

Vincent Munoz:

I think what we need to do is explain how our principles of free speech, free inquiry will help serve the cause of justice.

Betty Friendan:

The First Amendment, the constitutional freedom of speech and freedom of conscience that is the bulwark of our democracy.

Bettina Aptheker:

There was a passion in what was being said affirming what people considered a sacred constitutional right, freedom of speech and freedom of association

Michelle Deutchman:

From the UC National Center for Free Speech and Civic Engagement. This is SpeechMatters, a podcast about expression, engagement, and democratic learning in higher education. I'm Michelle Deutchman, the center's executive director and your host.

Almost every headline you read these days is focused on how generative AI, such as OpenAI's hugely popular ChatGPT will revolutionize the way we live, work, and learn. Today's guest, Matt Perault, a nationally recognized technology policy expert, will help us unravel the web of predictions about this new tool. Does ChatGPT signal the death knell for the college essay? Will classroom learning never be the same? Are we all going to lose our jobs to increasingly intelligent and personable chatbots? We will talk about AI's impacts big and small. But first let's turn to class notes. A look at what's making headlines.

In higher education we most often associate June with completion and commencement, but June is much more than the end of the year. It's also LGBTQ Pride Month celebrated in June to commemorate the Stonewall Uprising, which took place on June 28th, 1969 in New York City. For those of you who need a history refresher, the Stonewall riots were a series of spontaneous protests by members of the gay community in response to a police raid that began in the wee hours of the morning. Patrons of the Stonewall and other lesbian and gay bars in Lower Manhattan fought back when police became violent. These riots are widely considered to be a watershed in the transformation of the gay liberation movement in the US. Cities and campuses around the United States celebrate pride month with parades, picnics, educational events and more. It's not too late to grab a rainbow flag and mark the occasion.

In June, we also commemorate Juneteenth, the day in 1865 that enslaved people in Texas learned they were free. This news was delivered two and a half years after President Lincoln's Emancipation Proclamation became law. President Biden designated Juneteenth, also known as Jubilee Day, Liberation Day and Freedom Day, a federal holiday two years ago. If you're interested in learning more about this moment in history, I recommend On Juneteenth a short but powerful read by historian and law professor Annette Gordon-Reed.

But there is more. June also marks the end of the Supreme Court term when we play the waiting game. This year, the higher education community waits for what virtually everyone assumes will be the dismantling of affirmative action. Also on the docket is the matter of student debt relief as well as the critical speech issues in 303 Creative. If you didn't already listen to Lambda's brilliant Jennifer Pizer discuss this case in season one, episode six. Get up to speed before the opinion is released.

In the flurry of all that June has to offer, there are yet more incidences of bannings and attempts to censor ideas allowed in schools. In one notable example, a Florida elementary school prohibited students from reading Amanda Gorman's *The Hills We Climb*, the poem Gorman read at Biden's inauguration. In better news, however, the coverage of the banning of the poem led to an upswing in sales of Gorman's books. Want to be part of the movement that fights censorship of ideas found in literature? On your summer vacation, read something from the American Library Association's list of the most banned books in 2022. You can find a link to it in the episode notes. One of my goals is to check them all off. And if you're able support a local bookstore and buy your copy there. Happy reading.

Now let's turn our attention to our guest, Matt Perault. Matt is a professor of practice at the University of North Carolina at Chapel Hill's School of Information and Library Science, as well as the director of the School Center on Technology Policy. Before that, Matt worked at Facebook where he was a director on the public policy team and the head of the global policy development team. He covered issues including antitrust, law enforcement and human rights, as well as oversaw the company's policy work on emerging technologies like artificial intelligence and virtual reality. Matt holds a law degree from Harvard Law School, a master's in public policy from Duke Sanford School of Public Policy and a BA in Political Science from Brown University. Matt also serves as a tech policy consultant working with some companies that operate in this area. Welcome and thank you for joining us, Matt.

Matt Perault:

Thanks so much for having me on.

Michelle Deutchman:

So I think we're going to just do some quick level setting, right? Chatbots have existed in various forms for decades, but the energy around ChatGPT, the fear, the panic, the glee, the excitement has far surpassed how other AI advances have been received. And I'm wondering if you can share with us, why is there so much buzz? Right? What makes ChatGPT stand out?

Matt Perault:

It's a really good question and I'm not sure I know with any kind of definitiveness what the answer is, but I think people see in it something tangible and different. When we were at Facebook, a lot of the strategy around trying to get users to join the product had to do with how can you get them to see as quickly as possible something tangible and beneficial in the product? And we would talk about sort of an aha moment that people would have when they saw some connection, usually to a family event or some family occasion or something friend related, and then all of a sudden they would be interested in using the service. But you had to get them to that point where they saw some component of a value proposition in order to show them that the product was worth using.

