

Thank you, Morgan, and thank you all for being here this morning. I was going to talk about our Sourcery project today, which is an application to streamline remote access to archival materials for both researchers and archivists, but at the last minute I've decided to bow to the inevitable and talk about this [ChatGPT] instead.

I can almost feel the inner groaning from those of you who are exhausted and perhaps dismayed by the 24/7 coverage of “Generative AI.” I’m talking about things like ChatGPT, DALL-E, MidJourney, Jasper, Stable Diffusion, and Google’s just released, Bard. Indeed, the coverage has been wall to wall, and the hype has at times been breathless, and it’s reasonable to be skeptical of “the next big thing” from Silicon Valley. After all we’ve just seen the Silicon Valley hype machine very nearly bring down the banking system. In just past year, we’ve seen the spectacular fall of the *last* “next big thing,” so-called “crypto,” which promised to revolutionize everything from finance to art. And we’ve just lived through a decade in which the social media giants have created a veritable dystopia of teen suicide, election interference, and resurgent white nationalism.

So, when the tech industry tells you that *this whatever* is “going to change everything,” it makes sense to be wary. I’m wary myself. But with a healthy dose of skepticism, and more than a little cynicism, I’m here to tell you today as a 25-year veteran of the digital humanities and a historian of science and technology, as someone who teaches the history of digital culture, that Generative AI is the biggest change in the information landscape since at least 1994 and the launch of the Netscape web browser which brought the Internet to billions. It’s surely bigger than the rise of search with Google in the early 2000s or the rise of social media in the early twenty 10s. And it’s moving at a speed that makes it extremely difficult to say where it’s headed. But let’s just say that if we all had an inkling that the robots were coming 100 or 50 or 25 years in the future, it’s now clear to me that they’ll be here in a matter of just a few years—if not a few months.

It's hard to overstate just how fast this is happening. Let me give you an example. Here is the text of a talk entitled coincidentally teaching with primary sources in the next digital age This text was generated by ChatGPT—or GPT-3.5—the version which was made available to the public last fall, and which really kicked off this wall-to-wall media frenzy over Generative AI.

You can see that it does a plausible job of producing a three-to-five paragraph essay on the topic of my talk today that would not be an embarrassment if it was written by your ninth-grade son or daughter. It covers a range of relevant topics, provides a cogent, if simplistic, explanation of those topics, and it does so in correct and readable English prose. [I'll give you a minute to read it.]

Now here's the same talk generated by GPT-4 which came out just last week. It's significantly more convincing than the text produced by version 3.5. It demonstrates a much greater fluency with the language of libraries and archives. It correctly identifies many if not most of the most salient issues facing teaching in archives today and provides much greater detail and nuance. It's even a little trendy, using some of the edu-speak and library lingo that you'd hear at a conference of educators or librarians in 2023. [Take a minute to read it.]

Now here's the outline for a slide deck of this talk that I asked GPT-4 to compose, complete with suggestions for relevant images. And here's the text of speaker notes for just one of the bullets in this talk. [I give you another minute.]

Now—if I had generated speaker notes for each of the bullets in this outline and asked GPT's stablemate and image generator, DALL-E, to create accompanying images—all of which would have taken the systems about 5 minutes—and then delivered this talk more or less verbatim to this highly educated, highly accomplished, Ivy League audience, I'm guessing the reaction would have been: “OK, seems a little basic for this kind of thing” and “wow, that was talk was a big

piece of milktoast.” It would have been completely uninspiring, and there would have been plenty to criticize—but neither would I have seemed completely out of place at this podium. After all, how many crappy, uninspiring, worn out PowerPoints have you sat through in your career? But the important point to stress here is that in less than six months, the technology has gone from writing at a ninth-grade level to writing at a college level and maybe even more.

Much of the discourse among journalists and in the academic blogs and social media has revolved around picking out the mistakes these technologies make. For example, my good friend at Middlebury, Jason Mittell, along with many others, has pointed out that ChatGPT tends to invent citations: references to articles attributed to authors with titles that look plausible in real journals that do not, in fact, exist. Australian literary scholar, Andrew Dean, has pointed out how ChatGPT spectacularly misunderstands some metaphors in poetry. And it’s true. Generative AIs make lots of extremely weird mistakes, and they wrap those mistakes in extremely convincing-sounding prose, which sometimes makes them hard to catch. And as Matt Kirschenbaum has pointed out: they’re going to *flood the Internet* with this stuff. Undoubtedly there are *issues* here.

