Computer Science Future

15-110 – Wednesday 4/24

Announcements

- **Hw6** due **Friday at noon**
  - No revision deadline, no late submissions! Submit what you've got before the deadline, even if it's incomplete. Submit early and often, and read your feedback.
- Practice materials are up
- Final exam:
  - **Monday, Dec. 9th, 1pm-4pm**
  - Locations to be announced

Grade estimation

The last two checks, HW6, and the last few exercises are not yet in Canvas, but…

Assuming that your current category averages do not change much, and given that the final exam is worth 25% of your semester average, a **rough** (non-binding) estimate of your semester grade can be predicted for a given final exam score by:

```
(canvasCourseAverage*0.75) + (estimatedFinalScore*0.25)
```

Remember that if you demonstrate any effort on the final exam, you're guaranteed at least a 50.

See syllabus for more details on grade breakdowns!

Learning Goals

- Recognize and describe the impact of key future computing ideas, including:
  - **Deepfakes**
  - **Generative AI**
  - **Virtual reality**
  - **General AI Intelligence**
  - **Quantum computing**

Deepfakes

Editing Media

Photoshop and other tools are widely used to edit images, to the point that it's sometimes hard to tell if a photo is original or edited. Now, computers can automatically 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. For example, consider this 'video' of Mark Zuckerberg, creator of Facebook:

https://www.instagram.com/p/BypkGIvFfGZ/

(Note, that video is from almost four years ago)

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.

How it works: www.youtube.com/watch?v=ohmajJTcpNk

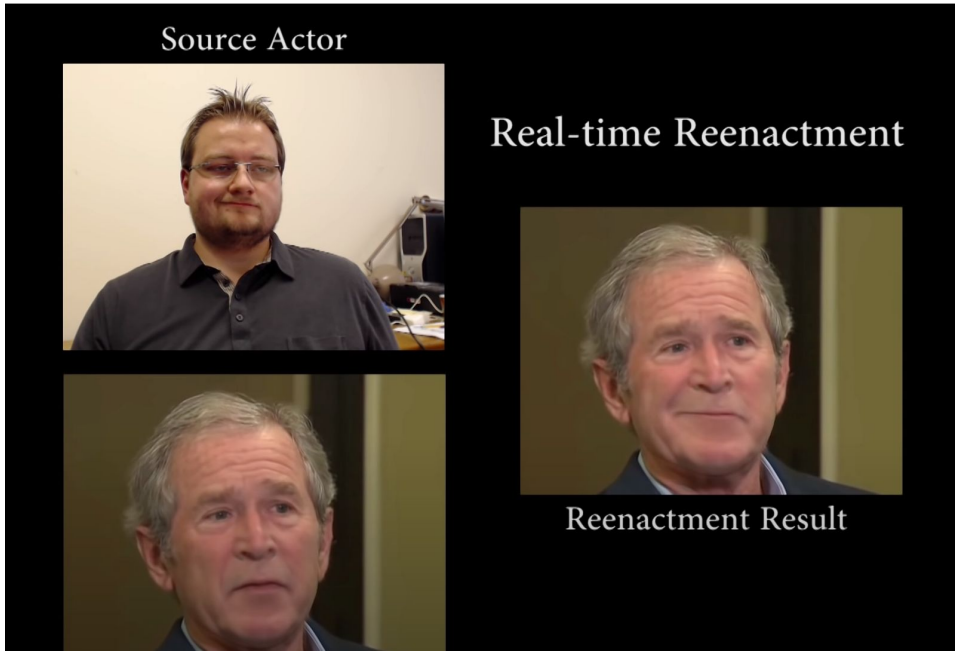

Intention vs. Potential Use

This technology is intended for post-production editing in film. For example, Disney has used deepfakes in several recent Star Wars movies to bring back actors who have aged a great deal or passed away since their original appearance.

Some people also use this technology for fun, turning all movies into Nicholas Cage movies or having Dr. Phil interview Dr. Phil.

There are lots of concerns about the potential repercussions of deepfakes when used for more serious purposes, like spreading fake news. This can be done by editing real footage or creating non-existent people to spread propaganda.

Audio Editing

Adobe (the company that created Photoshop) also showcased research on a system that could 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.

Learn more here:
www.wnycstudios.org/podcasts/radiolab/articles/breaking-news

Generative AI

Generative AI

Some AI models classify or cluster data (we've talked a bit about those). Others create new data – we call those **generative** models.

We briefly talked about large language models before. Here, instead of generating text, we'll talk about models that generate images.

The image generation technique that we'll talk about is called a generative adversarial network (GAN).

GANs are used to create the images on https://www.whichfaceisreal.com/

GANs

GANs have two pieces:

- A generator, which creates new images
- A discriminator, which classifies images as real or fake

At a high level, the generator learns to create good-looking images by trying to fool the discriminator into thinking the generated images are real.

Discriminator

The discriminator is just a classification model. It takes an image as input and estimates the probability that the image is real or fake (outputting a score between 0 to 1 for each category).

We train the discriminator by showing it a large set of real images (say, pictures of cats) and a large set of fake images (say, random noise), all labeled.

Distinguishing cats from random noise is easy, so the discriminator does well at first.

