CMPS 10 Lecture Notes: Lecture 8 (1-28-2016)

A Momentous Day in Computer Science: Featuring AlphaGo, Alan Turing, and Morphogenesis

Last Time

- We were talking about a newspaper article.
  - It said that there is one gene that seems to be connected to a neurotransmitter in your head that is associated with happiness, and that gene is more common in a certain demographic
- How do you measure happiness?
  - Well, humans are pretty good at self reporting. You are asked "on a scale of 1 to 10 how happy are you" and you answer
    - And enough studies have been using this method that people are convinced that this works out just fine.
- Why did teacher ask us these things?
  - Wants to see how hard he can push us.
  - And there is a connection to music
    - The idea that even though music has all of these things going on, but we can still manage to break it down to our Left and Right game.
    - We don't need to know it at every moment in time, we only need to know it at a sufficiently dense sampling of intervals.
      - And we use numbers to describe the path that leads to that interval. (i.e. there is a tree).
    - It's crazy to think that there is a guy just lifting his arms, you can completely replay the symphony on your own without any loss of information.
      - **And there is evidence that complex living organisms can be reduced down our similar language of 1s and 0s.**

How many people have heard of Alan Turing? And seen the imitation game?

- Alan Turing is the influential, intellectual hero of computer science.
  - We'll talk about one of his major contributions today.

In Biology there is an idea of **morphogenesis**

- Morpho means form, genesis means birth
  - So birth of form.
  - Wikipedia: (Morphogenesis) is the biological process that causes an organism to develop its shape. It is one of three fundamental aspects of developmental biology along with the control of cell growth and cellular differentiation.

Drawing of sperm and egg

- Genetic Material in Sperm, Genetic material in Egg
  - Once sperm goes inside egg, barrier around egg becomes impermeable.
    - And so now there are two pieces of genetic material there, and they combine to make a new bit of genetic material.
    - So, imagine a sequence of 1s and 0s from Sperm, and a different sequence of 1s and 0s from the Egg. And the two sequences are the same length.
      - And so you compare the two side by side. If one is 0 and one is 1, then there is a 50 50 chance that the new genetic code will be either 0 or 1. OR if they are both 1 then you get a 1, if both are 0 you get 0.

      - And then there is DNA, and the DNA have A, C, G, T.
      - However, the above is a little bit a cartoon version, because there is actually "coin flipping" happening. It's all happening in a big soup of water, and so this process of combining these strings of genetic code is happening in an environment that is shaking things, so in some sense it is a race.
  - It is also interesting to note that most humans share 99 percent the same genetic material.
    - Humans can be thought of as having the same subsections that are the same for all humans.
    - And then there are small subsections that are polymorphism and that explains all of the difference between us.
      - We are composed of parts that have come together, for some of those parts there are different versions, for other parts thats it.
    - Imagine that our parents are cars
      - We get windows from dad, tires from mom, etc.
      - BUT we definitely GET windows, and we GET tires. We aren't breaking it down into infinitely small "mass."
      - That is, pretty much all humans have two arms, two legs, a head, a brain, a liver, a heart, etc. And even those parts are mostly the same between all humans (our arms all have elbows, our legs all have knees, etc.). We are composed of relatively large building blocks that are identical for all of us; it is within those building blocks that there is variation.

Clicker question: I started life as a single cell: A for NO, B for Yes

- A 20 percent
- B 80 percent

Eventually we reach a point where we are several trillion cells.

- JOKE: Sometimes more cells than we would like! (I'm not sure if anyone laughed at this, but it was very funny).
- And it is clear that the cells that make up your teeth are different than the cells that make up your skin.
- But what does it mean to say that they are different?
- How do you go from one cell to many?
  - Cell division.
  - What is input and output of cell division?
    - Input is one cell, output is two cells.
  - Are these two cells the same?
    - Clicker Question: I was not taught cell division. A for Yes, B for No. 100 percent was taught cell division.
  - So again, are they identical or are they different?
    - We believe they are identical.
    - We have to ask ourselves, how could they possibly be different? If they were different, there would need to be some process to handle the versioning.
  - Remember that DNA is a double helix. And the two strands of the double helix are complimentary, i.e. you can reconstruct one half from the other.
    - i.e. given a C, you KNOW that the other helix is going to have G.
  - So why have these two strands, if one is redundant?
    - Because it makes it super easy to make duplicate! You have two strands, but when they get split up, you can now double it very nicely.
    - So the way that the scientific establishment believes is you start with a single cell, and it divides (using a similar process above) and again and again and again.
    - So why, when we look in front of us, we don't see a massive lump of protein, but we see something with eyes and ears and mouth and a liver.
      - We only know how to divide, but division should be perfect copies!

