This is supplementary material for Maggie Fernandes and Megan McIntyre's webtext, "Giving Voice to Generative AI Refusal in Rhetoric and Writing Studies," published *in Kairos: A Journal of Rhetoric, Technology, and Pedagogy, 29*(2), available at http://kairos.technorhetoric.net/29.2/disputatio/mcintyre-fernandes/index.html

Everyone's Writing with AI (Except Me!) Interview with Dr. Michael Black

In this episode, Dr. Michael Black spoke to Maggie and Megan about his 2022 book *Transparent Designs: Personal Computing and the Politics of User-Friendliness* and how writing teachers can respond to AI hype.

Megan: Hi, I'm Megan McIntyre.

Maggie: And I'm Maggie Fernandes.

Megan: And this is *Everyone's Writing with AI (Except Me!)*, a cheeky writing studies podcast about refusing generative AI in higher education and writing studies specifically. All of the articles and texts that we talk about today will be in the show notes. So, feel free to follow along.

Maggie: Today we're excited to talk to Dr. Michael Black. Dr. Black is an associate professor of English at the University of Massachusetts Lowell. Dr. Black specializes in studying the intersections between computers and writing. His 2022 book *Transparent Designs: Personal Computing and the Politics of User-Friendliness* documents the rhetorical history of usability and highlights how the concept of usability has been leveraged historically to smooth over conflicts between the rhetoric of computing and its material experience.

Prior to joining the faculty at UMass Lowell, Michael served as an associate director for the Institute for Computing and Humanities, Art, and Social Science at the University of Illinois at Urbana–Champaign. We wanted to talk to Dr. Black because of his work on rhetorics of usability and user friendliness, and the connections between the historical development of usability studies and our current generative AI moment. Welcome to the podcast, Michael.

Michael: Hey! Hey! Thanks for having me.

Megan: Thanks so much for coming on. So, I get to ask the first question. And this is a question we're going to start with in every conversation we have with very smart people like yourself, which is: what is your reaction to the AI conversation in writing studies at the moment? And what are the important pedagogical or political or ethical challenges that you see teachers, writing teachers specifically, facing at the moment?

Michael: I would say, probably my main reaction has been that I was just kind of stunned by the rhetoric of inevitability around generative AI. Whether you were talking about people that are really excited to use it in their research, or in their classrooms, or people that are skeptical or looking to resist it. There's this shared sense of, "it's here. There's no going back. It's only going to become a more and more, larger part of the writing process at all stages in all disciplines." I expected there to be a little bit more doubt, and maybe there is. It's just that certain voices are louder, getting a lot more attention. You saw the same thing around digital humanities 10 years or so ago. Lots of enthusiastic embrace. Lots of, I think, quiet criticism, some loud criticism. But I think a lot of it was quiet. And then here we are today, and not everybody's a digital humanist. But it's still around. But arguably, it didn't transform the humanities like a lot of the proponents inevitably said it would.

Megan: Yeah, it makes me think, too, of the MOOC [Massive Open Online Courses] moment: when we were all on campuses where they were looking to the Ivies who were putting together

these hugely well-funded MOOCs, and the MOOC was the future. And then it just sort of faded and fizzled in ways that really smart people at the time were saying, this is what's going to happen. But there was this drive of inevitability in that moment, too, like with digital humanities. But it seems like MOOCs had less lasting power than DH has had, so far at least.

Michael: The thing to keep in mind is that you see this a lot with technology generally. There'll be a big hype cycle that comes out. Silicon Valley is very good at pushing these narratives of inevitability to drive up investment. Just think about some of the stuff we've seen in the last 10 years, and kind of what's become of it. This morning I was thinking about cloud gaming, and they were saying how physical game consoles wouldn't exist anymore. We would all just play them on our phones. And where has that gone? It's found a niche. But it didn't transform the industry entirely.

You could probably say the same things about crypto or self-driving cars. They found a place. There's people that are really into them. But we're still driving cars. I'm still paying for things in U.S. dollars.

I think it's probably going to be the same with generative AI. There's going to be a lot of challenges in the short term. But long term, I'm not really sure where it's gonna go. But I would, we would be very surprised if it lives up to the fever dreams coming out of Silicon Valley that it really does replace all of our writing or become the norm for writing in every single field.