And I think with ChatGPT, and I should say other integrations of similar technology, people can see very quickly that there is something new and valuable, that you can ask it to generate a 500 word essay on a topic, and yes, there might be instances where there is inaccurate information or it's imperfect or you would write something slightly differently, but the fact that you can get in many cases something tangible and meaningful and valuable for you is really helpful. I have lots of examples of my own usage, although I think they're sort of silly and just related to my idiosyncrasies and the things that I am interested in. But I have had experiences where I think a search on a search engine pre November, 2022 would take longer to find a result that would be satisfying to me then it seems like you can get, at least in some instances in a generative AI tool.

Michelle Deutchman:

Okay. So sort of a combination of immediate gratification and also product. And hearing you talk about this is making me think maybe I should be playing around more with it, right? Because I think I've heard a lot of examples of sort of the fun funny, but I haven't given as much thought as I think I should to how it might actually enhance my work. And I don't know if I'm clearly in the minority because there's allegedly 100 million users, so maybe I need to get going on that.

Matt Perault:

Right? But I'm a user too and I'm not... And there's a difference between being... I mean, we can look at the ways that tech companies have traditionally looked at usage metrics, monthly active users versus daily active users versus using something once every six months or once every year. My usage, I can confess, is more like... I'm trying to think of the last time I was actively using a generative AI tool and it might have been weeks. And I'm actually sort of embarrassed to say that because I think there is something in me that is hesitant to find it as useful, to experience its utility at its maximum. And I think that's because for reasons, and I'm sure we'll talk more about this, but I think it's because it's coming for what I view as my expertise in the world. In my career, the thing that has I think enabled me to get various different jobs and do various different types of things is the ability to basically do a bad first draft.

I said to someone at Facebook, "I see my job as basically generating bad first drafts." And he was like, "Yes, yes, yes, yes." And then he said, "But could you do some good first drafts every now and then too?" And I think generative AI tools really take that off the table. It might not generate a perfect draft that is the kind of thing that you could... a law review article or an op-ed or something or a speech, but they can generate, they take you from a blank page to some form of a written product. And that ability to wrestle with a blank page and come up with something has always been a skill that I've prided myself on. I'm not saying I'm the world's best person at it, but that's kind of been the value proposition that I bring to the world.

This technology is not one that is other technologies coming for a part of the population that feels different or distant from me. This isn't displacing low wage workers necessarily or exclusively. This is coming for people who write stuff for a living and I am amongst those people who write stuff. And so I think the idea that I would start every assignment or every project with seeing what I can get through a generative AI tool is a very scary prospect for me. And my thought is if I continue to feel scared by it, I will be out of a job. Someone said to me recently you won't lose a job to generative AI. You will lose a job because you're not able to use generative AI effectively. And that seems really true to me. We're going to go into a world where you have to know how to use the technology and that means jumping in and figuring out what works and what and becoming a sophisticated user. And so I think that I'm saying to you today, I can't really remember the last time I was playing around with a generative AI tool is going to really be a mindset that will be to my detriment over time.

Michelle Deutchman:

I really appreciate that and I appreciate your candor, and in some ways it really sets the stage well, even though that might not what have been what you were intending in terms of why there is so much talk in academic and higher education circles about generative AI, especially this focus on classwork. Is the college essay dead? Are people going to be able to give exams the same way? How we're teaching and learning and research, which at its core is largely a lot of writing and thinking kind of going to have to be upended. And I'm curious what your thoughts on that are. And then a follow-up, which I'll give you in

advance, is what do we do as we think about, and I don't know if this is a real thing, but it's not media literacy, but AI literacy for people that are starting at least their higher education journeys?

Matt Perault:

Yeah. Well, I think what you just said will be the key thing. This is why I think bans and slowing down and trying to get it out of the classroom is not the right approach. I mean abstinence has never been a particularly successful approach to a variety of different issues, drugs, sex. I think it applies as well to AI trying to get people to forswear the technology is I think a long, brutal, losing battle. And also I think a disservice to students. I don't think the future of the workforce will be places where they say you can't use AI tools to generate something. The question will be how good can you make the thing, which will mean I think bringing a layer of skills and knowledge that sits on top of AI. And so I think the kinds of assignments that strike me as really interesting and important is you is actually encouraging or requiring even students to use the technology and then figuring out what's the learning experience that moves them beyond a kind of unsophisticated use of the technology to a more sophisticated one.

