Computer Science Future

15-110 – Wednesday 04/23

Quizlet9

Learning Goals

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

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 currency 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.

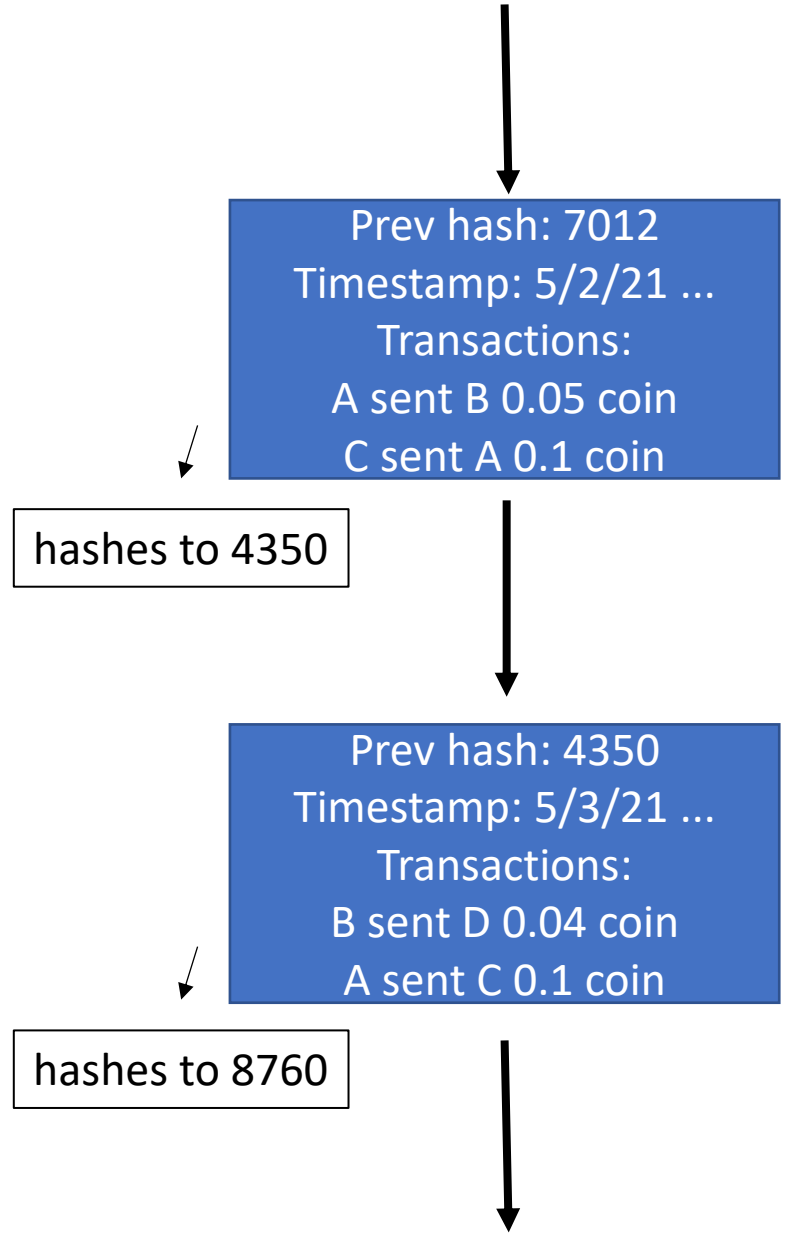

Bitcoin

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; a year ago it was worth $66k per bitcoin, and now it is $93k per bitcoin.

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

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.

