Introduction to GenAI and Chatbots for Healthcare Professionals

Presenter: Leigha Jansen, EdD, CPACC; Carolyn Smaka, AuD

- Hello and welcome to Introduction to GenAI and Chatbots for Healthcare Professionals. My name is Dr. Leigha Jansen and I'm the Director of Educational Technology at Continued.

- My name is Dr. Carolyn Smaka, and I'm editor-in-Chief at Continued.

- We are your course presenters today and we're excited to talk about a topic we are both very, very passionate about. So let's jump right in. But first and foremost here are our disclosures, presenter disclosures, content disclosures, and sponsor disclosures. Our learning outcomes today are as follows. After this course, you should be able to define AI and list current AI applications in healthcare. You should also be able to list uses of generative AI to increase the efficiency and or creativity of professional tasks. Also, you should be able to generate effective prompts to engage in meaningful and productive AI-driven conversations. And finally, you'll be able to explain the limitations of generative AI and strategies to mitigate risk.

This field is evolving quickly, so don't be surprised if some of this information has changed by the time you watch this course. We've tried to present the topics in a way general enough to give you a great foundation, even as some of the details change. And as an introductory course, we've equipped you with a comprehensive list of some of our favorite resources so that you can continue exploring and learning on your own as you'd like. You will find all of these resources in your handout. So disclaimer, generative AI is a rapidly evolving topic and everything is subject to change. I'm gonna turn it over to Carolyn to kick us off.

- Thank you, Leigha. Well, we'll start with a definition. Artificial intelligence or AI as we'll refer to it in this presentation is the development of computer systems that are able to perform tasks that typically require human intelligence. So tasks like pattern recognition, natural language processing, decision making, language translation and

more. And the term AI is not new. It was actually coined back in 1956 by computer scientist John McCarthy at the Dartmouth Conference, which is considered the seminal event that launched artificial intelligence as a field of study. One really important thing we want you to take away from today's course is that the key to AI is data. AI uses various techniques including machine learning and deep learning to enable machines, in most cases, we mean computers, learn from data and improve their performance over time.

The more data an AI system has access to, the more effective it is. Now, you are probably already using AI in your life. Virtual assistants like Siri and Alexa are all powered by AI. They use natural language processing, predictive text. So if you've ever been typing a text and your smartphone is suggesting the next words, those suggestions are generated by AI that has analyzed and learned from millions of common phrases and word patterns. Email spam filters have been trained on data, what is spam, what isn't spam? And they've learned to detect and block spam. Social media platforms like Facebook and Instagram use AI to personalize your newsfeed and recommend products, too many of which I've purchased from those custom newsfeeds and GPS systems like Google Maps and Waze use AI.

They predict future traffic based on all the data they have from historical traffic patterns. And on the right hand of your slide, you'll notice another use of AI. So I was recently holiday shopping on Amazon and I came across this waffle maker that, believe it or not, had more than 238,000 reviews. Now obviously I don't have time to go through all those, so if you notice there's a little arrow there that says this is an AI generated from the text of customer reviews. So AI actually went through all that data. There was over 238,000 reviews and summarized here. So I have an idea whether I should buy this waffle maker or not. Pretty cool. So for those who wanna delve in deeper with regard to exactly how AI works, we have in your resources two complimentary and insightful Coursera course links.

Many of the AI courses out there you'll find are primarily designed for IT and computer science professionals. But these courses that we've given you the links to are accessible, they offer a solid understanding of AI. And again, you can find them in the resources. But we did wanna go into these terms because most of the course today when I turned things over to Leigha in a bit, is gonna cover generative artificial intelligence or GenAI. So we wanted to put that in context. So if you look on the slide here, machine learning is the broadest category encompassing a range of technologies for machines, in most cases, computers to learn from data. Deep learning is then a subset of machine learning, which uses something called artificial neural networks.

This is the infrastructure of deep learning and it mimics the brain's neural networks to process and learn from complex data. Then we have GenAI. GenAI leverages deep learning in order to create new outputs based on learned patterns from the data. And ChatGPT is an example of GenAI and Leigha will cover more of that later. But again, I wanted to put it in the context of of the bigger AI. So I said earlier that AI is not new. The term was actually coined back in 1956. So you might be thinking, well, why all of a sudden in the last couple years has it seemed like AI exploded? And there's three main reasons for that. The first is that is because of data.

