

What We're Learning About Learning

A podcast by the Center for New Designs in Learning and Scholarship
Georgetown University

Season 3 Episode 5

Chat GPT and AI in Higher Ed

[MUSIC PLAYING]

EDDIE MALONEY: These tools are continuing to grow and change. What it will mean for us to regulate, adapt, or integrate will also need to change as well.

KIM HUISMAN LUBRESKI: That was Eddie Maloney, the executive director of the Center for New Designs and Learning and Scholarship, also known as CNDLS, at Georgetown University. Welcome to this episode of What We're Learning About Learning, a podcast about higher ed teaching and learning, created and produced by CNDLS. I'm Kim Huisman Lubreski.

JOE KING: And I'm Joe King. This podcast series has covered experiential learning, ungrading, concurrent hybrid practices, and now we are diving into the oceanic world of artificial intelligence tools and how they are affecting how we teach and learn. At a recent forum we hosted on AI in Higher Ed, faculty shared their excitement, concerns, and approaches to integrating AI into their teaching practice. In this episode, we'll bring you highlights from their conversation, as well as interviews and statements with other GU faculty.

KIM HUISMAN LUBRESKI: There have been so many ChatGPT articles and insights that it's hard to remember why it's dominating the news cycle. So we thought we'd start there. We spoke with Georgetown Computer Science Professor Grace Hui Yang about the magnitude of the development of AI, specifically large language models like ChatGPT and GPT 4.

[MUSIC PLAYING]

GRACE HUI YANG: This is not just another tool. ChatGPT is more like another kind of intelligent being.

JOE KING: Nick Lovegrove, from the McDonough School of Business, calls this the biggest story of the decade.

NICK LOVEGROVE: About five years ago, McKinsey set up a unit focused on AI and its implications, and it had 15 people in it. And I checked last week. It now has 1,500 people. And I think it's a reflection of two things. One is that this is the topic at the moment. And second of all, it doesn't seem to be replacing people. It seems to be creating opportunities for people, at least for people who are focused on helping figure out what to do with AI and its implications.

KIM HUISMAN LUBRESKI: The unusual speed with which this transformative technology has rolled out has also impelled tech leaders, educators, and others in March 2023 to call for a pause to its development and release. Georgetown faculty who followed this development closely are also pointing to the need for government regulation and clear policy making.

As faculty consider how to navigate AI with respect to their teaching and learning practices, the public discussion of the abilities and uses of AI has become increasingly complicated. Georgetown Computer Science and Linguistics Professor Nathan Schneider helped frame the questions we started with.

NATHAN SCHNEIDER: The discussion over these past couple of months, from my view, in NLP, is, people are asking, what is it good at? What is it bad at? What are its limitations? And what does it all mean? So ChatGPT is an instance of a large language model. And a large language model is about synthesizing text based on text it has already seen before. And it is trying to synthesize text that it thinks is plausible, that is likely to occur, as opposed to text that is unlikely to occur.

But what I want to emphasize is that there is no mind behind this system. It doesn't really know what these concepts are that you're mentioning. It just has access to a vast amount of text that it has been exposed to and also some interaction with humans who have been trying it out and giving it feedback, in some cases, on, was this a good answer, or was this a bad answer.

What's remarkable about this system is that it does synthesize really plausible sounding text to a lot of kinds of questions. The text is highly fluent. And both in the sense of grammatically, but also stylistically, it's able to imitate different styles that you ask it to produce. It's also really impressive to me how it handles the interaction, so the chat aspect of it, that it's not just responding to your latest prompt necessarily.

It has the history of the previous step turns in the interaction. It often is able to stay on topic, and you can give it a short prompt with the implicit understanding that has come from the prior context. And it is good at retrieving topical concepts that it has seen. So if you ask, what should I cover in my natural language processing course, it is able to list a bunch of topics that it has seen mentioned as topics in natural language processing courses.

It also can produce code and follow instructions about manipulating text in systematic ways. There's great potential, I think, for helping people explore topics and create content. It's obviously very exciting that we have these AI systems that are a lot more natural sounding than they used to be. It seems like it's good at creating boilerplate that you could then edit or to customize to your needs. So maybe we don't need to do as much writing from scratch in the future. But I think it's important that we are able to take responsibility for the content, the ideas that we're getting out of the tool.