I think there are zillions of ways that you can test that. I mean, this wouldn't work for every type of a class, but I sort of thought if you ask students to generate an original draft of something using a generative AI tool and then you graded their use of draft changes to improve upon it, that is something that I don't think AI can cheat that. And you can evaluate two things. You evaluate the sophistication of a prompt, which I think will be hard and will be something that students need to learn. How do I prompt the tool in the right way? And then you evaluate their ability to take a first draft to something that's really significantly better. That strikes me as the skillset for a lot of future jobs in the workplace, not generating the original draft.

Michelle Deutchman:

Well, and I think it's interesting because in some ways your focus is less on how we train students, I mean ultimately there. But ultimately faculty need to know how they can do the best job of helping students to utilize it in a way. So in some ways the training might have to be, wait, okay, how do people in different disciplines learn what kinds of prompts and assignments to give to students and others in order to enhance kind of day-to-day learning in the classroom?

Matt Perault:

I think that's totally true. And actually there, I don't know exactly know what it is, but there's something about how you're framing this that's making me think about my experience as a new academic watching the way, at least the institutions that I was close to, treated the pandemic. Which was sort of I thought, oh my God, we need to hold on to our prior model, our prior business model, our prior pedagogical model as tightly as possible because we might lose it and then we'll be dead. And that struck me as a major mistake. I think there needs to be significant innovation in how we approach education to keep pace with the evolution in our world. Holding tightly to a prior model seems absurd to me and seems like a real disservice to students.

I remember a conversation with someone where they were like, "Well, if you broadcast lectures on Zoom, then people will quickly realize that you could just go to YouTube and watch a lecture instead of coming to class and sitting in it." And I thought that begs the question about what you're doing with your style of teaching. The issue shouldn't be you got to come to class because if we put it on YouTube then our business model is significantly undermined. If we just commoditize education and everyone can access anything, then all of a sudden you wouldn't pay the tuition and that's a problem. That's not the right framing or rationale. I think the provocative thing that the pandemic asked of us, and I think

generative AI will ask of us as well, is how do we take light of the world that we're in and figure out how to provide a rich educational experience? And to me that means engaging with the technology and using it, not trying to get students to put it in a box and pretend it doesn't exist.

Michelle Deutchman:

Well, and I think it's so interesting because one of the reasons I like the work that the center does is because we're looking at higher education, but the idea being that higher education on campuses are just microcosms for society, right?

Matt Perault:

Yes.

Michelle Deutchman:

And already we're talking about AI and education, but we're also talking about all of these things, the privatization of education, the commodification of the classroom, that AI is just another window that we're looking at through, that how are we framing it, and how do we need, like you said, to have an evolution of education. I mean, I'd like to look back and think that there wasn't the internet, there wasn't search engines, at some point there weren't calculators and that education has been able to evolve and that is this just another step in that evolution or is it really a five alarm fire or is it just another step in the process? Or is it both? Maybe it's not either/or.

Matt Perault:

I think that's right, I think it's both. I mean, I guess my view is I would put much less weight on how harrowing is it? Is it going to reduce enrollment? Are we going to lose tuition? Those just feel to me to be the wrong questions to ask. I think the question is how rich, by which I mean intellectually and culturally, can this phenomenon be for the educational experience and for cultivating people to be curious and engaged and to live gratifying lives? And I think we spend a lot of time saying, well, let's categorize on a scale of one to 10 how disruptive this is and this is how we currently get the money that allows us to keep the doors open. And what if this undermines that? That is an approach that I think is not desirable, not sustainable, I think is a disservice to students. And my hope would be that if universities take that approach, that there's competition that kind of comes in and makes it harder for them to take that approach in the long run.

Michelle Deutchman:

No, I agree. And I think in some ways, unfortunately bad news sells. And so I think that might be part of why there has been, I feel much significant more focus on the bad stuff that's going to happen as opposed to potentially the good stuff.

Matt Perault:

It really is weird, isn't it? Just the number of things that I've read that kind of try to wrestle with the value proposition. The positive value proposition are so few and far between. And I didn't have this experience in school because I... I mean again, on the things I do relatively well versus relatively poorly, long list of relatively poorly. But on the list on the relatively well side, taking something in my brain and turning it into a written product was always something that was relatively easy for me. That is not easy for everyone. That is a thing that lots of people struggle with and there is a sort of skewed advantage for

people in the world who are able to do that thing well. And that is different from having good ideas. It's really a subset of it. It's having an idea that you can translate into a written product.

And that strikes me as just such a profound shift in the planet where there's much more of a meritocracy around idea generation that doesn't just privilege people who have certain writing and speaking capability. And that to me is just profoundly revolutionary, not revolutionary in the sense of, I don't think the way that it felt when I started at Facebook and there were people marching into our rear square that were pointing to Twitter and Facebook as the things that were generating pro-democracy movements in the Middle East. But it is very revolutionary in terms of leveling a playing field for idea expression that feels just unbelievably important to me. And there really aren't that many... there are people who are sort of talking about, but what if AI makes something up? And that feels like a sort of inevitable thing that just seems so much less consequential to me. I think. I mean, I shouldn't belittle it, but so much less consequential than this fundamental shift in the ability to express oneself.