This is the problem of morphogenesis

- The only thing we have at our disposal is take a cell and make a copy of it
- BUT that still leads us to organisms that are really super different.
- So to say, what does it mean for two cells to be different? What is the difference between a liver cell and a tooth cell?
  - They EXPRESS different genes.
- You can think of Dna as a giant recipe book. It has the recipe for making every single protein in your body.
  - And you can think of each cell as a restaurant. It knows how to cook french, it alien, mexican, american, etc.
    - But each cell wants to specialize, as specialization brings lots of benefits (e.g., efficiency/quality). It wants to be a world class French restaurant, and not just a generic diner.
    - Each cell knows how to make everything, but it actually only makes certain things.
  - And we are pretty sure that all of the recipes are inside the DNA.
    - Mitochondrial DNA, something you only get from your mom, no mixing.
    - For the purposes of our discussion, all information comes from DNA.
  - So what do we know
    - All cells in body are identical with respect to DNA.
    - Yet they are different because each cell doesn't execute all of the recipes that are available to it.
    - Furthermore, we know that if we take a cell of a certain type, you can TURN it into a cell of a different type, without changing the DNA. How?
    - To make it more complete: Let's say we go at the moment when you have 1,024 cells, already there is some differentiation there. Already things are a little bit different from cell to cell.
      - But if you take any of those two cells, you can in some sense start again, and cause the two of them to become whatever you want.
    - What mechanism is it that allow us to control what recipes each cell becomes.
      - It is based on the environment. There is a very strict notion that the cell changes based on what is around it.
      - Think of the environment as being the customer, making demands of the restaurant.
      - Specialization brings about all sorts of efficiencies.
      - All sorts of things need to happen, but your cells learn to cater. We have pushed the problem a little bit.
- So again
  - Every cell has the cook book for every "dish" (i.e. every protein, i.e. it knows what a "tooth" cell looks like, it knows what a "liver" cell looks like).
  - The thing that differentiates liver and tooth is which recipes the cell decides to use in the cook book.
  - And the recipes it decides to pursue are determined by forces coming from the outside.
- But we still haven't addressed the main problem: how does the symmetry breaking process happen.
  - What causes the divergence? We haven't solved or addressed that problem. How does the environment affect things?
  - The cells are perfectly willing to change based on requests, but where do these requests come from!
    - The environment is water and salt, basically! The environment is not very smart!
      - There is something called In Vitro fertilization (e.g. in glass fertilization) glass, salt, water, and a single cell, and it will grow a human.
    - It does turn out that cells not only take orders for food, they also give orders for food.
      - The cells tell each other what to do.
      - But still: the cells are identical, so we haven't broken the symmetry problem.
      - The short version is: We do not know the solution.
      - And this is morphogensis. This is the other half of what Alan Turing was thinking about, besides computer science.

Alan Turing

- Besides helping the Allies win the war and also make amazing contributions to computer science, he was also thinking about morphogensisis.
- And this was in the 1930s, before much research into DNA was around. Very impressive.
- Does anyone have any ideas about how to solve the problem?
  - One hint: if you try to do this in zero gravity, then it fails.
    - You put a single cell in zero gravity, it does divide, but it does not turn into a human, it turns into a monster. Something that is not sustainable and dies right away.
      - That is to say: the cells don't specialize as they should.
  - It seems that bottom cells are treated differently than top cells (i.e. if you have a bunch of cells attached to each other, some of them will be on top, and some of them will be on bottom), as they would be affected by gravity differently.
    - Also, eventually enough cells grow that you can begin to definitely speak of inside cells vs. outside cells, meaning that they are experiencing different environment, just by geometry.
      - The outside is exposed to salt and water.
      - So those outside cells get signals at a different, slower, rate than the inside cells. It might still be the identical signals, but this is enough to get the ball rolling on this differentiation.

