

[Mary] Hi everyone in-person, online, in the future. My name is Mary Ton. I'm the Digital Humanities librarian here at the University of Illinois.

It's my great pleasure to be joined by Rebecca Stover, and Sarah Christiansen and Cadence Cordell, for this presentation on Deepfakes: How Computers See and Mimic Us.

So today's slides are at go.illinois.edu/deepfakes, all lower case. This link is case sensitive, so that is go.illinois.edu/deepfakes. And the slides will be available after the workshop, so you can come back and refer to them.

Little bit of a - a spoiler alert. If you take a peek, you might see some things, but, um, you can refer back to those after the workshop is done.

So deepfakes. What I'm about to show you is a deepfake that was created at the University of Illinois, or an AI Avatar by our friends in CITL.

You're about to see a real Robert Brunner and an AI Avatar version of Robert Brunner.

And what we're going to do is, we're going to guess which one is the real one.

So I'm going to play you the video first, give you a link to the - the poll, and then come back to the video.

So I want you to - to watch the video and think, which one is the real one. And there's no sound with this, other than light music.

[light music plays over video]

Though, I am a little concerned because there should be some - some light music. Let me try...okay. All right.

So this is the poll. For those of you joining us online, there should be a link in the chat. I cannot see what you respond. So please feel free to answer.

The QR code for the folks in-person is here. So I'm going to give you a moment to, um...

Let's see if that Zoom panel goes away. No!

For our friends in person, if you want to use the QR code to participate in the quiz. It's there. Tada.

Alright. So looking for, who is the real Robert Brunner? And going back. Let's see. Alright. Here's the video one more time.

[light music plays over video]

Alright, so let's see what people are thinking. Okay! We have a strong sense of A, the person on the left. We'll come back to this particular recording at the end and we'll see if these responses change.

So our goals for today. We're going to describe how deepfakes work. We're going to consider current applications and implications of deepfakes. And then we're going to discuss how to spot a deepfake.

And if you need the link to the slides, it's there again, go.illinois.edu/deepfakes.

Okay. So the first bit is to talk about, how are deep fakes created, um, and how does that inflect the way that we spot those?

So what is a deepfake? Deepfake gets its - the 'deep' part of its name from the term 'deep learning', which is a method in computer science to study patterns in large quantities of image, sound, and video, and we'll talk about how a computer sees and hears multimedia in just a moment.

These tools, after they've studied large quantities of material, they apply what they've learned about how humans look and sound to create new arrangements of pixels that mimic the human face and voice.

So, we mentioned pixels. When we're talking about digital images, especially digital images of humans, we need to talk about the pixel, which is the basic building block of all digital images.

AI tools and deepfake tools, as a form of artificial intelligence, break down images into pixels, look for patterns in those pixels, and then use those patterns to generate new arrangements of pixels, AKA, a new image, or a new image of a - of a person that it has studied.

So when we're talking about that - that process, what it's really looking for as it - as it looks for patterns is, it's going to take the pixel, compare it to the ones surrounding it, and look for three things: red blue green color values, contrast for outlines, and contrast for texture.

Is it something that's solid like my shirt, or something that has a pattern? It's then going to use those patterns to identify segments.

So in this, um, image on the right, there are actually six segments: three for the bicyclists, three for the bicycles. And then it's going to classify those features.

I think that the - the thing that's in this box, based on the patterns that I'm seeing, is a human. I think the patterns that I'm seeing in this box is a bicycle, so on and so forth.

Which brings me to my absolute favorite, uh, problem in all of computer vision, the Chihuahua or Blueberry Muffin problem.

Computers have difficulty classifying and distinguishing the differences between these images of blueberry muffins and chihuahuas because, they have similar red, blue, green color values, they have similar outlines, and they also have similar textures.

They all look delicious.

If you've seen the animated version of the Mitchells versus the Machines, you've seen a version of this problem where the robots are having trouble classifying the family dog as, dog, pig, dog, pig, or loaf of bread. So you've seen this in action.

This inflects the ways that people use computers to study the human face and to make three dimensional models.

So it's not only something that is taken into account when people are building deepfakes, but things like when we are, um, generating images of hands.

So a computer understands what hands should look like, but not necessarily the underlying anatomy. It is a - It's trying to understand a three-dimensional object in two dimensions and failing miserably.

