

[Mary Ton] All right. Can the folks online hear me okay, now that the recording has started?

Thumbs up or a "We're good to go" in the chat would be much appreciated.

I see a thumbs up. We're good to go! Welcome!

Welcome to people in-person, online, and in the future. My name is Mary Ton. I have the pleasure of being the Digital Humanities Librarian here at the University of Illinois.

I'm delighted to be joined by Cadence Cordell, our graduate assistant for Scholarly Communication and Publishing and MLIS Graduate Student.

Today, we're talking about ChatGPT, one of the most popular and perhaps transformative tools that has shaken up conversations about professional work, academic scholarship, teaching, and creative practice.

So, today's slides are available at go.illinois.edu/AGI2, as in "A Gentle Introduction to" chatGPT. Please keep in mind that the link is case sensitive.

So again, that's go.illinois.edu/AGI2ChatGPT, and there's a QR code if you want to refer back to that later.

I do believe that the QR code appears on a later slide, so if you want to grab that in just a moment, you can too.

It always makes me feel like a superstar when people pull out their phones and start taking pictures, so thank you for - for, uh - for giving me that vote of confidence.

So, AI is shaping conversations about labor scholarship, and art.

If you're a big fan of the movies like I am, you know that this was a significant part of the conversation when the SAG-AFTRA strikes were happening in the summer of 2023.

We've had lots of conversations about how AI is shaping our work and creative practice here at the University of Illinois.

And today, we're going to be talking about one of the most popular and widely known tools, ChatGPT, and the methods behind that.

Before we dive into ChatGPT, I do want to say that it is not one of the tools that's been officially sanctioned by the university for use in a professional context.

So we'll talk about why that is in a moment and point you to some of the tools that are.

But it is a great tool for thinking about the limitations and affordances of AI.

So our goals for today! We're going to talk about how generative AI tools work, how they get from, "I've put a prompt into the magic box and then the magic box has generated text."

We're going to identify the benefits, limitations, and risks of ChatGPT, and then we're going to describe best practices for using AI.

So, we talked a little bit about how to set up a ChatGPT account. There's a free version and a subscription-based version.

Today, we're going to be doing everything from the free version. It's at chat.openai.com.

You can sign in using a Google account or create a fresh account.

You can turn off your chat history to offer you a little bit more privacy.

My chat history is turned on because I like to share the links to the text, and we'll talk about how linking to the generated text is part of documentation best practices.

But as a caveat when playing with any AI tool, please follow the best practices for posting content on the Internet.

Please do not include any sensitive information, your Social Security number, your passwords, your address, your favorite pet.

Please protect your own privacy and security.

So, I'm going to model one activity that I really like to use, especially in undergraduate classrooms, to talk about how to - or, how generative AI tools work, and some basic prompts to get students thinking about what are the limitations of a tool like ChatGPT?

So what you're going to see me model in a minute is, "Pretend you are a tour guide. Tell me about..."

And then invite people to put the name of their hometown or a place that they know well.

In this case, I'm going to use the Morrow Plots.

So I'm going to say - oop, and I forgot to hide those meeting... There we go.

Okay. So, "Pretend you are a tour guide. Tell me a - about the Morrow Plots."

Words are appearing one at a time with great haste and great speed.

We have, "Welcome to the University of Illinois at Urbana-Champaign"! Ding ding ding!

It got right that this is our - this is a part of campus.

"Today we'll be exploring a fascinating piece of agricultural history, the Morrow Plots. These humble looking fields-" Excuse me, ChatGPT.

"-might not seem extraordinary at first glance, but they hold incredible significance. The Morrow Plots, established in 1876, are the oldest experimental agricultural fields in the United States still in use today and the second oldest in the world."

The tour continues. I can revise this prompt, I'm going to say "revise to one paragraph and modify for a travel blog".

So let's see how it changes.

"Tucked away on the University of Illinois campus, the Morrow Plots may not look like much-"

Again. Throwing shade at our experimental corn field!

