Podcast Transcript: CDT Tech Talk – AI and Free Expression

Host: Jamal Magby

Guests: John Coleman (Legislative Counsel on AI and Free Expression for FIRE) and Becca Branum (Deputy Director of the Free Expression Project at CDT)

Audio Source: CDT Tech Talk - Grok AI / AI & Free Speech

Introduction

Jamal Magby: Welcome to CDT's Tech Talks, where we dish on tech and internet policy while also explaining what these policies mean to our daily lives. I'm Jamal Magby, and it's time to talk tech.

Today's conversation sits right at the heart of one of the most pressing debates of our time: how to govern artificial intelligence while protecting free speech rights. As AI systems become more powerful and more embedded in our daily lives, from content moderation to generative tools, they're raising complex questions about expression, accountability, and the role of both government and private platforms. Where should we draw the line between regulation and censorship? And how do we ensure innovation doesn't come at the cost of fundamental rights?

To dig into these questions, we're joined by two leading voices with thoughtful perspectives. John Coleman, Legislative Counsel on AI and Free Expression for FIRE, and Becca Branum, Deputy Director of the Free Expression Project at CDT. John and Becca, thank you so much for being here.

Becca Branum: Thanks, Jamal. Glad to be here.

John Coleman: Thank you for having us.

Jamal Magby: Of course. So, to kick us off, when each of you think about AI and free speech, what feels most urgent right now, and where do you see the biggest risks or misunderstandings in the public debate?

Public Misunderstandings of AI Regulation

Becca Branum: Maybe I'll start with what I think are some of the biggest misunderstandings when we're talking about AI, and I think that misunderstanding itself is pretty urgent. I think one of the biggest misunderstandings, in my view, is that when we're talking about AI regulation,

we're often not really thinking about what's really at stake when we're thinking about regulating chatbots and their outputs.

Just given litigation incentives and the like when we're thinking about where we assign liability, we're often thinking in terms of the developer's or the deployer's liability—the company—and who at the end of the day is going to have to sign the check to cover the bill for whoever is harmed.

What's more interesting and more urgent, in my view, is what that regulation and allocation of liability means for the rest of us—the people who actually interact with AI, whether willingly or not. Restrictions on the kind of information that chatbots can provide aren't really just regulatory requirements; they're restrictions on what we, as people, can read, can learn about, and can create with technology.

Too often, the regulatory debate, in my view, collapses into one that assumes people don't have choices or can't be trusted to make them. That assumption—that users lack autonomy—ironically hands over a lot of power, and a lot more power than is safe or warranted, both to AI companies and potentially to the government.

First Amendment Rights of AI Users

John Coleman: Becca, you make some really great points. I would just follow up by saying that AI is an expressive tool, and the people who use it retain their First Amendment rights. It's really important that we get the conversation right about this interaction between this technology and people, and that very much fits into the narrative about how we regulate it. We need to take into consideration people's rights, and that includes the right to receive and share information, the right to speak anonymously, and the right to be free from compelled speech. If we bulldoze those rights, we're giving the government free rein over this technology that impacts us every day.

Drawing the Line: Speech vs. Conduct

Jamal Magby: So, I want to pull that thread a little bit more because there's growing momentum for AI regulation at both the state and federal levels. John, I want to kick this to you first: where do you see the line between legitimate regulation and policies that could chill lawful speech?

John Coleman: Yeah, well, let's look at the environment right now. We have over 1,500 state bills that have been introduced this year alone. We have the Senate moving federal legislation as we speak. On top of that, the White House is doing its own consideration of AI policy in a lot of different ways. There's a lot going on, and there's a lot of issues at stake—from data centers to content-based restrictions.

FIRE is particularly interested where the First Amendment applies. That doesn't mean all regulation out there. If we focus on particular areas like access restrictions, content-based restrictions, and things like age verification, we have a difficulty drawing a line between what's

good and what's not. Those types of regulations are constitutionally suspect, and they're going to have an uphill battle in court if they're introduced and enacted.

Becca Branum: I agree. I think what will become increasingly important is drawing a line between what is speech and what is conduct—when folks are using AI for the purposes of communication and when they're using it, for example, through an AI agent to go out into the world and buy them plane tickets or reserve them a table at a restaurant for Friday night. Right, that's a really hard line to draw, but we have seen the Supreme Court increasingly make clear that you can't dress up speech regulations as conduct and get away with or get around the First Amendment. So, it's a line that's going to be extraordinarily important to draw.

