, and is denoted $A \cap B$. The *complement* of a set A, written $\sim A$ contains all the objects that A doesn't.

A set A is a subset of a set B if all the elements of A also happen to be elements of B. In this case, we write $A \subset B$. For example, $A \cap B \subset A$ for any sets A and B. Notice that the set $A \cap \sim A$ has no elements. It is called the *empty set* and denoted by $\emptyset$. For any set A, we have $\emptyset \subset A$.

Counting the number of elements in a set can be tough. For example, the set $\{A, B, C, A\}$ contains three objects. If a set is specified using a description, this may be even harder. For example, consider the set BBQ of US states where there are at least 10 barbeque restaurants. This set has somewhere between 0 and 50 elements, but it's hard to tell how many.

If we know how many elements a set has, we can figure out the number of elements in some other sets, or at least estimate it. For example, $A \cup B$ has no more than $\#A + \#B$ elements, and at least as many as the smaller of A and B. The general formula is $\#(A \cup B) = \#A + \#B - \#(A \cap B)$.

If the universal set has a finite number of elements, then we can figure out how many elements there are in $\sim A$. For example, if we know there are 13 spades in a deck of cards, and we know that a deck of cards has 52 elements, then the set of cards that are *not* spades has $52 - 13 = 39$ objects in it.

It's also useful to consider how many combinations one can pick from certain sets. For example, if a restaurant serves three types of burgers and two types of drinks, then there are 6 types of combinations one can select. We can write this out, or use the formula $\#(A \times B) = \#A \cdot \#B$. Here $A \times B$ consists of pairs: one from A and another from B. We concluded in class that if one wants to choose a card and a domino, one has $52 \cdot 7 \cdot 7$ options, since there are 49 dominos.

The *power set* $\mathcal{P}(A)$ is the set of all subsets of A. We guessed at the end of class that $\#\mathcal{P}(A) = 2^{(\#A)}$ for any set.

Lecture 2 Wednesday 1/27

We talked about a possibility for a semester project; namely, building and documenting an L-shaped billiard table. We then talked about using Venn diagrams to display data.

Lecture 3 Friday 1/31

Quiz 1. Then review of fractions, percentages, parts per million, etc.

Lecture 4 Monday 2/3

Discussion of probability. What are the chances of picking a certain number of balls out of a bag? Repeated random trials, with putting balls back or not.

Lecture 5 Wednesday 2/5

Testing theory: what is a test and what does it mean for it to be correct?

We first discussed various examples of tests through history, both bad (witch hunts, red scare) and good (gifted student testing, security, medical). We decided that a good test should have a low false-positive rate and a low false-negative rate.

We then focused on a specific example: testing for vampires using garlic. We estimated that the test would have a decently high false-positive rate (normal people that dislike garlic) and might have a false-negative rate (some unusual vampires that are immune to garlic). Using specific false-positive and false-negative rates, we worked through an example of a city with vampires and computed how much of the population would be diagnosed as vampires using the garlic test.

We noticed two things:

1. Restricting the population tested reduces false-positive cases (e.g. if all vampires are known to be tall).
2. After several years, the vampire problem might return (garlic-resistant vampires infecting new individuals). The garlic test will become completely ineffective.

Lecture 6 Friday 2/7

Quiz 2. "She's a witch" clip. A discussion of estimating outcomes of random events, focusing on expected value.

We first discussed a variety of cases where unexpected outcomes have to be analyzed (choice of major, buying a car, deciding to cheat on an assignment). We focused on the cheating example to figure out whether it was worth it to cheat: the chances of being caught are multiplied by the punishment, while the chances of not being caught are multiplied by the benefit. This led to the idea of *expected value*: the benefit of each outcome is multiplied by its probability, and the sum is the expected value.

We switched to a gambling example. Here, the expected value has a clear interpretation: if one gambles 1000 times with the same odds, one can expect to have 1000 times the expected value in winnings.

Lecture 7 Monday 2/10

Probability games. The class broke up into teams of 2-3 people and was provided with dice. Each person played 27 rounds of one of the following games (player's choice):

Game A: Roll a dye. Gain the number of points displayed.

