

Computer Science Future

15-110 – Friday 12/06

Learning Goals

Define a few relevant **buzzwords** of up-and-coming computer science trends.

Recognize the possible **implications** of upcoming technologies.

Cryptocurrencies

How does money work?

A dollar bill does not have any intrinsic value; it's only worth \$1 because we all agree it should be, and we all trust it will continue to hold that value.

This is how all currencies work! Usually we trust currencies because they are backed by a powerful system (a country or government), and we trust that system to not start printing a lot more money.

Cryptocurrencies are just like normal currencies, except that they are **independent**- they are not backed by a country. So how can we trust in their value?

Blockchain – how cryptocurrencies work

In most cryptocurrencies, the value of the currency is protected by the **collective** that uses the currency. The collective keeps track of who has how much money with a data structure called a **blockchain**.

The blockchain is like a series of records of how much money everyone has, and everyone in the collective can contribute computing power to keep their own personal blockchain. This is called **mining**.

Whenever someone makes a transaction, they send that information as a message to the rest of the collective; if a majority of people (weighted by computing power) accept it, the transaction becomes official.

Discuss: What happens if an individual controls more than 50% of mining?

Bitcoin



Created by a mysterious individual(s) called Satoshi Nakamoto in 2008. To this day, no one knows who this is...

The code was made open-source in 2009. Link here: github.com/bitcoin/bitcoin

Bitcoin started out as worth a few cents per bitcoin. It is now worth over \$7000! But the majority of the value increase happened in 2018- it's actually decreasing in value right now.

Learn more here [start at 2:41] :

www.npr.org/sections/money/2011/07/13/137795648/the-tuesday-podcast-bitcoin

Deepfakes

Editing Media

You likely already know about photoshop; it's used widely to edit images, to the point that it's sometimes hard to tell if a photo is original or edited. And we've reached the point where computers can even generate photos of people who don't exist:

<https://thispersondoesnotexist.com/>

Recent advances in technology are making it possible to edit other types of media as well:

www.youtube.com/watch?v=cQ54GDm1eL0

Video Editing - Deepfakes

Edited videos of this form are called 'deepfakes'. They work by mapping the facial motions of an actor onto the target's face.

Example:

www.youtube.com/watch?v=ohmajJTcpNk

This technology is intended for post-production editing in film. Some people also use this technology for fun purposes, like turning all movies into Nicholas Cage movies. But there are many concerns as well.



Audio Editing

Adobe (the company that created Photoshop) has also started work on a system that can edit audio, Project VoCo:

www.youtube.com/watch?v=I3l4XLZ59iw&t=58

Again, this technology is intended for post-production use, but could obviously be put to other purposes as well.

Deepfake Repercussions

Discuss: what effect will this have on communication? How can we prepare for it?

Learn more here:

www.wnycstudios.org/podcasts/radiolab/articles/breaking-news

5G

Cellular Networks

Our phones use **cellular networks** to place calls, send messages, and also download data from the internet.


Just like with the internet, these networks have different **protocols** that define how they work. Your phone probably runs on 3G, 4G, or maybe 5G. These are three different generations of wireless protocols.

Each new generation broadcasts data at a different **frequency**, and can usually transmit data more quickly than the previous generation. We usually get a new generation once every 10 years.

5G vs 4G

Many telecommunications systems have started rolling out devices that can work on the **5G standard** this year (2019). Many more will join in 2020.

How is 5G different from 4G? The primary change is that 5G will support **much faster download speeds** than 4G. While 4G can transmit megabytes per second (2^{20} bits), 5G can transmit gigabytes (2^{30} bits).



How 5G Works

This speed is possible because the electromagnetic waves 5G uses to transmit data are much shorter than 4G waves; that makes it possible to transmit more data in the same amount of time.

Because these waves are shorter, 5G requires **more access points** than the previous networks; in a crowded area, that could mean one on every city block. This infrastructure change will take a while to implement.




Illustration by WiredScore

5G Implications

Discuss: how will radically faster download rates change the way we consume information on mobile devices?

Learn more here: gimletmedia.com/shows/science-vs/j4h39x

Quantum Computing

Quantum Physics vs Quantum Computing

Quantum physics states that particles are not limited to being in only one of a number of states. A particle can also be in a **superposition** of states. For example, a particle that could be charged or not charged can also have those two states superimposed, and be both of them at the same time.

This same idea is used in **quantum computing**. A classical bit can be one of two states (0 or 1); a quantum bit, or a **qubit**, can be 0, 1, or the superposition of 0 and 1.

BIT

0



1



QUBIT

0



1

Quantum Algorithms

When we represent data using **qubits**, we can process it using **quantum algorithms** that test all the possible states of the data at the same time.

This produces a quantum result, which we then need to translate back into a classical (non-quantum) answer. This is done probabilistically, but certain translation algorithms (like [Grover's Algorithm](#)) have a high likelihood of success.

If this translation can be done quickly, it can lead to **huge efficiency gains**. For example, Grover's algorithm can take a $O(N)$ algorithm and solve it in $O(\sqrt{N})$ time.

Quantum Implications

Quantum computing may seem very theoretical, but it can have real impacts on the computing we do today.

For example, consider **integer factorization**. This is an intractable problem for classical machines, but [Shor's Algorithm](#) can solve it in $O(\log N)$ time, if N is the size of the integer. This has been successfully implemented with up to 21 digits.