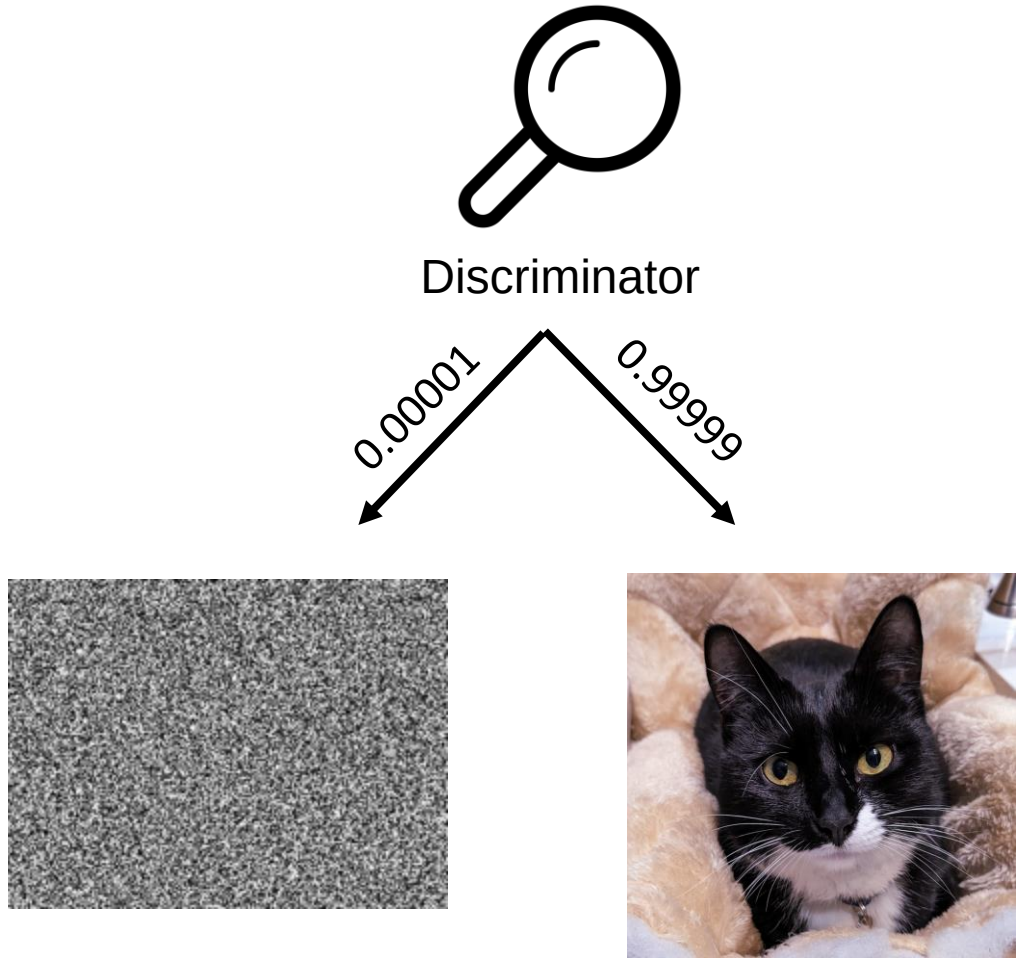

Generator

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

0.00005

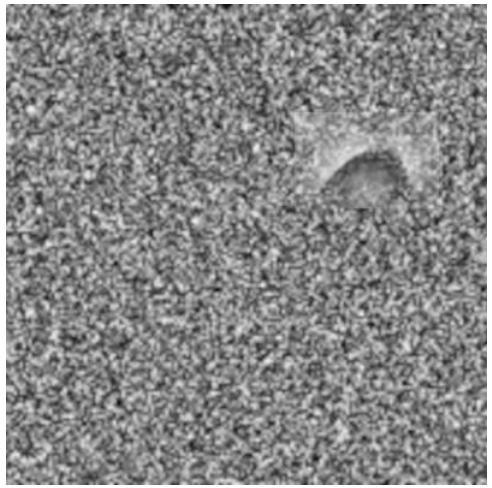

Adversarial

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.