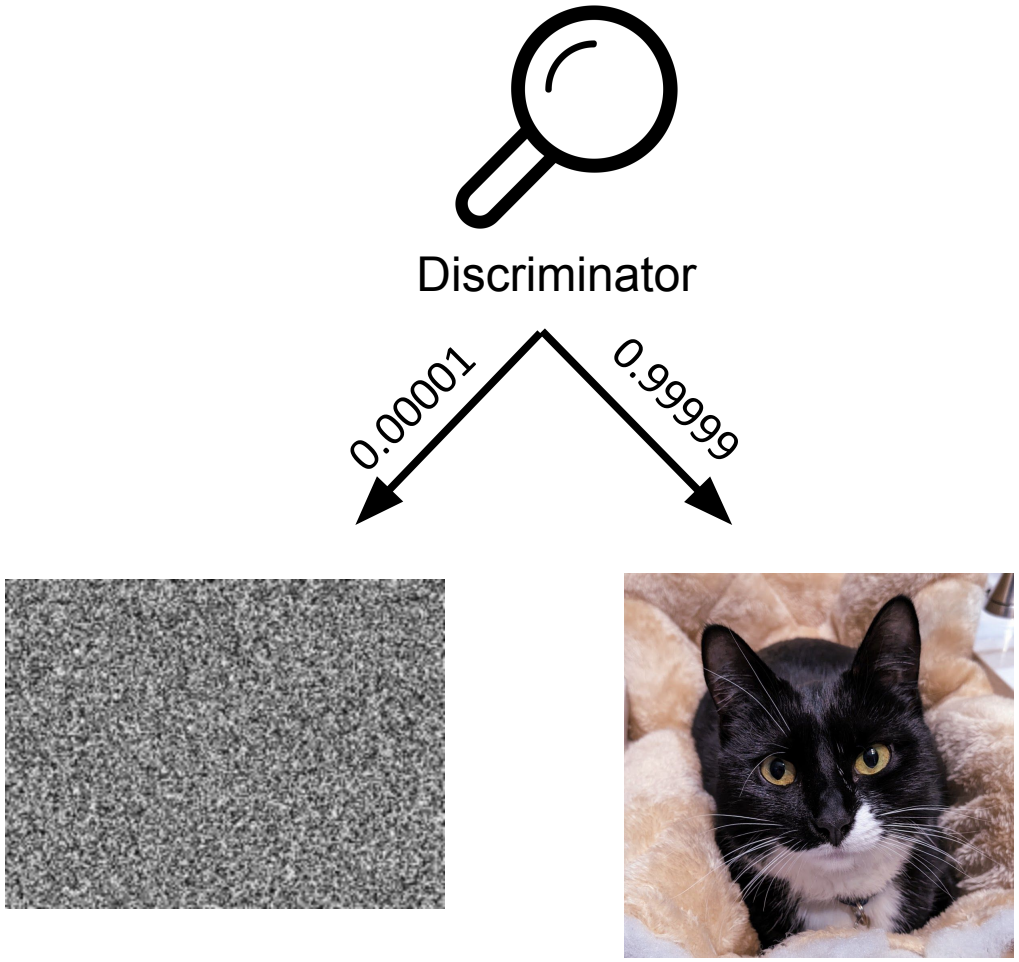

Generator

The generator is supposed to create new images (fake images) that fool the discriminator into thinking that they are real images.

At first, it has no idea how to do this. It generates lots of nonsense images that are just random noise. But some of those images are slightly closer to being cats than others. When the discriminator starts to think the generator's images could be real, that reinforces the generator's current settings. Over a long period of time, the generator gets better and better.

Discriminator

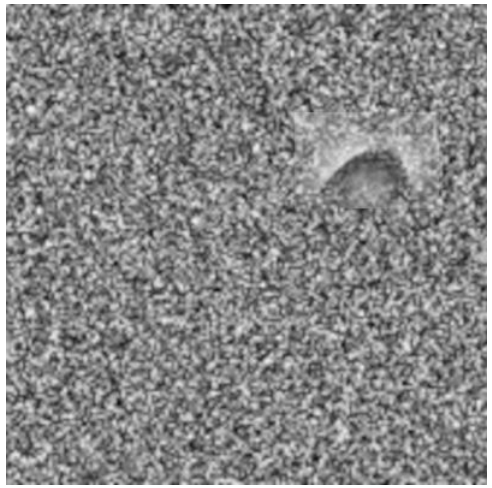

Adversarial

The “generative” part of “generative adversarial network” makes sense, but what does “adversarial” mean?

Eventually the generator gets reasonably good at its task, and can reliably fool the discriminator – and so, we update the discriminator! We can now give it the newly-generated images as examples of fake data, and the original real images as examples of real data.

After improving the discriminator, the game begins again! The generator is trained until it can fool the new, better discriminator.

Each time one model improves, the other is trained until it catches up: they are adversaries, each working to be better than the other.

The cycle repeats until the person training the models is satisfied with the quality of the images the generator can produce.

Learn more: https://machinelearningmastery.com/how-to-develop-a-generative-adversarial-network-for-an-mnist-handwritten-digits-from-scratch-in-keras/

Virtual Reality

Virtual Reality is Embodied Simulation

Virtual reality (VR) is a technology that lets you experience a virtual space as if you're actually there. You can see the world, hear your surroundings, and sometimes move around and interact with different parts.

At its most basic, VR maps your senses to computer representations. It uses a headset with many sensors to change what you see based on your movements, speakers to change what you hear, and controllers to let you interact with the virtual world around you.

VR has been explored widely in pop culture and is available commercially in several game consoles. But what can it actually do?

3D Rendering Translate a Space

We've shown how to draw two-dimensional images on a canvas to make pictures and simulations algorithmically. What changes when we move to three dimensions?

In three dimensions our view is based on a **camera** situated within a space. That camera has a position and a viewing angle that determine what it sees.

The world is then generated by taking the model of the 3D space and drawing each object in that space based on the geometry of how the object would appear to the camera.