So back in 1956, we didn't have the internet, but now huge data sets are available because we live our lives online in this digital age. These data sets are literally available at our fingertips. First reason. Second reason, innovations in chip technology. So the computational power inefficiency of these chips today, the is just grown tremendously in recent years and it enables the processing of vast amounts of data that's needed for these complex AI calculations. And then third, in November, 2022, ChatGPT was launched. It took ChatGPT, just five days to reach 1 million users and only two months to reach a hundred million users. So ChatGPT is one example of a GenAI chatbot. It's one of the fastest growing applications ever.

And now that it's in wide use, people are really starting to understand the enormous potential of AI and GenAI. So on the next slide here, we have put into context this explosive growth of ChatGPT as compared to other popular technologies. So if you look, you'll see Uber took almost 70 months to reach a hundred million users. And then if you look over, you'll see TikTok almost 10 months and ChatGPT, again, only two months to reach a hundred million users. Pretty amazing. So now let's move on over to discuss AI and healthcare. In healthcare, we have access to large data sets to some of which you see listed here. The electronic health record or electronic medical record represents one of the largest and most diverse data sets in healthcare.

So for example, my son was born in 2000. His whole life has been captured in his EHR, from the time he was born, his medical history, any diagnoses, any medications he's been on in the last 23 years, treatment plans, immunization records, allergies, radiology, images, lab test results, et cetera. These massive and complex data sets in healthcare provide the provide the foundation for AI. So now let's get into some specific applications of AI in healthcare. And this is by no means an exhaustive list or literature view. These are just some things that we pulled out that we find really interesting. So here is an application that applies both to physical therapy and audiology. The Mayo Clinic is developing an AI powered system to triage patients who complain of dizziness to get them to the most appropriate specialists and to support both outcomes.

And also in your resources is a brand new article that just came out in a PTA magazine about potential uses of AI in physical therapy. And how some physical therapy companies are currently developing AI solutions to assist with documentation. I'm hearing all your PTs groan out there, documentation, and help integrate clinical practice guidelines into treatment plans. So please check out that article. In terms of speech language pathology, there's a company called AI4ExceptionalEd. This is a team of

researchers, special educators, AI technologists and others who are developing AI solutions for school-based speech language pathologists, for things like universal speech and language screenings, to streamline treatment and to help customize SLP interventions for students. And you can check out more about that in the ASHA leader.

In terms of behavioral health, if you go to PubMed and you search the terms artificial intelligence and behavioral health, you will find more than 10,000 PubMed entries. There are just a lot of uses for AI in behavioral health. So I've highlighted two on the slides here. The first one is a fascinating project at Cedars-Sinai, where they're using virtual reality goggles plus an AI robot named Maya who's delivering AI expert psychotherapy. In your handouts in your resources, we've include a link to a 10 minute YouTube video where you can hear Dr. Brendan Spiegel, who directs the virtual medicine program there talking all about this project. How it works, you'll see videos of patients using it. You'll see Maya, and it's really fascinating.

Check that out. And then at the bottom of this slide here, I have a quote from an article by Graham and colleagues. The articles entitled Artificial Intelligence from Mental Health and Mental Illness, and kind of paraphrasing here, but to summarize that article, these authors say AI is poised to analyze and provide insights on the huge amounts of data from varying sources to help identify trends in mental illness and mental health and to develop better outcomes. Mental health is so multifactorial as you behavioral health professionals know, there's so much data out there from so many diverse sources. And AI is a perfect tool to help us really mine that data for better patient outcomes. And then in medicine, I highlighted this.

This study's been in the news quite a bit lately. There's was a study out of Sweden that shows the promise of AI to help read mammograms. And this study concluded that using AI and one radiologist was superior at detecting cancer in mammograms as compared to two experienced radiologists. And that's because AI and the radiologists

are looking at different things, so they compliment each other. Really amazing work being done in that area. So again, this was not a literature review, but rather just a few highlights of how AI is transforming the work we do as healthcare professionals. There's many other ways in which AI is being applied to healthcare and the other industries that touch us. And here at the bottom of this slide, this statement is adapted from the summary of a recent World Economic Forum article, where they interviewed tech pioneers on the subject and they said AI is poised to make healthcare more accurate, accessible, and affordable.

And the point here is not that AI will take your job. So the quote on this slide comes from the global head of medical Technology at Johnson & Johnson. "The most sophisticated thing in the clinical workflow is the human brain." In a recent article about AI and PT that I referenced a little bit ago, Ryan Eater, the CEO of Include Health, put it this way. He says, "While AI can augment and supercharge the capabilities of a physical therapist, it is not a replacement for the specialized knowledge and human touch of a trained medical professional." So this is a call to action. As you can see from these examples, healthcare professionals are already involved in AI projects. We wanna empower you as a healthcare professional to be part of AI conversations, to have a seat at the table when it comes to implementing AI in your workplace and in your profession.