Maggie: Yeah, I think that those comparisons are spot on. A lot of people who I hear subscribing to the rhetoric of inevitability around these things, they compare the GenAI moment to tech evolutions in spell check or learning management systems, which that's probably more similar. Or you know, we survived the word processor. And we're still here. And is it more like the word processor? Or is it more like Uber, which has disrupted an industry and without living up to the promises that Uber sold itself on.

It's always good to remember that a lot of the things that seemed like they were gonna take over, didn't. But also, this feels like the hype is really sustaining still, as we move into two years after the release. I don't know how to feel about that, though that might just be the August blues.

Michael: Yeah, the hype cycles can last for a while. I think what's really going to decide things is going to be, broadly how or whether these companies are able to monetize it. I don't really feel like they've really been able to figure that out yet. I think a lot of their biggest clients for companies like OpenAI or other tech companies that they're selling access to, who are using their API to create all these bots for various purposes. I don't know, I mean, in thinking about the challenges that we're going to face. I think probably one of the hardest things to think about is that I don't know that that these debates are going to be settled with arguments about the quality via zoom writing that it produces. I think it's going to be determined largely by the same old institutional pressures and administrative incentives that we've been dealing with in writing for a while now, and maybe a good tech to compare it to in that regard is Turnitin.

Maggie: Mmhmm.

Michael: I mean, there's been a lot of people who've written very critical articles, critical research on Turnitin. And yet it's available on most campuses, and from what I hear on my campus it quietly sees a lot of use. So why is that, right? We object to it, but it's still here. Well,

there's all these labor issues around writing, not just in English or writing departments but across the disciplines. It takes a lot of time to enforce academic integrity. It becomes very difficult to enforce academic integrity when you're a part-time faculty member who's teaching at multiple institutions, who's overloaded, so it quietly sees a lot of use because it solves a lot of labor problems. But if administrators really did care about writing, we would probably have more full-time faculty teaching it. We would have smaller class sizes, ways that we could really implement our pedagogy effectively. But instead, we're in this moment where it seems like there's more incentive to turn to technology as a quick fix. And we're seeing that, not just in universities, with generative AI. But you're seeing that in a lot of industries as well. There's a lot of people, or a lot of companies, who are downsizing, and they're pointing to generative AI as an excuse for that when it could just be because interest rates haven't fallen. Or, a lot of businesses saw a kind of a crush or influx of money during the Covid years because people were stuck at home.

Megan: Right.

Michael: So yeah, I don't know. I just don't think that. That's, I think, one of the weird challenges we're going to face is that we're not really going to settle this based on arguments about writing pedagogy. It's going to be settled at an administrative level and determined on other things, like labor issues and legal issues.

Megan: The labor question is so interesting to me, too, because to the example you were just talking about: Companies have laid off parts of their workforce, pointing to generative AI. And then we simultaneously have studies starting to come out about how much labor generative AI has added to, not just places where downsizing has happened, and so there are fewer people. But even when there hasn't been downsizing, the labor of onboarding yourself into this new space, this new way of interacting or just making your process that has worked for you or has worked for your company change in these ways that actually aren't productive because the company needs to justify whatever financial investment they've made in this tool, which isn't cheap, right? As you point out, they're selling to other tech companies, and that selling is sometimes free to begin with, and then gets really expensive very quickly. And so, I think that feels like an important thing for us in higher ed to be thinking about. To what extent are the labor questions we're asking or the labor impact we're seeing, actually adding a lot of labor? This is something Maggie and I've been talking about because I run the writing program here at University of Arkansas. And so, there's limited amounts of time at things like orientation, or in the pedagogy class, or in the professional development workshops. Do I really want to spend that time helping my faculty learn how to better use this tool? Or are there other ways that our labor is better spent? And I think those are questions that I want to be asking. But there hasn't been as much space to ask those questions as I was expecting in the disciplinary conversation, I guess.