Interesting Idea: The "point" of an organism at any moment in time is to sufficiently house the stuff for the organism to get ready for where it needs to be tomorrow.
For the above idea to work: If you have a target shape, you would somehow need to know what "half shape" (or "in progress shape" would help you to achieve the target shape?

- It's so easy to imagine this thing going wrong! There are so many intermediate phases.
- We know it to be a relatively robust process though! Most times it works, by virtue of the fact that there are so many humans, and most of the time we look pretty similar to each other!
- An incredibly elegant problem.
- And for those of us who are "hard core" whenever we turn on our computer, basically the same thing happens.
  - older computers have something called the BIOS
    - Really tiny program written on hardware. Electricity comes in into the chip and it makes things happen.
  - And the only thing the BIOS knows how to do is read more bits. And its in hardware.
    - But then it starts reading off of the disc the core part of the operating system.
      - And after a while you have the complete new layer that is the operating system.
      - And then you have applications that speak the language of the operating system (as opposed to the language of the BIOS)
      - There are multiple levels of abstraction in organization.
- Teacher hopes that one of us pursues morphogenesis as a living.

part of the reason why he is telling us this is as an introduction to something important that happened yesterday.

- how many have heard of this quote: (First They Came...)
- idea is the greatest evils happen because the majority of people sit back and watch.
- A quote about the Go boardgames from Facebook.
- Very few people knew about Go in the class.
- Yesterday, an AI beat the European Champion in Go.
  - Why is this a big deal?
  - they have scheduled a match against the world champion in a few weeks, who may be the greatest go player of all time.
    - What does this have to do with morphogenesis.

First, let's learn a little bit about Go.

- 19 x 19 grid.
  - Way that the game is played; rather than thinking of squares as important points, think of the dots between the square as important points.

  - If you manage to put white pebbels that fully surround a black pebble, then the black pebble is removed, and if you surround a large collection, then you get to take all of them out.
    * Each player is trying to remove opponent pebbles while protecting their own.
  - And if you reach a configuration where you own most of the board state you win.
    * Game is very old, 2,500 years
  - And for a long time, it is a litmus test for the strength of AI.
  - The reason why is the following:
    * In chess, the pieces are different. From the main point of view, if we've played a little bit, the number of legal moves available to each person is relatively small.
    * Once we've gotten relatively far into the same, there are a relatively small number of moves we can make that are legal.
      · And so we make our move, and then our opponent has a small number of legal moves as well.
      · And you can keep on doing that again and again, and you get a tree.
  - so let's say that, for both for us and for the opponent, every time you have to play, you had 7 moves available to us (not true, of course, but illustrative).
    * So if we were to play m moves each.
    * Then the number of possible situations would be 7 * 7 .... * 7.
      · Then the total number of possible situations is 7^2m
      · And so this number goes up really fast! if you play 6 moves each, you get 13 billion possibilities
      · As large as that number is, though, your iPhone can consider this in just a second. Your phone can beat you in chess while also showing you tube videos.
  - In chess, there is a standard system for the value of pieces in chess.
    * All your computer has to do is go down all the way and say, in this position, what is the value of opponents pieces
      · And figure out if this is a favorable situation or not.
  - First 10 moves in chess are literally known by heart.
    * If one person opens in this way, then you are supposed to follow a certain pattern.
      · It is only after the first 10 moves that humans stop playing like humans.
  - And then also you can make yourself faster by pruning off certain parts of the tree.

So in Go, instead of there being 7 possible moves, we are instead dealing with 200 possible moves (technically even more, since you put stones on the 'dots' instead of inside the grids). It makes it go INCREDIBLY high.

- A really really really big number.
- Computers can't just do lookahead (like they could with chess)
  - The tree grows so fast that there is no chance of figuring out how to play go by looking ahead 6 or 7 moves and choosing what seems to be good.
  - People recognized this, and appreciated that Go would not be won by brute force. That something exciting would have to happen.
- AND top chess players DO play by lookahead.
  - And because they are super good they don't consider all possibilities, they only look down certain paths that they know are going to probably be good.
    * A good chess player only has to decide between 2 moves generally, OR there is one clearly best move, because all the rest is crap.
  - But if you Ask Go players to do the same, they can't do it.
    * They rely on something called intuition. It is because it felt right.
    * Up until about a year ago, the best Go computers would lose to an amateur human who had only been playing for a few months.
      · And yesterday, it was announced that they beat the European champion.
    * This is a watershed moment. And they did not do it by brute force. Everyone expected it couldn't be done by brute force, and that was proven correct. It was done by something far more sophisticated.