The best way to do that is to get informed and perhaps the best way to get informed is to get hands on experience with AI tools, specifically GenAI, as it's extremely easy to use and effective. And you can't learn to drive and park. So let's shift into drive as we move from the big picture AI to the practical, and I'll turn it over to Leigha.

- Thanks, Carolyn. With new technologies, it's normal, it's natural to be skeptical of new technologies. As humans, we know that skepticism and even some fear and anxiety is a natural response to the unknown. Just a few headlines of late from the

"New York Times" include "The race to dominate AI." "Inside open AI's crisis over the future of artificial intelligence." "One year of ChatGPT, how AI changed Silicon Valley forever." Our goal is to take everyone beyond the headlines and increase your practical knowledge regarding AI, specifically generative AI. Knowledge and experience are the best tools to reduce fear and anxiety around AI, but maintain a healthy skepticism regarding the future. So let's dig in because using AI will teach you how to use AI.

As Carolyn mentioned earlier, generative AI is a subset of artificial intelligence. Generative AI uses massive data sets like the healthcare example, neural networks or banks of computers that mimic the neural structure of the brain and pattern identification and prediction to create new media, which includes text, images, audio, and even video. The heart of many of these tools is a large language model, also referred to as an LLM into which we'll dive deeper shortly. Generative AI tools include ChatGPT, which celebrated a one year release anniversary on November 30th, 2023. ChatGPT has led to the creation of a $207 billion, yes billion dollar industry according to Tech Republic. Generative AI is often collectively referred to as chatbots, computer programs that are designed to simulate conversation with a human user.

They're designed to interact like humans and have conversations. Of note, the images on this slide were created using chatbots. The top one was created using DALL-E, and the bottom one was creating Duet AI for Google Slides. Our goal is to teach you how to use these tools in a very general sense so that you're comfortable getting in there and playing with these tools. But before we do that, let's talk a little bit about the large language model or the LLM. Large language models are highly technical and purely beyond the scope of my understanding, but we can unpack the basics of how an LLM functions in order to better understand how a chatbot works. As Carolyn mentioned, it's all about data sets and data collection.

Generative AI is fed tons of examples like photos, articles, or music to learn from. Just remember that more high quality data equates to a smarter AI, think the entire internet. The design of the model is also important. The model has a structure of different layers to process all the data. Typically, GenAI is a deep learning model that utilizes those neural networks that mimic the neural activity and setup of the brain. These models are responsible for finding patterns. So the model analyzes the examples to detect patterns, features and relationships in unstructured data, all on its own. Natural language processing or an NLP is a common component. Often there's specialized training and that's when the model is trained further with labeled data for a specific task, such as generating text, generating images, audio or video, et cetera.

There are even training for chatbots to learn computer code. Finally, there's the output. After training, the model can generate new images, text, et cetera, similar to what it learned from. So although an LLM is a complex mathematical system, it can learn skills by analyzing vast amounts of data. Think digital books, Wikipedia articles, online bulletin boards, basically all things on the internet. It identifies patterns and all that text, and the LLM learns to generate text or language all on its own. So yes, it can write term papers, generate computer code, or even carry on a conversation. To summarize, think about an LLM as taking the entire internet and tossing it into a bowl of alphabet soup and using that DF data to create a term paper on the critical elements of persuasion.

Generative AI can create new realistic outputs after learning the essence of huge data sets without needing every example labeled. So what can generative AI do for you? Think of generative AI as a thought partner. Generative AI can be effective as a writing partner, as a web search partner, a creativity partner. But like any partner, generative AI has its strengths and its weaknesses. The greatest of these weaknesses is hallucination, generation of content that is partially or wholly inaccurate, but with a very high sense of confidence. So think about generative AI as something that can generate

new content, it can simplify content, it can explain content, it can synthesize content, and yes, it can hallucinate. So let's take a little bit more time to talk about hallucination.

To exemplify hallucination in practice, I asked a simple question of our team members here continued about their experiences with generative AI and hallucinations by asking what is the craziest or most bizarre hallucination you have experienced with GenAI? The first one had to do with grammar. Someone was taking a Japanese class and found that the generative chatbot hallucinates constantly about grammar. It can be completely wrong, but it sounds so confident. In the case of legal cases, ChatGPT found a number of cases with actual docket numbers and judge assignments that looked completely legitimate. However, that person was unable to verify the cases and confirm that they were all complete hallucinations. Interestingly, it made the person think that AI wanted to come up with an answer so badly and please the user that it just made everything up.