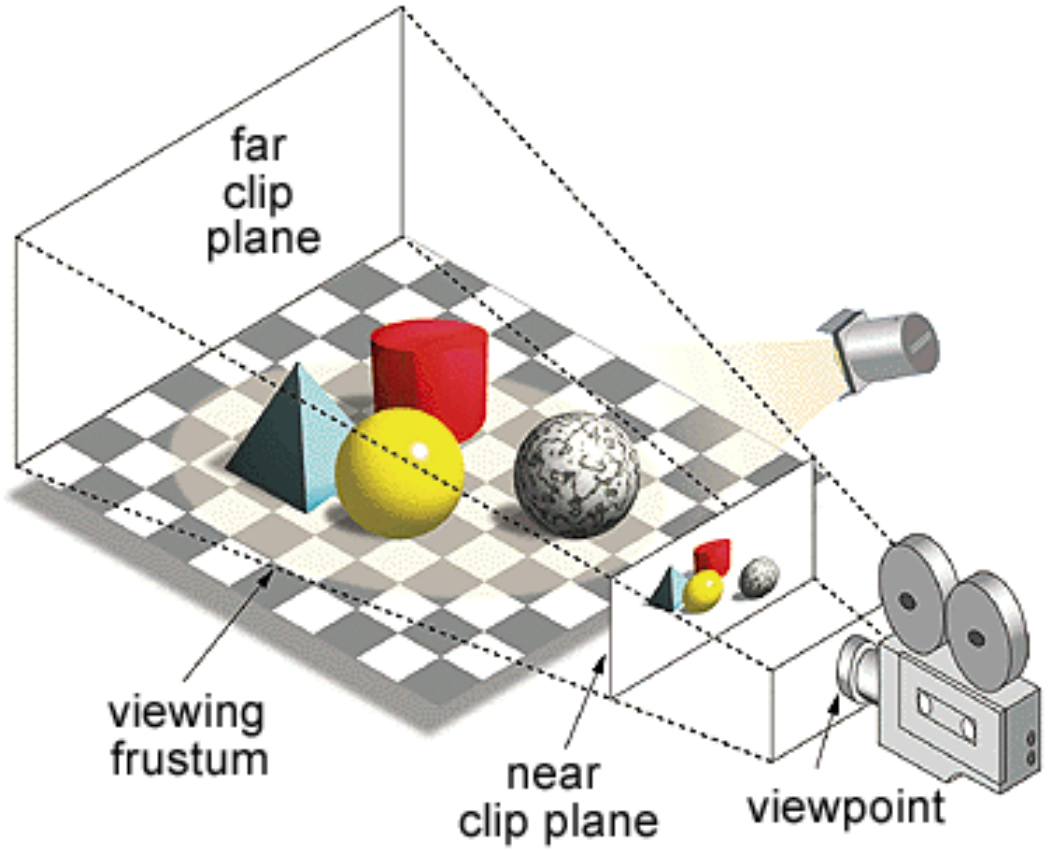

Virtual Reality Adds Complications

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.

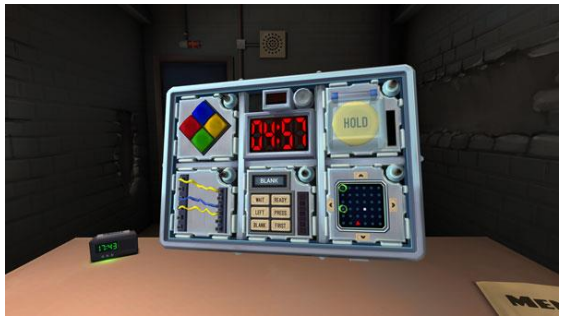

Recent developments have used hand tracking (with controllers) and head tracking to allow basic **interaction** with the VR environment.

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.

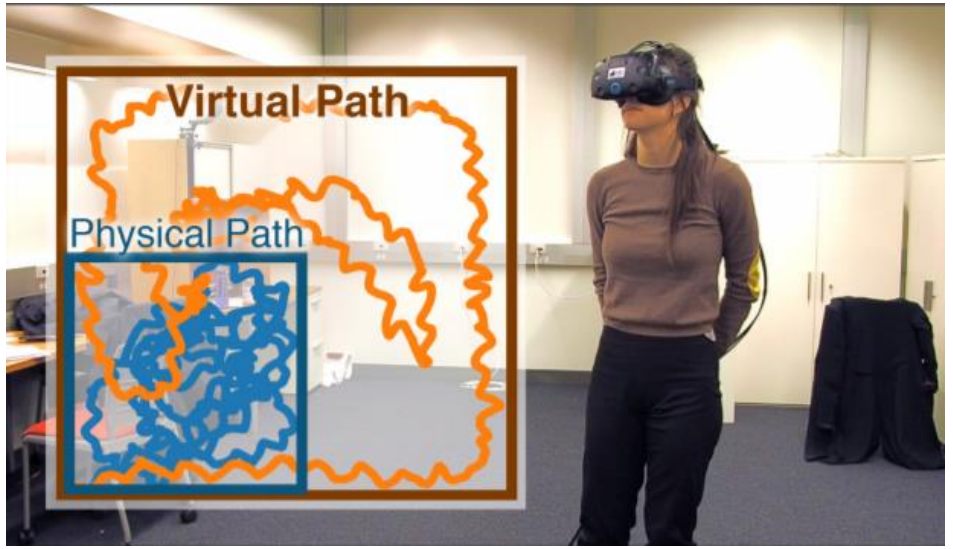

Sidebar: Augmented Reality

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.