Michelle Deutchman:

Well, let's start with this kind of fundamental shift because I think it's, again, an interesting take on something that I've read about, which is how will generative AI impact equity issues on campuses? And again, what I've mostly read is, is AI going to merely augment preexisting socioeconomic inequities or not? And again, I'd like to hear your response to that question writ large, but the example you just gave is actually talking about a way that it could level the playing field, an inequity that maybe people aren't always thinking of when they think about who has access to technology, computers, stuff like that.

Matt Perault:

Yeah, so I'm not an expert on this and I don't think I'm a particularly thoughtful voice on it, but a friend of mine wrote a piece a while ago, there was a law review article and he also published it in the Lawfare blog called Robophobia. And it was about how people have sort of disproportionate fears when it comes to robots. I think one of the examples in the piece was if you have human delivery of medical information and you have robotic AI delivery of medical information, that people prefer the human delivery even if the information is shown to be less effective or impactful. So essentially they have a preference for worse outcomes if it's delivered by a human over a robot. And his point was we're sort of disproportionately and unstrategically... we have a preference for human stuff over robot stuff.

And one of the other examples I think he gave in the piece was related to the issues that you're describing, which is I think there's understandable and appropriate attention on AI bias, but it's not an issue where humans have a long track record of excelling. So I guess my thought about it is we should be conscious of bias, but we should look at it as an opportunity, I think, to improve upon this aspect of humanity that has a relatively poor track record to date.

And I don't know how optimistic we should be about that. I mean, maybe there are reasons to be very pessimistic, but when you look at... There's been a lot written about the use of AI in bail determinations, I think, that might be an issue that you know... I don't know if you've tracked and maybe can correct me if I say anything that's off base here, but my understanding is the AI has been biased in bail determinations. But again, humans are not good at this and it seems to me like maybe there's some potential to turn over some part of a process like that to an entity that might have an opportunity to exhibit less bias than humans do in some situations. And again, that's a thing that I would be excited about rather than fearing.

Michelle Deutchman:

I definitely think I need to read that piece by her friend. I might be one of those people that prefers human contact. And I think it presses me to think about why that is. And I also, again, I think in today's world, optimism is in such short supply, so I'm sort of happy to have a guest who's sort of acknowledging that there are likely going to be inaccuracies and problems and bias, but that you're able to flip it and say, but again, you're seeing it through the lens of opportunity and benefit, which I just, again, I think is in sort of short supply.

And I kind of wanted to ask you, I'm just picking, there's been a lot of open letters by people and researchers who've helped develop this and other technology. There was one recent open letter, I think it's like 23 words and it said, "Mitigating the risk of extinction from AI should be a global priority alongside other societal scale risks such as pandemics and nuclear war." And I think I want to ask you again, do you think that message is sort of accurate/necessary, and/or is it sort of hysteria? And I think it gets more credibility, and again, I'm not an expert on this, because it comes from people who are in the field. And I'm just curious if you have thoughts on that?

Matt Perault:

Yeah, I would say I don't have strong thoughts on the merit of the message, because it seems like there is some existential risk and we can say is that a 1% chance or a 50% chance or an 80% chance or a 100% chance? But even if it's a 1% chance, I mean the end of humanity is a significant thing. We should probably be concerned if there's a 1% chance of the end of humanity. But I'm less focused on that and more focused on what does that mean in terms of what we do?

And when people say, well, let's slow down development of the product, that to me feels like a pretty concerning remedy because slow down just punishes good actors who follow the rules and lets other people continue developing the technology in ways that would be problematic, I think. And it doesn't seem to me we're going to know anything in six months or a year or two years, I don't think, that's going to mean now we can bless the technology as safe. So unfortunately I think we're in a position of we need to continue developing the technology and then ideally learn as we go. So to the extent that perspective results in let's cease or abstinence is the best approach, those things I think are deeply problematic. Or if it provides more momentum to bans and that sort of thing, I don't think that's the right approach.

I also think it's problematic in terms of thinking about governance mechanisms because I think it puts us in a position where we start moving up the spectrum in terms of the level of stringency of governance. And some people have suggested something modeled on atomic regulation for instance, as the right model for thinking about AI. And it feels to me like that is a very strong view that maybe it's right, but feels likely to me to overstate the risks, that it would be dropping a bomb that would wipe out a city or kill hundreds of thousands of millions of people. It doesn't strike me that it is likely at that same risk level. And nuclear energy is not known for being a field of immense innovation and sort of dynamism. It is tightly regulated because it's so dangerous. And part of the result of that is it really hinders the pace of innovation.