So when - generative AI tools have difficulty representing hands because they're treating something as a two-dimensional image rather than a three-dimensional object, which brings us

to techniques that are used to bring two dimensional images into a three-dimensional digital model.

So what you can do is take photos in 360 degrees around an object and then use computer vision to compare those photos and say, okay, if this is the same image from this angle, and then this is the same object from this angle, how can I stitch those two images together in a three-dimensional space?

This particular technique is called photogrammetry, and it's used frequently in the cultural heritage sector, museums taking photographs of their objects to create three-dimensional objects.

In the entertainment industry, you've seen this in action if you've been to the movies.

One of the breakthrough movies that developed this - this kind of computer graphic technique and this special effect was the Curious Case of Benjamin Button with - starring Brad Pitt.

The VXF team used Paul Ekman's facial action coding system to record 70 expressions, filmed those from multiple angles, then used those to create a digital model of Brad Pitt's head. And there's a link in the slides to hear more about this process.

In the early days of motion capture, they had to paint him this rather spooky color of fluorescence just so that it would be easier for the computer to track patterns and movements on his skin.

But this was able to create a far more detailed three-dimensional model of a human head in real time than what they had been able to do previously.

And so this particular technology has come a long way from the Curious Case of Benjamin Button.

If you saw Free Guy, you know that there's the - the moment where they digitally - uh, Ryan Reynolds has a beefy doppelganger in this movie, and it's hilarious.

So you see the actor on the left, who's playing the body double. You see Ryan Reynolds above recording the motion capture performance of his face. And then on the far - the bottom right, you see his head digitally put on top of the bodybuilder's.

If you've seen the most recent Indiana Jones movie, the Dial of Destiny, they use similar techniques to capture Harrison Ford's performance and digitally de-age him.

Okay, so that's a little bit about how - how people create, you know, um, representations of human faces in three dimensions and can account for some of that movement. But what about sound?

Surprisingly, sound is also an image recognition problem.

So to talk about how it's using image recognition, we need to talk about two specific features of sound first.

There's the frequency, the pitch. So something that has a high frequency has a high pitch.

[singing a high note] So it would be something up like here.

[voice dropping very low] And if it has a low pitch it's down here.

And then we need to take into account amplitude, intensity. Usually we tend to think about this in terms of loudness and softness. So if something is very soft, it will have a very low amplitude.

Okay, so this is a visual representation of the phrase "nineteenth century".

What it's doing is it's representing pitch, whether it's something - a sound that's high or low on the, um... the vertical axis. And then it's representing intensity through color.

So we have nineteenth century, You can see where the T's are coming into play, just for where the hotspots in the soundscape are.

And from there, it's a classic image recognition problem, using one pixel, comparing that to the surroundings, looking for patterns, and once those patterns are identified, then recreating those in a - a soundscape and using audio equipment to translate that digital signal into something that we can hear.

So image recognition for soundscapes can get pretty complicated and really exciting.

This is a - excuse me, a soundscape of Mount Rainier or rain-ier, depending if you live in Seattle and are constantly soggy.

But you can see how it can use image recognition techniques to identify the frequencies and amplitudes of things like insects, a marmot, a jet, a distant jet, birds, et cetera, et cetera.

So the audio equivalent of the AI image-generating hands and having trouble, um, is garbled audio.

So it can understand and - and replicate pretty well the sounds of things like instruments and human voices. But sometimes there's a little bit of noise or a gap between, um, what the human voice will actually sound like and what it's representing through that generated two-dimensional image.

So I'm going to play for you a song, um, that is generated using a tool probably like Suno. And the things I want you to listen for, it's catchy, but it's fairly generic. The voice sounded garbled, and the instruments are a little bit too accurate. They lack that kind of nuance that a human performer would have.

The sound - song starts off a little quieter. So hopefully - yeah, there it is.

[music playing]

[Mary, voice muffled] The voice is a little hidden in the instruments, but pay attention to when he goes up [indiscernable].

[music continues playing, includes garbled lyrics].

[Mary] I'm going to pause it there. [music stops]

And for those of you who are online, you can't see all the people in-person kind of squinting. That's not you having trouble hearing. That's actually the garbled audio.