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.

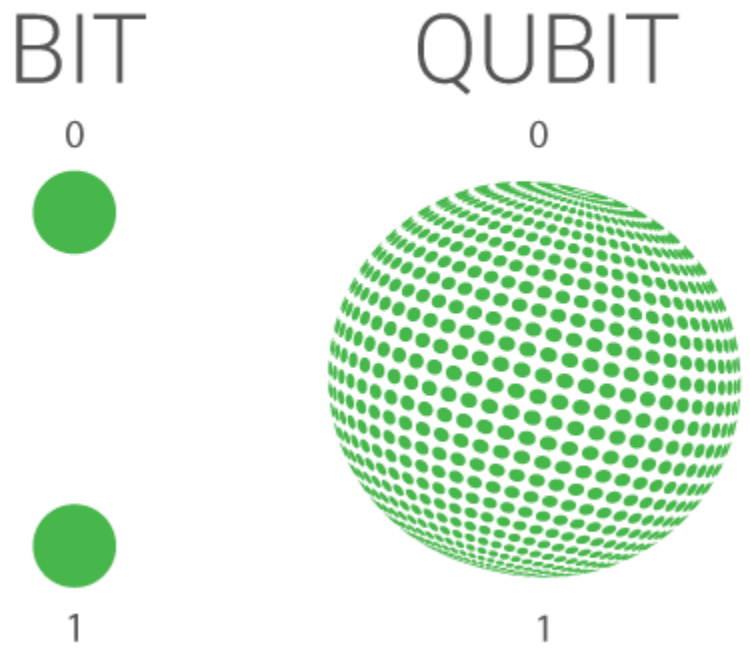

Entanglement

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!

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

Learning Goals

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

Other Future Topics

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

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. For example, consider this 'video' of Mark Zuckerberg, creator of Facebook:

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

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

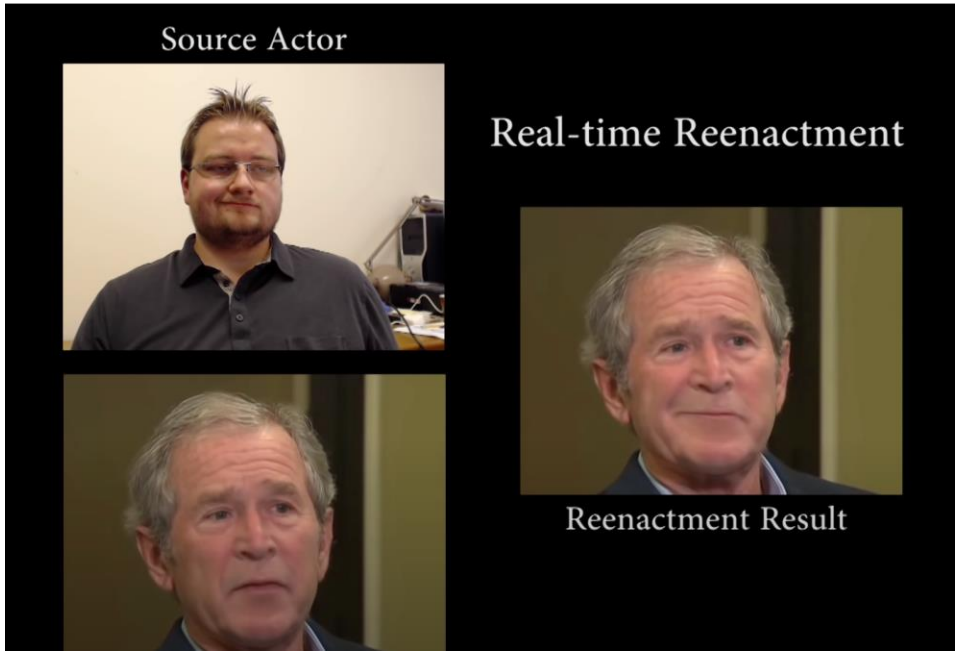

Intention vs. Potential Use

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