

Episode: AI Hype, Classroom Experiments and an Educator's Dilemma

Series: [AI in the Classroom](#)

Produced by: The Digital Futures Institute, Teachers College, Columbia University

Featuring: [Alex J. Bowers, Professor of Education Leadership, Organization & Leadership Department](#)

When you have machine learning, the machines are going to learn from the data that exists, that we give them, unless we do something about it. So they will learn our biases and all of our problematic behaviors and serve them right back to us. I think that we need to be very cautious about: Who's been included? What data have these models been trained on? Which, given the current corporations, we don't know for most of OpenAI and ChatGPT, they won't tell us. And I think that's highly problematic.

[series graphic and brief futuristic music interlude]

In looking at the surveys about the use of technology and advancing technologies like AI across multiple industries, education has shown continuously that managers and educators in education have a lot of interest and they really want to talk about AI, but its deployment across education lags behind most of the other major industries. There's a lot of potential for what the media likes to call as “artificial intelligence”. We don't have artificial intelligence, but I think there's a lot of work yet to be done for education.

In education over the last 30 years, through a lot of really hard work across a lot of different research domains and practice domains such as learning analytics, as well as the education technology domain, there's been a lot of really interesting progress in using educational technology to help teachers really be able to teach in stronger and more effective ways. Where the technology is used to help solve the problems that they may have in the classroom and be able to help learners have fun, engage more, to provide useful data back to the teachers. Over the last few years, when I talk to teachers about their use of technology in the classroom, when it comes to things like instructional tutors as well as engagement detectors and many other kinds of technologies, teachers continuously tell me that it helps them do what they came to teaching to do more effectively.

That the technology itself that comes from a lot of the research that we do here, such as a Teacher's College, is helping them understand what it is about their day-to-day teaching practice that's working with students, what needs to be retaught, where students may be engaged or disengaged, and how to help understand their own practice in a way that helps them not have to worry about writing curriculum, writing tests, so they can work on really what they came into teaching to do, which was teach students in real time and help the students understand that curricular item and that issue right now in a way that works with the students and what the student knows, and does, and helps the student succeed.

And all the while many of these technologies are also based around educational gaming as well as the idea of serious games and education. But why shouldn't kids be able to come to school

and play games, but learn faster than they ever have, with a teacher who really has the information coming to them that helps the teacher understand exactly where each student is? And so rather than have to pitch to the middle, the teacher can differentiate in really strong interesting ways and use the technology to help them do that.

So here at Teachers College, I teach, and I'm a professor of Educational Leadership and a lot of my research and instruction is around helping school leaders use the data that we collect every day in schools more effectively. And so a lot of my research is on data dashboards and technologies that can help us understand the patterns that are in the data, but then at the same time also, how do we talk to practitioners and talk to each other about the evidence that matters in our organizations?

And so technology in my class here at TC really not only includes Canvas, and has for many years. I launched the online non-credit course here at TC where we offered the first digital badge here at Teachers College so that it could be an open class globally, so students didn't have to enroll in at TC, they could be administrators, teachers, I've even had high schoolers take the Leading With Evidence course and using these kinds of technologies.

But also in many of my classes, I use open access and open source kinds of coding environments for when it comes to statistics and data visualizations, things like R, the integrating these kinds of technologies into the classroom is really helpful. Yeah, I'm the faculty advisor to the Teacher's College Student Data Science and Education Association. And I've been very interested in these new large generative models that have come out, where the media has been talking about them as "AI", or the term I heard last week, the "AI hype bros"?

With the death of most of the cryptocurrency, the hype had to go somewhere. So it's come into now AI. So there's a lot of hype, because the technology is fascinating and I'm fascinated by it. And so with the release of many of the open source models such as the AI image generators, these large models that generate lots of different image . . . that can make these images.

It was fascinating seeing not only the hype but the development of the use of these kinds of large language models as well as the generative image models. And so immediately in my classes, I stream ChatGPT right into the class to show everyone what it was. But I tried to show them the many different things I was learning about what these models could do, some of the really weird things they could do, but also the idea of: What happens when you put the entire first assignment of the class into ChatGPT and ask it to write the paper, what does it look like?

And immediately in one of my classes, students asked, "Are you showing us how to cheat right now?" That was one of the first things they said. And so we started to have some of that conversation.

For spring semester in my Advanced Methods course, I included a statement in my syllabus that said that students are encouraged to use these large generative models, but that they would take responsibility for the fact that these models lie, that they're biased, that they can create

offensive materials and just give incorrect information, but in a very convincing way. And so I included that statement in my syllabus to see what would happen because this is here and the models are going to get better. They're not going away. So I figured I needed to start having students encounter this kind of work and see what happens.

And so in that research methods course, there's been times I've got a certain set of prompts that will make funny outputs around statistical concepts. And so I implemented that in class and had the students all open up ChatGPT, and they all had an account, which was part of the check because I'm like, "Okay, everybody open up ChatGPT," and they all opened it right up.