Game B: Roll a dye. If it is even, gain 10 points. Otherwise, lose 3 points.

Game C: Roll a dye. If it is even, gain 3 point. Otherwise, gain 4 points.

Each student kept track of the rolls, points gained on each turn, and the running score. Each group graphed its running score and reported the final scores to the class.

We then added a line of slope 3.5 to the graphs, which allowed us to compare the actual outcomes of the games with the expected value of the games (as it turned out, all three games had the same expected value of 3.5). Some of the graphs matched very well, while others didn't.

To figure out how likely a game was to be far from the expected value, we used a new tool: the *variance*. One might expect the variance to ask "how far is the outcome of the game likely to be from the average"? Instead, it's a little more complicated: for each possible outcome, we compute the difference from the expected value (3.5 in our case), then square this difference. We then add them up as before, multiplying by the probability of the outcome.

So for game B we get that the variance is $50\% \times (10-3.5)^2 + 50\% \times (-3-3.5)^2 = 21.125 + 21.125 = 42.25$, which is huge compared to the variance of the other two games (calculate the variances of those to compare). So game B is very unpredictable compared to the others.

Lecture 8 Wednesday 2/12

We discussed ways in which data is presented: various charts (bar graphs, histograms, pie graphs, line graphs), and basic statistics (minimum, maximum, mean, median).

Lecuture 9 Friday 2/14

Quiz 3. We then discussed the notion of data collection: when one requires studies, and how one generally goes about setting them up. We focused on the following test case: suppose you decide students are doing badly in Chicago schools and would like to make them do better. We brainstormed various reasons why this could be happening and discussed ways to test the hypotheses. These came down to three methods: experimental (trying to control various factors in students' experience), interactive (surveying the students to see what they think), and observational (seeing what correlates with good performance).

We discussed drawbacks of each approach. Experimental studies in certain contexts are unethical (e.g. putting good students in bad schools). Intractive studies presuppose that people actually know what is happening around them (which in certain cases such as child development is false). Observational studies use correlation as a primary tool, and correlation does not imply causation.

Lecture 10 Monday 2/17

We discussed experimental design, starting with a webcomic showing off bad experimental design, watching videos of the marshmallow test, and then focusing on the question on child language development. Specifically, we discussed the question of how students begin to understand the connection between verbs and nouns, as far as singulars vs plurals go. We brainstormed some ideas for analysing such a relationship. We then went through some experiments in child language acquisition, talking about how to design a good experiment, and seeing how one goes from data to final results (via data coding and statistics).

Lecture 11 Wednesday 2/19

With the test coming up on Friday, we reviewed the content of the previous lectures.

Lecture 12 Monday 2/24

Idea of storing data. Alphabets vs pictographs, basic "computers" such as player pianos and music boxes. Encoding data on bits. How high can you count with x fingers, or x bits? Base 2 numbers.

Lecture 13 Wednesday 2/26

More base-2 numbers, followed by base 16 and 8.

Lecture 14 Friday 2/28

Storing data as actual bits. Convering text to ASCII codes, and further into shaded boxes – and back.

Lecture 15 Monday 3/3

Checksums: basic, $3-1-3-1$, etc.

Lecture 16 Wednesday 3/5

Review of some previous topics. Notion of a prime, looking for them using the Sieve of Aristothenes. Factoring composite numbers.

Lecture 17 Friday 3/7

Quiz 5, followed by some awesome activities and discussion of cybersecurity.

Lecture 18 Monday 3/10

Discussed requirements for step one of the semester project. Discussed modular arithmetic: the fact that it works very nicely mod a prime, and not as well mod a composite number. Quadratic formula still works mod a prime, but not every number has a square root, so there aren't always solutions (just as there aren't in usual arithmetic).

Lecture 19 Wednesday 3/12

Introduction to logic and some discussion of its use in computer programming. Boolean statements. And, or, parentheses. Truth tables. Translating things into pure logic.

Lecture 20 Friday 3/14

Quiz 6, followed by more logic: turning different phrases into logical statements (more or/xor, and the many versions of "and"). We then went through the process for negating a sentence: breaking it up into the basic logical pieces, then using the negation rules to reverse and simplify, then writing it in English again. We wrote down the rules for "not (A and B)", "not (A or B')", and "not (A xor B)".