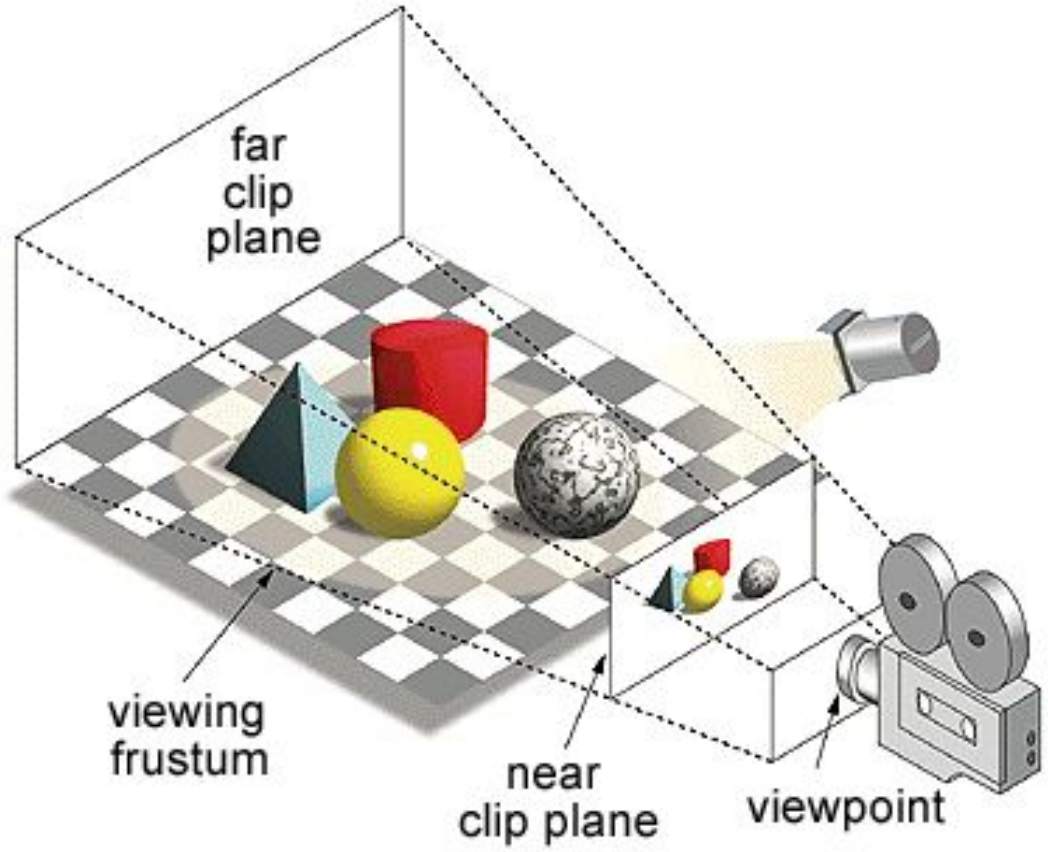

Virtual Reality Adds Complications

How is VR different from 3D rendering? The main difference is the **distance from your eyes**.

Most VR systems have a headset that puts a screen very close to your eyes. This makes it harder to trick your brain into believing that it is perceiving the world.

This disconnect can cause motion sickness, especially when there is a low refresh rate, or when the visual display does not track head movements perfectly. Companies are still trying to figure out how to reduce motion sickness to make these technologies more widespread.

Virtual Reality Capabilities

Most VR kits available for commercial purposes today focus on the headset. That means that common VR applications are mostly focused on the visual and auditory.

Recent developments have used hand tracking (with controllers) and head tracking to allow basic **interaction** with the VR environment.

Currently you can play games, watch documentaries, and participate in social experiments through VR.

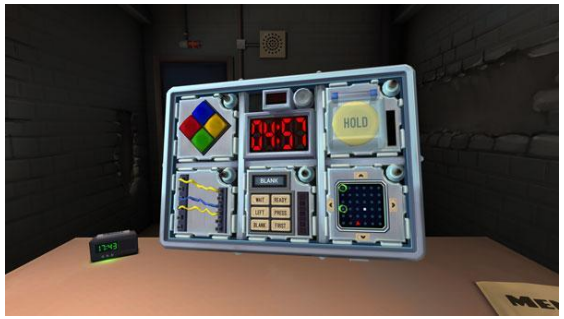

Virtual Reality Limitations

Though you can see, hear, and interact with objects in a virtual reality, other senses – touch, smell, taste – are much more limited. Controllers may be able to vibrate, but they cannot replicate most touch sensations.

Additionally, any simulation that involves moving across space has a simple limitation- the size of the room. Developers can use clever design tricks to make a virtual space seem bigger than the physical equivalent, but design only goes so far.

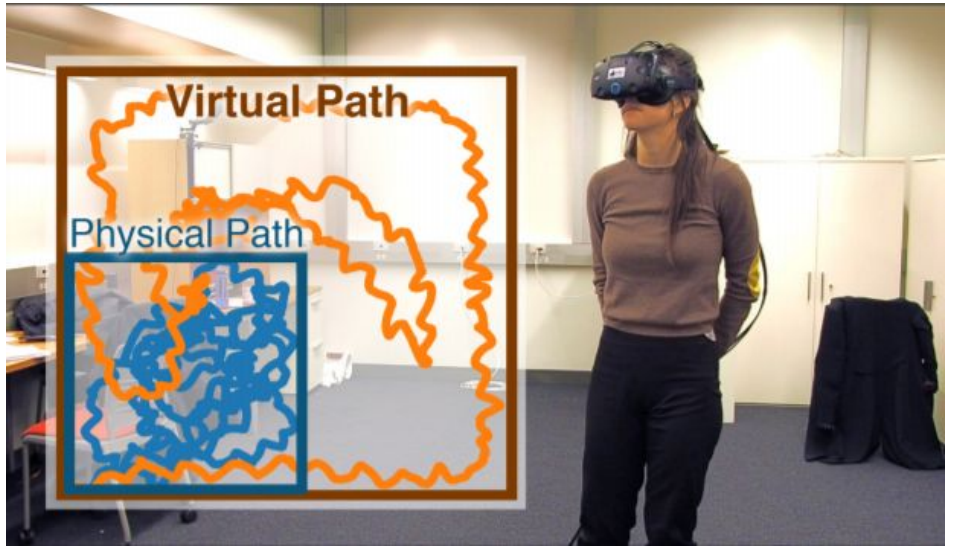

Sidebar: Augmented Reality

You may also have heard of **augmented reality (AR)**. This is like virtual reality, except that the virtual components are **overlaid** on the real world instead of taking the real world's place.

AR is also used in widely, in games and applications.

General AI Intelligence

Unstoppable AI Intelligence

The idea of 'The Singularity' is often used in science fiction to describe the point at which AI intelligence grows so spectacularly fast that it outpaces human growth. What that leads to – the end of humanity, or a new age of prosperity? – is speculation, as we can't guess what would happen next.

Many popular figures in science have expressed concerns about the growth in artificial intelligence, warning that intelligent AIs could endanger humanity. Other figures have claimed that it will take a long time to develop the field to this point, if it ever gets there at all. There's no real consensus at this time.