KIM HUISMAN LUBRESKI: When we consider how capable AI will become in a very short time, we can understand faculty whose first instinct might be to ban ChatGPT's use altogether, which some schools have done. Instead of forbidding its use, however, we might investigate which questions AI poses for us as teachers and for our students as learners.

JOE KING: How can we think anew about producing and sharing ideas? How might the task of communication change and grow? As we mentioned, CNDLS recently hosted a forum on this subject. And at the start of the event, Eddie Maloney offered some framing.

EDDIE MALONEY: I think there are three ways we can think about responding to this right now. First is really to regulate, and you will have seen that that is starting to happen in places. You get policies either at an institutional level or at the course level that can try to limit what students are doing with these tools to make sure that we have in place a set of rules people who are writing and producing to adhere to as we think about using tools that can create text for us.

It was our response to the word processor. It was our response to the internet. We often get into that place where these tools come out, these things become available to us, and we need to figure out how we limit them because they challenge some assumptions that we have about the work that we do in our relationship with our students.

The second approach that I think we move into and maybe move into relatively quickly is to adapt to this space. So to create composition assignments that cannot easily be, at least right now, replicated by AI or by some of the tools that we have out there right now. Likely, that will change and the things that we think that we can do to adapt will probably be adapted to by the AI tools. But it's one approach, right?

And then the third stage of this is really figuring out how do we integrate these tools into our space in the same way that we had to integrate the calculator or the spreadsheet or the internet, how we had to integrate online into the environment, even though, 12 years ago or 10 years ago, we were afraid that online was going to take over our jobs and the institutions that we care so much about.

So we probably get to a place where integration fairly soon will feel like the norm, but it's probably worth also noting that as you probably have been reading, these tools are continuing to grow and change as well. And what it will mean for us to regulate, adapt, or integrate will also need to change as well.

KIM HUISMAN LUBRESKI: You can check out the full recording of the forum event on our website, cndls.georgetown.edu. Let's listen to some of the panelists' thoughts at that time, as well as conversations with other faculty, some of our own discoveries, as well as students' experiences with ChatGPT. We hope this helps further your consideration of how you're interacting with this phenomenon in your teaching practice.

JOE KING: These developments are sure to transform many processes and tasks across industry, government, and of course, education. We were especially eager to hear faculty hypothesize about possible complications of the tool and others like it. Here's Grace again, followed by Andy Zeitlin, from the McCourt School of Public Policy.

GRACE HUI YANG: We really should get our faculty and our students more involved into this discussion. It's not only like AI and my job. It's really AI and America's future or even humanity's future. I do think Georgetown faculty and student should be more actively involved and think for mankind. It needs regulation.

ANDY ZEITLIN: In the public policy problem area, I think there are two big buckets of the types of problems where we want to be thinking about this. The first of this is to help our students to think about the application of-- or how these models and their output is being used to guide decision-making. And there are analogs of this in the use of, for instance, predictive models in lots of other domains. And then a second is, actually, how this relates to the incentives for the production of knowledge when people will no longer go to your favorite professor's website to get the answer for how to do something, but instead can get those things off the web or through these tools in ways that don't pass credit on to the people who generate the knowledge in the first place. What does that mean for the incentives to be the person who you go to for the answer to problem X when that's always going to be some chat tool for the future?

One obvious application of that is in the writing of code where I think these things really are productivity enhancing, potentially, for a lot of people. And you've seen my syllabus language around the kind of skeptical use of that.

KIM HUISMAN LUBRESKI: Check out our show notes to see Andy's syllabus language and that of a few other professors, including Nathan Schneider's, whom you heard from before.

JOE KING: Nathan Schneider shared some limitations of implicitly trusting this technology.

NATHAN SCHNEIDER: So it'll pick up some concept and realize it's related to some other concept, but it will get the details wrong. It may invent something that's totally false, but sounds plausible. So if you ask it

to give citations for the things it's saying, it can make up fake bibliography entries, and it'll be maybe mixed with correct bibliography entries. So, again, it's going for this plausibility, but it has no capacity to tell us when it's guessing versus when it's correct. So that's something to be concerned about.