So I will confess, I was totally fooled by this song when I heard it for the first time. I had YouTube music playing as I was driving around. I was listening to some music to inspire me for Dungeons and Dragons campaign. And this song came on. I'm like, "Man, that nails the vibe. What is this? Who is the artist? This is great. But there is a problem with my car speakers. It just sounds really weird."

And so, I was - curiosity got the best of me, and I was like, what does, what do the lyrics mean? This is really cool that it's in Old English. I want to know. So I'm Googling it.

And a Reddit thread told me, ah yes, this is from an account that cropped up once AI-generated song tools became popular. It's - it's probably generated by Suno.

And then listening back to the song, I was like, oh, that wasn't my, my car speakers. That was actually the song.

So I mentioned this story too, just to flag: we've all been tricked by AI. And one of the things that we're - we're trying to help you today is just to practice spotting things, 'cause we've all been fooled by it before.

Okay. So, use cases. There are some really exciting applications for deepfake technologies in the entertainment industry. It is being used to help tell a more vibrant range of stories that weren't possible before.

I think one of the most exciting uses of generative AI is to - and of deepfakes, is to make the lip synching process more seamless.

So I'm going to play you a clip from a movie. I will say that there is an expletive, but it's been beeped over. And it's going to walk you through how they match the movement of the lips to the dub as they are taking it from an R-rated movie to a PG - PG-13 rated movie, and then into other languages.

[video clip] Now, we're stuck on the stupid [bleep] tower in the middle of [bleep] nowhere. And I don't blame you, and now we're stuck on this stupid...this stupid freaking tower in the middle of freaking nowhere. And it's all my fault.

[Speaks in Japanese, followed by Spanish]

[Mary] Pretty awesome. Pretty awesome. So, gone are the days of someone's lips moving in and out of sync when we're - we're watching a movie in translation that's been dubbed.

[video clip] Now, we're stuck on this-

[Mary] Oops. Okay. The other exciting thing is the way that AI Avatars, or authorized deepfake versions of ourselves can help us communicate.

So there's a tool called HeyGen that will help you create your own AI Avatar, and you can add text and generate an AI avatar based - or, um, generate a video based on that text.

What I find is really cool about this tool is that it can help you speak a foreign language that you probably - you may or may not know how to speak yourself.

It can also be a useful way to maintain anonymity. So if you feel uncomfortable being on camera, that this is an opportunity to use an AI avatar to convey your - your ideas in a way that you feel more comfortable.

Now, while these applications for entertainment and for communication are really, really exciting, unfortunately, they are a relatively tiny, tiny, tiny, tiny fraction of how deepfakes are being used.

So we've seen instances of generated images being used for political motivations. If you've been following Taylor Swift, you know that she came out with a statement about who she was voting for because there had been an AI image, a deepfake of her circulating, and she posted her support for Harris because she wanted to be clear that it was the real her who supported a particular candidate.

We've also seen Russia use deepfakes of Ukrainian President Zelensky, to um, to advance Russian propaganda.

[short pause]

Of all the statistics related to deepfakes, this is the most unsettling and terrifying for me.

So it's estimated that 98% of usage of deepfakes is to generate erotic content of people mostly without their consent.

And it's really harrowing to see women who are Twitch streamers and their images being used without their consent um, for purposes without their consent and the - the harm and trauma that that's enacting.

To talk about how the legal landscape is responding to the rise of deepfakes and to this particular situation, I'm going to turn it - things over to Rebecca.

So excuse us while we switch the mic for just a moment.

[long silence during mic switch]

[Rebecca] Can you hear me?

[Mary] Yep!

[Rebecca] Okay. So the legal landscape.

[short pause]

[Rebecca] Um, just something that I'll quickly say is that next week, there is a Copyright, Text Mining, and AI [workshop].

So we're going to address text mining and AI, like, a little bit, but it'll mostly be addressed next week if you want to go to that Savvy Researcher workshop, and we're mostly now going to be concentrating on the legal landscape of deepfakes in particular.

So...one kind of, like, framework of thinking about how AI in general is able to be created, generative AI, is, um, fair use, which allows... [pause to adjust mic]

Okay, um...Oh, okay. I'll just hold it up here.