And I'm concerned that the more stringent the mechanisms are, the more you cut off competition, makes it harder for startups to compete, the more you harm innovation, I think in some ways you sort of freeze the technology and you learn less about it. So I feel like a lot of what we've talked about so far is kind of like there's a lot that we don't know, we should learn as we go and ideally figure out ways to mitigate harms and maximize positive use cases. I'm not trying to minimize the likelihood that there'll be harm. It's just I think you want to learn as you go and develop better and better mechanisms for governing the technology as you know things. And right now I don't think we sort of know enough to say we should regulate this like nuclear.

Michelle Deutchman:

Well, I think this is, I don't know if you're anticipating where we're going to go next, which is about regulation. And I'm wondering, of course there's this idea, should we use 230, something like 230? But this is where I would love to play on your policy expertise. If you could take us through some of the alternative proposals that are out there and maybe even opine on what you think has the most likelihood. I don't want to say success because that is not the right word.

Matt Perault:

Yeah, I don't know. I mean, it seems right now that everyone knows this is a thing that's important and everybody knows this is a thing where something should be done. And I haven't seen a lot of evidence of this is the thing, this is the right way forward. I'm obsessed probably to a fault with this idea of regulatory experimentation, that typically we regulate from a position of certainty, like we know how it will be and thus we need nuclear style regulation and we will enact it tomorrow and then it will be on the books in perpetuity. That strikes me as the wrong model for tech for a lot of reasons.

And actually one thing that I think is a little counterintuitive about it is that it actually makes it hard to initiate lawmaking because people are fearful that if you put a bad law on the books, it'll just be there forever. And so it actually creates a kind of higher bar for getting something done. And the result is, as we've seen in the federal government and tech policy, essentially nothing happens. Really nothing. And have my opinions about how various different proposals might play out in practice, but lots of those opinions are probably likely to be proven wrong when we see how something actually plays out in practice. But we don't pass laws so we can't actually see that.

Generative AI I think is a perfect area for experimentation, which people are very hesitant to experiment in tech policy, which I don't understand. We have COVID vaccines because we experimented with vaccines on humans on a life or death issue. Clinical trials in medicine are explicitly about experimentation to try to develop things that are safer and better and improve humanity. That certainly strikes me as something that's more important than artificial intelligence or other areas of tech policy. So I think an experimental model where we kind of try out two different things, one is different types of products. So there's experimentation on the product side. The second is experimentation on the policy side. What are different types of governance regimes that might be desirable? That feels to me to be a really helpful way to move forward so that if we do something that has adverse results, we will learn about that and correct it the same way if you do a medical trial and you learn about side effects of a certain treatment, then you try to alter the treatment so as to avoid the side effects.

So that's the kind of approach that I like. That approach, I guess as a good academic, my ideas have no traction in the real world. I would not say that there are lots of policy proposals that sort of mimic that approach. You could make the argument, I think, that there's been much more of it that's happened at the state level with different states passing different things in tech policy reform. California's been a leader in that in many ways. And so that gives us a little bit of data that we can compare across states to see how different things work in practice. The federal government has done very little experimentation and my guess in generative AI is that we won't see it.

Michelle Deutchman:

I am fascinated by the analogy to clinical trials in medicine, because when you mention that, that seems like so natural even though there's so many unnatural things about it, which is experimenting on people who are alive. But it makes sense and it does make me kind of see through a different vantage point of



why aren't we doing experimentation? The focus is like what are we going to do on the federal level? What's Congress going to do? It's going to apply everywhere kind of full stop, stat.

Matt Perault:

Yeah, right? It just strikes me as so bizarre. And there are laws that, I mean one is reforms to Section 230, which is, maybe all of your listeners know this, but it's basically a law that provides immunity to platforms that host content. And so if you wanted to file a defamation case for content posted on Twitter about you, you could file it against the poster, but Twitter would likely not be liable. And that has enabled platforms to host content because they can host content without fear of getting hauled into court every time someone posts something that might be deemed to be illegal. You can go after the individual just not the platform.

And so there's a law passed a few years ago called SESTA/FOSTA that was aimed at sex trafficking and protecting sex workers. And almost immediately after it was passed, the sex worker industry started campaigning for the law's repeal. And I think in a world of experimentation, we would not fault the law's drafters for that. We would say that seems like a noble cause, that's a noble objective. And you can try many different ways to try to achieve that objective. And this particular one seems like it has not worked. And in a world of experimentation, just like with clinical trials, you would be generating evidence that hopefully would see the light of day, so you could actually evaluate in a more nuanced way how it actually has played out in practice.