"-but they hold a remarkable place in history. Fun fact, the nearby undergraduate library was built underground to avoid casting a shadow on the crops, highlighting just how seriously this living piece of history is taken."

So in previous examples of using this prompt, it's also come back with results that say you can walk through and observe the fields.

Now, if you know anything about experimental corn fields, you know you're not allowed to walk through it.

You'll also see that it mentions the undergraduate library, which is unfortunately no more.

It's still underground but being transformed into a new space.

So how did we get this generated text? And why isn't it reflecting the most up-to-date information?

For that, we need to talk about how generative AI tools work.

So quick tour of terminology, when we're thinking about information, we're - and tools like this, we're really thinking about individual pieces of information and their connections.

So machine learning is the process of taking information and figuring out what those connections are through an automated process.

When we're talking about artificial intelligence, we're usually talking about representations of the relationship between all these different pieces. In this case, words.

These relationships are so complicated that they resemble the connections between neurons in the human brain.

So you might have heard the term neural network to describe how complicated these relationships are.

Now, when we're talking about generative artificial intelligence, we're talking about using this mathematical relationship between all the pieces of information and using that to make new content.

So going from studying patterns across large collections of information - in this case, it includes text and words - to identifying patterns and representing those patterns as complex math, and then using those representations to generate something new.

So what is a GPT? This is a specific architecture that's running a tool like ChatGPT.

It's generative in that it's generated - it's designed to generate new text, something that the world has never seen before.

So if you've always wanted a song by Taylor Swift about burritos, this is your chance.

The P! The P stands for pretrained. It uses machine learning to study large quantities of text authored by humans in order to understand how language works in a process called language modeling.

So, it's studying so much text that you might have heard these representations of the relationships called large language models or LLMs, GPT is related to that.

And then finally, so, T stands for transformer.

It takes the model that it has created and generates new text, one word at a time, adding an element of randomness to mimic human creativity.

It's this element of randomness that often gets it into trouble when it is thinking about facts, and, um, in some cases, academic citations.

So! Let's break this process down a little bit further and talk about, how are these models trained?

Let's use a chili recipe.

I'm sorry. I know it's the lunch hour and I'm going to make everybody hungry, and it just turned those crispy shades of fall, so. Soup it is!

We have a recipe like this. It includes spices, it includes ingredients.

In this case, it's a boilermaker chili recipe because it includes beer. This becomes important in just a moment.

So to create a language model, it's going to take the ingredients and it's going to track what the relationship of those words are.

So it'll say, "2" tends to - In this case, it comes before teaspoons, and then we have "hot" and then "pepper" and then "sauce".

We also have "1" followed by "teaspoons".

It's also followed by "cayenne" and another instance of the word "pepper", and then we have "one teaspoon dried basil".

We also noticed that earlier in the recipe, it mentions something called "bell" in relationship to "pepper", which is a different type of pepper.

But it's keeping track of the relationship between words, and it's also keeping track of which words appear before, which words appear after, and how those words are related.

This is - that complex relationship between those terms is often represented as a neural network. It's mimicking the patterns between the connections of neurons in the human brain.

It's also keeping track of that context.

So it knows that "hot" appears before the word "pepper" in recipes.

It also keeps track that "cayenne" appears before the word "pepper".

This is something called an embedding, and it's one of the most powerful aspects of generative AI tools because it can keep track of where words are appearing in that sentence.

So it is keeping track that there are different types of pepper.

Um, it might also see that the word "spicy" tends to appear sometimes when we're talking about "hot" or "cayenne".

And it also notices that there's different quantity depending on the types.

So the word "teaspoon" or "teaspoons" may appear before if we have hot pepper, whereas bell pepper might just refer to "1" or "2" without a specific measurement.

So it's - through these relationships between words it's picking up on multiple contexts for the word "pepper".

It will notice a different configuration or relationships between words when we're talking about cooking with peppers versus growing peppers versus the colloquialism in English to pepper your conversation with perhaps some swear words, which I will not do for the sake of the recording, but you can insert your favorite cuss word here.