Also, plagiarism not just in the sense of students copying from ChatGPT, but in the sense of ChatGPT copying from all of the data that it has seen in its training. So we don't know if ChatGPT utters a phrase or a sentence. We don't know if it has synthesized that from scratch or if it's repeating that verbatim. OpenAI built in some attempts to protect against it spewing toxic language or expressing harmful opinions, but there are ways to circumvent that if you really want to. And there may be some more subtle kinds of biases that could be harmful in terms of what it actually says. So those are some concerns. And remember that there's no mind behind the curtain that's actually making it even necessarily consistent, let alone correct.

KIM HUISMAN LUBRESKI: Andy Zeitlin from the Policy School highlighted the equity problem that the technology presents.

ANDY ZEITLIN: As we see some of these models go under subscription models, as opposed to advertising models, I actually think we might worry, on the other hand, about the question of equity of access among our students to these tools, and not just think about the problem of restricting that access.

JOE KING: Eric Saldana, a graduate student representative from Georgetown's graduate student organization GradGov, expressed serious concerns about students with more financial resources having more access to higher quality AI tools. He worries about a not too distant future where there is an increasingly big divide between low quality AI and high quality AI that perpetuates inequities.

KIM HUISMAN LUBRESKI: With an awareness that we can't solve for every issue, at CNDLS, we've mostly focused on how faculty can design assignments along a range of AI awareness, from banning it for a reason-- for example, an oral exam in a second language would preclude the use of ChatGPT-- to designing an assignment to interact with ChatGPT intentionally, as Nick Ludgrove describes here.

NICK LOVEGROVE: I'll give you what the assignment that I set my students yesterday in a management course. I set them to write a paper about the respective business models of two businesses that compete, Zara and H&M, for those who follow the fast fashion arena. This is what I was going to ask you to do. What I'm actually going to do and what I have done is I've put all the questions into ChatGPT. This is what it's come up with. I want you now to critique what they've said. Identify where you think they're wrong, identify where you think they're right, and improve upon what they've done. That's my first experiment with trying to be additive to what's available, what they could easily do themselves.

JOE KING: Andy shared how he was thinking of some of his pedagogical considerations.

ANDY ZEITLIN: Thinking about how to teach our students to leverage these tools to do those core functions of their jobs more effectively may be an important set of skills that involves things like the engineering of good prompts, so how to ask the right questions in order to get the right answers, and knowing what the right answer is when you see it and learning to detect problems in these things. And then lastly, on the kind of pedagogical problems, I think these are not unique to the Policy School, so I'll be brief. I think these are most challenging for us when we think about how do we make sure that we are giving students a context in which we can encourage them to really think about both the development of knowledge and critical thought in ways that can be baked by the output of these systems pretty well. So how do we make sure that students really know how to recognize when the fundamentals are right and wrong?

On maybe the more positive side, I think it may well be the case this is a rapidly evolving landscape, as we've heard. I think part of that evolving landscape are tools for detection and ideas of watermarking and other things that are part of this as well.

KIM HUISMAN LUBRESKI: Our panel also featured students, including Camber Vincent, a current junior here at Georgetown in the School of Foreign Service, where he studies Science, Technology, and International Affairs. He also has minors in Environmental Studies and Sociology and serves as the current president of the Georgetown Student Association. Camber spoke on behalf of the student body. At the time of the presentation, he was in the middle of completing an assignment that asked him to use ChatGPT directly, as he describes here.

CAMBER VINCENT: As was spoken about in the presentation, as this tool continues to develop, it is going to become reality of how students learn, whether you like it or not. Thinking about ways to integrate it meaningfully into your practice is something that I think is really important. One of my professors, just to highlight her because I'm doing this assignment actually this weekend, so I have to turn it in on Monday. My professor for Life on the High Seas, so Professor Rebecca Helm, she's new in our Earth Commons Institute, wonderful professor. Our assignment for this week is that we have to think of a high level critical prompt in the field of what we've learned about the high seas. We asked ChatGPT to create a four-paragraph essay at the collegiate level on this topic, answering the question that we posed to it. Once we generate that response, we have to go back through, critique where it's wrong, find out more depth, find out more information, and basically prove why it's a very poor tool to use as an essay writing prompt. I already generated the prompt, and I've already gone through it. And just in the first paragraph, I have like 18 different notes on where it went wrong in five sentences. So it's not a phenomenal tool for producing this high level content, but it's interesting for her to demonstrate to our students hands-on how the AI generative model works. So we actually had a 15-ish minute session in class discussing how it generates answers to the questions you pose it, and then going through and having our students actually directly critique that model and understanding where they can find that information and create stronger prose.