Instead, we are in this bizarre universe where that law's passed. It seems like the community it was designed to protect dislikes the law and would like it removed. And yet politically it's very hard to repeal a law that's aimed at protecting against sex trafficking, that is challenging politically to take a vote on that issue. And so it'll just be on the books in perpetuity even though it makes sex workers' jobs less safe and secure.

I think of that's less about a problem in the substance. I mean, you can say that the drafters of the law should have anticipated that result, there are a lot of people who are advocating on it who anticipated that result. But I think it's more of a problem in just the way that our policy world is set up, that we have this particular way of doing things that makes it very hard to do anything. I mean this law is one of the few laws that actually has reformed Section 230. There have been dozens of different proposals introduced this law actually went through. And then it's not a particularly good one I think and it's just on the books forever.

Michelle Deutchman:

I mean it's the law of unintended consequences, that you would like to think that if there are unintended consequences, then we can make some changes to have the consequences that maybe we intended. And what you're saying that for many reasons, including that everything right now politically is challenging, that there isn't room to do that?

Matt Perault:

Yeah, I think it's a lot of weight to put on a lawmaker on day one. That is a big weight and I don't envy them. It is really hard mean. Our center's about to release a report on age verification, age assurances, so a lot of activity right now in online child safety. And if you're focused on protecting children online, which again, is a very noble cause, then you have to know who a child is. And it turns out to know who a child is, you also have to know who is not a child. And so if you are really taking age verification seriously, meaning you want to accurately know every person on the internet who is or is not a child,

then that has real privacy consequences. I mean that requires a massive amount of data gathering and data knowledge by companies that many people want to know less about them.

So in the course of this report, we're very focused on trying to develop recommendations that policymakers can use, but it is really hard in age verification to have good ones. It is just unbelievably hard because every direction that you want to go in has big costs. And so it's more about figuring out what is the path that minimizes those costs or has costs that you're prepared to bear. And that again, just seems to me that it's just ripe for experimentation to really learn more about how this plays out in practice. What are the privacy consequences? How do kids experience this? How do parents experience this? But we don't have a framework that enables us to do that experimentation and so-

Michelle Deutchman:

I can tell you as the parent of a 13 and 11 year old who are, we're starting to already work on some of these things, like, I want to read that report. And also just so you know, this kind of conversation for me is why I want to go get my MPP, because I think the policy implications of things are so fascinating. But I don't want us to get too in the weeds. And I do want to pull back a little bit, not just to sort of the regulation piece, but to the kind of speech aspect of it. And not even so much does AI have speech protections, but more this question of when I think back to social media and the ways that it really changed the landscape of expression both in the US and globally. And again asking you just to spitball on this, it's a big question, but do you anticipate that generative AI may have similar impactful results on speech and expression? And you can kind of take that where you want.

Matt Perault:

Yeah. I mean, like I said earlier in the conversation, I think it has a massive one in terms of being able to literally translate ideas to a written product. And that just feels to me it has massive expression possibilities. It is, I think, helpful to keep in mind that in the two main implementations of the product to date, either an open text field where you get a result or a search engine, which I guess is an open text field where you get a result, neither of those really are distribution platforms. So the kinds of things that many people have been concerned about when they express concerns about social media aren't present right now for generative AI in that people wouldn't care a lot less, I think, about content on Facebook or content on Twitter, or content on YouTube, if there was no distribution, if it was just you and a computer. Where people are concerned is that there is information that gets distributed broadly.

Certainly we could be headed in that world. You could imagine generative AI helping you to compose tweets, which I would like because I struggle in tweet composition. Or helping to compose a Facebook post or an Instagram post or something. Certainly that world could be coming, but it's not where we are right now. So I think that does sort of shift the dynamics around it.

And I wrote a piece in Lawfare in, I think it was February, about whether ChatGPT and other general AI tools would get Section 230 protection. So would they be considered, the technical legal term as an interactive computer service or an information content provider? Interactive computer service is a host information content provider is the creator. Hosts are protected, creators are not. And my view in the piece was that in most cases they probably are not going to be protected. They probably are not going to be considered hosts because they do actually generate the content. So the language in the statute is responsible in whole or in part for developing the content. So if you develop the content in whole or in part, then you would be considered an information content provider. It strikes me as difficult to argue that a generative AI tool does not at least develop the content in part.

Since I wrote that piece, I think more thoughtful and better legal minds have weighed in on that question. So Eugene Volokh has written about it in the defamation context. Derek Bambauer has written about it. Jess Meyers has written about it. And I think all of their arguments are incredibly nuanced and thoughtful. So I would suggest to your listeners, if they're interested in this issue, check out what they've said about it.