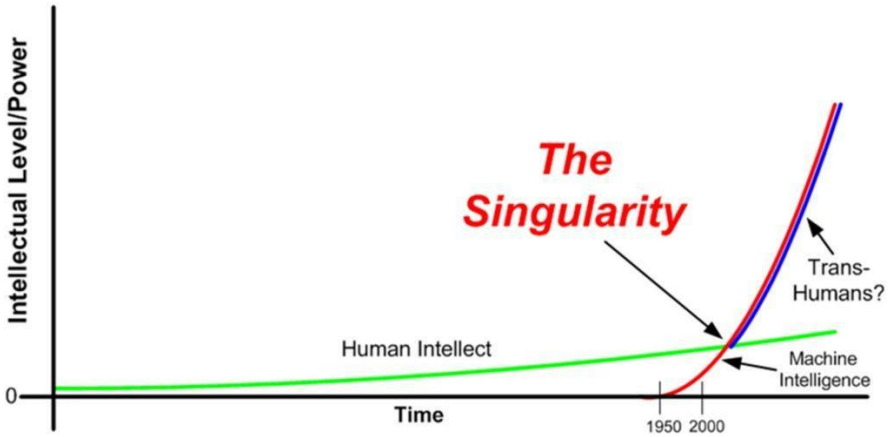

General Intelligence in AIs

Before artificial intelligence can outpace human intelligence, we need to develop algorithms that demonstrate **artificial general intelligence**.

A general AI agent should be able to learn anything that a human can learn. In other words, it should be hard to distinguish from a human being.

Even AI agents that can do a large range of things- like Siri, or Watson, or chatGPT - are not *generally* intelligent; you can't teach them an entirely new field.

How can we tell if an AI *is* generally intelligent?

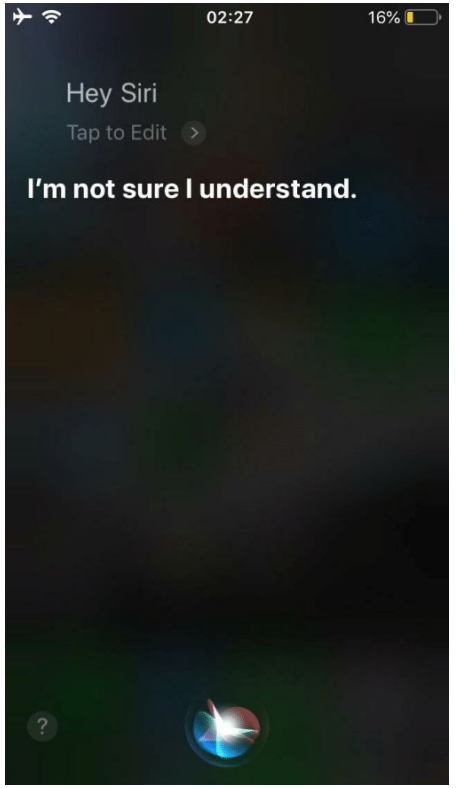

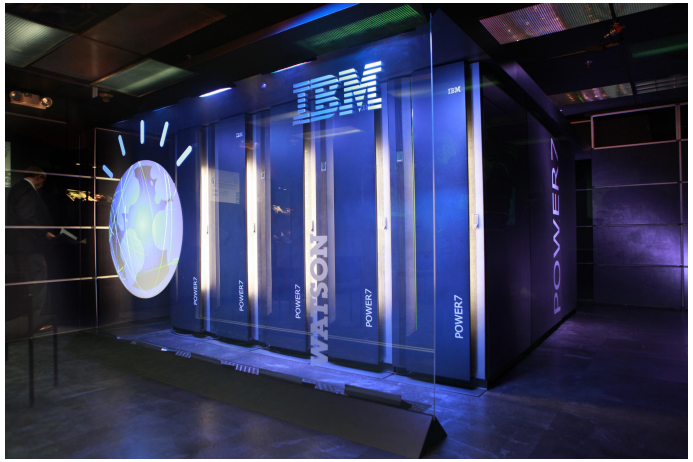

What is Intelligence? – Turing Test

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.

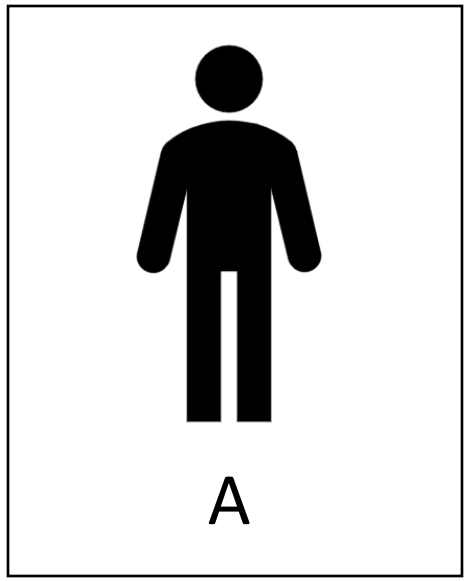

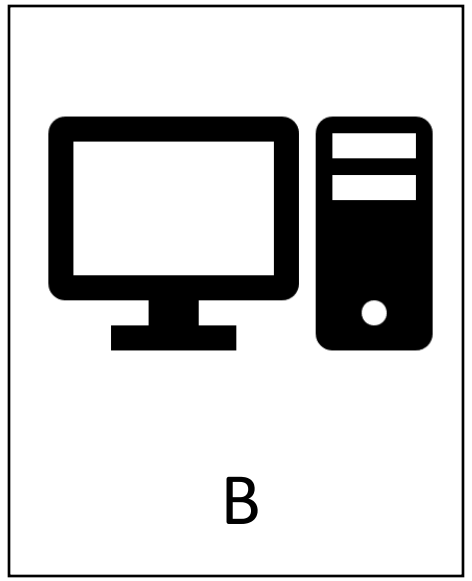

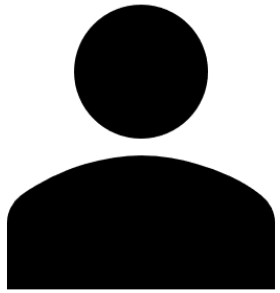

Turing Test is Hard to Pass

Turing thought that by 2000, machines would be able to fool 30% of human judges after a five-minute conversation. We use 30% as a benchmark to see whether an agent passes the Turing Test.

Some algorithms managed to 'pass' the Turing Test in past year, but they all did so by employing conversational tricks:

- The program ELIZA in 1966 simulated a therapist by looking for keywords in the typed input and sending them back to the judge as a question
- The chatbot Cleverbot in 2011 was rated as more human-like than not in a competition by recycling messages inputted into it by humans in previous conversations
- The chatbot Eugene Goostman in 2014 fooled 33% of judges in a Turing Test competition by pretending to be a 13-year-old boy from Ukraine

Much more recently, ChatGPT passed the Turing Test, which opens up many new questions about the intelligence of AI agents!

Counter-Example – Chinese Room

Even if an algorithm authentically passes the Turing Test, we still might not call it generally intelligent. Consider the thought experiment called the Chinese Room, which was introduced by John Searle in 1980.