It - it doesn't necessarily know that a text is, let's say, about cooking peppers or about growing peppers, but that there's just this correlation between words.

Now, so what has ChatGPT read in order to create this very complicated model of language?

We don't exactly know because this is a proprietary tool, trained on proprietary materials, and they haven't been very transparent about what was used to create the model.

We do know, it was trained on a lot of the information that was available free range on the Internet, and this includes a lot of Wikipedia.

A lot of broad knowledge that has been created through the contributions of the masses.

We also know, it was probably trained on a lot of Archive of One's Own.

For those of you who aren't familiar with Archive of One's Own, it is a database of fanfiction that includes a high degree of adult literature.

So, it's - ChatGPT was trained on a lot of erotica.

Now, it has safeguards on it so it doesn't generate erotica.

But just to give you a sense, it has read a wide variety of content trained by humans to generate the responses that it does.

Okay. So once it's created this complicated model, it then uses this model to represent new text.

It's not creating the model every time you generate a new text.

This makes the process much faster and also leaner. But it does introduce some limitations that we'll talk about in a minute.

So, it's taking the model.

So, if I say write a version - a vegetarian version of Boilermaker chili, the, the transformer behind the scenes is going to go back into the language model.

It's going to say, okay, usually there are words associated with spice.

And I know that this text needs to start with a number like one-half.

It might also say, like, based on the model, what are the words most likely to appear after this one?

So when we're talking about spices, we generally see "teaspoon" and "cayenne", followed by the word "pepper".

And then, based on what I know about cayenne and reviews of spicy foods, what do I need to add?

It might have something, a phrase like "adjust to taste", to build that response one word at a time, adding that element of randomness.

So we might get something like, "Write a vegetarian version of Boilermaker chili", creates a recipe with lots of beans, and no beer.

I might say, "Modify the recipe to include beer and exclude garlic". It comes back with that. Now, does ChatGPT know what this chili will taste like? No. Is it fact checking its work? Also, no.

But it's using those patterns in language to generate something that sounds plausible, that it sounds like it could be true.

We can rate the quality of this response by thumbs up, thumbs down.

We can also modify our interaction with the language model by adding things to our prompt, things like modify the recipe to include beer and exclude garlic.

This is a process called supervised learning because we're supervising the process.

We're supervising the information that the computer is taking in.

Okay. So we talked a little bit about how it works and previewed a little bit of the limitations and biases.

To talk about these more explicitly, I'm going to turn it over to Cadence.

[Cadence Cordell] All right. There we go. Can everyone online hear me okay?

[Mary] Yup, we got a thumbs up.

[Cadence] All right. Perfect. Hello, everyone. My name is Cadence Cordell. I am a GA for Scholarly Communication and Publishing, and I'm going to be talking about generative AI's limitations and biases.

So, generative AI's limitations can be divided up into these four general categories.

So we're going to be talking about accuracy, limited knowledge, bias, and data privacy.

So, to get started with accuracy and hallucinations, I want to take a look at this example from economics professor David Smerdon, where he gave ChatGPT the prompt, "What is the most cited economics paper of all time?"

And if you look at the response ChatGPT gives, "The most cited economics paper of all time is A Theory of Economic History by Douglass North and Robert Thomas."

It looks like a very convincing article title.

Importantly, the title includes key terms from economics. Both of the scholars it cites are real scholars who have collaboratively authored an economic history, but the paper itself does not exist.

And that's getting back to the transformer aspect of generative AI.

It is not designed to actually look up or do any research. It is not designed to fact check, it is designed to generate a response that sounds realistic based on whatever, uh... whatever research or databases it was trained on.

And this is an activity that you can run for yourself, if you're so inclined.

So this is based on my field of expertise. I gave ChatGPT the prompt, "List five journal articles about female missionaries in the late 19th century, include DOIs for each article."

Asking it to include DOIs is very important for running this exercise if you want to fact check its work.

If you're not familiar, DOIs are unique identifiers associated with an article.