Michael: Yeah, I agree. You bring up a good point about how it would require us to rethink how we do all this training, how our classes are set up. But let's just assume for a moment, as a thought experiment, that we could figure that stuff out, that we could implement generative AI in the ways that administrators seem to want us to do. One question that I would ask is if you look at people who are advocating for prompt engineering to replace writing, and you actually look at the prompts that they're engineering, they are very sophisticated descriptions of writing process, rhetoric, audience, purpose, situation, motivation, intention. And the question that needs to be asked is, how do you get students there when you can't even get them sometimes to do really

just basic stuff on their own? I think it takes years for students to become that level of expertise. I mean, I don't think I reach that level of comfort with my process and self-awareness till sometime in graduate school. And so, this idea that we're going to train people to be process engineers, even if we could figure out all the labor issues around retraining and all that. I just don't know a lot of people that are advocating, for it seems like they're already very skilled writers and just assume that because they can do this, that everybody else can be brought up to that level. Editing work is in some ways a lot harder than writing, and I'm very skeptical of this assumption that we'll all just become editors of AI work when even editing human work can be very difficult, especially when you have somebody to talk to and try to figure out what their goals were.

Maggie: That is so true. Because I've seen the prompt engineering work that people are doing. And it's compelling in it. I point to it when I say I'm not necessarily against the concept of this technology writ large. I have other objections to it. I understand that we write with technology, but it's hard to imagine getting students to that point without teaching the way we teach right now, and having the time with them that we would need to get them to that point. Like you, I think I got to that point in graduate school, and I don't always feel like I'm at that stage of writing with every project I'm in until closer to the end. Part of me thinks that there are students who are going to be really interested in that. But if universities aren't that concerned with good writing in the majority of classes, and workforces aren't concerned with good writing, why would students even engage the technology in that way? That's part of how I'm sometimes thinking about it. When I think of my students, who oftentimes are rushing between jobs to park on our extremely inaccessible campus where they have to walk uphill for 10 min to get to the main campus, I'm thinking about how fast they're doing assignments. And when I imagine teaching them prompt engineering, I'm thinking, the more time that I ask them to spend with this tool, I'm not sure that it helps them spend more time on their work

Michael: Yeah, I mean, I was thinking, we're all skilled, accomplished writers. I was looking at some of the prompts that have been engineered, some of these articles on it, and thinking to myself that by the time it took me to write all of that out and perfect that set of instructions I could have just done it. I could have just edited the document, or I could have read back through it and made notes to myself. I don't know.

I mean the other thing to consider, too, and this comes up in some of the publishing around Generative AI. Specifically, I'm thinking of the article on stochastic parrots by Emily Bender, her collaboration with data scientists, they talk a lot about how these large language models, they're basically just pattern matching systems at a large scale (Bender et al., 2021). If you have a lot of patterns that you can match to, you can produce some very sophisticated niche stuff. But it's important to remember that they don't have an embodied sense of context, and they don't know what our intention is, and they don't have intention themselves other than to match a pattern. So, when you're trying to do editing, you have to figure out what the author's intention is or what your own intention is. You're trying to sus out what their experience of the topic is, whether it matches your own, and trying to leverage your own experience to give them feedback. I don't know.

I just wonder sometimes when people are promoting prompt engineering what they understand editing to be. They ever run a peer review and tried to work with students to get them to provide substantive comments? It's a very hard skill to teach, and I have to confess, as someone who's been teaching writing for over 15 years now, I think that's the hardest part for me is getting

students to do self-assessment, to do peer assessment. Once they get over the hurdle of the first draft, I think they start to get a handle of what they want to say. But then, thinking about the mechanics of actually saying it in later drafts can be very difficult.

Maggie: I struggle with that in my classes as well. And when I think about teaching students peer review with, or self-assessment or peer review, like peer with heavy quotations, with an LLM, I think about the time I'm sacrificing. It's the time that they could spend thinking about their own work somewhere; the time has to come from my class. And so, if I'm asking them to practice that skill with ChatGPT, I wonder would it not be better to have to force us all to struggle through the effort of trying to talk to another person about it. It's really hard work to do. I find it hard in my own practice, oftentimes to articulate to another person what I want to gain from peer review, or when I'm assessing my own, my own writing. But that's where that context driveness comes from. We think that matters as a discipline. I guess it's hard to imagine giving that up or supplementing it with something that doesn't really understand.

Michael: I mean, I think it's important to a lot of disciplines, not just writing.

Maggie: Yeah.

Michael: You know, in tech a lot of times they have post mortems, which are just a big reflection exercise. Businesses have their quarterly reports, where they have to reflect on what's working and not working in their strategy, put that in words. And I think that whether you're doing that collaboratively or just by yourself, having to think through it and find the words is very important. And I worry about short circuiting that process by having a tool that can generate a first draft. Even if you're feeding it ideas to generate that draft, coming up with the individual words is very important, and gonna have important consequences in many legal contexts, too.