    - First day of class we talked about teaching machines how to paint photos in the style of a painting
        * That same idea is how they did this Go problem.
        * The fact that computer can do this same thing to solve two such different problems is a little scary.

(**AlphaGo Video**)

- Beating happened with no handicap. "Even Stones"
- Demis Hassabis, CEO Google DeepMind
- David Silver Google DeepMind
- Game of perfect information
    - Refers to two things: no element of chance, and no secret information.
        * E.g. when you play cards, you don't know what your opponents cards are so that is a game of imperfect information.
- Toby Manning (British go Association)
    - AlphaGo is name of their program.
        * They use deep neural networks under the hood.
        * First neural network is policy network, second is value network.
        * first just considers a handful of moves, then decide if you like that board state.
            · Less brute force, and more like imagination.
        * Going to be playing against the world champion.

Mastering Go Through Deep Neural Networks. Deep Learning doesn't ring a bell. Not seen in popular press.

- Deep Learning was something proposed about 25 years ago.
- Thought to be computational intractable.
    - Sort of says something like this:

Idea is we would like to build a circuit that we can think of us a function that can do handwriting recognition.

- The US Post office was interested in being able to automatically classify mail.
    - Be given an envelope, take a picture, and be able to read the handwritten address.
        * People would be happy if you could have computers sort mail according to post code.
    - So one of the first applications was trying to read these postal codes.
        * Imagine that we had a 10 by 10 grid of pixels, and for each pixel, you ask Is there any ink in there yes or no.
            · so you have a single array that contains 100 bits, each of which is zero or one.
            · Now even if the shape is the same, if you push it a row up or down, or shit if a column left or right then it also moves a lot.
            · Not to mention that different people write two differently!
    - But HUMANS find this task to be pretty trivial.
        * So goal is to take a machine that takes our array of 100 bits, and connect it to a 100 circles, and a pattern of 0s and 1s manifest there.
            · And then there is another layer of boxes each box takes two inputs (two bits) and if both are 0 or both are 1, then it returns 1, otherwise 0.
            · And so now imagine we have 5 layers of these boxes,
            · And in the last layer there are only 10 boxes, from 0 to 9, and we are guaranteed that by the end only one will fire.
            · So all that is being described is the possibility of an abstract promise based on digitizing the image, and then feeding it through layers of processing, and by the end, it will do the job.
            · We are told that IF we wire it up correctly then we can do it.
            · Given that promise, we can attempt to DISCOVERY that structure.
            · Learning refers to this process of discovery.
            · How do you discover the circuit

      - You take millions and millions of hand written digits that have been annotated by hand what digit it is.
      - So you have millions of examples of twos and threes and fours etc.
      - And so then you start with a random wiring, feed it all the labelled data that you know.
      - You measure the quality of the circuit by how much error it makes on the training set. Start with a random circuit how well does it do?
      - And then you say, Okay, let me take a single edge, and yank it out. If I commit to removing that connection I now need to make a new connection. And let's say that I only have 10 options for that.
      - And so then you do the same thing for each of those 10 options, and you pick the one that minimizes the error!
      - Metric is only how we minimize error.

How does a human player become better? By playing! You play, you lose, you introspect, and then you modify.

- So the more examples of hand written digits the better, the more computation you have to do the search the better.
- And we've reached a point where we have enough computation for it to work.

The other major revelation is having multiple layers

- They form, in the final result, successive higher level of abstractions. By which we mean in the case where we were doing image recognition, the first few layers is merely searching for stray lines.
  - You first do the dumb blind search.
  - And you get a circuit that does incredibly well.
  - and it then recognizes lots of different levels of abstraction of vision.
    - And that is the connection to morphogenesis: successive layers of abstraction giving rise to multiple layers of added complexity.