

"AI Goes To College: Large Language Models and the Teaching of Writing"

Transcribed by Eric Detweiler

[*Rhetoricity* theme plays]

Eric Detweiler [D]: Hello hello *Rhetoricity* listeners! This is Eric Detweiler, here to kick off a slate of new episodes that will be dropping this fall. This first episode is a part of The Big Rhetorical Podcast Carnival 2023. The theme of this year's carnival is "Artificial Intelligence: Applications and Trajectories." If you're interested in other podcasts that are participating, check the episode description for links and information. As for me, I'm bringing you a roundtable discussion featuring members of the MLA-CCCC Joint Task Force on AI and Writing. In July 2023, that task force put out a working paper identifying key issues, principles, and recommendations for writing teachers trying to navigate the wave of generative artificial-intelligence tools that have arrived on the scene over the past few years. In this episode, task force members Antonio Byrd, Holly Hassel, Sarah Z. Johnson, Elizabeth Losh, and Anna Mills elaborate and reflect on that working paper. Among other things, they address the labor implications of AI, the importance of precision in the language we use to discuss different artificial-intelligence tools, and how AI fits in with and diverges from the long history of moral panics around new writing technologies. Before we get rolling, I also want to give a quick shout-out and word of appreciation to the task force members who weren't able to participate in this conversation: Leonardo Flores, David Green, Matthew Kirschenbaum, and A. Lockett. And with that, I give you this roundtable on writing and AI.

[transition music: Smoked Meat Fax Machine's "Artificial Problems," an upbeat electronic song]

D: All right, my name is Eric Detweiler, and I'm here talking with the MLA-CCCC Joint Task Force on Writing and AI. And I really appreciate all of you taking the time, especially in the middle of summer, to have this conversation. I wanted to start by just asking each of you to introduce yourselves to listeners and talk a little bit about your background with and interest in artificial intelligence and its relevance to the teaching of writing.

Liz Losh [L]: Sure. I'm Liz Losh. I am the cochair, with the fantastic Holly Hassel, of this enterprise. I started writing about text generation in the mid-2000s; I was really interested in kind of web 2.0 cultures around text generators that were kind of seen as novelties at the time. So things like the [Postmoderism Generator] that Andrew Bulhak created, was one of the first text generators that I read about. But I've also been—I was on the executive council of the MLA and organized their "AI in the Humanities" executive council session at the MLA convention, particularly focused on members of Black in AI talking about their concerns about AI technologies, and then have been part of a couple of big research groups working on things like facial recognition technologies or simulation technologies. So that's sort of how I got to be part of the task force

through my work on the Executive Council, thinking about these technologies in the humanities.

Holley Hassel [H]: I can jump in. This is Holly Hassel. I'm the cochair with the fabulous Liz Losh, representing the CCCC side of things. It might be fair to say that I'm not at all interested in AI, actually, [laughs] but I am really interested in literacy, and college writing classrooms, and how students engage with language, and, you know, just sort of what happens in our first-year composition classes. You know, I've also sort of, I'd say, been fairly technologically involved having taught online courses for 20 years, way before they were kind of common, at least at four-year institutions. And so for me, I came to this in my elected role at the time that the task force was getting assembled. I was chair of the CCCC—the conference, you know—as more sort of stories and technology hot takes and whatever came to our attention. It became clear that this is, as someone really interested in writing pedagogy and engaging students as writers, you can't not care about or be interested in this technology, if that's what you're interested in. So, for me, it was partly just a kind of governance priority, recognizing that as the organization that represents college composition in the United States we had to get involved and take on the issue.

Sarah Z. Johnson [J]: Hi, I'm Sarah Z. Johnson and I am a faculty member and Writing Center Director at Madison College in Madison, Wisconsin. Like Holly, my background is not necessarily in digital rhetorics or history. I haven't been doing this for a long time, like Liz Losh has. But I was chair of the Two-Year College English Association, and as part of that, also a member of the CCCC executive committee, that is how I at least became aware of the work of the task force. And I am particularly interested in the work that we've been doing because of my role at my institution as academic integrity officer. And while my scholarship is not in that area, because of that role, I'm a member of the International Association of Academic Integrity, and I listen into those conversations quite a bit. So for me, that's been really interesting to hear the overlap in how we're talking about AI, generative AI, and academic integrity issues.

Anna Mills [M]: I'm Anna Mills. I've taught writing at community colleges in the San Francisco Bay area for over 17 years, mostly City College of San Francisco. I wrote an open educational resource writing textbook, *How Arguments Work*. And I've been involved in the OER world. I got involved in AI in part because it's sort of native to me: I grew up in Silicon Valley; my dad was a machine learning and expert systems engineer. And I did some web development and coding and technical writing time before I became a teacher. But I became really interested in how I could sort of support the discourse and started a resource list through the Writing Across the Curriculum Clearinghouse when I got obsessed with AI in June of 2022. And I did some consulting with OpenAI on GPT-4 before it was released. And so I've been writing and kind of in that larger open educational practices collaborative discourse. I feel like it's been great to be part of that.