Megan: We also just know so much about how important those opportunities are for things like knowledge transfer, and that transfer in and transfer out process are essentially reflective processes. It requires reflection and lots of practice with that reflection before learners can do it most effectively. And so, I also worry about the reps. Are we giving students enough opportunities to try hard things in places like first-year writing, which, at least on our campus, is intended to be a place where you can mess it up and try again. That's part of the writing process, that there's room for failure without the kind of consequence that is life altering and with support along the way. And so, it seems like if we give up that time for something else, we really are creating a situation where students aren't getting the time, space, freedom to fail that really benefits them as lifelong learners, as college students, as you know, future writers in other contexts. All of those things that we care about, that lots of disciplines, I think, care about those sorts of things. And so, it's tough to think through why I would want to make so much space for this piece if it means I won't have enough room for the things, the reflective things, that the research tells me that benefits students from our classes.

Michael: Yeah, exactly. Really. I mean, so much of writing is just tacit. I always tell my students that I could explain to them. I could give them an outline of all the steps they have to do to write their paper, but they just have to kind of do it over and over again until they figure out a procedure that works for them. And even then, they're probably going to struggle to explain it to me. But that's okay.

Maggie: That's a good jumping off point to talk about. How are you approaching generative AI and perhaps resistance, refusal, rejection in your classes? How are you approaching these tools in your classes as a teacher?

Michael: It depends on the class. I teach writing at multiple levels, everything from first-year, writing to senior-level seminars on digital rhetoric and technical writing. But I guess there's really two big things. The first is that I really think it's important that students don't anthropomorphize the technology. So earlier, Maggie, you were talking about peers in scare quotes. I don't want them to refer to it as a collaborator or a partner. It's a tool. It's a tool that can do certain things very well. It also has very specific limits. And they need to understand what those are. The other thing that I think they need to realize, in terms of why it's bad to anthropomorphize it, is that a lot of times the meaning they see in it. If they say, "Oh, that's, that's what I was trying to say." They're projecting that sense of what they wanted to say in the moment; they're deciding when they read it that, "Oh, that's what I was thinking." You know, I've done a lot of work with machine learning. Before I joined the faculty here, my job was basically to find humanities and social science professors and get them to use the high-powered computing systems on Illinois' campus. So, a lot of those projects were machine learning because that's how you justify using those big systems. And one of the things I learned about it is that, in that process, a lot of these tools will find patterns in almost anything, tons of patterns. If you feed it data, you tell it to spit out 50 patterns, it'll give you 50 patterns. But you really have to take a lot of time post hoc to figure out what those patterns mean. Some of them are meaningful; some of them aren't. And the thing about large language models is that because they're grammatically correct sentences that resemble genres that they may be familiar with, or think they're familiar with, it seems very tempting to be like, "oh, this is writing." "This is responding to me," when really what it's doing is it's generating text to match the pattern that it thinks is going to successfully reflect the pattern that you requested with your input.

That being said, I think it's really hard to explain that to students. I won't even say just in first-year writing. I've had trouble explaining that to students in upper-level classes. So, one of the things I've been trying to do is develop some handson exercises that get them to recognize what I mean by pattern matching. So, usually I like to devote a day. I've only done this a few times but devote a day to simulating large language model algorithms by giving them a set of sentences that they can't read. I usually just give them printouts of some short sentences in Windings, and then I ask them to study the patterns and see if they can produce a new sentence using any of the symbols they see on the page. They usually can't produce a grammatical sentence, or if they do, it's nonsensical. There's a subject and a verb, but they don't really match or reflect anything in reality. But it's interesting because I think they start to see the limits of it, and they start to see, "well, hallucination just happens." Why, it's almost impossible to eliminate in these systems and opens up conversations for other aspects of the writing process.

Megan: So, if we could switch gears just for a second, and talk about your book. So, first of all, thank you. The Coda, in particular, was incredibly helpful for giving me some vocabulary that I did not have before. So that, first of all, but in that Coda to *Transparent Designs*, you say, quote, "if we are to imagine an unfriendly future, we must begin by recognizing that user friendliness helps designers to conceal, not just the technical aspects of computation, but also their constant interventions in digital culture." So that was really helpful for me in understanding one of the discomforts I was feeling with conversations around generative AI and the efficiency and "easy" discourses that feel really embedded in the conversation. So, do you see similarities between the histories you trace in your book and current conversations about generative AI?