Uh, so fair use allows for a limited use of a copyright material without the permission from a copyright holder.

Courts have generally found that text and data mining are quintessential fair uses as AI takes this, um, body of data and then transforms it.

So, transformative use is an element of fair use. And because AI transforms the work, it is considered a fair use. However, we're not entirely sure if training AI is going to totally be fair use, or it's an ongoing question of how legal it is, or I guess, will be.

So, in the case of the New York Times versus Microsoft et al, which is a currently ongoing case, the New York Times is suing Open AI and Microsoft for using New York Times' data, New York Times' articles, et cetera, both without their permission and also claiming that it infringes on the New York Times' ability to sell newspapers, to sell their online apps, et cetera.

And that's an ongoing case, um, that will be decided at some point, presumably within the next few years.

Another ongoing case is Sarah Andersen et al versus Stability AI.

So it's a group of artists who are suing the Stability AI, um, tool. So there is this AI tool called Stable Diffusion where you can put in a text prompt and that creates an image for you.

This group of artists contends that the data that went into creating this tool was mostly taken from artists without their permission and is infringing on their copyright.

It is also an ongoing case, but most recently, the court declined to dismiss these infringement claims and in fact, moved forward all of their copyright claims, suggesting that they think that this is a valid case that should move forward.

And once that case moves forward and is hopefully eventually settled, we'll have a better idea of what the legal ruling on, um, using this AI set is going to be.

Can you use this video to train AI? Maybe. We're actually not sure yet. So we might think, yes, but there also might be other measures preventing you from using this text or this video.

So examples such as terms of service, licensing agreement, et cetera.

If you use the platform DeviantArt, it says in one of their Terms of Services that you can use no AI, and that won't be included in any kind of mining set. And DeviantArt is actually one of the, um, companies, wings, that is actually included in that lawsuit that I just mentioned of artists suing the Stability - Stable Diffusion platform.

Another example is Terms of Services.

So under TikTok's Terms of Services, it says that, "You or the owner of your user content still own the copyright in the User Content sent to us," so you still own, kind of, the official copyright.

However, "by submitting the User Content via the Services" - by posting your TikTok - "you hereby grant us an unconditional, irrevocable, non-exclusive, royalty-free, fully transferable, perpetual, worldwide license to use, modify, adapt, reproduce, make derivative works of, publish and/or transmit, and/or distribute, and to authorize other users of the Services and other third parties to view, access, use, download, et cetera, et cetera, make derivative works, et cetera. On any platform, either now known or hereafter invented."

So, that's, um...a pretty infinite amount of use of what happens to your TikTok after you post it.

Um, so looking at these kind of terms of services, especially when thinking about how these things might be used is, I think, super important and also, this scares me of TikTok personally.

So this is going to be, like, a few second video of a voice artist whose voice ended up being used for TikTok's Text to Speech app without her permission. And just as an example, we're going to watch it real quick.

[woman in video] Anywhere you hear a voice and don't see a face. It's kind of what I do.

[voiceover] And now it seems she does a lot of this as well.

[TikTok] Hi, TikTok. My name is Misty. How text message-

[Rebecca] Yes. So that's an example of TikTok using this woman, Bev Stan-

[video] Anywhere you hear-

[Rebecca] Sorry.

[video continuing] -voice, anywhere you hear a voice and don't see a face-

[Rebecca] There you go. So in Bev Standing versus TikTok, so she made recordings, um, for the Chinese Institute of Acoustics Research for a Text to Speech tool, and this was used by TikTok without her consent and without compensation.

She sued them for damage to her brand, and this was settled out of court, which means that there was no legal ruling made out of this case. They just probably gave her a hefty lawsuit, and they now use a different Text to Speech voice, presumably with someone that they paid, we hope.

Um, but this is just a case - one such case of how these kind of things are settled. Um...

So currently, there isn't any federal law on deepfakes. Most recently, in July of this year, the Office of Copyright published, um, a report, part one of this bigger, um, thing. But part one, which was about digital replicas, so –

And the Office recommends that Congress enact a new federal law to protect individuals from the knowing distribution of unauthorized digital replicas.

It's a really fascinating report. They go through the different instances where, um, they see it's been harmful to people and go through what the current legal landscape looks like and how Congress could enact new laws on how to protect individuals' images.