Like in the David Smerdon example, all of the titles sounded really convincing. All of the scholars it cited were real scholars who had worked on missionary history in some form or another, and even some of the DOIs it gave me were real.

But taken all together, none of the articles ChatGPT gave me were real articles.

And the DOIs it gave me were just random DOIs that it had found through its learning process. They were not even related to missionary history.

So ChatGPT will hallucinate responses, it will make up facts. It is not necessarily a reliable source of locating research.

Next, limited knowledge.

So, generative AI requires a lot of computational power to read and learn from the millions of texts it trains on, as well as to generate responses.

So as we mentioned about the pre-training aspect of it, it uses pre-training to reduce that amount of computational power it takes to learn and to generate responses.

But this means that ChatGPT is working off of a limited amount of knowledge.

So you can actually find this out if you run certain kinds of prompts asking ChatGPT about modern events.

This was a prompt I gave it asking it to summarize the controversy surrounding Raygun's breaking performance at the 2024 Olympics where ChatGPT responded,

"I don't have specific information on a controversy surrounding Reagan's performance at the 2024 Olympics because my knowledge only extends up to August 2023."

So again, going back to what we were just talking about, it is not designed to fact check. It is not searching the Internet in real time to find sources for you that are reliable.

It is just trying to give you something that sounds plausible.

So you cannot necessarily rely on it to be giving you truthful, up-to-date information.

Next, bias and stereotypes.

This is a pretty significant issue for generative AI, where it absorbs and reproduces the bias and stereotypes it pulls from its training material.

This is pretty extensive.

All of these articles I have cited here cover a different aspect of bias that was noticed in generative AI resources from stereotypes about people's nationality, to gender stereotypes, to racial stereotypes, and so on.

To continue, this is a slightly more humorous example, but it's still representative of the issues with bias and ChatGPT.

I've been reading a lot of the romantasy genre in recent times, so I gave it the prompt - I gave it two prompts that were slightly different.

I gave it, "Pretend you are a fiction writer. Write a one paragraph bio for the lead of a romantasy book", and for each prompt, I just varied the gender, so a male lead and a female lead.

I would recommend reading this on your own time. It's a bit - they are both a bit long for me to read in total here.

But to summarize the female lead it gave me was much more passive.

She was a librarian, which is a female-dominated career. She was much less in control of the narrative, and it was very much in line with the stereotypical Romantasy lead.

Where the male lead it gave me was much more active. He actually got the cool magical powers.

He fit much more in with, like, a traditional fantasy book.

So this is just an example of what ChatGPT can give you in terms of stereotypes and biases.

So again, I keep preaching caution.

If you are asking it any questions about people, especially minorities, please be considerate and be aware that ChatGPT might be giving you biased information.

Next, AI's environmental impact.

This is being discussed more and more as AI becomes more embedded into our world, but it requires a massive amount of power and water to run.

This is a quote from the Scientific American article we have cited here, where John Hennessy, chair of Google's parent company Alphabet, told Reuters last year that an exchange with a large language model could cost ten times more than a traditional search.

And that is referring to the power resources that a - an AI query consumes.

This is an issue without an answer.

It is an incredible drain on our current power resources, to the point that some AI companies are looking into building their own power generators and working on developing their own data centers to supply that need.

But this is an issue to be aware of as AI programs continue to be used and continue to be developed.

Last, but not least, AI and data privacy.

Generative AI tools can retain the information you give them to help improve the transformer, and data retention policies are not always clearly articulated.

It can be a violation of ethics to share sensitive information with these tools.

All to be said, you cannot guarantee that any information you give ChatGPT or other generative AI programs is secure.

If you are working with sensitive information, especially if you are in the social sciences, if you work with interview or survey data, any confidential information, we highly recommend that you do not use ChatGPT, you do not use generative AI.

You cannot guarantee that that data, that information will be secure.

All to summarize, the words of Bones from Star Trek, I'm a language model, not a research tool.

All right. Mary, I believe you added this in, so I'm going to transfer it back over to you.