Michael: Oh, yeah. If I had written the book this year it would probably be a little bit different. I mean, I think the main arguments would be the same, but the emphasis on how I got there would be different. I don't know. I could talk about this in a couple of different ways, historical connections and thematic connections.

First, artificial intelligence really did have an early influence on thinking about usability and personal computing. A lot of early thinking about how to make computers more usable was heavily influenced by the field of cognitive psychology, who in turn were heavily influenced by the field of artificial intelligence. Ulrich Niser, in his book *Cognitive Psychology*, proposed that the human mind was an information processor, basically a computer. You could break it down to subsystems. And if you understood how they work and how data was flowing through it, then you had a sense of psychology. And so, his work informed Donald Norman, whose work informed a lot of the Macintosh UI design and user friendliness more generally, as it's practiced now. The information processor model also influence more quantitative approaches to usability, like the goals, operators, methods and selections framework.

So artificial intelligence was there. It was kind of behind the scenes, but it was strongly influencing a lot of those theories, but maybe more directly. I don't think I quoted from any of these in the book, but Bill Gates gave a ton of interviews and submitted a lot of articles to magazines in the early 1980s, where he suggested that the future of UX would be artificial intelligence that use the computer on our behalf. Lots of other people talked about how the future of usability were autonomous agents, that we would talk to, and they would carry out stuff on our behalf for our computers. So, what's interesting is a lot of people thought that that would be realized by the late 1980s, and it wasn't. And I think computers were very frustrating. I guess they still are, but very, very frustrating for people that were into them through most of the nineties. I don't know you. You always see tech come back. It happened with virtual reality. It's happening with AI and autonomous agents now. So those are some historical connections. But yeah, if I were writing it now, I probably would have had a whole chapter just on those early discussions of AI and usability.

Megan: So, you also talk about an unfriendly future. And I wonder if you can talk a bit about, what does an unfriendly future look like in terms of generative AI? Also in the coda, you say, "software designers structure our relationship to their software to maximize momentum and minimize our agency." How do we create the friction or the unfriendliness in this moment, as we were talking about, that has been so hyped, and really relies on these conversations about ease and efficiency?

Michael: When I was writing that I was really concerned about being critical of the idea of friendliness because a lot of that rhetoric proposed making computers more accessible, in the sense that you had to have a computer science degree – that's what they said – to use them before the Macintosh and similar interfaces came about. But one of the things that's interesting about friendliness is that it assumes that certain things aren't worth knowing. And if we think about how that maps onto generative AI, there's this assumption that certain aspects of our literacy are not worth our time. It's not worth our time to struggle through a rough draft. It's not worth our time to rewrite completely into a second draft, or do other things like that. It's not worth our time to go talk to somebody else, or it's not worth their time, so we'll just ask the computer to give us revision. And that's a pattern that happens a lot through history. So, friendliness is, as I

saw it, was about concealing the tech and automating on our behalf, and unfriendliness was about exposing how it works and making more room for decision making.

In terms of what an unfriendly future might look like, hard to say because I think that at least generative AI, as it's discussed now, is only friendly. I can see some uses for it, but it doesn't really lend itself to unfriendliness. It's all about automating and concealing parts of the writing process. I don't know.

I mean one way to do it might be to take the logic of Silicon Valley to extreme and imagine what a hyper, friendly future might look like, and then try to think of ways to avoid that. So, let's take some of their narrative at its face value and think about what that future might look like. Let's say, we reach a point where human writing isn't really something that has to be done anymore. So, think about all the writing that we have to do, not just as academics, but just as professionals generally. All that email that people complain about having to spend so much time reading and writing. Well, let's say that now AI reads all of our email but also writes all of our email. So, at that point, what is even the point of having email in the workplace? It just doesn't exist. It's not a part of professional culture anymore. You see all these ads now for AI assistants that will attend meetings for you. Okay? Well, what if it's one person speaking to eight AI, or, better yet, what if they trained an AI with a voice synthesizer to lead the meeting, and they're all just AI talking to each other at this meeting. Why, even have meetings? So, to me it just doesn't make sense if you try to game out what they want. That seems like there's a lot of outright contradictions and just areas of work just won't exist.