But based on the trajectory of improvement, do we really think these problems won’t be fixed? Let me give another couple of examples. Look at this chart, which shows GPT-3.5’s performance on a range of real-world tests. Now look at this chart, which shows GPT-4’s improvement. If these robots have gone from writing decent five-paragraph high school essays to passing the Bar Exam (in the 90<sup>th</sup> percentile!!) in six months, do we really think they won’t figure out citations in the next year, or two, or five? Keep in mind that GPT-4 is a *general purpose* model that’s engineered to do everything pretty well. It wasn’t even engineered to take the Bar Exam. Google CEO, Sundar Pichai tells us that AI computing power is doubling *every six months*. If today it can kill the Bar Exam, do we really think it won’t be able to produce a plausible article for a mid-tier peer reviewed scholarly journal in a minor

sub-discipline of the humanities in a year or two? Are we confident that there will be any way for us to tell that machine-written article from one written by a human?

(And just so our friends in the STEM fields don't start feeling too smug, GPT can write code too. Not perfectly of course, but it wasn't trained for that either. It just figured it out. Do we really think it's that long until an AI can build yet another delivery app for yet another fast-food chain? Indeed, Ubisoft and Roblox are starting to use AI to design games. Our students' parents are going to have to start getting their heads around the fact that "learning to code" isn't going to be the bulletproof job-market armor they thought it was.)

There are also some big philosophical issues at play here. One is around *meaning*. The way GPT-4 and other generative AIs produce text is by predicting the next word in a sentence statistically based on a model of drawn from an unimaginably large (and frankly unknowable) corpus of text the size of the whole Internet—a "large language model" or LLM—not by understanding the topic they're given. In this way the prose they produce is totally *devoid of meaning*. Drawing on philosopher, Harry Frankfurt's definition of "bullshit" as "speech intended to persuade without regard for truth", Princeton computer scientists Arvind Narayanan and Sayash Kapoor suggest that these LLMs are merely "bullshit generators." But if something meaningless is indistinguishable from something meaningful—if it holds meaning for us, but not the machine—is it *really* meaningless? If we can't tell the simulation from the real, does it matter?

These are important philosophical, even moral, questions. But here we are. As a purely practical matter, we need to start preparing our students to live in a world of sometimes bogus, often very useful, generative AI. The first-year students arriving in the fall may very well graduate into a world that has no way of knowing machine-generated from human-generated work. Whatever we think about them, however we *feel* about them (and I feel a mixture of disorientation, disgust, and exhaustion),

these technologies are going to drastically change what those Silicon Valley types might call “the value proposition” of human creativity and knowledge creation. Framing it in these terms is ugly, but that’s the reality our students will face. And there’s an urgency to it that we must face.

So, let’s get down to brass tacks. What does all this mean for what we’re here to talk about today, Teaching with Primary Sources?

One way to start to answer this question is to take the value proposition framing seriously and ask ourselves, “What kinds of human textual production will continue to be of value in this new future and what kinds will not?” One thing I think we can say pretty much for sure is that writing based on research that can be done entirely online is in trouble. More precisely, writing about things about which there’s already a lot online is in trouble. Let’s call this “synthetic writing” for short. Writing that synthesizes existing writing is almost certainly going to be done better by robots. This means that what has passed as “journalism” for the past 20 years since Google revolutionized the ad business—those Buzz Feed style listicles (“The 20 best places in Dallas for tacos!”) that flood the internet and are designed for nothing more than to sell search ads against—that’s dead.

But it’s not only that. Other kinds of synthetic writing—for example, student essays that compare and contrast two texts or (more relevant to us today) place a primary source in the context drawn from secondary source reading—those are dead too. Omeka exhibits that synthesize narrative threads among a group of primary sources chosen from our digitized collections? Not yet, but soon.

And it’s not just that these kinds of assignments will be obsolete because AI will make it too easy for students to cheat. It’s what’s the point of teaching students to do something that they’ll never be asked to do again outside of school? This has always been a problem with college essays that were only ever destined for a file cabinet in the professor’s desk.