Why does this matter? RSA- the algorithm that provides encryption on the internet- depends on integer factorization being hard to do!

Recent Quantum Breakthroughs

Very recently (in October 2019), Google announced that it had created a quantum processor, called Sycamore, that could represent 53 qubits.

This processor could solve a task that would take a classical computer thousands of years in only 200 seconds. IBM later claimed that it would only take a classical computer a few days, but this is still a huge improvement.

Paper here: www.nature.com/articles/s41586-019-1666-5

Learn more here: www.youtube.com/watch?v=lypnkNm0B4A

Automation

Automation in the Workplace

If we shift the scale of time to look at how computing will affect the world in the long term, we need to consider how **automation** will affect the world economy.

Automation isn't a new process. Since the Industrial Revolution, humans have been finding new methods to solve problems that require less human labor. This means that the skills expected of workers are constantly changing.

Check out the most common jobs in the United States over the past forty years here:


www.npr.org/sections/money/2015/02/05/382664837/map-the-most-common-job-in-every-state

Automation by AI

More recently, however, the success of artificial intelligence has led to concerns that common jobs we previously thought couldn't be automated will be replaced.

Specifically, jobs that are **routine** are more likely to be replaced by algorithms. This includes certain expected jobs (like manufacturing and data processing), but also more surprising jobs, like truck driving.

2.1.9 Impact of automation on jobs



Note: Percentages are from Frey and Osborne (2017) estimates on probability of automation. Framework is based on Acemoglu and Autor (2011).

Concerns about Automation

The rise of AI automation has led to some concerns.

First- what will happen to people who are displaced? Will they be able to find a new line of work, or will governments need to turn to universal income? This is a policy question, but one that technology companies need to be involved in.


Second- how can we make sure that the AIs taking over jobs are well regulated? Consider self-driving cars. The AI for a car won't work 100% of the time; it will need human intervention in rare circumstances. Should the person in the car be responsible, or should a remote monitor step in?

AI + Human Collaboration

On the other hand, there is a great deal of potential in using AI to **support** people in their work, instead of replacing workers outright. This is possible because AIs are really good at solving specific tasks, while humans are good at generalizing.

This is already done in a number of industries- autopilot controls for airplanes, for example, and [AI-assisted grading](#) for professors. In 2014, a human-AI chess team (called a [centaur](#)) beat an expert computer program, even though the human was not a chess master!

In the future, AI will be able to provide even more support. This is predicted to drastically impact the fields of healthcare, law, and accounting.



Potential Impacts

Discuss: will AI automation and support have a more positive or negative impact on the world?

Learn more here:

www.npr.org/sections/money/2018/10/24/660297140/the-rise-of-the-machines

The Singularity

AI Intelligence

Finally, with all this talk of technological improvements, we need to address a big question: will we someday reach a point where technology starts rapidly improving by itself?

This could occur due to a number of possible scenarios, but let's focus on one: **sentient AIs**. Is it possible that AIs will one day obtain human-level intelligence?




Modern computer scientists have a range of opinions on if and when AIs may become sentient. Here's what we know now.

What is Intelligence? – Turing Test

First, we need to address a simple question: what would it mean for a machine to reach human-level intelligence?

In 1950, Alan Turing introduced the idea of an 'imitation game', now called a Turing Test. In this game, a human volunteer is asked to hold conversations with two different entities, A and B. One is a human; the other is a computer. The volunteer is supposed to guess which entity is the human.

If the volunteer cannot consistently tell the difference between the human and the computer, we say that the computer passes the Turing Test.





What is Intelligence? – Chinese Room

One popular counter-example to the Turing Test is the thought experiment called the Chinese Room, introduced by John Serle in 1980.

In this experiment, a person (X) is kept in a room with a large book of Chinese characters that maps questions to answers. X cannot read this text, but if someone sends them a piece of paper with Chinese characters on it, they can look for a matching set of characters in the book, write down the corresponding answer's characters, and send the response back out.

Now imagine that you are the person outside of the room. You send in a question, and the room sends back a response. You might reasonably assume that X is communicating with you. But they don't understand the text they're writing; they're just following pre-prescribed rules.



Implications of the Turing Test

Turing thought that by 2000, machines would be able to fool 30% of human judges after a five-minute conversation.

Technically, a machine called Eugene passed the Turing Test in a competition held in 2014. However, the machine was pretending to be a 13-year-old boy from Ukraine, and only fooled 33% of the 30 judges.

We still have a long way to go to create machines that can truly pass the Turing Test.

What's so hard? Open Research Areas

We have at least two huge open areas of research.

Edge cases. In driving, for example, cars are pretty good at driving in reasonable conditions, but if you add construction or snow or even a lot of traffic it is incredibly hard to test what will happen outside of simulation.

General AI. How do we go from using many specific AI systems to one general one? Can we transfer knowledge about one task to another? How does a person do that? We don't know how to make a system that's good at a lot of things all at once.

When we can solve these problems, we'll come closer to AI Singularity. Most recent estimates from top AI researchers average around 81 years from now, in 2099!

Learn more here:

www.theverge.com/2018/11/27/18114362/ai-artificial-general-intelligence-when-achieved-martin-ford-book

Learning Goals

Define a few relevant **buzzwords** of up-and-coming computer science trends.

Recognize the possible **implications** of upcoming technologies.