Antonio Byrd [B]: My name is Antonio Byrd, and I'm Assistant Professor of English at the University of Missouri–Kansas City. I do research on digital literacies and race. And right now, I'm especially focused on computer programming and how Black adults use coding to obtain liberation for themselves, depending on how they define liberation. And for me, I like to joke that I am the professor that students can come to when they don't want to write essays anymore. So I teach a lot of classes in professional and technical communication, multimodal composition, and digital rhetoric. So a lot of the genres are not always recognizable to our English majors. So last spring, we made board games and talked about the rhetoric of the design of board games and instructions. And that is kind of like my interest with artificial intelligence is thinking about these different tools and technologies that students can use for their writing. So in the fall of 2021, I was revising an undergraduate course in technology and literacy studies and I came across GPT-3. And I thought, this is a great way that we can use this tool to talk about originality, remix culture, copyright, ethical uses of AI. And then our main task was to use GPT-3 as a way to revise our essays. So it started off as kind of like a teaching subject. And since that class, I've been thinking about next steps in terms of my own research and how this seems to be really the next thing that I should really step into, especially as it's coming onto the scene for a lot of our writing classrooms and really other areas of our society as well. Holly had sent out an email to the executive committee asking for representation. And I thought, well, this is perfect. I've played around with GPT-3 a little bit. I'm interested in digital literacies and the writing classroom. And so that's how I joined the task force.

D: All right, well, thank you all for those introductions. So before we go any further here, I wanted to ask a question about vocabulary to frame the rest of this conversation. As you all note in the working paper, there are a lot of distinct terms at play in the conversation around what we sometimes just refer to as "AI": you've got generative artificial intelligence, artificial general intelligence, large language models or LLMs. And as you point out, even the more everyday terms we use to describe the workings of AI carry really important rhetorical connotations. You know, does AI "think," or "invent," or "lie," or is that just us projecting, you know, human activities or ways of understanding that don't really apply here. So to help listeners appreciate those distinctions as the rest of this conversation unfolds, can you all tease out some of the key terms that you use to talk about AI?

L: I can start. But I'd also like to hand it off to Anna Mills for all of the work that she's really been doing, about making sure that our various constituencies understand these terms a little bit more. One of the things in the working paper that we did is really used that LLM term to emphasize the fact that these technologies are statistical models, and that there's a tendency to invest them with sentience or sort of supernatural powers that are really misplaced given the ways that these things operate. And because they are statistical models, that means that we have to think about their biases. And we also have to think about the ways that statistical models can operate through generalities in

ways that ignore evidence and facts. A term like “AI” can sometimes invest these technologies with more sentience or consciousness than they really have.

M: Definitely, and I think Eric mentioned AGI, artificial general intelligence. So there's this sense that we use even the term “AI” very advisedly. Kate Crawford has written a really important book, *Atlas of AI*, that attempts to sort of contextualize current AI systems within sociopolitical, economic, and environmental structures. You know, the argument has been made by Crawford that AI is neither artificial nor intelligent, that it's very much based on human data. And it doesn't have, certainly, the kind of intelligence that we recognize as such, or that's represented in science fiction. And that there's this long documented tendency to project onto these systems since Joseph Weizenbaum coined the term the Eliza Effect. Weizenbaum was a computer scientist who made one of the early chatbots, which was called Eliza, and it was sort of a Rogerian therapist. It wasn't anywhere near as sophisticated as what we have now—it was spitting out kind of rote answers. But even so, the people who worked with it, Weizenbaum noticed that they were tending to—even though they knew it was software and it wasn't that sophisticated—they developed feelings about it, they felt they had some kind of relationship with it. And they started to see it in this way that was very much anthropomorphizing, and they tended to be emotionally invested in it. So we really tried to avoid those verbs like the system “understands” because we want to dispel that notion that there's real understanding or consciousness or communicative intent there. We want to help people understand that it is a statistical model. And that's part of why we started the introduction with kind of a framing of what these systems, how they're structured, and what that means for the language they produce.

L: And I would say that emotional piece is important for students as well. You know, we're interested in seeing our colleagues do more research into how students are using things like ChatGPT because there's a lot of anecdotal evidence that they're actually using ChatGPT for a lot of things that aren't even related to their academic work. And so I think that in terms of a whole repertoire of literacy practices, we need to learn more about, you know, that guy who's asking ChatGPT for help with picking a birthday gift for his girlfriend. And that, you know, his focus on interacting with ChatGPT isn't cheating. It is, you know, getting essentially relationship advice, much like those early users of Eliza in the 1960s in the Weizenbaum laboratory.