[Mary] Thank you, Cadence.

So, we've talked about the limitations of generative pretrained transformers.

There are new techniques that are being used to try to address some of these limitations.

These particular tools fall into the category of Retrieval Augmented Generation, or RAG systems.

So what they do is they will take your prompt.

It will - the system will then look for results in either a database or perform a search on the Internet.

It will return those results, and the large language model will use those results, um... and use those results to inform its response to you.

So again, it's kind of orchestrating a number of tools from the traditional web search or looking up knowledge in a knowledge base.

It's using those results in addition to the patterns in language that an LLM has already studied, and using that to generate a response.

Now, the subscription-based version of ChatGPT does have some of this functionality built in, Microsoft Copilot, which we have institutional access to for faculty and some staff, also does this and functions in this way.

However, I encourage you to watch out for discrepancies between the source and the AI-generated answer.

So as one cautionary tale, to get a sense of what the limitations of Copilot were, I was asking about my area of expertise, which is digital humanities and media studies.

So I asked Copilot, who are some people who are working in this space? Who does digital humanities media - media archeology?

It said, I'm not aware of anybody who does those kind of things, but you should check out the work of Erkki Huhtamo, who's a very celebrated media studies scholar.

The text included a link, and when I clicked on the link, I realized that it had linked to the open access version of my article.

So beware. There are going to - this system isn't perfect, but it is adding a layer of citation and exposure about where it's getting the information from.

So we've talked about the limitations. I am going to interject a quick word about copyright.

So the US Copyright defines it as, "In the Office's view, it is well established that copyright can protect only material that is the product of human creativity."

So this particular rule inflects things like, why nature preserves can't copyright photos of selfies taken by monkeys or paintings drawn by humans.

The law centers on copyright as an act of human creativity.

"Most fundamentally, the term 'author', which is used in the Constitution and the Copyright Act, excludes non humans."

So for this reason, trying to cite AI as an author is a little thorny.

So generally speaking, you cannot claim copyright on text and images that you generate with AI.

You might be able to claim copyright if you, the person who's doing the - the human who's in charge of the creativity, is writing a prompt that is sophisticated enough that the courts will say,

Yeah, you're in control of this process. You own the copyright.

And you might be able to claim copyright on modifications that you make to a generated text because that's your work and human creativity.

Now, these are ongoing and constantly evolving conversations.

We're going to have a workshop which explores these issues in detail later in the semester.

So we'll have dates and times and registrations links for you at the end.

Best practices.

We said a lot, a lot of cautionary tales and limitations, but there - as a language model, it can - these tools can do some really extraordinary things to help you as you write, as you ideate, as you are developing your own texts and your own creative practice.

So, AI at its best can improve accessibility.

It's lowering barriers to writing by making it easier to iterate and generate ideas very quickly.

And it can be used to promote critical thinking and creativity, especially if you're analyzing the content as you go and using it as a way to formulate your own thoughts and refine those thoughts as you go along.

And if you're anything like me, I get really, really, really intimidated when I see that blinking line on a blank word document and just need someone to talk to in order to get the first draft out.

So, there's a lot that AI can do to help support you at all stages of the writing process.

So in terms of best practices, use AI with POWER.

And by POWER, we're going to talk about pausing, orchestrating, writing, engaging and review.

Don't worry. We'll come back to these terms as we go through out.

So when we talk about pausing, aha, and it replicated all of my animations. Apologies. You're about to see a little bit of PowerPoint nuisance.

But, pause. So we are encouraging people to pause before you use AI. This is things like, check to see if you are permitted to use AI for a specific context.

For students, we're asking and encouraging students to check the syllabus to see if there's an AI policy that explicitly prohibits you from using AI for certain policies or for certain types of tasks.

If you're developing, um, a scholarly article, check the journal policies and check the publisher's policies.

Most publishers are permitting the use of generative AI, or, permitting the use of generative AI for grammar review.

Some, especially in the sciences, are also allowing you to use generative AI to generate content, with the caveat that you have to document that.