One of the things that you see if you look back at that early period, I guess you see it now, too, around usability and user friendliness: tech is always presented as a solution to problems that tech creates. So, Google was considered to be groundbreaking because instead of having to do all these manually curated search engines, now, Google could do it all automatically on our behalf, and it would somehow know what you wanted through its PageRank algorithm. Now, we're at a place where there's all these literate demands on us in the workplace. And I think, and this is something I've been thinking about, too, that is in some ways depressing, but also liberating to realize. So, I don't think a lot of people like writing, even if it's a really core part of their job. They see it as something that gets in the way of what they really want to do, even though, arguably, they have to carry out those actions through their writing. I mean, an unfriendly technology version of large language models would be one that, I think, exposes the sources of the text that it's generating, and its goal would be to direct users to those sources rather than stand in places of them. But as it's currently used right now, almost all of them are trying to block users from going to the source and just give them this simple summary instead.

Megan: It makes me think, too, about – and this was a conversation a year ago and that conversation sort of coming back around a bit, at least in my little corner of academic Twitter – about how AI outputs as training for future AI models creates really unstable and far worse AI models. And so, there's this situation where for over a year, we've been talking about, there are limits to this. You're gonna need new human input that is actually human generated for AI models to advance in the ways that people like Sam Altman are promising they're going to. And yet the discourse around the generative AI models in educational contexts and professional contexts is all about eliminating those kinds of human output that would train the model. And so, it also feels like it's not actually possible to reach what they're arguing for, in part because their models aren't built to pull from only AI generated materials.

Michael: I think it in one of her recent newsletters, Annette Vee (2024) referred to human writing as the new oil. We're always going to need sources of human writing to keep these things moving and updated. Even if we were to assume that all human writing would stop now, then writing would basically be locked into the style and discourse of 2024, which wouldn't work. Things would break down. The world changes, right? We would need new approaches, new ways of thinking, new ways of phrasing. Genres change because situations change. And we can't just have these large language models spitting out these stale genres as the world changes around them.

But I think you bring up a good point about it being trained on AI output text. I think there's other issues around AI output text, too, becoming dominant, especially legally. What I think is interesting is that at my institution there's this push from above for everybody to reskill and think about how to incorporate AI into their teaching. Very big, lofty stuff. They gave us this book called *Teaching with AI* (Bowen & Watson, 2024) that takes in that rhetoric of inevitability, and says, here's all these things you can do to enhance your teaching with AI. But quietly, before they gave us access to an enterprise ChatGPT license, they sent us this little—not that little, it's 3 or 4 pages—user agreement authored by the IT department and the legal team, which makes some things very clear. Number one: currently, according to the U.S. Copyright office, AI text cannot be protected as intellectual property, and that matters in a lot of areas of professional writing. So, no company is going to want to put stuff out there that people can just wholesale steel and slap their name onto that. The other thing they warn us about is there's the potential for these generative AI tools to unintentionally spit out copyrighted text. So, you could publish something and then be accused of plagiarism, which is a big, big deal for academics, career ending if it becomes public. And then all these other privacy issues. If we want to put any sort of protected data into these AI systems, we have to go through an audit with a university team. And there's all these constraints that are in the fine print around computer generated text. So even if we were able to produce all this stuff out there, even if the problem of model collapse wasn't a thing – AI feeding on AI generated text – there's still these little legal realities of writing that would really prevent it from becoming the norm in a lot of contexts.

Megan: I was talking to a lawyer on our campus not that long ago about employment law. And making sure that we're in compliance with things. And he was saying, "yeah, I could theoretically generate some of the text that I write with generative AI. But the reading I have to do with them is backed up by so much knowledge, genre"—I mean, he didn't use the word genre—"but genre information, understandings of complexities of purpose and audience and secondary audiences, and all of those things, that it's easier for me to write it myself." And also, all of that knowledge, even if ChatGPT could spit out the kinds of letters that he writes on a daily basis, it's built on other people's texts, and then he has to bring his human knowledge to that because that's the only way the institution stays protected from his perspective, and also because we know we can't trust it. There are nuances that he recognizes that the machine can't. And so, it's a big, wide thing, and it's an erasure of all of the labor of learning, going to law school, learning all the things he needs to know to tell what the text is and isn't, what it does and doesn't do.