And finally, someone asked for 10 scholarly references when writing a blog post on learning, and every single one was bogus. It looked real. It had dates, issues, authors, dates of publication, page numbers, links, and when you went to click on the links, nothing existed, right down to the APA formatting being correct. So what's common across all of these hallucination experiences? Is that the chatbots exhibited confidence and produced responses that appeared to be legitimate at first glance. So here's a real life example. As I was creating this presentation, I was trying to generate an image using DALL-E that would illustrate the five functional steps of how an LLM works. If you take a look at this slide, you'll see that even though I kept saying You're wrong, you're wrong.

It kept generating images with more than five steps or more than five in images. And it also apologized for the lack of confidence that I had in its responses. It said, I will ensure the illustration only includes these five steps. Here are the images that

accurately depict the five stages of generative AI in the LLM process. And if you look at those images, there are more than five images and there are more than five steps described in the images. So again, that confidence of being correct comes through in the language and the interaction with the conversation in the chatbot. Just to give you an idea of what hallucination rates look like, Edelman published these rates from the common chatbot just a couple months ago.

Like all things AI these values are subject to change as dataset grow, models evolve and new chatbots enter the field of play. But the key concept to understand in this case is that chatbots can and will be confidently and convincingly incorrect, even with hallucination rates of three to 27%, most of these fall between five and 10% when it comes to hallucinations. So check your data. This is an example of the need for a human brain to complete the process. With that being said, it's just as important to understand the limitations of generative AI in how humans differs from chatbots. So remember that generative AI cannot understand abstract concepts, they cannot reason about ethics or morals. Generative AI cannot perform original research, nor can it provide lived experience, especially of note, it cannot fact check.

Neither can it ensure quality. So remember that generative AI is a language or an image tool. It's not a knowledge tool. Critical thinking is a human process and is still required. Let's move on to consider some practical applications of generative AI as professionals. Like we talked earlier, think of generative AI as a partner, a thought partner that can help you think, brainstorm, create, develop, write, translate, proofread, summarize, think about it being a fresh college graduate. What could you assign as a task to a recent college graduate with sufficient instruction for them to complete the task? Kind of like having a generative AI partner. So you can think about as a professional, you could use generative AI to help you write your resume or update your resume as we'll see in my case later on.

You can use generative AI to help you write a cover letter if you're applying for a job. If you job involves communicating with insurance companies often in a repetitive nature, you could have generative AI create those letters for you and those communications. Letters of recommendations for students who are applying to graduate school. Writing performance reviews is also something that you could use the help of generative AI with. Graduate school application essays and job postings are just some of the ideas that you could use to implement the use of a chatbot. If we think about it for business, generative AI can also be a business partner for you and help develop things like patient handouts, in services or professional development learning experiences, help you plan community events, create images for marketing community events, and I'm sure we've all encountered the customer service chat bot where it's getting a little bit harder to decipher whether you're talking to a chatbot or you're talking to a human.

And finally, generative AI can be used for business and marketing. In a recent working paper published by the Harvard Business School, an esteemed group of authors investigated the effects of AI on knowledge worker productivity and quality in their fields. Their data suggests that for tasks which represented four different task types, creativity, analytical thinking, writing proficiency, and persuasiveness within, and this is an important designation, within the current capabilities of AI, knowledge workers completed tasks 12% fast, they completed 12% more tasks, completed those tasks 25% faster and produced significantly higher quality results compared to a control group. Also, the experimental group that received a prompting overview, which we'll talk more about shortly, consistently demonstrated a more pronounced positive effect on the output.

As Carolyn mentioned earlier, the role of AI is to augment a human decision or action in a way that improves speed, quality or both. And the early data of Dell 'Acqua et al is definitely showing that that is the case within the workplace. In summary, using these types of examples, we know that chatbots can enhance productivity, efficiency and

quality for knowledge workers. But what about your clients and your patients? With the rise in the use of chatbots, remember our race to a hundred million users, ChatGPT won that race in just two months. I'm confident your clients and your patients are also playing with chatbots for education in a wide variety of subjects, including your area of expertise.

In this case, as the healthcare provider, it's critical to differentiate between disinformation and misinformation, especially when considering the hallucination rates of chatbots. Misinformation is false or inaccurate information. It's simply getting the facts wrong. Chatbot hallucinations are considered misinformation. Disinformation is another story. Disinformation is false information, which is deliberately intended to mislead, intentionally misstating the facts. Remember that chatbots cannot reason about ethics or morals. At this point in time, deliberate and intentional misrepresentation of information is human behavior, but it's also important to remember that disinformation lives on the internet and internet feeds our chatbots with data. So kudos to you for taking this course to educate yourself and in turn support your clients to dispel chatbot misinformation and the perpetuation of human disinformation.