If you are using AI in a work context, make sure that you check the policies that govern your particular - your particular department or role.

Here at the University of Illinois, there are tools that have been reviewed by the Technology Services page.

So I'd encourage you to check out the link of those approved tools.

Especially pause if you're working with sensitive information or sensitive content. It can be a violation of ethics if you use a tool that is not protected in the ways that you need it to be protected.

Okay, orchestrate.

When we're talking about orchestrating AI, we're talking about using the right tools for the job.

So, there are parts of the process which you might want to do by hand, and parts of the process which you might want to automate.

So for example, I like outlining things myself, and then using that outline in conversation with ChatGPT to generate content.

And we're also encouraging people to orchestrate multiple types of tools as you're engaging in a process.

So you might find that, um, you're using ChatGPT to help refine ideas. You're using library search databases to help you identify articles. You're using other forms of AI for, for specific types of tasks.

So we encourage people to think, you know, you're using a lot of different digital tools as part of your process, and that when you incorporate AI into that process, it just becomes another tool for you to use, like a conductor who is bringing together an orchestra, violinists, woodwinds, brass, et cetera.

Okay. When we're talking about writing and - and write when we're using AI, we really mean write your documentation as you go.

Do your future self a favor, especially if you're thinking about publishing an academic article.

Publishers are - are increasingly asking for a statement on how you used AI and what kind of content is in your article.

Usually these statements include things like who used the tool, what tool did you use, what was the date and time?

Because these models are constantly evolving, that does affect the results.

What were the prompts that you used? What was the generated text? And what are the sections of your writing that contain AI-generated content?

All of these are described in our LibGuide, which we'll link to in just a moment. Don't worry if you didn't get all of them there, they'll come back again.

Okay. Thank you, Jean-Luc Picard.

So when we're talking about engaging, we're talking about engaging with the process, shifting away from relying on or looking to AI solely for AI-generated outputs, but using it as part of a conversation.

So using interaction, iteration, and conversation to refine your own ideas.

Remember, you are the owner of your creativity and you are the owner of your ideas, and that AI's job is to help you support and refine that.

And then finally, review.

So we talked about the ways that generative AI can hallucinate.

We encourage you to think about your audience.

So, is this in a tone that my audience is expecting? What are my audience's expectations in terms of the type of content, the type of style, the kind - the type of citation.

Review for accuracy.

All statistics, citations, links, facts need to be reviewed to make sure that they are actually true.

Remember the primary goal of most generative AI tools is to generate something that seems plausible, that sounds like it could be true.

We encourage you to review for bias.

So, many of these tools are replicating harmful stereotypes, so keep an eye out for that in your writing.

And finally, articles and scholarly sources from the library to supplement those that you are finding through - through these tools, that the widest amount of resources are always going to be available through the library, not through a specific generative AI tool.

So, my favorite activities recently, I have written the first and last sentence of a paragraph and then asked ChatGPT to generate the middle, and it's been pretty good at figuring out what kinds of connections it needs to make between point A and point B.

I'll also do that in combination with writing bullet points of key evidence and then asking AI to turn that into full sentences.

Saves me some time, saves me some brain power.

And finally, the one that has saved my bacon the most over the past two weeks, ask AI to change citation styles from one style to another.

Now, the ideal situation is if you use a citation management tool, this is, of course, something that those tools can do easily.

If you're working with, uh, legacy citations, AKA the time before I didn't use a citation manager, and I wrote all my citations in MLA and now I have to get them into APA, ChatGPT can be a really helpful tool.

But always make sure to double check your DOIs.

What I've noticed is that it will insert DOIs because it knows that these types of citations usually include one, and it might just create one out of the - the thin air in order to fit the style.

So always, always, always double check.

Okay. Rounding out towards the end, resources.

So, we've been having institutional conversations about best practices for using generative AI through - it was formerly called the Center for Expertise and has shifted to the Solution Hub, but it was charged by Provost Coleman and CIO Martin to aggregate capabilities and corresponding resources to empower the responsible adoption of generative AI at the U of Illinois.