Michael: I imagine that or somebody like him, he's not writing every single document from scratch, poring over all these years of knowledge. He has templates for certain situations that he reuses. Most people do.

Megan: Right.

Michael: I don't open up a new word document for an expense report. I pull up last year's and I change the stuff around. I think that's what most people do. Programmers do it. They reuse their code all the time. There's formal systems in place for code reuse. So, when people talk about it as giving a jump start. I'm like, well, I guess if you have a blank page, if it's something you've never written before. But in those cases, I tend to turn somebody to somebody who has a lot of lived experience in that genre or writing towards that specific question and be like, "Hey, can I see an example? Something you've written and talked to you about how you responded to that situation?" I don't think these large language models, as good as they might be at spitting out examples that look good, I don't know that they can necessarily explain the context of their creation.

One of the things I'm worried about really long term is the loss of expertise if large language models become the default. I don't think they're going to go away. But I just don't think they're gonna be as common as folks want them to be or believe them to be. I think they're going to serve as interfaces to help access databases of human written text. I was reading about this technology recently called retrieval augmented generative systems. Basically, if you have a set of documentation, a knowledge base. Right now, if you go to a knowledge base, you have to sift through this hierarchical listing. Like, "where's my question? Or is this similar enough to the problem I'm having?" What these retrieval augmented generative systems promise is that you could type in your question, and it would search through all this human written documentation and try to find something that it thinks matches your question, give you a summary, and then, ideally, although I'm not sure it does in every implementation of it, link back to that document so that then you can read it and get the full explanation. I've heard that even those, even though they're working on a very specific corpore, can't prevent hallucinations.

Megan: So interestingly, we have a library integration tool that is pulling just from the database of materials owned by the institution, digital materials owned by the institution, that is claiming to be able to do some of the things you're talking about. And what we were talking about, in response to them beta testing this tool, I said, one, this summary is terrible. But two, that's not how it gets used by students. This is the wrong use case. A tool like that for somebody who is looking to build expertise in the topic, sure. But this will be my 18th year teaching college writing, and I don't know more than 2% of the first-year students I've worked with who are going to do the six clicks it takes to get back to the original source. They're going to read the summary. They're gonna see that there are sources. And for most of them that's going to be the end. And so, if that's what we're enabling with this tool, is that what we want for students? And everybody was like, not really. I was like, then why?

Maggie: Especially when it's sponsored through the library. You know, a lot of classes with a research component will have a library day. So, a librarian is coming and telling them that this is how you can do your research. And everyone wants their students to go to the library more. And so, there's an extra layer of authority that's embedded into the process. So, if you're not caught, if you're not always worried about hallucinations, if you're trained throughout your education that going to the library is encouraged more than using Google, why would you not trust it? Which is especially problematic on our campus, and probably other campuses where the academic integrity policy has been updated. It's essentially at the teacher's discretion how to implement enforcement against AI writing. Whereas, if it's okay in my class, and I'm not going to police it because I don't really feel the need to do that. But if someone else really doesn't want their

students to do that, the library has become an access point to an academic, an academic integrity violation.

Michael: I think in instances like this, the people promoting these tools are really placing a huge emphasis on product. They think the point of doing these research papers is the paper itself. No, it's the going back and reading all the stuff, figuring out what's important about it. I mean, I think that's one of the challenges, too, for writing instructors, and not just writing instructors but any class that involves written writing as a form of assessment, whether it's a philosophy exam, philosophy paper, history paper. They're going to be repeating a lot of the same arguments semester after semester, year after year. But the point isn't necessarily those arguments. It's the process of getting there, of internalizing all that stuff and learning to think through it. When you look at a lot of the discourse around academia and people getting upset around large language models, I think there's a lot to be upset about. But I think there's also this – I don't want to sound too critical or cynical – but I think there's a lot of this emphasis of, yes, we do a lot of writing in the humanities, but for a lot of classes writing is a vehicle for assessment, but it's not the object of assessment. A lot of times when I teach students outside of our major, they talk about how they get comments from their professors. But a lot of times it's just, "Yes, no, correct. Don't understand this." But not really engaging with the style or structure of the writing itself, and that's a problem. Why not turn to these tools which could give you that style and structure that you're struggling with, that you can't figure out, that nobody's really explaining to you or giving you a chance to fail and figure out what you did wrong in a safe environment? I feel like every good writing class should give you that sense of "it's okay to fail because you have time to revise and fix it." And that's normal.