H: Yeah. And I'll just put in a plug for the Quick Start Guide page that's on our task force website that Anna and I have just actually been updating in the last couple of days. If listeners are interested in, you know, more examples of the kinds of uses that that Liz mentioned.

L: ChatGPT is really good at writing dating profiles, by the way.

D: [chuckling] Oh, interesting.

L: If you have any single friends.

D: Good to know.

J: I'll just add, we were bringing different areas of expertise. You know, some of us have not been longtime experts in digital rhetoric, literacies. I think that actually helped because there were many in the group who were very familiar with the terminology and the conceptual frameworks. Others weren't. And so having those conversations, I think helped clarify for everybody what that terminology was and where we might need to make sure we were being more clear.

D: In the working paper, you all discuss some of the institutional processes, exigencies that brought this task force together. And that's something you all have touched a little bit on in your introductions in some cases. But regardless, I wanted to ask if you could share a little bit more about what the day-to-day, week-to-week work of the task force looked like, particularly because I think this kind of collaborative organizational writing is immensely important, but, you know, documents like this working paper can sometimes feel to folks like they sort of come out of nowhere. So I'd love to hear just a little bit more about the mundane process of deliberating, drafting this document—kind of what that looked like for you all, and how you worked through the process of putting this together?

L: Well, I'll turn most of that over to Holly since she was really the one herding cats a lot of the time. But I would say that one of the interesting things about our task force is even though we have a wide range of areas of expertise, including electronic literature, digital rhetorics, computers and writing, a lot of us have been WPAs or are WPAs right now—writing program administrators. And so we were really concerned about the institutional effects that are happening right now in real time. And I think one of our big concerns is that this is a dramatic shift that requires a real pedagogical response because these tools are going to be integrated into commercial writing platforms, like Google Docs and Microsoft Word very soon. We've seen demos of this, right? So we know this is coming. And so I think that we really felt, like with COVID, when there was a major pedagogical pivot, institutions understood that they needed to do more to support their writing faculty. But I think there's been just a lot less administrative and institutional support for, you know, this new state in which we're going to be teaching writing, composition, and literacy.

H: Yeah, I'd say, you know, the group kind of initially got formed when Paula Krebs, who's the executive director of the MLA, reached out to me—just because we've been working on a couple of things together in my role, as you know, at the time chair of CCCC—and shared a problem statement that had been written by some MLA members. So MLA and CCCC were kind of looking for, you know, "We have a problem statement, we have a problem, like, what's the solution?" What are our next steps moving forward,

because this is clearly something that our two organizations are going to be affected by and are uniquely positioned to take the lead on, you know, sort of setting professional standards and providing professional resources. We each kind of as an organization had to determine or identify, like, what does this group look like? Is it a task force? Is it a committee? You know, MLA and CCCC, these are both language, literacy, literature organizations, but we have very different structures and very different resources and very different kind of governance processes. So, ultimately, we moved ahead with this kind of joint task force model, you know, and once that group had been sort of populated and charged, it hasn't been easy to sort of figure out, okay, what do people need? What are our next steps? How are we going to get this moving forward? Not because, you know, we're confused, but because there are so many needs. I guess that's the part where if I feel overwhelmed, ever, it's just this kind of accelerating and shifting landscape constantly. You know, so there's linguistic justice, there's academic integrity, there's pedagogical issues, there's policy, and sort of how do we put our time to use in a way that's gonna pay the right kind of dividends for people in what they need right now? And so I'd say our process was pretty collaborative. And it was iterative. We came together every couple of weeks, identified what our goals were for that week, tried to assign specific kinds of tasks, whether those were writing or research or, you know, getting our web presence set up. Then we would alternate sometimes between a formal agenda where we were gathering, either getting information or sharing information, or sometimes just working, writing and drafting time often, which I think we all found kind of helpful to make sure that, you know, it was a sort of coherent kind of piece.

J: Holly started to touch on some of the things that I was going to talk about too. As many of us who are writing teachers, I found the process fascinating because we really did the things that we tell our students to do in terms of, we tried all kinds of different approaches when it came to collaborative writing. You know, we wrote in pairs, we wrote individually, we drafted together, we met. Sometimes we put a meeting on the schedule and then just promised ourselves that we would write, and I think a lot of us did, and then came back together, held ourselves accountable. But we were really productive, where at the end of a meeting, we're like, "Yeah, we got a lot accomplished."

M: I would just add that it didn't feel like the document was disproportionately produced by like one or a few people because we were dividing up into pairs and generating, but then we were also constantly commenting both in the meetings and asynchronously on the document and editing it.

L: And we shared drafts of the documents with our various committees so that they had an opportunity to give feedback, we surveyed our memberships and invited interested

people to kind of comment on these pressing issues. A lot of people have volunteered to participate in the writing process.