More to come soon from Carolyn on this topic. I leave you with this quote from the APA on this topic, just about the power and spread of misinformation and disinformation and the fact that it's prevalent throughout our society. It's our responsibility as care professionals to promote the science as a means to protect ourselves and the clients in our care to the best of our ability. Again, we're gonna talk more about strategies for safe use of generative AI later in this course. So let's chat, shall we? Let's move on from chatbot background information to chatbot use or conversations with a chatbot chat. Remember, we talked about how chatbots are designed to simulate conversation with a human user, but how?

How do we do this? And the good news is that it's really pretty easy. It's kind of like a human conversation. Everything starts with a prompt, similar to when you introduce yourself to a new human. The purpose of a prompt is to communicate your query, request or instructions to the GenAI platform to obtain specific and relevant responses from the chatbot. Thanks setting the topic of conversation and turn taking with another human for further information, clarity or confirmation. So let's begin with a prompting model from Bliss 2023, who use the acronym CREATE. Character, request, examples, additions, type of output and extras. So when you think about a character, think about who do you want your ideal expert to be that you're chatting with, or who do you want your chatbot to emulate?

When it comes to requesting, be specific, the specific task that you want your chatbot to complete. Be sure to describe that. If you can, provide examples of the desired output desire or example information to better shape your output. Just like human conversation, there is always room for interpretation. If the output isn't what you expected, tell the chatbot how it should be different with additional information. As part of your request, specify the desired output. Do you want your output in paragraphs? Do you want bullets? Do you want numbers? Do you want a form for people to fill out? Be sure to tell the chatbot what you expect from them and think about any extras that you might add to the conversation once the chatbot gives you from some output.

My favorite is, is there anything else I should know about this? You might be surprised what a chatbot comes back with. So here's an example of the CREATE model in sort of real life. I'll give you a moment to read the example and then I'll share how the create model fits in. Our character in this case is a seasoned occupational therapist who mentors students considering the profession. Our request is for the chatbot to write us a letter of recommendation. For examples, we included details about the student in order to highlight their experiences within the letter. Where in terms of additions or adjustments In this particular prompt, we started out with a lot of detail. We may not

need to adjust much, but based on the output, be prepared to specify or ask the chatbot to confirm details or ask if there should be anything else added.

In the example we specified the type of output being a letter to not exceed one page. And if there was something extra, maybe I would ask after I saw the output, maybe it would be, is there anything else I should include to meet the requirements of an NYU OTD application? You would be shocked at how quickly the response comes back and how detailed the question might be to that. But be sure and fact check it. What I did is then I then answered this prompt exactly as you see it into Claude. And Claude responded in 11.18 seconds. After reading the output, it was a typical experience in that it was about 80 to 90% towards a final version status, which is pretty typical.

It's important especially when generating text to add your own voice, your thoughts, expertise, and experience are critical to achieving an optimal result For exposure, here's a list of the current landscape of chatbots. Examples include Bard, which is a language chatbot. Bing, which can also create images as well as language. ChatGPT, which is one we've talked about a lot. Claude, which is I would say probably anecdotally, the second most popular chatbot that folks are using for language. And then for images, we have DALL-E, Midjourney and Synthesia, which I always struggle pronouncing. Note that this table was created in the last week using ChatGPT. I asked ChatGPT to give me a list of the most popular chatbots, their type and a brief description of each one.

But like all things AI, changes are inevitable and everything is subject to change. But note the different types of chatbots. We have language, image and video represented, and also chatbot models are merging to support multiple modalities. For example, ChatGPT and DALL-E are integrated in order to support both language and image creation. Midjourney supports both still images and video. Moving forward, expect to see more and more support for multiple modalities within chatbots. So we know from

Dell 'Acqua et al 2023 that AI can boost quality, productivity and efficiency for knowledge workers. And that an overview of prompting had a significant positive impact on the output compared to the control group without a prompting overview. And as we have seen, prompting can be driven by a model or conversation, which is often more natural and dynamic.

And we look at these two comparatively and they often achieve nearly the same result. Prompting can be easier to get started and allows you to transition to more of a conversational mode with the chatbots. It's nice to note that conversation often can create some unexpected useful results. Wise words from Ethan Molik at the Wharton Business School. "The lesson is just using AI will teach you how to use AI. You can become a world expert in the application of AI to your own domain by just using it a lot until you figure out what is good and what is bad, what is it good at, and what is it not so good at." He also notes that evolution will likely be at play with prompting.