But at least we could tell ourselves that we were doing something that simulated the kind of knowledge work they would so as lawyers and teachers and businesspeople out in the real world. But now?

(Incidentally, I also fear that synthetic scholarly writing is in trouble, for instance a Marxist analysis of Don Quixote. When there's a lot of text about Marx and a lot of text about Don Quixote out there on the Internet, chances are the AI will do a better—certainly a much faster—job of weaving the two together. Revisionist and theoretical takes on known narratives are in trouble.)

We have to start looking for the things we have to offer that are (at least for now) AI-proof. We have to start thinking about the skills that students will need to navigate an AI world. Those are the things that will be of real value to them. So, I'm going to use the rest of my time to start exploring with you (because I certainly don't have any hard and fast answers) some of the shifts we might want to start to make to accommodate ourselves and our students to this new world. I'm going to quickly run through 8 things.

1. The most obvious thing we can do it to refocus on the physical. GPT and its competitors are trained on digitized sources. At least for now they can only be as smart as what's already on the Internet. They can't know anything about anything that's not online. That's going to mean that physical archives (and material culture in general) will take on a much greater prominence as the things that AI doesn't know about and can't say anything about. In an age of AI, there will be much greater demand for the undigitized stuff. Being able to work with undigitized materials is going to be a big "value add" for humans in the age of these LLMs. And our students do not know how to access it. Most of us were trained on card catalogs, in sorting through library stacks, of traveling to different archives and sifting through boxes of sources. Having been born into the age of Google, our students are much less good at this, and they're going to need to get better. Moreover,

they're going to need better ways of getting at these physical sources that don't always involve tons of travel, with all its risks to climate and contagion. Archivists, meanwhile, will need new tools to deal with the increased demand. We launched our Sourcery app, which is designed to provide better connections between researchers and archivists and to provide improved access to remote undigitized sources before these LLMs hit the papers. But tools like Sourcery are going to be increasingly important in an age when the kind of access that real humans need isn't the digital kind, but the physical kind.

2. Moreover, we should start rethinking our digitization programs. The copyright issues around LLMs are (let's say) *complex*, but currently Open AI, Google, Microsoft, Meta, and the others are rolling right ahead, sucking up anything they can get their hands on, and processing those materials through their AIs. This includes all of the open access materials we have spent 30 years producing for the greater good. Maybe we want to start asking ourselves whether we really want to continue providing completely open, barrier-free access to these materials. We've assumed that more open meant more humane. But when it's a robot taking advantage of that openness? We need a gut check.
3. AIs will in general just be better at the Internet than us. They'll find, sort, sift, and synthesize things faster. They'll conduct multi-step online operations—like booking a trip or editing a podcast—faster than us. This hits a generation that's extremely invested in being good at the Internet, and, unfortunately, increasingly bad at working in the real world. Our current undergraduates have been deeply marked by the experience of the pandemic. I'm sure many of you have seen a drastic increase in class absences and a drastic decrease in class participation since the pandemic. We know from data that more and more of our students struggle with depression and anxiety. Students have difficulty forming friendships in the real world. There are a growing number of students who choose to

take all online classes even though they're living in the dorms. This attachment to the virtual may not serve them well in a world where the virtual is dominated by robots who are better than us at doing things in the digital world. We need to get our students re-accustomed to human-to-human connections.

4. At the same time, we need to encourage students to know themselves better. We need to help them cultivate authentic, personal interests. This is a generation that has been trained *to write to the test*. But AIs will be able to write to the test much better than we can. AIs will be able to ascertain much better than we can what *they* (whomever *they* is: the school board, the college board, the boss, the search algorithm) want. But what the AI can't really do is tell us what we want, what we like, what we're interested in and how to get it. We need to cultivate our students' sense of themselves and help them work with the new AIs to get it. Otherwise, the AI will just *tell them* what they're interested in, in ways that are much more sophisticated and convincing than the Instagram and TikTok algorithms that are currently shoving content at them. For those of us teaching with primary sources this means exposing them to the different, the out of the ordinary, the inscrutable. It means helping them become good "pickers" – helping them select the primary sources that truly hold meaning for them. As educators of all sorts, it means building up their personalities, celebrating their uniqueness, and supporting their difference.
5. I think we also need to return to teaching names and dates history. That's an unfashionable statement. The conventional wisdom of at least the last 30 years is that that names, dates, and places aren't that important to memorize because the real stuff of history are the themes and theories—and anyway, the Google can always give them to us. Moreover, names and dates history is *boring* and with the humanities in perpetual crisis and on the chopping block in the neoliberal university, we want to do everything we can to make our