B: Yeah, I just wanted to kind of echo what Liz said: that it kind of gave a lot of confidence, at least for me, to hear comments from other committees, from folks thinking about contingent labor to thinking about technology and having their own thoughts and ideas that kind of shape their own thinking. And so that meant that we knew that we were really on the right track with producing a paper that could have interest to many members of both MLA and CCCC, and even people who are outside of those two organizations teaching. One other thing that I was also thinking about is, so this is actually the first time that I have worked extensively with such a large group of other writers. And so it is kind of like getting a sense and feel, like, "What's everyone's writing process like?" So when we were in person together, writing on the document, I noticed I'm kind of the type of person who's just like, "Make a decision, and then let other people think about it." Like, let's just kind of move forward. So I like to get words down on the page, and then kind of worry about it afterwards. And so it is really good to see how other people's writing processes complement how I approach it, like, "Slow down, maybe you should think about the language." Like, that's a really good idea! I should. It's been a really helpful moment to think about how collaboration really works. And it is about kind of bringing the best parts of your writing process together.

D: Great, thank you for that. We've been talking a lot kind of about the background of this working paper and the task force so far, and I wanted to now jump into the paper itself. And one of the things you note in the introduction is that writing itself is, of course, a technology even though we sometimes forget that, and the other specific technologies that facilitate it have always been an inescapable part of how writing gets taught. So given writing's rich and, as you all point out, inherently technological history, could you talk a little bit about what you see as historically familiar about the panic and the enthusiasm that's happening around AI right now, and what if anything you see as significantly different here? In other words, you know, what have we seen before here, and what is actually, maybe, potentially new?

L: I think that, you know, a lot of us have studied rhetoric, and we know that some of these debates about how writing functions as an assistive technology actually kind of go back to Plato and Aristotle in some ways and that, you know, we'll hear these kinds of concerns about the ways that the technological constitution of writing are things that will inhibit human capacities and abilities or encourage certain kinds of deception that we should all be concerned about. And I think those fears continue to be real fears. But I think that what we're concerned about is this connection between writing and learning, and also writing and public participation. And I think that, you know, we are at an important inflection point. And that's one of the things that we tried to sort of say in our introduction, where we're both talking about the sort of long history of writing technology and assistive purposes, and the ways that writing facilitates civic

participation, it facilitates metacognition, it facilitates reflection—like, all of these things that writing does that can't be outsourced to a machine we really wanted to highlight in framing the work. And the other thing we really wanted to highlight is our role in facilitating student development and why, you know, thinking about this just in punitive or corrective terms seemed to us like such a wrong road to go down. Really, it's about our ethics as writing teachers and what we hope to achieve, not only with our students, but just in interacting with any kind of public, how do we see the value of human-centered writing.

J: I also want to add—and this is more in my role as a writing center director because I do a lot of WAC, writing across the curriculum work, and then in my role as academic integrity officer, I see how writing is used in other disciplines. And I feel that actually, MLA and CCCC are much closer to finding very productive ways to think about and incorporate generative AI into our pedagogies, and I think other disciplines are going to need our help as they think about completely overhauling assessment practices and how they're going to be using writing. And so while we are struggling with this, I also feel like we are much closer to figuring out really excellent ways to use AI to help teach AI literacy. But I also think that we're further along the path because we've always been focused on process rather than product we've always been focused on, as Liz was saying, how do we connect writing with learning?

M: I'm hoping people will correct me if I'm wrong because I'm my background is not in the history of writing technologies. But it does seem like this is different because it mimics thought and learning and understanding in a way that other technologies haven't. So, you know, it sort of passes the Turing test of, you know, you can't necessarily distinguish if the words come from a human or from an automated system. To me, that's significant.

J: I agree. And I mean, I think that there has always, in terms of writing as a technology, I mean, anytime a technology like this comes along, we're afraid that the technology will make us dumber. But I agree with you too, Anna. And I think the task force has talked about this as well, you know, how do we make sure that our students have the skills that they need to critically interact with these AI platforms.

M: But maybe like previous waves of fear, I mean, there's the fear that it will be used to replace thinking. But then there's a recognition that's beginning that it can be used in a lot of complex ways, it can be used to stimulate thinking, that we get to choose how to use it and how we want it to be part of a thinking process. So that, to some extent, maybe counteracts the panic, which is a bit reductive. It's not just a phone-it-in tool, as I think Liz has said. Use it to push yourself.