A person (X) is in a sealed room with a large book of Chinese text that maps questions to answers. X cannot read this text. If someone slips a piece of paper under the door with Chinese text on it, they can look for a matching set of text in the book, write down the corresponding answer text, and send the response out.

Imagine that you are outside of the room. You send in questions and the room sends back responses, forming a conversation. You might reasonably assume that a person who understands what they're writing is communicating with you. But X doesn't understand the text they're writing; they're just following a pre-determined set of rules. Who are you really conversing with?

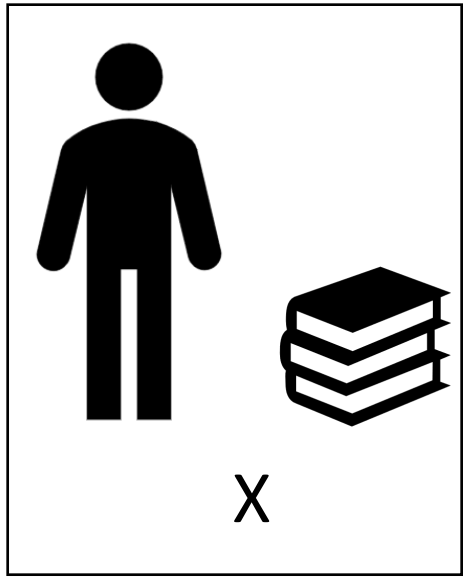

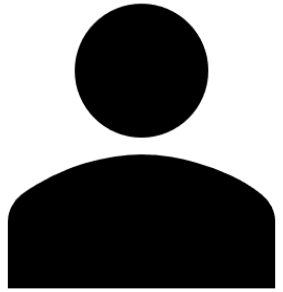

Long Way to Go

Is an AI agent like ChatGPT generally intelligent? This is a hard question to answer! Previously top AI researchers thought it would take until 2100 until we could accomplish this goal – now we're not so sure.

This is an area of open debate for AI researchers in the current day.

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

Quantum Computing

Quantum Physics vs Quantum Computing

Quantum physics states that particles are not limited to being in only one state at a time. A particle can also be in a **superposition** of states.

For example, a particle that could be in the "spin up" or "spin down" states could also have those two states superimposed; then it would be in some mixture of both states at the same time.

This same idea is used in **quantum computing**. A classical bit can be in one of two states (0 or 1); a quantum bit, or a **qubit**, can be 0, 1, or a mixture of 0 and 1. The mixture percentage is a real number, not a binary value.

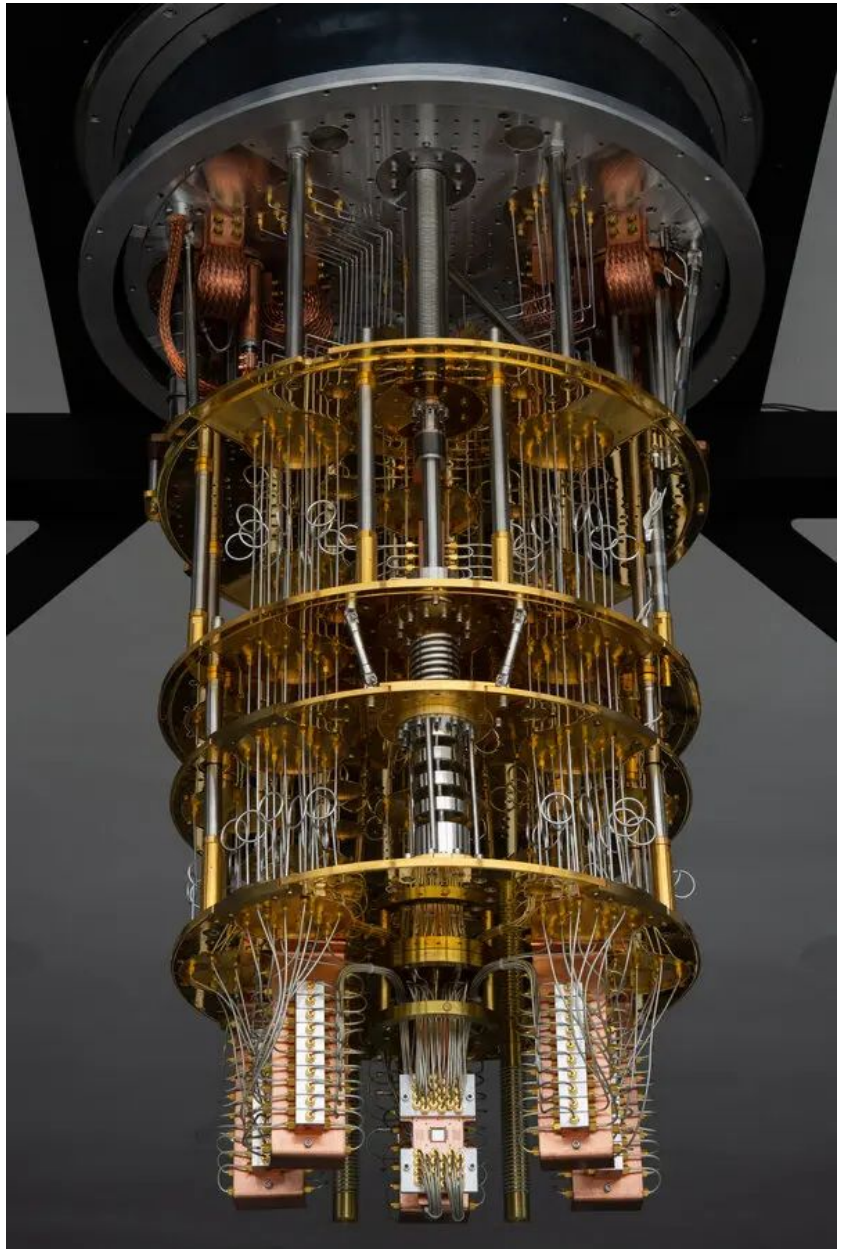

Quantum Physics vs Quantum Computing

The diagram to the right represents the Bloch sphere, which represents the state space of a qubit.

The state of a qubit can be represented as a specific point somewhere on the surface of the sphere, where “north” represents a classical 0, and “south” represents a classical 1.

Points on the hemisphere represent states where the qubit is both 1 and 0 with equal probability.