Nothing is constant in AI but change on the daily. "Having to worry about prompts in all could become a very temporary state of affairs," Molik states. "As AI systems improve, the need for prompting could very well decrease because the AIs themselves become really good at figuring out what you might want." So just like a human, you have to talk to someone to get to know them. Using AI how you prefer, whether it's using a prompt model or just conversation, get in there and use it and you'll continue to get better and better. So why not ask a chatbot how to best engage with the chatbot? This list was generated from Claude AI using the following prompt.

What are the best practices for engaging in conversation with a chatbot? The rules of engagement according to Claude, and you'll see there's a lot of human quality to these recommendations. Be clear and concise. Chatbots have limitations in understanding complex sentences or ambiguity. Keep your requests and questions simple and to the point. Ask one question at a time. Don't overload the chatbot by asking multiple

questions together or may get confused. Ask your questions one by one instead. Be patient. Chatbots require some processing time, especially for more complex requests. Allow a few seconds for the chatbot to respond after you type your question or request. Speak in a natural tone, don't feel like you need to use stiff robotic language, chat casually while still being polite and respectful.

Using natural language helps the chatbot understand you better. Provide context when needed. If you find the chatbot is not understanding your request, politely rephrase while providing more background context to clarify meaning. Avoid inappropriate questions. Chatbots are not intended for controversial or insensitive questions. Keep the conversation productive and positive. Be willing to rephrase. If the chatbot provides an unsatisfactory response, politely rephrase your question instead of repeating the same question verbatim. Different phrasing can elicit better responses. And finally say, thank you. Simple courtesies like thanking the chatbot or saying goodbye help provide social cues to indicate positive interactions. And Claude closes by saying, following these basic etiquette rules helps ensure the best conversational experience for both humans and chatbots.

The key is to engage respectfully and effectively. Note the human tone to all of these recommendations, respect, clarity, and appropriate communication and positivity. To wrap up this section, I'll share some of my chatbot world pearls of wisdom. The first is the 80% rule. AI or a chatbot will get you 80% of the way to a finished product. You still need your voice, you still need your brain to make it yours. Hallucination happens, hallucinations do happen. So check the facts and check the links. If you're not happy with a response, rephrase it to refine the output. Repeating the same thing again and again and expecting different results, well, we know what people call that. Don't be bashful. Provide your chatbot feedback.

Know about knowledge cutoffs. Some data sets have limitations. They may have, let's say whatever data set they're searching, maybe their data is only as current as let's say June of 2023. That means that nothing is put into that dataset after June of 2023. Also, there are length limitations or character limits, and although they are quite large, they can limit your input and your output. And finally, typos matter. As you'll see in my two images here that I tried to create, I typed in pear of wisdom instead of pearl of wisdom. And the images that I got were very, very different. So like in all things typos matter. Let's move on to an activity. Let's pause and play and have you choose a chatbot.

So what I would like you to do is pause the recording and choose a chatbot. ChatGPT, Claude, start with Google to get to the right place and create an account for yourself. Start talking to the chatbot. And here's a list of ideas to get you started with a generative AI request. But don't be afraid to request with purpose and creativity to suit your own needs. Use the create model or simply engage in conversation. Come back when you are ready to continue the course. Welcome back. I hope your chatbot experience was enlightening and exciting. As an example, I used Claude to provide suggestions for updating my resume. It's been a while. Claude provided eight specific recommendations that are relevant and reflective of changes and resume formatting of late.

Claude was also very kind, stating that overall it has great content and achievements are captured. I started with a very simple prompt, what are your recommendations for improving the attached resume? And yes, that means that you can attach documents to your prompts. I adjusted the prompt to get further detail on how to best list the 20 plus courses I've taught as an adjunct faculty member. The options and explanations were sensible, relevant, and helpful, and they condensed the information effectively. What surprised me most about the experience was the accuracy of the recommendations to improve my document. I'll definitely take Claude's advice. So

what information did you discover that was missing from your prompt? And how did you adjust your prompt to improve the output?

And finally, what surprised you about the chatbot experience? Now that we've increased our familiarity with chatbots and their function, I'm happy to pass the baton back to Carolyn to discuss risks, limitations, and safe AI practices.

- Thank you so much, Leigha. So as we've already touched on, as Leigha sort of already touched on with all technologies, there are some risks and limitations to be aware of. And once you know, you'll be more skillful as you use these tools. So I've summarized the risks and limitations of GenAI and of AI in general into these four categories on the slide. And I have a slide build in here too that you'll see. I've also added an umbrella here with the word ethics on it is I see a lot of overlap between these risks and what we think of as ethics in the healthcare realm. So let's dive into the first one. Leigha already touched on misinformation, disinformation.