But I think it's important to do so because, too often, I hear from my friends in the AI safety space concern that if we think about the First Amendment, then that means that all regulation is off the table. That's just simply not true. The First Amendment is important, and I spend most of my day working to defend it in the tech space, but it's not absolute.

There are really important spaces where we can focus, whether it's civil rights enforcement against discriminatory uses of AI, other types of conduct rather than speech that's regulated, ensuring that there's human oversight, and redress for people who are harmed by automated decision-making. There's a lot of ground that can be covered. That doesn't mean that the First Amendment isn't implicated at all, but it's a lot different than a regulation that says what you can read, what you can share with your friends, or what you're allowed to think about in the privacy of your own home. Drawing those lines appropriately, I think, will allow us in the tech policy space to both have our cake and eat it too, where we can have access to tools that allow us to create and learn, while also not ignoring the very real harms that can come from this technology.

Historical Parallels and Legal Precedent

John Coleman: Thankfully, this is not the first time courts have had to deal with a new technology. In fact, it's not the first time they've had to deal with a communicative technology. We even have recent decisions that have shown how the courts handle social media regulation, and we can borrow from that precedent to understand how courts are going to apply the First Amendment to AI. It's very informative when it does come to issues like output restrictions, or even how developers design their AI models. So, if you want to go back even further, you can look at how the court has dealt with video games, or even newspapers. It's very interesting to see a lot of the parallels that can be drawn from these past decisions and apply them here, and I think we're going to see courts willing to do that.

Becca Branum: And if I can nerd out for just a second on the video games point, I think that's a particularly good analogy to draw into the AI context. There was a Supreme Court case that considered restrictions on kids' access to video games, and it was either the main opinion or a concurrence by Scalia that honed in on the interactive nature of that medium—of the video game medium—and what that meant for First Amendment protection. That interactivity was actually quite important and was every bit as protected as a more passive type of communication.

AI chatbots, in particular, are very much of that same kind of technology and have that kind of interactivity that I think is both important legally, but also really worth protecting. When you sit down to a chatbot and are requesting information, you are imparting something to the chatbot, and then the chatbot's programming is giving something back to you. But I think it's really important for First Amendment purposes to keep in mind that that interactivity matters for doctrinal purposes, but also for what it is that we're trying to protect, which is people's autonomy to engage with the information that they care about. I think the video game line of cases will actually end up being pretty important in the AI space.

Separation of Design and Output

John Coleman: Interactivity is definitely a good point to raise here. Becca, you wrote an article recently that kind of got at a different misconception, going back to the original question that you asked, Jamal, about how we should view AI's separation between design and its ultimate output.

I've heard criticisms that because there's a gap between the designer's intent and then the ultimate output that corresponds to a user's prompt, that there's no expression at all there. But Becca, you raised a lot of great examples in your article, the first one being *Art by Telephone*, which was the title of your piece. I thought that article—just to nerd out a little bit with you—was so great, and it's something that I reference all the time talking to people: that we can draw from the art world and see that even when there's a disconnect between the artist and the ultimate work, there's still expression happening.

Becca Branum: Yeah, thanks for raising that. So, a bit of background for folks who are tuning into the podcast: there was an exhibit held at a museum in Chicago called *Art by Telephone*, and it was a really interesting, sort of modern art exhibition where artists called curators and fabricators at the museum and gave directions for the pieces that they wanted made, and then the fabricators went forth and created them—whether it was music that they played or sculptures that were made.

The artists obviously gave direction, but the fabricators and the people making the art don't have complete knowledge of what the artist wanted, and so there was, I guess, a bit of intermediary interpretation that happened between the giving of instructions and the art that was actually created. It was a really interesting piece and exhibition because it demonstrated that creativity happened all throughout that process—both in the giving of instructions by the artists themselves, the creation of the art, and the translation of those instructions into the ultimate piece.

When we think about chatbots in particular, you don't have to assign rights to the chatbot itself to recognize that when you give a prompt—asking it to write a... what I do now, I had a baby a few months ago, I ask it to write bedtime stories for me, right—the chatbot itself is not someone that needs to have rights, and it doesn't have expressive intent itself, but it is a facilitator of my creativity and the creativity I'm asking it to produce on my behalf.