B: Yeah, and in preparation for writing the working paper, I did some reading in histories of computers and writing, and I read about how there was a lot of excitement for word processors and computers in the 80s and early 90s, if I remember right, and

there was this hope that it would improve student's writing. And then when scholars in the field did research, they found that it did not improve student's writing at all. And so I use it as an example to kind of think about how there's, whether that's hype or there's peril, there's a tendency for reality to kind of settle in. And then you come to realize, "oh, this is how it is being integrated into the writing life that we have." That there is more of a desire to do collaboration with other types of writing tools, as opposed to giving them the most meaningful work and letting them actually do that. So some other examples I was thinking about is, like, how the internet will replace the university, and it didn't. We actually created tools for properly navigating the internet, especially when you're doing research. As an MA student teaching my first composition class, Wikipedia was actually a big conversation. And granted, there's a dedicated team of editors—most of them are probably white men, but nevertheless, Wikipedia is a starting point, but not like the final conclusion for all of the research that you're doing. So I think that we're at a moment where we're trying to come to understand what is the reality of using these AI tools, and AI is always developing in different kinds of ways. So we're always going to be trying to keep up with what's happening.

D: Yeah. I wanted to ask one quick follow-up here because you all have brought up some really great sort of ethical considerations here, some historical comparisons, but some of you already mentioned, of course, you have used or experimented with these tools in your classes. So I was curious for you all, are there concrete use cases, I guess, you could kind of talk us through here to think about what this has looked like for you as teachers as well?

B: Yeah, I've thought about a couple of things. I did write a piece that is for a WAC Clearinghouse publication that is basically like a collection of assignments [ed.: this collection is *TextGenEd*]. And those assignments come from people across disciplines. And it's set to be published in the fall, sometime this month or later on in the semester, where people from different disciplines can use those lesson plans. The one that I wrote about was how to use ChatGPT or other large language models for peer review, where you're giving that large language model, you know, the same kind of questions that you would give to a human peer about thesis about organization, about the tone or the language that's being used. And then from there, almost kind of like bringing together the perceptions of the human being with the feedback from the large language model to kind of help you think through the revision. So it actually builds on the initial revision assignment that I created back in 2021 with GPT-3, where at that time, you could give it some text, and then they would kind of fill in the rest of it afterwards. And then from there, I would have students kind of think through, what are the decisions that you're going to make for the revisions with that essay, this one is kind of like thinking of how large language models can be used as a peer reviewer. And just another couple of examples: One is a student in my spring semester class who could not come up with an idea for her ethnographic qualitative study. And we had talked about ChatGPT, just as a way to talk about technology and how it changes writing. And

the student used that a few weeks later and said, "Hey, this gave me some great ideas for my paper," and just the burden was lifted, the floodgates opened, all the ideas started coming from just from that list, from ChatGPT. And that was not with my, any kind of prompting or anything like that. The student just decided to use it. And this is the argument for folks in the first-year composition program here at my institution, is thinking about how does this kind of integrate throughout the writing process, kind of like assisting you in the process of writing something.

J: I'll add something to kind of go along with that. Talking about using the technology, using an LLM, in different parts of the process, one of the things that I've done this summer, what I've had my students do is create prompts about topics that they are really expert in. So you know, whatever viral TikTok trend or whatever, so that then they can do the fact-checking, so to speak. And they've found that kind of fun because, again, they get to be the expert. And it's also a great way to teach them that these LLMs are not perfect, that they are not thinking. And so that's one idea as well.

M: I love that. It reminds me of how I would teach students to work with a grammar checker by highlighting places where it makes a bad recommendation on my own writing. And so highlighting the mistakes, I think, is empowering to students in giving them chances to practice identifying those mistakes. It's empowering to them in terms of finding their own voice and seeing their own capacities as writers. And if you want to make the argument that we need to prepare them to work with AI in the workplace, that's also what that looks like, is that level of critical thinking and expertise where you can find the mistakes, you can find the problems in logic. And you can sort of cut through the mystique of the academese that it speaks so fluently, and students can see that it's sort of producing these gorgeous academic sentences that are maybe really empty, or just complete nonsense.

L: And I think because these technologies generate results so quickly, and in real time, it's a useful way to model different kinds of things in the classroom very quickly. And it makes it possible for students to respond quickly as they're sort of seeing. You know, one of the terms we haven't talked about yet is prompt engineering, which is the ways that you often—in order to get the kind of output that you want, you actually have to put a lot of thought into the input, into the instructions that you're actually giving an LLM like ChatGPT. And so it's interesting. You know, say you're talking about invention and talking about genre. This is an opportunity to really help students understand nuance in different genres, right? The fact that there are often subgenres to be considered when we're talking about the kind of output that we want. And as everybody has said, it's not just about using ChatGPT for invention. It's also about using it or other LLMs for revision as well. You know, there are opportunities to talk about things like tone. You can have students essentially generate a bunch of different outputs very quickly with different tones. And it gives them that opportunity to sort of

talk about, you know, nuance and connotation and language use in a way that can be really productive in the classroom setting.