Currently, in the House of Representatives, there is a No AI Fraud Act, um, which would provide for individual property rights to likenesses in voice.

This was introduced in January of this year, and it is bipartisan sponsored. I believe it is currently still under review, but we'll see as it moves forward, what happens.

Currently, as you may know, some states have moved forward with Deepfake legislature, but right now it's a state by state process.

So, this is all to say that the legal landscape is ongoing and ever changing. And if you come back, probably, if we give this talk again in the spring, it may have changed since then as well.

So protecting your image and voice.

[muffled voices during mic switch]

[Online participant] Do you mind if I listen to this? Okay.

[Sarah Christensen, laughing] It's like, no, not at all. Um...okay. Can you hear me okay? Bump it up a little bit. I think on backwards... Okay. So if you're...I'll just hold it.

If you're a content creator like Bev, there are ways to protect your image and voice.

So a few things out of University of Chicago. They have a Glaze tool, and now they have a web Glaze tool, so you can use it online without having to download anything, that introduces noise to your image so that it will disrupt the diffusion process and not compromise the appearance to... Yes. [laughs]

Um, so it won't compromise the appearance to a human viewer. But when AI reads it, it's like, wow, what's that?

Um, nightshade is a similar thing. The difference is more, Glaze is more of a defensive thing. Nightshade is more of an offensive thing, where it will like, kind of like, poison the image. So when AI goes to create something, it's all - it's all janky and messed up.

So we have - I can walk around here, right? We have over here some examples of what Glaze will do.

So this is original art pieces by Karla Ortiz up here. This is without Glaze, what AI has created from it.

[muffled] I'm just going to - I'm not good at microphones. And then, can you still hear me okay?

Okay. And then, this is what AI has generated with the - the Glaze on it. So you can see it, like, doesn't really look anything like her original art, which is really cool.

And then Antifake is a Python script that will introduce noise into audio files to prevent deepfakes. So if you're an aspiring voice actor, that's something that you should look into.

Um, okay. So how do spot a deepfake? Okay. I added - I [laughs]. I added this just for a friend who will remain nameless. I send stuff like this to her all the time, and she's like, "What are you doing? This is AI."

And I'm like, "Let me enjoy things. Okay." Like, I want this giant cow lounge. Totally completely realistic. I think.

Um, so we've been talking about deepfakes, but I also wanted to introduce the word "cheapfake".

So this is a whole spectrum that Britt Paris and Joan Donovan came up with. And a few, I - I linked to the report and the - the graph online.

Um, so we can see a lot of what Mary talked about, real, like, really fancy Hollywood, like, not super accessible tools that the average person can use, but we do see a lot of suspect things on social media nowadays, you know, like ten fingers and really warps anatomy, things like that.

Um, but a lot of it is also what she calls cheapfakes, what they call cheapfakes.

So, some examples of those are like recontextualizing things, like taking an image from something completely out of context, and then, like, adding a caption that's like, "Eugh, look what's happening here", you know.

And we can see just like speeding and slowing can really affect how we look at videos and the whole context behind that.

These are all pretty accessible tools that we can all use, and it kind of, like, goes up over here.

Here we have, like, you know, the Snapchat filters and things like that. Someone did this, like, rotoscoping, which is like tracing over action to- I mean, it takes a while, but you can do it with things like Adobe and whatnot.

And then these are more, like, the deepfakes with the more - more expertise and technical resources are required to create those, um, which, you know, there's still obviously a problem, and they're going to get better and better and easier to - to use. And these are - these are just a few examples that I found.

The first was this - Hurricane Helene, which is super recently. In Florida, this - this video was actually from, like, a storm in Iowa in like, 2018 or something like that. Um, you can see they've dated it September 25.

Um, and over here, this is something that went around on WhatsApp. It was actually a - a study that this company was doing to see how people would react in natural disasters, or - not a natural disaster. Just disaster in general.

So they - they created this whole video and they put it on YouTube, and they, like, said it was fake all over it, but then someone downloaded it and put it on WhatsApp, which can add because WhatsApp is like end to end encryption. It can add a little bit of authenticity to it, or like Facebook, it's, you know, we're just sharing things, and Whatsapp seems like slightly more credible somehow.