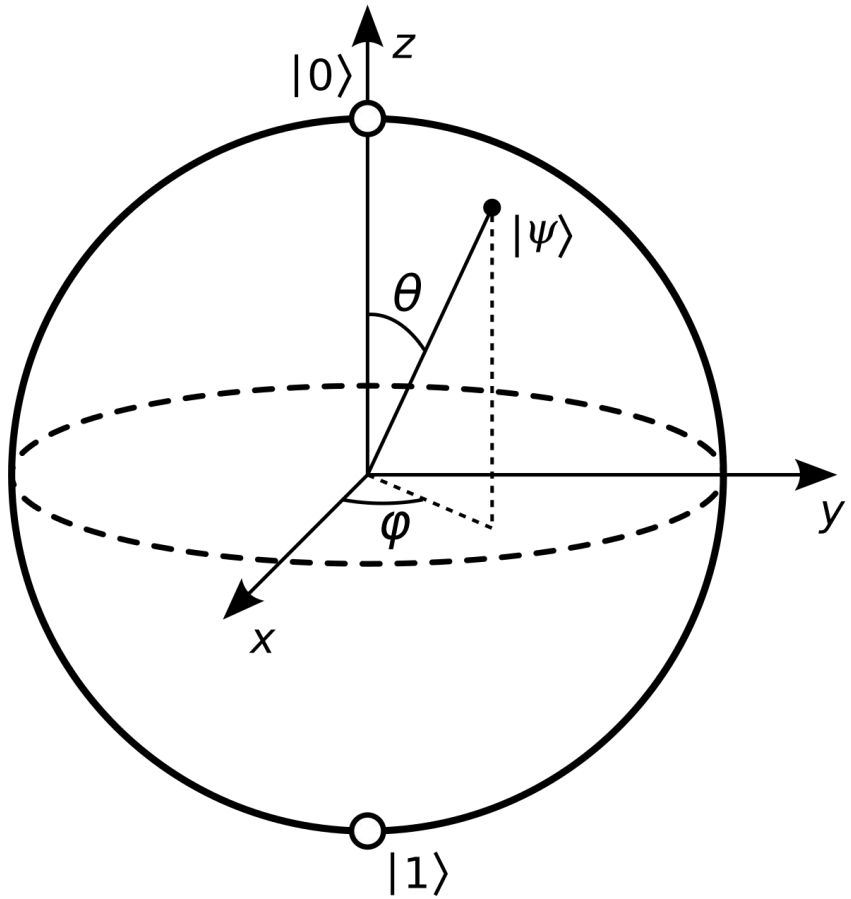

Entanglement

When qubits are represented by physical phenomena such as particle spin, it's possible to connect the states of several qubits together. This is called **entanglement**.

If a system of N qubits is fully entangled, we can represent $\mathbf{2^N}$ possible data values **simultaneously** instead of just representing one value. For example, two entangled qubits can have a state that is a mixture of 00, 01, 10, and 11.

This is what makes quantum computers so powerful – they rapidly speed up **efficiency** by playing by non-classical rules.

Quantum Algorithms

When we represent data using qubits, we can process that data using **quantum algorithms** that operate on 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 $O((\log N)^2)$ time, if N is the size of the integer. This has been successfully implemented only for tiny integers so far - up to the number 21.

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

Why is it hard to build a practical quantum computer?

To solve real, practical problems (ones for sufficiently large input sizes) we need *many* qubits working simultaneously.

Qubits are extremely sensitive and decohere (fall out of superposition) easily. The more qubits we entangle, the more likely our results are to decohere.

Also, many of our old tricks (i.e. error correction algorithms, and even some simple things like the OR operator) don't work with qubits in superposition).

This means it can be difficult to translate certain tasks into quantum logic.

Quantum Breakthroughs

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

More recently, Atom Computing created a 1180-qubit computer in late 2023

Learning Goals

- Recognize and describe the impact of key future computing ideas, including:
  - **Deepfakes**
  - **Generative AI**
  - **Virtual reality**
  - **General AI Intelligence**
  - **Quantum computing**

Other Future Topics

Bonus slides from prior semesters – material will not be tested on the final exam

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** and **decentralized** – they are not backed by a country. So how can we trust in their value?

Collective Accounting

In most cryptocurrencies, the value of the currency is protected by the **collective** that uses the currency. Every person in the collective keeps track of the financial record of the system **independently**.

Whenever someone makes a transaction, they send that information as a message to the rest of the collective. Everyone in the collective can contribute computing power to verify the transaction on their own personal record. This is called **mining**. Occasionally, miners receive a bit of bonus bitcoin for successful verification; this helps incentivize the distributed work.

If a majority of people (weighted by computing power) accept a transaction, it becomes official and is added to the public record.

Everything is public, but you can only spend money from an account if you have that account's **private key** – it uses asymmetric encryption!

Example Transaction

Suppose Soren wants to send Ariana 0.05 coin.

Soren posts a transaction to the collective (using his **private key** to encode) saying he wants to send 0.05 coin to Ariana's account (using her **public key**).

Individuals among the collective use computing power to verify that both account numbers are legitimate and that Soren has enough money to make the transaction.

When 50% of the collective has verified the transaction, it becomes part of the official record.

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

Blockchain

The collective ensures that no one spends more money than they have through a transaction, but how does it ensure that no one modifies the record? They use a **blockchain**.

A blockchain is just a data structure used to store records over time. It is literally a chain (list!) of **blocks**, where each block contains information about the state of the financial system at that point in time (by listing transactions that have been made).

The blockchain is **secure** because every block contains a hash of the block that came before it. This means that no one can go back into the blockchain and edit one of the older records to give themselves money; this would break the hashes on all the subsequent blocks, and people would notice.