H: Everybody kind of said all the really great smart things. If I worry about things related to you know, ChatGPT or whatever LLM it is, I worry about the disparate harm and help it provides to students, and I think about, you know, that if we don't make it an explicit part of our writing courses, then, you know, you're missing opportunities to help the full range of students think about how that this is a tool that will or will not work for them. And I think about just, like Sarah, also having taught at open-admission two-year colleges, you know, so many students who have had traumatizing academic writing experiences, right, and who have felt their whole lives like they've been told they're bad writers. For those students, there's a lot of work that has to happen in the classroom to help them feel confident and to help them feel empowered. That's something that I'm spending a lot of time worrying about: I think, you know, what is going to be the disparate impact on—I don't know who it's going to be, you know, probably the same students who are vulnerable or marginalized. And so I think we have to be really smart and forward-thinking and intentional about making sure that whatever we do in our classrooms, and whatever we do as professional organizations, are like a testament to that.

D: Yeah. Thank you. Thank you for that, Holly. And I mean, speaking of those kinds of disparate helps and harms, potentially, of AI, I mean, one of the things that comes up time and again in the working paper that I think resonates with that is issues of labor. And I will admit, that's one of my, you know, biggest hang-ups about generative artificial intelligence: at what points have or will these tools displaced, devalued, supplemented, or even supported the labor involved in writing—more specifically, in this case, the labor of writing teachers and students? So I'm curious, for you all, what you see as kind of the most salient and pressing labor issues that are raised by the current wave of generative artificial-intelligence technologies, whether those are potential threats or benefits or kind of messy mixes of the two.

L: Yeah, I just finished writing an essay for *Critical AI* that'll be out, hopefully, in a little while, where I really focused on these questions of labor. And I actually was drawing upon work that Holly has done in academic labor because I think that often, in the university, we have a tendency to make labor invisible, and often sort of devalue labor in favor of talking about things like, you know, intellectual property or ethics or, you know, these other kind of more intangible qualities than this question of investing hours in learning to do something, and investing hours in figuring out how to respond to student writing in ways that are productive. One of the things that's sort of come up in the MLA side is concerns about actually the academic labor that's performed in scholarly communities being outsourced. And so what does it mean when faculty are using things like ChatGPT to write recommendation letters, or write peer reviews for journal articles. And I think that, you know, talking about labor is often hard for the sort of elite

professional practices that a lot of academics are socialized with. But it's really key, I think, to have honest discussions about labor. And one of the things I try to do with my students who are using these technologies is I try to get them to send me everything, right? I want to see the prompt you give it, I want to see the output, I want to see the whole, you know, paper trail of how you're interacting with these technologies so I can better understand the labor that you've invested in our sort of joint enterprise of, you know, this cocreation of knowledge that we're doing together in the classroom. You know, it's interesting to look at things like the WGA writers strike right now, where there's a real attention to labor politics. I think we need to think about these questions of solidarity around labor, and, you know, obviously, there's real pain points for the most contingent employees of college campuses. This is going to disproportionately hurt writing teachers and language teachers who are teaching world languages, who are already the most precarious employees who are often working on contracts where they have too many students and not enough time to give feedback and are already overwhelmed. They're anxious about things like health insurance, and this is just going to make their lives worse. And then some institutions are expecting them to police student behavior in a way that's completely unrealistic. I mean, that's just, you know, playing whack-a-mole of the worst kind. So, I mean, I think that figuring out how to center labor in this conversation is really important.

J: To speak to this, Antonio and I were on a panel together earlier this summer for the MLA. It was with language department chairs, and one of the things that I found really interesting was something that we as a group have talked about a little bit already, but the idea of flattening of discourse: that by simply allowing these LLMs to reify the systems, the discourse, the language that exists already, we're no longer pushing against ourselves, or encouraging our students to push against the systems. And so while it does allow more students—and again, I'm thinking about this as a teacher at an open-access institution—allow students access to the language of power, also stopping the questioning of the language of power is an important aspect that we as a task force talked about quite a bit.

B: Yeah. Part of the way I'm thinking about it is, like, what makes up the corpus of data that these AI actually run on? And reading one article that's well known about stochastic parrots—

D: Is that “On the Dangers of Stochastic Parrots”?

B: Yes, that's it. It's that, you know, what data is actually worthy to go into the large language model, what language is filtered out? And so some of the languages that are specific to marginalized communities, Black people, LGBTQ people, is kind of filtered out because it's considered offensive, or “that's not a word that we really want” even though that's been reclaimed by the marginalized community, it means that they're not going to appear. Or they may actually be flagged, as, you know, something that's incorrect. So one example that someone posted on Twitter was asking ChatGPT, “What