Um, but BBC had to, like, issue a report, and like, they had to take this video down, and it was - it was a whole thing.

So this is a - a really fun video, which I would encourage you all to watch, "In the Event of Moon Disaster". And MIT created this and it goes into how deepfakes are created and a little more about AI, and, you know, they - they're using a few different techniques in this, so you can go through and try to spot all of those, and then read more about it on the website.

Where is, like, half of...? Oh. [laughs] I was like, why is it not? Um, there we go.

Okay. So, these are - and you should be using all of these when you're looking at what, um, trying to spot deepfakes.

So there are some visual cues, um, in the Robert Brunner video earlier. You know, we could look at gestures, hands, anatomy, um, background, camera movement, green screen artifacts. Those are hard to remove. Ask me how I know. Um...

Also for audio, like, there's sound cues, lack of sound, voice, tone, the amplitude.

Metadata can be tricky. Usually, you know, photographers, you'll have, like a photo fingerprint in it, and you can extract metadata using, um, tools like EXIF, which is freely available online.

It doesn't - it can be hit or miss if you're just downloading things from the Internet, like, your download date will often be like the date created? It will usually tell you the source, but if you downloaded it yourself, like, you know, you can see it's from Facebook. But, um, using more advanced tools like EXIF can help you put all these clues together.

So the content, we talk about this a lot with like phishing scams, too, the - the tone of how it's written doesn't match like what you would think, like for example, the BBC would put out. Is their layout the same, is it, um, the same, uh...authenticity, I guess I'll say.

Um, and you can also, you know, kind of fact check and corroborate, like, is the date the same as what they say it is? Are they citing any sources? You know, try to find different news sources to say, oh, okay, I guess this might be true.

And then there's a lot of forensic analysis that's happening now too, and there's a lot of, um, tools online that you can just use from home, which is really great. But a few different things that they'll look at is, like, the compression of JPEG specifically. When you edit them, they - they get like, compressed a second time, a third time, a fourth time.

Um, localization kind of goes over the pixels that Mary was talking about more in depth to look at patterns and how they might be kind of messed up. If you - if you've done Photoshop a lot, you can see there are usually traces of things left behind from like, clone stamp or blurred or like, just not quite the right hue. A lot of that.

And then for audio, you can, like, feed it through, like, a forensics thing. Sounds really official, "forensics thing". For oscillation signals, and they call this, like, the Gibbs effect and I don't know if that's Barry Gibbs or, um, a forensic scientist named Gibbs. I want to say Barry.

Um, there's also the SILL method, which is, you know, stop, investigate, look and listen. Like, the - yeah, yeah. Count me and Mary.

So just kind of, you know, without - there's different - different approaches to identifying deep fakes.

A lot of companies and people are favoring the technology side and developing all these tools where you can just, like, run it through a program and be like, oh, no, it's a deep fake.

But you can also look for yourself, stop, investigate the source, like that, um, the cow couch from earlier is from an account called, like, inspirational design or something, which, uh, it kind of implies it's not real.

Um, but also looking for artifacts and unnatural movement, just doing kind of a deep dive yourself.

And then, um, also listening for garbled sound, which I - you know, even though you're telling me that song is fake, I still - I'm like, it's just old English.

Also, like, TinEye is a really great reverse image search that you can use to identify the original source of something. Um, Google Lens, and there's a few others, like Amazon has one now too. And then there's one that begins with a Y that's also new that I forget. Okay.

[Long silence during mic swap]

[Mary] One last sound check. So hopefully the folks online can still hear me. If not, then I'll see some waving from Cadence in just a moment.

But we're going to revisit the video of Robert Brunner, and the things that I want you to keep in mind. So these are general guidelines that are not necessarily always true, and there's - especially when we're working with industrial strength deepfake technology. Some of these, uh, techniques can be accounted for - or these - these features can be accounted for.

But generally speaking, when there's an AI image, start being suspicious if it has a very simple background. If there's no trace of a green screen.

If someone has very limited hand gestures and - and the hands aren't visible, they're kind of below the podium, be suspicious.

If they have a fairly limited voice range and don't really have much intonation, that can also be a sign of an AI-generated audio. Um, as well as garbled audio, that's created as part of the image generation process - or the sound generation process.