Lecture 21 Monday 3/17

Guest lecture by Noel DeJarnette.

Quick review of the last two days of logic for warmup (for example: negate A="Bob ate either a hotdog or a hamburger, but not fries" by writing it as "(B or C) and (not D)", negating it and rewriting in English, perhaps with truth tables for both A and not A).

Main topic: $A \Rightarrow B$, $B \Leftarrow A$, $A \Leftrightarrow B$, with corresponding truth tables and examples. Start off with a basic example (e.g. "If your car breaks, then you have to stay at home.") with the corresponding truth table. Somewhat strangely, $A \Rightarrow B$ is true if A is false; similarly to "innocent until proven guilty". Do a couple of examples of writing sentences as implications, with different phrasings becoming $A \Rightarrow B$. Then introduce $A \Leftrightarrow B$ and provide an example or two to differentiate it from $A \Rightarrow B$.

Finish off by talking about negating implication (mention counter-examples) and concluding things based on known implications (modus ponens).

Lecture 22 Wednesday 3/19

Review for Friday's test.

Lectures 23 and 24 Monday 3/31, Wednesday 4/2

Introduction to non-Euclidean geometries. We started by inflating a balloon and thinking about familiar notions on it: straight lines (we switched to the term "geodesic"), triangles, circles. We noticed that these don't behave the same way they did in Euclidean geometry.

We went on to more exotic spaces, all of which we drew in the plane. We had the taxicab metric, for which one has to always travel horizontally or vertically; the Paris metric in which one can only travel along the straight lines coming out of the origin; and the hallway metrics in which one could move around as usual, but had to stay inside some "hallway" region.

In each of these spaces, we explored distances, geodesics, triangles, circles, and disks.

Lecture 25 Friday 4/4

On the weekly quiz, the new Polish Comb geometry appeared, in which one can only move along the vertical lines in the plane or along the x-axis.

After the quiz, we began talking about graphs, which are geometric spaces made up of points (vertices) connected by line segments (edges). We drew some graphs, counting edges and vertices.

Lecture 26 Monday 4/7

We reviewed what a graph is, and talked about some ways the notion is useful in the world (road networks, internet, etc.).

We asked whether, given a graph, it is possible to visit all the edges exactly once, while coming back to the initial starting point. A trip like this is called an Euler cycle.

After playing around with several examples, we decided that the existence of an Euler cycle is based on how many edges each vertex sees (the *valence* of the vertex). Indeed, Euler's Theorem states that an Euler cycle exists in a graph only when each vertex has even valence — and that if the valence of each vertex is even, then a graph is guaranteed to exist.

We also briefly discussed the related Traveling Salesman Problem, in which one wants to visit every vertex in the graph and get back home, with as short a trip as possible. Since the Traveling Salesman Problem is officially super-super-hard, we didn't try to solve it.

Lecture 27 Wednesday 4/9

Often, the edges of a graph don't have equal length. For example, if the graph represents a rail network, the vertices would be train stations and the edges would be the rail lines connecting the stations — with some longer than others. Our goal during this lecture was to compute distances in a graph with labeled edges.

The main algorithm for this is Dijkstra's Algorithm, which works as follows:

- Suppose you want to compute the distance from a vertex Start to the vertex Finish.
- We know that the distance from Start to itself is zero, so we label that vertex with a 0.
- We put a check mark next to Start to indicate that the zero next to it is the actual distance from Start to Start.

- Suppose then we have some vertex that is connected to Start with an edge labeled 2. We label that vertex as 2.
- After we are done labeling all the vertices next to Start,
- We now find the vertex with the smallest number next to it (but without a check mark). Let's call it NowVertex. We put a check mark next to NowVertex, since the distance fro Start to that vertex has been computed.
- Suppose NowVertex has a neighbor vertex Neighbor, connected by an edge of length 3. We update Neighbor's label with 3+(NowVertex's label), unless that's more than what's already written there.
- Once we've updated all of NowVertex's neighbors, we are done with it; and choose the next NowVertex by picking the vertex with the smallest number next to it.
- Once Finish has a check mark next to it, the number next to it is the distance from Start to Finish.