are some examples of incorrect sentences?" And sometimes the list that it gives are sentences that are actually Black linguistic practices, but it's counted as wrong. And so that kind of builds up this bigger question: what happens when Black students have access to something like ChatGPT if you are a teacher like myself, who really wants students to think about linguistic justice, to let them use the linguistic practices that they have, and they're using something like ChatGPT, you know? Then there's kind of like a conflict. And there's a larger conversation to be had: how do we put some type of critical frameworks on these technologies where we know that we can use, in my class, at least, these practices, but it's not represented in the large language models that you're using. So there are two sides to it. We often think about, like, what is the output, but at the same time, we don't have access to the data that they're using. There are some large language models that are trying to be more open, like Hugging Face, for example, where you can get access for research purposes to the data corpus that they use for their large language model. And that could be another level of critical AI literacy for students: here is a technology where we can actually look at everything that they use that's put into it. And now we can get a sense of how these things are made and why sometimes we are not found in those tools. I also think about, like, the implications for, say, using AI tools and kind of passing the responsibilities of teaching over to the companies that actually provide them. And it's almost like universities no longer become partners with these companies, but the universities are the companies themselves if you rely on, if you want to shift the labor from our contingent faculty over to those companies. And it brings up questions about, like, what kind of decisions are you outsourcing to those companies to let them decide for your institutional context? And do you really want them to be making those decisions? Are some of the assessment models going to be heavily biased against marginalized students if there is an approach to assessment that's supposed to show which students are very successful? What's going on behind those AI tools that make those kinds of decisions? And there's no kind of transparency on how those assessment models may work, so I'm just kind of thinking about how one implication of shifting labor away from people who understand the local context to understand their students, their colleagues, who really understand how the discipline is taking shape, in ways that AI tools may not be able to keep up with. And, you know, what does it mean to say, this company can now take over the majority of that labor because they have these great tools for us to really use? It might kind of be a little bit messy.

L: And I don't want to make this an advertisement for Kate Crawford's *Atlas of AI*. But one of the things she points out is how much labor is invested in producing these AI technologies. And there was a recent article in the *Wall Street Journal*, actually, about Kenyan labor groups protesting the ways that workers were being abused while training the algorithms for ChatGPT in order to deal with violent and sexual content and hate speech, and so those workers, you know, much like content moderators in the Philippines who are doing work for social media companies, really were being abused at sub-minimum wage in order to make these AI technologies appear friction-free and easy to use and inoffensive. You know, they were the people investing labor in that.

M: I was just going to add, on another note, another document we worked on this summer alongside the working paper was a comment to the White House Office of Science and Technology Policy. They had a request for input on AI. And one of the things we brought up was just this question of the labor required to adapt, to build critical AI literacy in faculty so that we can guide students around it, and just the investment that's needed that may not be feasible at individual institutional levels. So really looking at, you know, where should that investment come from? You know, does it make sense for the companies just to be making money off the product while the teachers and the institutions and the taxpayers have to pay for the adaptation of education to it? Or how else can we support that labor that supports student learning, and that supports a flourishing democracy, theoretically? And the other thing that comes up in relation to labor is this question of transparency around what is AI and what is not. So, you know, given the context in higher ed, we might be driven sometimes to feel that we need to use these systems to help us grade student work. But then, do students know that? If they do, how does that affect them? What are the privacy concerns around submitting their work to AI? Is that legitimate? Just big questions that touch on, you know, is it going to be disclosed, and I feel pretty strongly that we need to disclose when AI is being used and when it's not. That's actually a core tenet of this White House Blueprint for an AI Bill of Rights—that you should know when an automated system is being used. And Kathryn Conrad has built on that in this Blueprint for [an AI Bill of Rights for] Education, where she's saying that students have a right to know when a system has been used to evaluate them and to ask for a human evaluation if they want. And similarly, you know, we might also ask, do we want to spend our time grading ChatGPT—brings up those questions around, do we know what's AI and what's not in student work? And transparency around, you know, whose writing is in the dataset that ChatGPT is trained on. That's also something we called for in our comment to the White House office—just to *know*. I think that question of labor touches on everything.

D: So finally, I wanted to ask about something that you will note in the working paper, which is that this is just a starting point—essentially, you know, that large language models raise questions about, and I'm quoting here, “ethics, citation practices, linguistic diversity, reading, pedagogical implications, language instruction, and others.” So I'm curious to hear whether as teachers, scholars, citizens, plain old people, what do you see as the most pressing concern related to LLMs at this particular juncture?

L: For me, I've been thinking a lot about what fellow task force member Matt Kirschenbaum has called the textpocalypse, which is the possibility that we're going from a read-write culture—and, you know, specifically a read-write internet—to one in which we have a technology that can mimic human discourse and generate a lot of it, and can also then generate stuff that it then gets trained upon, so that we are going to

potentially face a situation in which human communication is just not something that we can count on in any way. And for me, I'm interested in things like trust and intimacy, and I think that the possibility that, you know, we can't feel the same kinds of feelings of trust and intimacy if we're not engaged in face-to-face communication—like, that possibility is a really disturbing one. It has all sorts of ramifications for our democracy, for our public life, you know, for the future of education. And so for me, that possibility, that there could be unintended consequences of this rapidly developing technology that's better and better able to mimic human discourse, just drowns out human discourse entirely.