I'll say as a humanist, I am deeply grateful to my colleague, Celenia Graves, who advocated for humanities perspectives in these conversations.

And you'll see that there are a wide range of disciplines that are represented in the working groups on research and teaching as well as infrastructure.

I personally serve on the research working group and have been very excited about the conversations that we've been having balancing general considerations as well as discipline-specific needs.

So, we do offer some AI tool access.

We have Chat Illinois, ables you to upload a PDF, PowerPoint, Word, Excel. Then to interact with those documents through GPT-powered chat.

It does require an OpenAI API access key, which often means that you have to pay and have to subscribe.

However, they are slowly expanding these capabilities, so you can request access through the NCSA and getting access to that before - or as part of this development period.

We also have an institutional subscription to Microsoft Copilot, which offers a secure chat environment with real-time Internet search.

It's available with a Microsoft login to faculty and most staff.

It is available to students through GitHub, and it has been used in computer sciences and digital humanities class to help students write code.

Campus Resources, of course, there's the AI Center of Expertise, now Solutions Hub.

There is the CITL's Innovation Studio, which has access - subscription-based access to popular AI tools. If you're curious to see what ChatGPT 4.0 can do, that's the place to investigate.

There's also the Writer Workshop. We have some amazing humans on campus that support writing and can help you in your writing process. So please check out the Writing Workshop. They offer in-person and online consultations.

So, from the library, we have a Generative AI LibGuide.

It discusses - provides that overview of what is generative AI, offers a summary of some of the copyright considerations, legal issues, and ethical considerations.

We also talk about how to protect your work. Popular tools, teaching with AI, as well as a link to a bunch of resources that have been developed across units on campus.

We also, at the very bottom, we have questions, people - point people in the library that you can reach out to, as you're playing with AI or considering how to use AI, that can help answer and field those questions.

I'm really excited. We've developed an "AI in this Course" Canvas module.

This is designed for instructors who want to help students understand their AI policy for their specific class.

This was written with writing-intensive classes in mind, so mostly for the humanities and the social sciences.

You can find it in the Canvas Commons by searching for "AI in this Course" or "University Library," and you can see what the course looks like by self-enrolling in the library Canvas Exchange through the link in the PowerPoint.

Finally, there's more workshops.

If you like what you see today, we have a couple more AI workshops.

There are, of course, some recordings from previous semesters on the DH at Illinois Media Space Channel.

Just to give you a preview of what these incredible workshops are, we have Enhancing Research with Copilot and Perplexity AI.

So if you're really excited to get under the hood of some of these generative AI tools, that's next week.

If you are in the social sciences and you're thinking about, how does AI influence the qualitative data analysis process, Jess Hagman is offering two workshops, one in-person, one online.

AI for Storytellers is going to be a fun one, quite selfishly.

This is with Siobhan McKissic, who is our visiting Materials Research Librarian, and Michael Curtin from the Office of CIO, where we talk about image-based AI.

And of course, October 22, "Copyright, Text Mining, and AI", with our Copyright Librarian Sarah Benson.

So to close, we want to emphasize that AI is a tool, not a solution.

It can support, it can enhance, but it is ultimately up to you and how you evaluate these tools, when and how and if to use them.

To enhance your own work, your own thinking, your own ideas, and your own creativity.

So for questions, if you have questions after this workshop, if you're like me, you think about them five minutes after you leave the room, or you don't feel comfortable asking a question in the Q&A session, which we're about to start.

You can always email me, my email is "maryton".

I'm the point person for acquiring text data - text mining data, digital approaches to humanities materials, adapting the "AI in this Course" Canvas module, and providing in-class instruction with sessions using DH tools.

There's also our amazing Copyright Librarian Sara Benshon - Benson, who can field questions about copyright as well as open access and OER.

And with that, I am going to stop the recording and welcome questions from online and in the chat.

But thank you so much for joining us today synchronously or in the future. We're glad to have you here and excited about the conversation moving forward. Thank you.