disciplines more attractive. But memorized names, and dates, and places are the things that allow historians to make the creative leaps that constitute new ideas. The biggest gap I see between students of all stripes, including graduate students, and the privileged few like me who make it into university teaching positions (besides white male privilege) is a fluency with names, dates, and places. The historians that impress most are the ones who can take two apparently disconnected happenings and draw a meaningful connection between them. Most often the thing that suggests that connection to them is a connected name, date, place, source, event, institution, etc. Those connections are where new historical ideas are born. Not where they end, for sure, but where they are born. AI is going to be very good at synthesizing existing ideas. But it may be less good at making new ones. We need students who can birth new ideas.

6. Related to this is the way we teach students to read. In the last 20 years, largely in response to the demands of testing, but also in response to the prioritization of “critical thinking” as a career skill, we’ve taught students not to read for immersion, for distraction, for imagination, but for analysis. Kids read tactically. They don’t just *read*. In many cases, this means they don’t read at all unless they have to. Yet, this is exactly how the AI reads. Tactically. Purely for analysis. Purely to answer the question. And they’ll ultimately be able to do this way better than us. But humans can read in another way. To be inspired. To be moved. We need to get back to this. The imaginative mode of reading will set us apart.
7. More practically, we need to start working with these models to get better at asking them the right questions. If you’ve spent any time with them, you’ll know that what you put in is very important in determining what you get out. Here’s an example. In this chat, I asked GPT-3.5, “How can I teach with primary sources.” OK. [Take a minute to read it.] That’s not bad. But then in another chat I asked, “Give me a step-by-step plan for using primary sources in

the classroom to teach students to make use of historical evidence in their writing” and I followed it up with a few more questions: “Can you elaborate?” and “Are there other steps I should take?” and then “Can you suggest an assignment that will assess these skills?” [Take a minute here too.] You’ll see that it gets better and better as it goes along. I’m no expert at this. But I’m planning on becoming one because I want to be able to show our students how to use it well. Because, don’t fool yourselves, they’re going to use it.

8. Finally, then, perhaps the most immediate thing we can do is to inculcate good practice around students’ use of AI generated content. We need to establish citation practices, and indeed the MLA has just suggested some guidance for citing generative AI content. Stanford, and other universities, are beginning to issue policies and teaching guidance. So far, these policies are pretty weak. Stanford’s policy basically boils down to, “Students: Don’t cheat. Faculty: Figure it out for yourselves.” [Show Stanford policy.] It’s a busy time of year and all, but we need urgently to work with administration to make these things better.

I’m nearly out of time, and I really, really want to leave time for conversation, so I’ll leave it there. These are just a couple of thoughts that I’ve pulled together in my few weeks of following these developments. As I’ve said, I’m no expert in AI but I think I can fairly call myself an expert in digital humanities and the history of technology, and I’m convinced this new world is right around the corner. I don’t have to like it. You don’t have to like it. If we want to stop it, or slow it down, we should advocate for that. But we need to understand it. We need to prepare our students for it.

At the same time, if you look at my list of things we should be doing to prepare for the AI revolution, they are, in fact, things we should have been (and in many cases have been) doing all along. Paying more attention to the undigitized materials in our collections? I’m guessing

that's something you already want to do. Helping students have meaningful, in-person, human connections? Ditto. Paying more attention to what we put online to be indexed, manipulated, sold against search advertising? Ditto. Encouraging students to have greater fluency with names, dates, and places? Helping them format more sophisticated search queries? Promoting better citation practice for born-digital materials and greater academic integrity? Ditto. Ditto. Ditto.

AI is going to change the way we do things. Make no mistake. But like all other technological revolutions, the changes it demands will just require us to be better teachers, better archivists, better humans.

Thank you.