H: Yes to everything Liz said. I think the things that I've been thinking a lot about are the classroom experience for our students. You know, I want students, if they don't love writing—I'd like them to love writing, but if they don't love writing, at least see the value of writing for doing things in the world, whether that's, you know, as local as just as a form of expression, journaling, diaries, internet blogs, whatever. Or, you know, as a tool that you can actually use to accomplish things, whether it's an email to your kid's teacher, or a letter to the editor, or a proposal to the school board, or, you know, advocacy for a candidate. The authentic writing tasks that have a purpose and an audience and a goal are meaningful, you know. And so for me, just with that as a core belief about my writing classrooms and about what I'd like to see for writing classrooms, and for students and teachers, the development of AI and how we use it, and how we talk about in our classrooms, you know, the extent to which it may threaten that sort of real goal, we need to dedicate our attention to making sure that whether that's like a critical AI literacy framework or pedagogical resources for students and teachers, you know, in terms of the task force's work, that's where I expect will be a big part of the conversation as we kind of identify what we want to do next.

J: I will say my biggest concern going forward, not as a member of the task force but in my everyday life as a faculty member, and mostly as an academic integrity officer, is the wide range of knowledge and attitudes in full-time faculty and part-time faculty when it comes to the issues of generative AI—that students can't really count on knowing exactly what their teachers' policies are, that colleges themselves are still figuring out what policies are and that it's changing so quickly. That those policies, and I'm in the middle of writing one right now, you know, might become outdated. And so I think the practical aspects of how to nimbly but also thoughtfully respond is a challenge. And we're back to this issue of labor. Because it takes time and it takes thought, and these are things that I think a lot of us are trying to do, on top of everything else we were doing, and weren't thinking about this a year ago. And yet here we are, with this brand-new load of information that we are having to incorporate into our pedagogy, into our curricula.

M: One thing that struck me when we had our retreat and we heard from Google was that the representative was sort of asking *us* how to maintain information integrity and how we could carry out that role in society, and saying there were limits to what he thought Google could do. And I think that connects with this concern that, you know, a number of people have mentioned about how fast this is rolling out, how it's really on tech companies' terms, but there needs to be more democratic oversight. And I think, you know, as teachers of rhetoric as writing teachers, we can see ourselves actually as experts who are needed in this domain to help shape how it rolls out, how it is regulated. We've also seen leaders of many of these tech companies, like Sam Altman, run to government saying, "Please regulate us," right? So there's an opportunity for some shaping of this that we could participate in. And, you know, we have this expertise in—they're calling it prompt engineering, and I think that it's mostly rhetoric, is what it is. It's not coding, it's rhetoric, you know? So I hope that we can feel empowered to see our own expertise and the need for it in this societal discourse, and as input to these companies. As the technology develops, how do we want to shape it to shape the future of human communication, to preserve that understanding of writing as a space for thinking, and for relationship, for mediating human relationships? And those are things we know about, and we teach about already, that we can bring to the discourse.

B: So I don't know if it's a concern; maybe it's a curiosity. And maybe this kind of digs into the benefits section of our working paper. But my mind immediately goes back to the brawl that many of us have been talking about in Montgomery, Alabama, that has turned into a very internet-memeable event. And I'm just thinking about so many Black people on what is now called X, the memes that were created that kind of speak to the cultural racial messaging that comes out of that event in Montgomery, Alabama. And so I'm kind of thinking about, you know, what are students doing that's really interesting or creative when they're using these technologies? And how can that inform the frameworks that we have, our expertise and rhetoric at the same time? And so that's where my mind is right now, is that with the fall semester coming up, I do have a couple of lesson plans in place where we're kind of like experimenting, playing around with the possibilities of large language models while also thinking about the ethics, concerns about climate change. And I'm really interested to know, what are the use cases from our students, and how is that going to really help us come to the realities of what these tools really mean to us, as opposed to the peril and the hype that we're in right now?

D: Thank you all one last time, once again, for not only your work on this working paper and on this task force, but thank you for taking the time for this.

[roundtable ends, drumbeat fades in]

D: Thanks so much for listening to this episode of *Rhetorcity*! I'll be back soon to bring you more episodes featuring Jennifer LeMesurier and Keith Gilyard. In the meantime, if you want to find out more about *Rhetorcity* or any of my other projects, you can visit RhetEric.org. That's r h e t e r i c dot o rg. And a reminder that you can check out the episode description for links to other materials from The Big Rhetorical Podcast Carnival as well as links to some of the sources and resources mentioned in this episode's discussion. Till next time, this is Eric Detweiler, trying to dodge the Eliza Effect for one more semester.

[*Rhetorcity* theme plays]