We talked about hallucinations. So this is very easy with GenAI. Another word I'd like to introduce here is deepfake and deepfake is actually a combination of deep learning and fake. So again, convincing, but fake content made using advanced AI technologies, particularly deep learning where a person's likeness is used in a false and deceptive way. So beware of this just as a consumer of AI. I've included a link here on your slide to a YouTube video that's a deepfake of Morgan Freeman. Lots of fun and fascinating, so check that out. There's also questions around data transparency in regard to where some of these inaccuracies are coming from. And the other citation here is from a study by Menz and colleagues, they highlighted how powerful, and as Leigha said, convincing health misinformation and disinformation can be quickly and easily created with GenAI by regular people.

We're not talking about hackers or computer programmers. So first risk, misinformation, disinformation. And then moving on, we have privacy. Privacy, another area of concern with current AI systems. So as we know, AI uses vast amounts of data. So one question is where did the data come from? Also, was it scraped from the web? Was it stolen? Did people give consent for use of their data? Who's responsible for data breaches? Accountability is an issue. And one other thing you should know about privacy that we hear of is that the training data, if it includes private information, that information can also appear in the outputs of GenAI. So if we go to the next slide, you will see right in the terms of use on Claude, it tells you personal data in your prompts, really any data you put into your prompts, they will collect the information and the information may be reproduced in the outputs.

And it doesn't only refer to your outputs with the GenAI chatbot, but with anyone using Claude, and you'll see her citation on the slide. Carlini and colleagues performed an attack on an earlier version of ChatGPT, and they were able to extract personal information from the training data. Things like names, addresses, medical information. So again, what you're putting in, could appear in outputs. Keep that in mind when it comes to privacy. And then moving on, copyright. So because AI may have been trained on data sets with copyrighted works without permission or remuneration, this is an issue for you to be aware of. Also, the same with privacy in that what you're putting into the AI may appear in outputs.

This, if systems were trained on copyrighted data, copyrighted data may appear in outputs or aspects of it. So there have been a lot of class action lawsuits in the news recently with regard to copyright and AI systems. In September, 2023 the Authors' Guild that that represents several well-known authors like John Grisham alleged that Open AI, which is the company behind ChatGPT, copied their works without permission or consideration and fed these copyright works into their large language models. Another lawsuit that's been in the news a lot, Getty Images, you've probably

heard of them, they own images and videos that you can purchase and license. They are suing the GenAI maker Stable Diffusion. They're alleging that stable diffusion used 12 million copyrighted Getty works in its training data without permission or payment.

And if you Google this lawsuit, Getty versus Stable Diffusion, you can see some images created with Stable Diffusion that look eerily similar to Getty images, and they even include a blurred Getty watermark. So as these lawsuits, they have not been resolved yet, but it'll be interesting to see how they play out and what changes they drive, if any, in terms of regulation and requirements around AI. And then we have bias. So this is really important for us as healthcare professionals to be aware that GenAI, AI systems may be biased. So their training data, for example, may not be inclusive. It may reflect wider societal biases or inequality. The algorithms themselves may be biased, particularly in the way that they correct for outliers in their data.

And then the way these systems are used may be biased, the decisions that people make based on the outputs of these systems. So how we understand statistics, how we interpret statistics, how we use statistics may be biased. And this on the slide you'll see Lee and colleagues, this was a very interesting study. It is available free and open access. It's a newer study. They looked at automated pedestrian detection systems. So we're not talking about GenAI, we're talking about Computer Vision, another form of AI. But these pedestrian detection systems do just that. They detect pedestrians for things like self-driving cars. So the researchers looked at the detection accuracy, and they found for some systems it varied by age and skin tone.

So for example, adults were detected more accurately than children, and light-skinned individuals were detected more accurately than dark-skinned individuals. So obviously there's enormous implications here regarding this bias for safety, for researchers, for computer engineers building these systems and for policy makers. So just again, an example of bias and something for us to be aware of as we use these systems. So we

know there are risks, with all technology there are, and being aware of these risks will make us more skillful as we use this technology. One of the ways that we're going to get to safer GenAI systems is through regulation. Regulation is gonna be one way that boundaries and guardrails are put into place.

And regulation is now in its infancy stages, but it's definitely emerging. We have four tiers of regulation. At the international level, we saw the UN recently put in place an AI advisory body. And the European Union just recently passed the AI Act, which is a global benchmark for AI regulation. In the US at the national level, President Biden issued a US executive order on AI safety in 2023. And then at the next tiers we have industries will be setting their own regulations and guidelines, our professional associations hopefully coming forth with guidelines and regulations for us. And then individual workplaces, individual companies will also be setting forth policies. And it is incumbent upon us as healthcare professionals to ensure that AI aligns with the values and standards set forth by our codes of ethics.