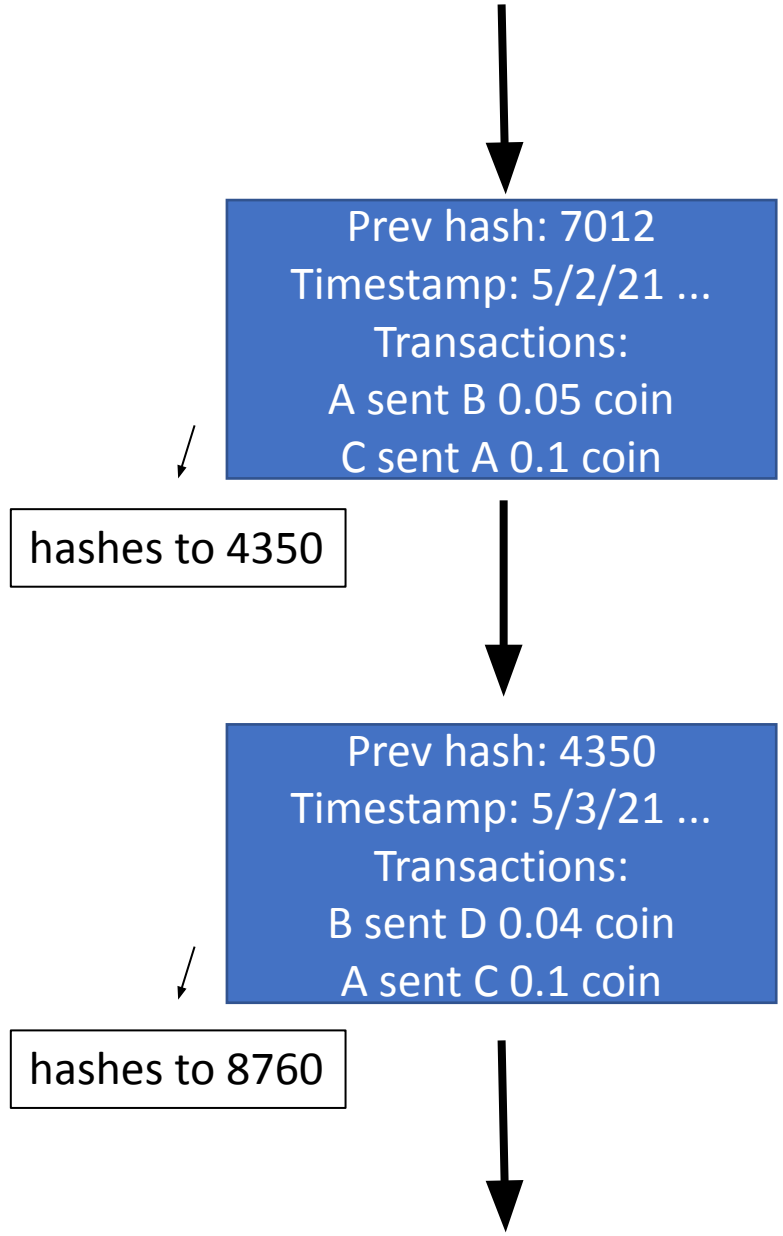

Bitcoin

Bitcoin is a specific type of cryptocurrency. It was 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 has gone up and down a great deal over time; two years it was worth as much as $64k per bitcoin, a year ago it went as low as $16k and now it is $44k per bitcoin.

Learn more here [start at 2:41] : www.npr.org/sections/money/2011/07/13/137795648/the-tuesday-podcast-bitcoin

NFTs

NFTs are Ownable Digital Items

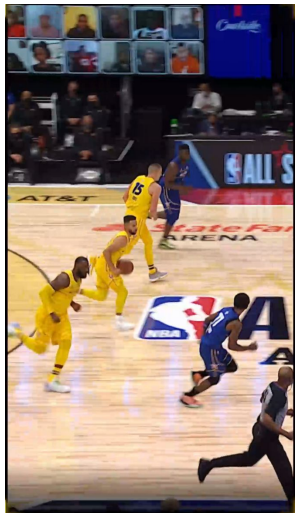

An NFT (Non-Fungible Token) is a digital item of some sort that is **owned** by a particular person. NFTs were highly popular from 2020-21, though interest has died off in the past year. NFTs are used to let people purchase digital items like sports gifs or digital art.

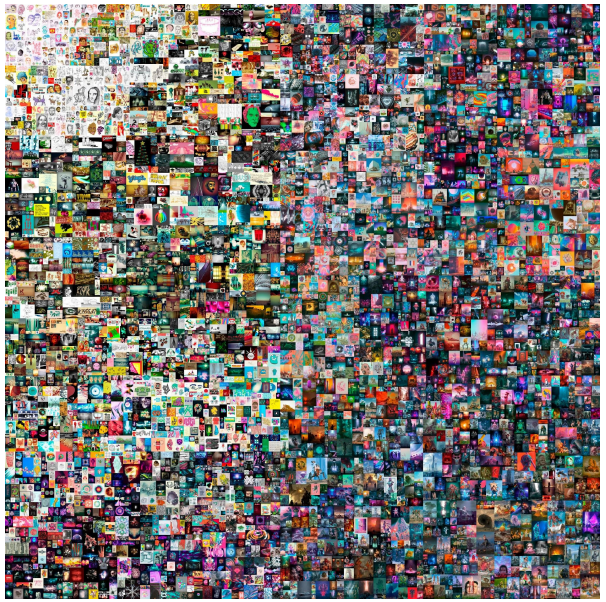

The digital item can still be posted publicly for anyone to see, but there is **digital certificate** that marks the item as officially owned by a particular person. Importantly, the certificate cannot be copied; it is unique to the owner.

How is this possible?

Blockchain Again!

The digital certificate associated with an NFT uses a **blockchain** to track the chain of ownership.

You can often copy the digital item itself, but if you don't have the certificate, you don't actually 'own' it. It's like printing a copy of the Mona Lisa – just because you have a copy of the painting doesn't mean you own the original.

This is similar to how collectible items are verified in real life with a paper certificate.

Difference from Cryptocurrencies

NFTs are different from cryptocurrencies because they are **non-fungible-** one NFT is not equivalent to another. For example, two NFTs of different NBA shots might be worth extremely different amounts of money. This is different from Bitcoin, as one bitcoin is worth exactly as much as another.

It's like comparing cash to collectible trading cards. You can exchange a $20 bill for another $20 bill, but two collectible trading cards may be worth very different amounts, even though they're printed on the same cardstock.

In fact, collectible trading cards are an excellent analogy for NFTs. They're really just collectible items on the internet.