JOE KING: Remember the graduate student representative Eric Saldana we told you about? He shared an example of how students might design their approach to assignments using the tool, which happened in an education neuroscience class, for which students have to write summaries of the class readings that aren't graded on writing, but rather mastery of content. These are then posted for the whole class to see. He wanted to ask if he could use ChatGPT to, quote, "punch up" his own summary, which he did and he felt that the quality of his writing improved. But the unclear expectations around the use of such a tool in this way caused Eric a lot of anxiety. At the very least, being open about expectations makes academic integrity a lot more doable.

KIM HUISMAN LUBRESKI: Camber, the undergraduate student government representative, talked about the ways students are actually using the tool and why faculty should make it clear how ChatGPT can be used in their classes.

CAMBER VINCENT: What students are using it for right now is as a generative tool for ideas. They will ask ChatGPT to structure outlines of possible arguments. They'll ask ChatGPT to give starting points when there's very open-ended questions where a professor basically gives students free leeway to decide a topic inside their course.

Students may choose to use ChatGPT to help narrow down on a specific idea for an essay subject. I know that I myself in late night study sessions in Lau have used it to think of specific words when I can't remember the exact word I want to use. I'll ask ChatGPT something general about the topic and hope it pops up with the right vocabulary that I'd like to use in my essay. And it's this kind of refinement tool. Students are very concerned about the implications for this, for the Honor Council system, myself included. We also run a student advocacy office that works with students who are in violation of the university's student code of conduct or the Honor Council. And students, by and large, are very concerned about what this could mean for their grades if they're accused of using ChatGPT, if they're caught using ChatGPT, what it means in their sense as students.

JOE KING: Nick Lovegrove, the business professor, builds on this concept to formulate the question of what students will need and not need to do now and in the future.

NICK LOVEGROVE: It seems to me this gets to the really fundamental question that maybe is implicit in everything that we all have been saying, but maybe this is going to require us to think harder than we have for quite a while about how do people actually learn, and what mechanisms, what process do they need to go through in order to learn. And I think we would probably all hypothesize they need to have done the reading, they need to have practiced the writing, they need to have gone through a learning process.

What these tools will do will supplement, and in some cases, offer the opportunity to replace some of those activities. And we absolutely have to worry about whether there'll be some diminution of learning. And that seems to me the fundamental thing that we're going to be grappling with, and I suspect we're all going to be grappling with for the rest of our careers, really.

KIM HUISMAN LUBRESKI: Many educators are reflecting on the depth of knowledge that may be sacrificed with this new technology. For example, a student uses ChatGPT 4 to generate a summary of an assigned text and then writes an essay in class based on the AI-generated summary. The student has missed out on a certain kind of depth provided by reading the book. Grace Hui Yang provided an example.

GRACE HUI YANG: If our students are not going through this process of laboring in learning, they're not going to get it. When I was young, my dad gave me the calculator [INAUDIBLE] from Japan. Well, I can tell you what. I never good at adding big numbers because of that, because of technology. And even I'm a professor in computer science.

I'm pretty good in other math, but not in that part, not that, because every time you're adding a big number, I will use that. And if you ask me to just do it by hand, I'm not good. And I'm just never good. And some other friends of mine, they're actually very good because they force them to use their brain, their human brain.

JOE KING: It's complicated to consider assignment design with a tool that is constantly and very rapidly evolving, not to mention the need to balance concerns about academic integrity when the very notion is being redefined, when we must prepare students with the tools of the world they live in and will live in. We should consider how to use these tools to generate syllabi or eventually provide feedback or even grades to students, which brings us to the pedagogical consideration of identifying what we want to assess that is the product of the student. Nathan Schneider points out how AI can level the playing field for speakers of different languages.

NATHAN SCHNEIDER: There are different possible use cases, and one of them is polishing writing, which is especially useful if you're writing in a second language, and you're therefore at a disadvantage, even if you have equally good ideas. So I would say that this should be considered as one of the possible use cases that is not necessarily invalid, but depending. And so I think instructors should think about this for their syllabus.

So in mine, I say that it's OK to use tools like Grammarly for polishing. I forget how I worded it exactly, but I said don't use ChatGPT just to generate the entire thing. But it's OK to use it as a language aid, essentially. Of course, that's because for my purposes, I'm more interested in the content. Teaching a second language, of course, there would be different considerations there.