The algorithm is a little complicated, but is very similar to how you would try to compute distances by yourself — it just provides a systematic way of making your intuition work out.

Lecture 28 Friday 4/11

After the quiz, we shifted gears again. Noel DeJarnette brought in paper cutouts that can be folded up into Platonic solids and discussed the topology of these assembled shapes. In particular, we noticed that $V - E + F = 2$ for each of those shapes. We called the number $V - E + F$ the *Euler characteristic* of the shape, and noticed that doing weird things like gluing two shapes together along a vertex changes the Euler characteristic.

Lecture 29 Monday 4/14

We talked about the notion of computational topology, where instead of finding familiar things like distances and angles, one is interested in things like "is this point inside our outside the loop I just drew". We talked about this specific problem, learning that squiggles are hard to work with, but a trick is available for figuring out if a point is inside or outside the squiggle: draw a line segment all the way out of the squiggle; if the number of intersections with the squiggle is even, then it's outside; otherwise it's inside.

We then continued the discussion of folded-up shapes, but this time without folding anything up. Instead, we drew shapes on the chalkboard, and used arrows (or colors) to indicate how we *would* fold them up and which edges or points would end up being the same in the folded-up shape. We then used Euler characteristic to figure out what the shapes actually looked like. Specifically, the formula $2 - 2g = V - E + F$ allowed us to figure out the *genus* g of the assembled surface — that is, the number of "holes" it has. For example, the shapes from the previous lecture are deformed spheres, while a shape with genus $g = 1$ would like like the surface of a doughnut (maybe also quite deformed).

Lecture 30 Wednesday 4/16

We then reviewed the previous two lectures, and kept working with folded up shapes. We figured out how tic-tac-toe and chess would work on a folded-up game board, and also talked about how one would compute distances in a folded-up shape. In particular, as long as one is gluing together parallel lines, the geodesics in the resulting space look like a bunch of line segments.

Lecture 31 Friday 4/18

After the quiz, we talked about fractals. We learned to draw the Von Koch snowflake and the Sierpinski gasket and carpet. We talked a bit about what fractals are good for, and about the idea of fractal dimension.

Lecture 32 Monday 4/21

We started with an example of a random Von Koch snowflake, where spikes were put in randomly rather than systematically. This gave us a pretty weird fractal, but is representative of how many fractal-creating processes actually work.

We then focused on ideas from Vyron Vellis's thesis. We asked: suppose you are in a fenced yard surrounded by zombies and want to stay two feet away from the fence, where in the yard can you actually walk? We drew pictures to figure this out.

Lecture 33 Wednesday 4/22

We discussed the idea of finding efficient travel paths in a space where travel difficulty varied from point to point. For example, this could be due to speed limits, traffic, or just different terrain (like sidewalks vs. grass on the quad). We talked about how calculus is the standard tool for figuring out such things, but in order to be able to do some problems switched over to a simplified version of the problem. Namely, we drew a big region tiled by smaller squares that were lighter or darker, indicating travel difficulty. We then computed lengths of various trajectories.

We then noticed that a trajectory could be easy to travel on, but have lots of turns, which is inefficient. So we included a cost to turning, and again figured out which paths were more or less efficient.

Since we were now looking not only position but also angles, we were now thinking of not just points in the plane, but points "facing in some direction". So we had to keep track of three coordinates: x, y, and the angle. What's more, we couldn't just move in any of these three directions, since walking sideways is pretty weird: to get around we have to either walk straight forward or change the direction we are facing.

We then talked about how this sort of geometry where you have very few things you *can* control but want to control lots of things, is pretty standard. For example, parallel parking requires you to control the angle and position of both the front and back wheels, but all you can control directly is the speed and angle of the front wheels.

The geometry that studies such setups is called *sub-Riemannian* and is the topic of my own thesis.

Lecture 34 Friday 4/24

Last quiz, followed by review for next week's test.

Lecture 35 Monday 4/27

Watching everyone's semester projects.

Lecture 36 Wednesday 4/29

Review for Friday's test.

Lecture 37 Monday 5/5

Review for Friday's final.

Lecture 38 Wednesday 5/7

Review for Friday's final.