AI images also tend to be - deepfakes tend to be a little too perfect.

So things that are possible indications of a real live human, it might have a complex background.

Again, industry and the entertainment industry has some tech - some pretty sophisticated technologies to overcome this, but, um, human-authored or, like, real human content, very complicated background in - and camera movements, because that's really difficult to deep fake.

Green screen flickers. So looking at the glasses, a head, or if someone has the crooks of their arm, you can see little bits and pieces where the - the video editing tool is having trouble distinguishing between the figure and the background behind them.

Hand movements. I like to talk with my hands, and it's in part an effort to show that I am human. Range in tone and voice, clean audio, blurred movement.

So these are things, again, generally speaking, be suspicious when you see the things that are under the AI column.

The - the things on the right hand side can be indicators that it's a genuine live human.

So with that, I'm going to put on the QR code. I want you to have that poll queued up. There should be, if - if I push the buttons correctly, it all should be working.

So take a moment to - to go back to that quiz, and I want you to be thinking about those things that we just talked about as you watch the video again.

Alright, I'm gonna play the video one more time, and then I'd like for you to cast your vote.

[Jazzy music plays during video]

[Mary] So things like green screen artifacting, probably human. Lots of gestures, probably human. AI, it's a little bit too clean.

So if you haven't already, please cast your vote at this poll. Who is the real Robert Brunner? Is it Person A or Person B?

And I'm going to go to our poll. Hopefully everything has worked.

So we have ten votes, and we have more folks who think that B, the Robert Brunner on the right, is actually the real Robert Brunner.

So here's the reveal. I'll say this has - the reveal, this video has a different real human Robert Brunner, but the AI avatar is the same.

[Robert Brunner] Hello, fellow disruptors. I'm Professor Robert Brunner, and it's my pleasure to -

[Mary] Oops.

[Robert] Hello, fellow disruptors. I'm Professor Robert Brunner, and it's my pleasure to welcome you as a real person to our groundbreaking course on navigating the disruptive landscape of emerging technologies here at the Gies College of Business.

[AI Avatar] Hello, fellow disruptors. I'm Professor Robert Brunner, and it's my pleasure to welcome you as an AI generated avatar to our groundbreaking MOOC on navigating the disruptive landscape of emerging technologies here at the Gies College of Business.

[Mary] So, the video was originally created for a massive online course or MOOC, um, that's through Coursera. You can take the course on disruption theory. But the AI avatar is the one on the left with a pin.

So when I first saw this video, I was - I was thrown off because I thought, Hey, that's a really nice detail. It has a U of I pin. And who's this guy on the right? With no pin, no cred!

But as you look at the - the video on the right a little bit more closely, you start to see the telltale signs of a human.

So if you look at his glasses, you'll see that there's some flickering from the green screen, also a little bit of flickering in the crook of his arm, and the figure on the right is much more animated, tends to use his hands. His hands are visible for parts of the video, whereas the - the AI avatar on the left has a very limited range of movement.

But don't be discouraged if you were fooled, we all were fooled, and it takes practice.

So looking at the - the "In the Event of Moon Disaster" is a great way to practice.

I'm going to speed through a couple of these resources before turning off the recording and heading into Q&A. So resources.

We have, um, future Savvy Researcher workshop series. Those will always be on the Savvy Researcher calendar, go.illinois.edu/sr.

All of our previous workshops on AI have - are posted to the Digital Humanities channel on Mediaspace. These include things like a closer look at how image - images are generated, and also how does AI generate text.

The library has also created a generative AI LibGuide. This has general guidelines for copyright, best practices, et cetera, et cetera.

I'd also like to highlight if you're interested in playing with generative AI and creating your own AI avatar, CITL's Innovation Studio has access to HeyGen. They have open walk-in hours every Monday and Tuesday from 12-4.

They also have subscription access to Midjourney, which is a really popular image gener - image generation tool, as well as, um, things like ChatGPT.

As an institution, we also have access to Adobe Firefly, which is an image generation tool. You can sign in using your @illinois.edu email address.

Before I, uh, turn things over to the Q&A, I just want to emphasize that we are always glad to field questions. If we don't get to your question today, please feel free to shoot us an email. We're happy to - to answer those as they arise, and we deeply appreciate you joining us today. Thank you.