But this has even come up in our discussions about policies for conferences and conference publications. The computer science field is heavily dominated by English as a second language speakers. And there's certainly an equity issue if they're being asked to write in English, and they're not being allowed to use any sort of technological aid for that.

Part of what I think will be asked of us as faculty members from our students is how to ask better questions of the tool. And so teaching to that is one way of thinking about integration. How do you create questions that are more meaningful? That's not dissimilar to our librarians, who wonderfully help us figure out how to ask better questions of the research tools that we have in place, or if you're as old as I am or older, you probably remember with the internet, when it started, the kinds of questions you first asked of the internet in 1993, are not the questions you would ask right now.

JOE KING: We've been thinking about corollary moments in history, like after the Industrial Revolution, when art and design changed to reflect intricate and curvaceous forms as a display of what humans can do that machines couldn't. We see today's transformation as a similar moment.

GRACE HUI YANG: First of all, don't be afraid of the change. And I guess, in the very beginning, when we say, OK, well, the student, you cannot use ChatGPT for homework or exam. We are not accepting any answer generated by GPT or something like that. So I think that's the very beginning of people's attitude. And very soon, you'll realize that you cannot resist that, basically, right?

And also changes are constant in human history, like television, car, camera, Netflix. You see why we're having this conversation. We, as a human race, will be pushed. I think we will become smarter. It will push us to think how to be different. You have to evolve. So I think that's good for us in the long-term.

KIM HUISMAN LUBRESKI: If AI is so capable, as Grace describes it as this other being, then the questions arise-- how are we going to stand on its shoulders? How will it catapult us into new capabilities?

GRACE HUI YANG: First of all, we need to understand these new dynamics. Now we not only have the professor and the student. We will have AI TA can be the AI private tutor. So it's this newcomer. This new type of being will be among us. So I think we should really think how can we welcome a newcomer and how to build this collaboration. Each of us will become a superhuman with our own team of AI friends. Instead of doing our job, you have to be a team leader, right? You have to think about this upper level architecture. So I think push us. Each of us have to think from this management point of view. That's something definitely will make each of us a super human being.

JOE KING: Grace emphasized that students will still need to know how to write, or at the very least, they need to know the characteristics of effective writing. All the more so because tools vary in what they are good at. For example, Google is a pretty precise search engine, while ChatGPT is more creative and thus

more unreliable. But it will only continue to get better. As a being that won't die, it has an infinite capacity to improve.

GRACE HUI YANG: AI can do exactly what we can do, with some certain exceptions. We can multiply. We can have babies later on, right? We have emotions they don't. And we have intrinsic motivations. That kind of intrinsic motivation, I don't think the AI has. No, they don't. They don't have the [INAUDIBLE]. So whatever reward we give it to it, it's like we are feeding them. They're trying to mimic human curiosity.

KIM HUISMAN LUBRESKI: Here at CNDLS, we wondered, with the immense capability of AI tools, whether the fundamental relationship between teacher and students would change.

JOE KING: We were encouraged and somewhat comforted by the range of faculty and student responses that all boiled down to, yes, and change is OK. Ultimately, these tools will help us be better at teaching, at learning, and at asking questions. There's no doubt that this conversation will continue to change, as the technology and our use of it changes. We'll continue to keep our eye on this space and hope you'll join us for future iterations.

We also hope you found some good food for thought in today's episode of What We're Learning About Learning. This episode was made possible by many people at CNDLS, including Molly Chehak, Eddie Maloney, David Ebenbach, Sophia Grabiec, Eleri Syverson, Noah Leiter, and Stefanie Chae. And a big thanks to our faculty who contributed to this episode-- Grace Hui Yang, Nick Lovegrove, Nathan Schneider, and Andy Zeitlin. Thanks also to Georgetown students Camber Vincent and Eric Saldana. Thanks, as always, to Milo Stout for creating original music for the podcast that has impacted our society more than any AI system ever will.

For more information about our podcast series and our guests, check out our show notes, where you'll find links to previous episodes, information about how to share your thoughts and ideas with us, our website and blog, and other resources. Again, I'm Joe King.

KIM HUISMAN LUBRESKI: And I'm Kim Huisman Lubreski. Thanks for listening.

[MUSIC PLAYING]