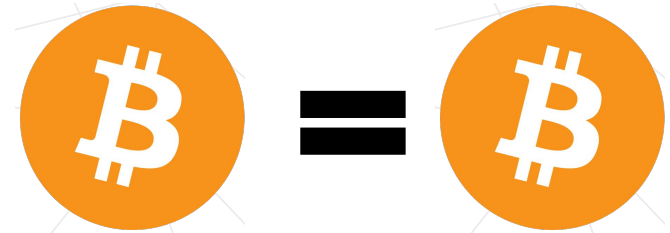

Ups and Downs of NFTs

With both cryptocurrencies and NFTs, there is a major concern around the **carbon footprint** of these digital items. Verifying transactions on a blockchain takes a great deal of computing power, and each time an NFT is sold, it adds another transaction to the blockchain. There is also a great deal of concern about the potential for scams.

On the other hand, NFTs have opened up new possibilities for artists. Some NFT blockchains can be set up so that every time the NFT is sold, a commission is sent back to the original artist. This could move more of the wealth associated with the fine art market back into artists' hands. Of course, that assumes that the art is not turned into an NFT without the artist's permission, something that has happened before.

Learn more here:
https://www.npr.org/2021/03/09/975450173/the-200k-nba-nft

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 started rolling out devices that work on the **5G standard** in 2019. If you have a new-ish phone, it's likely 5G-compatible now!

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

Read more: https://www.forbes.com/sites/forbestechcouncil/2019/09/13/what-effect-will-5g-have-on-our-world/?sh=73543da46de9

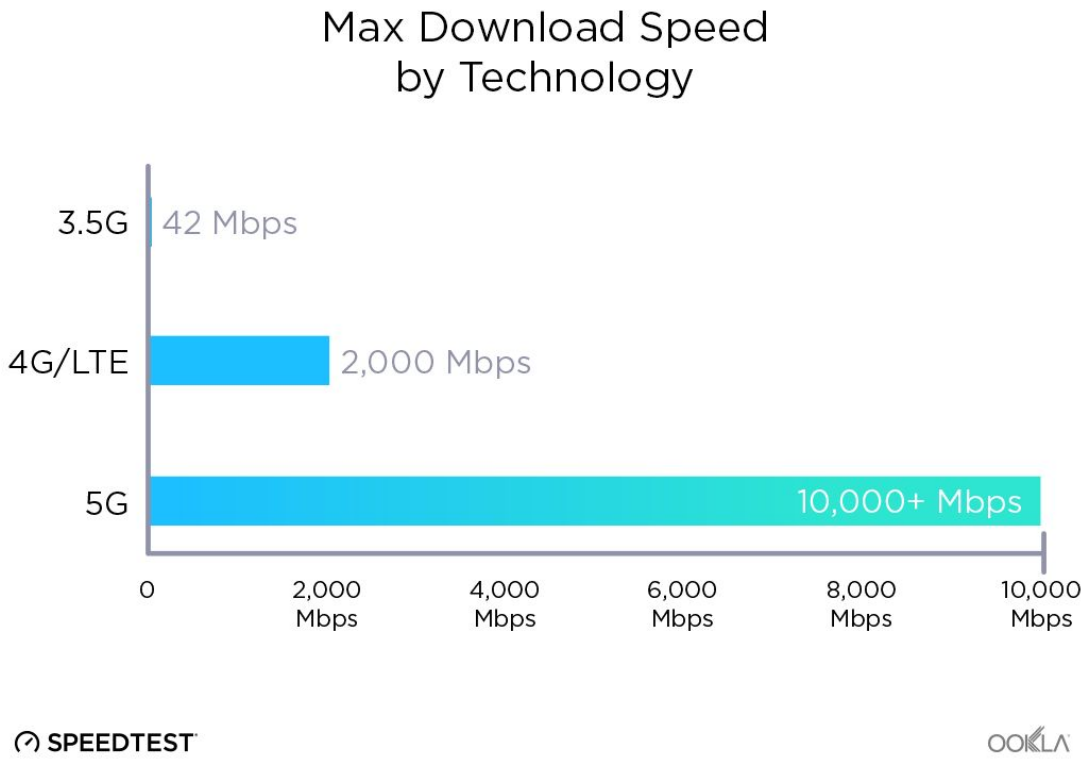

How 5G Works

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.

[Pittsburgh has already had some new access points installed.]

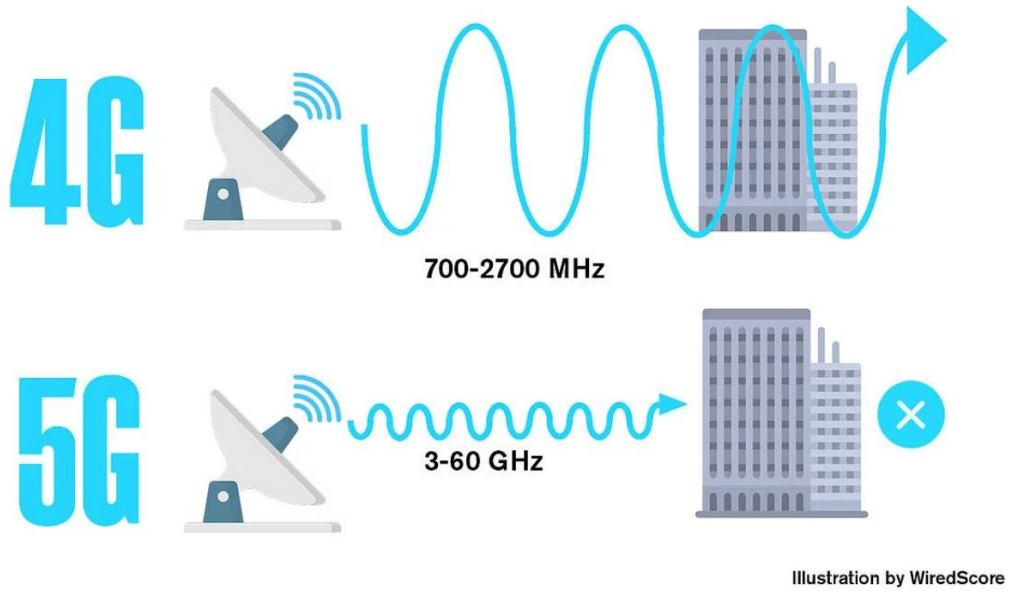

Health concerns?

Some people are concerned that increased use of 5G could lead to health risks, but the science doesn't back this up.

The electromagnetic waves generated by 5G are smaller than previous generations, small enough to the point that they aren't strong enough to break chemical bonds in cells. In fact, the waves can barely get through your skin.

It never hurts to do more research, but the current science does not suggest any serious risks from 5G.

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