Megan: Yeah, absolutely. It feels to me like it's even more important to me, at least on our campus, that I do everything I can as the writing program director to try to protect that space and time in our classes. Because WAC programs are amazing. We don't really have one on our campus. But I've been on campuses with really robust and lovely, wonderful WAC programs that have really done a great job of working across the curriculum to bring that sense that we should teach writing, and we should talk to students about writing, and we should give them time to practice. But I think that's pretty rare. And so, on a lot of campuses, our first-year writing classes may be the only place that it's still possible for a student to experience that space to fail and to do the repeated practice that we've been talking a little bit about. And so, what can I do as program director? What can writing programs do to protect that? That seems to become an even more precious thing the more these systems show up on our campuses.

Last question we wanted to ask was, what are you reading right now that we or our listeners should read?

Michael: I can name a lot of stuff that's helpful for thinking through this topic. I guess, on the tech side of things. I've been thinking a lot about Shoshana Zuboff's (2019) book *[The Age of] Surveillance Capitalism*. If my book covers very early inklings of this in the late '70s, early '80s, I would think of Zuboff's book as Part 3, the '90s, and then the post dot com bubble into the present. She really looks at how machine learning became built into a lot of our tech data collection. And those logics are really what power these large language model systems, the economic logics behind them. There's also a lot of really stunning quotes in that book from people that she interviewed in the field about how, in a sense, they recognize that taking all of this data that people are producing maybe isn't the most ethical thing, but also, they're making money off of it. And it's not really that illegal. So, they're just going to keep doing it. I don't know.

I feel like there's a lot of that sense in today's conversations around large language models and especially the sense that they are ingesting all these copyrighted works without permission. And is that infringement? Or is it just unethical, but not illegal? So, I would highly recommend everybody read that. It's long, but it's really, really good and written very accessibly.

Back in writing studies. I've really been thinking a lot about Deborah Brandt's (2014) *The Rise of Writing*.

Maggie: Hmm.

Michael: Particularly, how a lot of people that she talks to do writing as part of their job. It is one of the main things that they do, but they don't think of themselves as writers, and I think that's really helped me think through a lot of the attitudes around generative AI, particularly the excitement. A lot of people are happy to hand off writing to something else because they really want to be in the lab or whatever else it is they do for their jobs. I think another book that I've been thinking about a lot lately has been Amy Wan's (2014) *Producing Good Citizens*. She talks a lot about how literacy both as a concept and literacy instruction structure relationships between individuals and institutions. And I think we're at a moment when generative AI could, if we cede ground to technologists and cede our authority on writing away, restructure the relationship of our students to the working world, to ourselves, and our institutions, etc, so that one's really worth reading.

I guess if you want some really historical criticism of AI, Hubert Dreyfus's (1978) *What Computers Can't Do*. It's a book version of a report that he delivered to the Rand Corporation in the '70s, reviewing a lot of early AI technology. And I think a lot of his criticisms still hold true today. The biggest one is that it's okay to acknowledge the technological leaps that they're making but also be very skeptical and critical of the social, political, and cultural claims that they're making about their technology.

If people want something more contemporary, just real quickly, I'd say look up 404 Media. They're a small tech journalism outfit, and they've been doing a lot of really cool reporting on generative AI.

Maggie: Well, thank you so much. That is a lot for us to dig into. *Surveillance Capitalism* is one of those books I love to assign for my grad students, because it's really like this. Thank you so much for talking with us, Michael. This has been really wonderful.

Maggie: Thank you for listening to this episode of *Everyone's Writing With AI (Except Me!)*. We want to thank Dr. Michael Black for being so kind and joining us on the podcast.

Megan: You can follow the podcast on Twitter for updates at @EWWAIpod (that's E-W-W-A-I pod), where we'll share info about forthcoming episodes and news about AI hype in higher ed. And you can follow me @RCMeg.

Maggie: You can also follow me on Twitter @magsfern. You can find these handles, as well as links to all of the texts discussed today, in the description of this episode.

Megan: If you'd like to talk to us about refusing gen AI in your work or teaching or anywhere else, please fill out the form, also linked in the description. Bye, y'all!