And we'll wanna be part of shaping standards and regulations when possible. Most of us probably won't have that opportunity at the international level, at the national level, but we can certainly get involved in our professional associations and in our workplaces to shape those standards for AI. But what can you do in an individual level to mitigate these risks? Well, the first thing as Leigha mentioned, verify and fact check always your GenAI results. At the end of the day, AI is a tool and you can't blame the AI if you use results that are inaccurate. You are responsible for how you use it. And I love this quote by Werner Vogels. He's the global CTO of Amazon, and I had the opportunity to hear him speak at the World AI Summit in Amsterdam in 2023.

This really stuck with me. He said, "AI makes predictions, humans make decisions." Again, it's up to you on how you use AI. These chatbots will tell you to verify their output if you read their terms. So this is from the chatbot Claude, and it says, right in

the terms, "You should not rely on outputs without confirming their accuracy and outputs may contain inaccuracies even if they appear accurate." Again, these are very convincing what you read, you have to check it. And then on the next slide, we'll see that ChatGPT, right on its homepage. I put a little red box around it in my slide build here, says ChatGPT can make mistakes. Consider checking important information. Well, we wouldn't say consider, we would say always check important information.

And then we have Google Bard, which is another chatbot. Also, if you read the fine print says right here, may sometimes give inaccurate or inappropriate information. And don't rely on its responses for medical, legal, financial, or other advice. This is something you can show to your patients. Don't rely on a chatbot for medical advice. That's why I'm here as your healthcare professional. All right, in terms of other strategies, when it comes to privacy and copyright exercise caution with regard to what you enter into these chatbots, as we've covered, any info you enter may very well be used in the training data for these systems. It may appear in other people's outputs and hackers may be able to extract it.

So don't enter PHI or personal information for yourself or others that you wouldn't want shared and used in this way. And consider how you're using the information you get out. For example, if you're publishing something or using something commercially, you're obviously gonna be more cautious with the information. Like say an image, you're publishing a book, you're looking for a cover for your book. Maybe consider using a proprietary system like Getty AI. So Getty Images has its own AI system that was trained on its own data sets and has clear terms of use rather than using an open access chatbot. And speaking of terms of use, on the next slide, just wanna say that the opt-outs, so if you think you're opting out of giving these chatbots your information, they are clear as mud.

Sometimes even when you think you're opting out, if you really read the fine print, it appears, at least to me, that your data is still being used. Opt-outs in some systems make them much clunkier to use and they even defeature it. So you can't use all the features if you don't allow it just to take all of your data on your inputs. So maybe down the road, as regulations are developed and enforced, we'll have clearer and easier ways to protect our data. But at this point, use a conservative approach to what you're putting into chatbots and how you use what you get out of them In terms of policies, see resources, we have an information in your resources about setting up a workplace policy for the use of AI.

So establish or advocate your employer for a policy in your workplace and advocate your professional association to lead us all to educate us and inform us and regulate AI in our profession. Always check with your employer before using GenAI for business purposes. And if you're your own employer, consult your business attorney if you have any questions. So I like this example from the American Psychological Association. They know their professionals are already involved in AI projects and they've developed these considerations to guide AI policy development, to ensure systems are ethical and effective. And I'd love to see all healthcare professionals really taking the lead on AI and how it's used in healthcare. So in summary, there are a ton of uses for AI and Leigha covered many of them for GenAI.

Keeping in mind these risks and limitations will help you determine when and how to best use these amazing tools. I keep my ChatGPT open on my computer all day, and I consult it for most writing tasks, even some emails. Just follow these simple guidelines. We've created this acronym for you, SAFE AI, security awareness, fact check, education, keep up on AI advancements. We've included a lot of great resources for you where you can really keep up on AI, kind of the big picture happening in just a few minutes a day. Accountability. And at the end of the day, integrity, you are responsible for how you use AI. So thank you so much for your time and attention.

And Leigha, I'll turn it back to you for any final words.

- Once again, I'll emphasize the role of experiential learning. To learn to use AI, you have to use AI, of course, with safety in mind, this is today my favorite image that I've generated using Midjourney. Orange cats learning together, using technology over coffee. It just makes me smile. My parting advice is to have fun experimenting with these technologies and share your experience with others. The more you play, the more you will learn, leading to increased comfort with the technology. Thanks again for your time today, and we hope to see you again soon.