And I gave them the prompt and then told them the different differences to put, the different stems to put in it, to make them different funny prompts around statistical concepts, like the difference between different kinds of T-tests or P-values. And we played with it to really show them that I was okay with it. Let's see how this goes. I asked the students to review it to say, and if they're going to use it in an assignment to say that they used it.

And what I've found so far, and we haven't gotten to the end of the semester yet, and I plan to ask them at the end of the semester how they thought it went, but each time they've used it, so far, it's really problematic for the students. It hasn't really helped because like ChatGPT, which is built on GPT-3 Turbo, if you don't pay, if you pay you get GPT-4, but the free version, the outputs come out very flat. And I can see the issues right there.

And if the students ask very specific questions, especially for methods, it gives the kind of response you'd get if you just read Wikipedia articles, it can't really do anything better than that. Now that's going to change and it's already changing, and students will get access to that kind of ability soon. But for now, I think most of the students in my class have learned that it's just quicker to learn it yourself.

The free version of ChatGPT is based on GPT-3 Turbo where if you pay, you get access to GPT-4, which is now the more advanced version. But we already know from the research literature around GPT-3 Turbo that it's only about 63% accurate across a variety of tests. It's better at some things, worse than others. But GPT-4 is better than GPT-3 Turbo and soon we'll have even more.

Students so far, are finding it sometimes helpful? It's easier I think for technical subjects for students to just search online like you normally would. And then you don't have to parse whatever it says, and if you ask it for references, it just makes up DOI links to articles. They're completely random. They'll lead to something, but who knows? And plausible sounding article titles with plausible sounding authors.

I'll give you an example, so in March I gave a talk for the Student Data Science and Education Association on large language models and GPT. It was really well attended, and the talk is available open access online, and I publish it online under a creative Commons license, so anybody can take it and mash it up.

But in that talk and in those slides, I show the different kinds of prompts that you can create to get ChatGPT to not only do things it's not supposed to be able to do and force it to do things that the Open AI says should not be allowed, but also get it to just write your papers and write them well, rather than very flat. Or to get better at getting the facts correct rather than lying, or what the AI hypos would call "hallucinations". That's a marketing term. When you hear that the AI hallucinates, that's a marketing term for lying and making up facts.

It can be a useful tool, but you have to know how to train it, how to give it a specific set of information to help build out a logical path that will get it to the answer. And lots of times it's just easier just to write it yourself.

I have multiple instances of students now using ChatGPT for their assignments, which is allowed. For the bigger assignments, I allow Revise and Resubmit, so I'll give comments and then provide an opportunity for students to revise or resubmit within two weeks.

When they've said they use ChatGPT, the assignments just aren't great. It seems like they're using it more right now because they run out of time on the assignment and they're hopeful that they can then use ChatGPT to just fill in some of the blanks and just get the paper done and submit it. And I don't know if that's better or worse than what I think they used to do, which was just write until they ran out of time and they maybe give me a hot mess of an outline and then give it to me. (laughs)

But this is going to change as these large language models get better and better at this. And as students get better at writing prompts. For example, the way to get ChatGPT to actually write a good paper is don't just ask it to write a paper because you have to train it. You have to teach it. And so I show in those slides for the DSEA talk, what I do is I'll give it the assignment, the whole thing, copy, paste, but I'll ask it at the beginning of the prompt, write a rubric to grade this assignment, so it ingests that information and then I ask it to write the assignment, once it's written that rubric. Then I ask it to grade what it just wrote, given its rubric, given the assignment. And it usually gives it itself a "B" because that's the level that ChatGPT writes at.

But then you ask it to give comments about how to improve the writing of the assignment and give specifics by sentences. And this is where, actually, I think these large language models are a really interesting way of expanding access to writing instruction. So ChatGPT is a fairly good 10th grade writer. It writes a very strong structured five paragraph essay. And for students who need that kind of support, it's now democratized for them. Anyone has free access to that kind of writing instruction where you can put in the assignment, you can ask it to create a rubric, then you can load in your assignment, have it tell you, have it grade it, which will be very close usually to the professor's grades. And then tell you where not only which sentences are problematic, but why and the reasons why you would restructure those sentences or the structure of the logic of the paragraphs, and what those rules are in good writing instruction.

Now, yes, is it easy for then a student in the next prompt to say, "Revise the paper given above?" Yes. But at the same time, I think for many early stage graduate writers, it is an addition

to the writing center, something like ChatGPT is this additional resource that I think can be very useful. Now, whether or not it should be encouraged, I don't know. I don't think anybody knows yet, but students are going to use it just like they use Grammarly and just like they use spell checkers.

Reading the research and comments from the natural language processing research literature, and people who do natural language processing where they pull information from the web and what people write on the web. And one of the things they talk about is that over the last five years in writing on the web, grammar mistakes have disappeared. They're not there anymore.

And it's not that people became better writers, is that everybody wrote uses Grammarly and spell checkers. And I think we're now going to next move into the next era where perhaps ChatGPT is your first round of edits. It's an editor on your writing first to tell you, this doesn't make logical sense. You could revise over here, you could make a stronger argument over here. Now, whether or not you then have the language model rewrite it your itself or do it yourself is, I think that's where this split comes from.