One sort of flippancy summary of it is that I think I didn't give enough deference to, attention to in the Lawfare piece is that it really is product specific. So the way a generative AI tool surfaces information will have an impact on whether it gets 230 protection. So that is my view, that it won't get 230 protection in lots of cases. And that will result in companies that use that technology significantly limiting use cases because they will fear legal liability. And I think that's a net negative. So even if judges side with generative AI platforms in some cases, I think there will be enough cases where they don't, where there will be legal liability that will stand in the way of innovation.

And it also seems to me that if the legal guidance is you're totally fine, as long as you don't do anything that's really generative, then what we're essentially doing is taking the most interesting use cases off the table. So I think, I mean, this is in the land of it's nice in theory, but will never happen. It would be useful I think to have an intermediary liability regime that does protect generative AI tools, at least somewhat. Maybe you make it conditional on certain things, maybe it's time limited. There are lots of different sort of asterisks that you can place next to the immunity to try to address potential concerns. But I think without it, we're headed to a universe where, in many cases, general AI tools won't get 230 protection and that will significantly impede innovation.

Michelle Deutchman:

Well, it sort of reminds me of what I say about free speech cases, which is the devil's kind of in the details. I mean, I think people want an answer like, oh, is it protected or is it not protected? And it's like, well hold on. Sometimes I can just say it is or it isn't, but sometimes it depends on the context and where was it said and who said it.

Matt Perault:

Yes.

Michelle Deutchman:

And how was it said? And that to me seems potentially analogous to kind of the products. Of course part of me is like I don't want to be talking about the presidential election because I feel like the election cycles start way too early, but I'm going to do it anyway and ask you a little bit about this because people are already starting to say things like, "This is going to be the AI election," and this idea of generative AI will be part of increasing misinformation, whether it's misinformation of audio or graphics. And I'm just curious what you think about that idea and whether generative AI is really going to be making misinformation and disinformation so much worse? Or again, that's just sort of the negative, one of those phobias, I think it's robophobia you said?

Matt Perault:

Yeah.

Michelle Deutchman:

Did I say it right? Is it robophobia or is it something that's real?

Matt Perault:

Yeah. Well I think it's real. I mean, every election there's a lot of incentive, financial and otherwise, for people who are active in an election to use whatever tools they possibly can to try to advance their interests. And the incentive is very, very strong to do that. Whether you're a candidate or you're like a communications agency trying to support candidates, if you are a political advertising service trying to support candidates, or if you're a foreign government attempting to disrupt the election or a hacker trying to disrupt the election, you would use whatever tools are at disposal. So I would be surprised if AI wasn't actively used in the election. I guess the question is, well, what are the use cases that are particularly harmful and what does that tell us about how we might govern the technology? And my hope would be that are very open-eyed to the harms and that we develop thoughtful policy tools for cabining those harms.

But I think that means, again, not blocking the technology. I mean, one thing that the slowdown movement does is, I think you asked about this earlier, to the extent the project is developing literacy, we need to have exposure. We cannot be in an abstinence mode because that just stands in the way of learning. And so we need to learn and there are downsides to being in learning mode. I mean, I think that there are things that will happen that will be scary and horrible. And my hope is that we are clear-eyed about that and can incorporate those learnings into whatever future governance regime we have.

One dynamic that I find to be very disingenuous and frustrating is political candidates saying one thing and doing something very different on tech policy issues. So there are a lot of candidates, for instance, who talk extensively about the importance of privacy and then purchase voter lists, and then share that data across a whole bunch of different jurisdictions with a whole bunch of different people. And at least for me, I tend to experience direct consequences of that. So whenever I'm asked to give money to a political campaign, one of the reasons I'm hesitant is that giving money to a representative running for office in North Carolina does not mean that in two years I'm going to want to get spam text messages from a candidate running in Wisconsin. And I think for anyone who has ever given any amount of money in a political race, they know that that sort of thing happens.

And I find that to be deeply frustrating, both because I don't like the spam from the random candidate in some other place, but also because if you are doing that with data, then it suggests that you don't think the right regime is certain types of data privacy regimes that you might be advocating for. And that disconnect I just find to be incredibly frustrating. We need a regime that works for a variety of use cases and it should work for political candidates and it should work for consumers. The idea that you're going to advocate for one thing, in theory, on behalf of consumers and then treat consumer data very differently when they happen to donate money to you or sign up for a list for your campaign frustrates me because it suggests a sort of disingenuousness in the governance that I find frustrating. And I concede the same thing emerging in AI where people actively use AI in their campaigns and then give speeches where they talk about how terrible AI is. And I just hope that there would be mean, we know how this is going to play out.