In graduate school, all of us need to start with derivative, terrible writing, because that's how you learn. We all start there. And if we skip that step, I don't know what that's going to mean for our writing down the road. We'll see what happens over the summer when this doesn't happen in the cloud anymore. It can happen just on your computer or in your phone. You don't have to pay for it.

And I think over the summer, students are going to get a lot savvier at how to use these models. I just don't think students have had enough time to process what this is. They've been trying it and it doesn't quite get them where it needs to go when they just ask it to do a thing. I think some of them are getting better and better.

And then the question becomes, should we teach how to use it better and to help them so that they build these skills early, so that TC is out in front on these kinds of technologies and it's our students who are leading? Or do we just not talk about it? Or do we do what I'm doing in my class, which is just play and let all the flowers bloom perhaps and see what happens?

And so I have the course in the fall, Applied Data Science in Organizations and Leadership, it's a class that we launched just a couple years ago. It's quite popular. It's a lot of fun to teach. There's a lot of coding in R, and between the large language models and coding environments like Copilot, I don't know what I'm going to do yet for the fall. I have applied assignments where they bring in their own data sets and they do their own coding, but just to learn how to do each of the initial applications is very much, copy this code, apply it to the same data set, get the same result.

It's very much, I treat it like a cooking show where we bake the cake and I show you what the cake looks like and try to replicate that, just so they can figure it out. But when you can with just a sentence, just talk to the computer and ask it for the thing, and it writes the code for you, and

this is how and where coding is going, and I'm trying to help our students learn how to do this work in an applied manner so they can go out and get jobs and do the work.

I have to think very deeply about what I'm going to do in the fall. I'm not sure how to change that yet because I can't make those decisions now, because who knows where we'll be by August and September? And so I think I'm going to have to be a conversation and I'll probably going to lean more into, there's already some co-constructed curriculum in the class itself, and I may need to open that up even more so that we have three partners: myself, the students and the models.

I guarantee you that 95% of our students at the college have tried it and have access to ChatGPT, have an account. But, that's where I'd put it is maybe 95%. If it's less than that, I would be very surprised. And so are students using it in their classes? Sure. Are they telling their professors? I don't know. I'd love to know what the conversation is at the college.

And I think it would be really useful for us as a faculty to come together around just what is the conversation we'd like to have at the college with our students, with these new technologies that are coming out that our students are definitely using? They're definitely interested in them. And incorporating them into our work I think is vital because they're here and they're not going away.

In thinking about the changes that AI is bringing about, I worry about the hype and that... I think the best joke I heard was on social media where a natural language processing researcher who's been doing this kind of work for decades put on social media that his aunt texted him and he's like, it's as if you're an ornithologist and your aunt texts you, "I heard they solved birds." And it's like, I just found that hilarious because it's kind of where the hype is at. It's like AI's here, it's solved all the problems. And it just is not that. And it's a lot of conversations to have around how these models are trained, the data they're trained on, they're highly biased.

We live in a highly segregated society that's dominated by corporations in the western world. And these technologies are trained on the data from the Western Internet. And so I talk about in my work, when you have machine learning, the machines are going to learn from the data that exists that we give them unless we do something about it. So they will learn our biases and all of our problematic behaviors and serve them right back to us. And I think that that conversation is a core conversation that needs to be had, because there are voices that are excluded. There are communities that are not privileged in these data sets. And if these models are to be used in a good way, in a way that's very positive and can help us really integrate these kinds of ideas and go to whatever that next step is, I think that we need to be very cautious about who's been included.

What data have these models been trained on, which given the current corporations, we don't know for most of OpenAI and ChatGPT, they won't tell us. And I think that's highly problematic, and especially when it comes to, if we begin incorporating these technologies into their courses, not having that information then bakes those biases into our courses if we use these technologies.

I think there's a lot of research on this right now and the bias across these models, and I think we'll know more, hopefully over the next easily six months, the year, two years. But I think this is a question that needs to be at the center of the conversation.

There are models, they don't seem to work as well yet as ChatGPT that we know more about what the training data are. But even the idea of these opaque models where we don't really understand how they work.

So, not knowing what they were trained on and then not understanding exactly how they work is highly problematic. So in my work, I have worked with the OECD as well as others to talk about how, especially in public schooling, if there's a technology in public schools that makes a recommendation on a student or a decision on a student, the algorithm must be open, public, and auditable. To do otherwise is unethical. A taxpayer paid for it.

Just like I can ask the teacher of my child, "Why did you make that decision?" Or I could hire a CPA to audit the budget. I should be able to hire a data scientist to audit the algorithms to make sure that they're not biased. Because we know through where we are with many of these predictive algorithms already across many industries that they're highly biased because they've learned from our bias system and they just serve it right back. So I think that this is the central issue going forward, especially if we incorporate these technologies into our courses and encourage our students to take them out and into schools, for instance.