Michelle Deutchman:

I mean, hypocrisy in campaigns and politics unfortunately, but-

Matt Perault:

I mean, I think you're thinking about this from a student perspective, what can people learn? And I think it is confusing and disingenuous to have people hear how terrible the technology is and then see people actively using it and try to figure out what does that mean for how I go about things? When I have

taught, I have really tried to put students in the position of being decision makers, sometimes in places that they would never in theory want to be or say they don't want to be. You're sitting on a public policy team at Facebook or you're a Republican and you're a staffer Elizabeth Warren, or you're a Democrat and you're a staffer for Ted Cruz. And part of the reason to put students in that position is like you have to be in a world where you make those decisions about trade-offs in a true way.

And I hope that we don't teach students to be the type of person who preps a speech on how awful AI is and then goes back to the office afterward and tries to use AI to generate all of the senator's future speeches or future targeting lists or whatever it is. And I anticipate that we will see that in this election and that will be frustrating.

Michelle Deutchman:

Well, I mean, I don't in any way want to simplify what you're saying, but I think to me it speaks to role modeling and it's kind of the same in issue that we see in deliberative dialogue or dialogue across difference. You can talk about how everyone should do that, but then when you turn around and are using inflammatory ugly rhetoric in the next breath, I think that's very confusing and it doesn't send a message of... it's like, do what I say, but not what I do.

There's been so much that we've talked about the big landscape of how much more there is to learn and that in order to regulate appropriately, we have to learn more. And one of the ways that we always like to close the episode is asking guests if there's a tangible action that listeners maybe can take? I mean, it seems like this is a moment, no matter who you are, whether you're a student or an administrator or faculty member or policymaker, we're on kind of a precipice of how this technology is going to be used. And I guess my question for you is, is there something, even something really small that listeners can be thinking about or preparing for as AI becomes more commonplace in society?

Matt Perault:

So I think that the number one thing that I would encourage students to do, and this goes back to the early part of our conversation, is to use the technology. I think it's a really exciting moment to become a sophisticated user. And I felt this pretty strongly, again, going back to how things played out in the pandemic when lots of people were sort of saying, "We need to be in the classroom. Or if you're on Zoom, you need to conduct yourself in x, y, Z way." And I felt like that was really a missed opportunity because if you think that you are going to go to school and when you graduate, you're never going to have to absorb information via a screen or deliver information via a screen and be effective, you're not well positioned for the workforce.

And the pandemic gave us this moment to really develop a skillset there that I think would set students up well for the future in a way that without it, I think if everything's just in person, you're not developing that skill. And that's a really valuable skill and isn't it better to be developing it while you're in school and you can experiment and fail and learn than doing it in a workplace where you might get fired or not promoted or feel very alienated from your job? And I think we're at a similar sort of moment where it's important for students to actively understand how to use the technology. And I think that probably means more in a daily basis kind of way than a once every six months kind of way. And that doesn't mean the technology is perfect, that doesn't mean it's low risk, that doesn't mean there aren't major downsides. But I think seeing those things and experiencing them intimately and becoming sophisticated users is really important.

Michelle Deutchman:

I think that's great. And I think it applies to any constituent in higher education. And I'm already thinking that maybe at our next center team meeting we should be starting to dedicate time regularly to experimenting and using it. You have been so generous with your time and your expertise and in my mind you should be running AI policy, so we're really grateful to you. I want to give you a chance to add if there's anything else that you don't think we were able to touch on. And then of course, reserve the right to have you back maybe a year from now to see where we are.

Matt Perault:

Yeah, that would be wonderful. I mean, to profess an interest in experimentation, I think that means you have to have a commitment to being wrong and... or you have to have a commitment to examining when you are wrong and looking carefully at how things have played out. And my expectation is that a lot of the stuff we've discussed today, maybe we had predictions that turn out to be off base. And my hope would be those would be interesting and fruitful areas to understand and learn from.

Michelle Deutchman:

And I think that you are role modeling intellectual humility, which is something that we talk about so much in the speech and dialogue space, and I don't think is often role modeled enough in academia, which is that part of learning is failing and getting it wrong, and then you just do it again. And so kind, I think we'll end on that very optimistic note of which is just to fail is to live and then to do better. And again, it was great to have you.

Matt Perault:

Thanks so much.

Michelle Deutchman:

I don't want to head into summertime without offering congratulations to all of the recent graduates who are heading into the world beyond college. Kudos on this significant achievement. With the Supreme Court term wrapping up in 10 days and most state legislatures out of session, join us next month for a conversation about the impact of court rulings and newly minted laws. Talk to you then.