To your point, John, art and the art world and that type of direction, I think, is a good analogy for us to think about in this space, and also, I think it's a way to make some of these theoretical tech policy and speech conversations a little more concrete. Sometimes we can forget when it's just typeface on a screen that we're talking about human communication at the end of the day, and sometimes art and physical art can make that a little more real for people.

Corporate Responsibility and Algorithmic Moderation

Jamal Magby: I love this, and I hate the fact that I have to shift gears, but I want to make sure we keep pushing forward. So, Becca, I want to ask this question to you first: because much of the AI ecosystem is driven by private companies, how should we think about their responsibilities when it comes to moderating AI-generated content or setting guardrails on these tools when we know these tools can sometimes, what some would say, cause harm?

Becca Branum: Right. I think it's really the space that has the most hope for this technology and for keeping people safe is in the design of these products, rather than a top-down approach from the government telling companies how they have to behave or how they have to design their products. And so, there's tons of room here, and we have an entire AI Governance Lab that is hyper-focused on the types of design and safety mitigations that are really important to ensure that people can use this technology safely.

I think we can actually look to the social media and the broader internet ecosystem for how this could potentially work. We've never assumed, or we've generally been skeptical of the idea, that there is a single way for people to talk to each other on the internet, or a single way for software to be designed, but rather we have trusted that market incentives and sort of the diversity of thought on the internet will allow people to pursue what is best, what they are looking for, and what they are hoping to get out of technology. So, through that diversity of approach and design, people can find what it is they need.

There are certain corners of the internet that I don't care to go to, but I'm glad in some ways that they exist because if they're over in that corner of the internet, they're not in the corner of the internet that I occupy. Right, and so that diversity can actually end up keeping people safe in certain ways.

That being said, there are very clear harms that I think all responsible AI developers should be pursuing, whether it's related to non-consensual intimate imagery or other clear harms, and other clear guardrails that not even just a legal matter, but as an ethical matter, these companies need to be considering to ensure that people are not victimized predictably through the technology that they're putting out into the world.

Diversity in AI Model Design

John Coleman: Yeah, and we've kind of seen that reaction from companies that have had that problem in the past. I think of Grok, for example, getting into hot water publicly, maybe even

legal trouble—I don't know at this point—but I know that they've since fixed it. But they had a problem in the past of their system generating CSAM and that creating problems.

But Becca's right, we can look at the broader universe of content out there and hope for diversity when it comes to all the various types of AI that could be created. Already, we have examples like Anthropic, for example, relying on a constitution that structures how its system responds to prompts. On Grok's website, they say their model was based on *The Hitchhiker's Guide to the Galaxy*. I don't know how that approach necessarily differs from a constitution, but I'm sure there is one—one is more based in fiction, but that could produce interesting results, and one on a constitution that's formalized and structural also sounds great too. What we need to make sure is that government regulation doesn't narrow this field to give us only a few AI systems to choose from. The more we have, the more that we can engage in truth-seeking that is absolutely important for our self-governing system.

Addressing Bias and Viewpoint Discrimination

Jamal Magby: One thing we haven't touched on that I'd love to hear your thoughts on are biases in AI systems, because they often come up in these conversations. John, I'd like to pitch this to you first: how should policymakers and companies address claims of viewpoint discrimination without overcorrecting in ways that limit speech?

John Coleman: Viewpoint discrimination can be defined in a very simple way: when the government favors one viewpoint over another in its policy, it commits what's known as viewpoint discrimination, and it is presumptively unconstitutional under the First Amendment. Courts are very wary of this because of the fact that you're favoring one side over the other. So first off, when we're talking about policymakers, they should be very careful that they're not engaging in viewpoint discrimination.

That came up recently in xAI's challenge to the Colorado law, where they alleged that Colorado engaged in viewpoint discrimination by favoring ideas like diversity and historical discrimination—redressing historical discrimination—over the other types of bias that could come up in AI use. The problem there, zooming out, is the fact that these risks of bias attach so upstream that it would affect how these systems were created.

Just as we talked about in regard to the previous question, we want to give this ecosystem as much of a chance to survive and thrive as we can. So, we need to be very careful about regulating bias where there's no harm attached yet. There is a place for anti-discrimination laws as AI is deployed across various sectors, but we need to be careful about taking it so far upstream that it impacts how these systems are designed.

That said, you also mentioned how companies can address it. That's going to be a real friction point to figure out how best to do that. I know companies right now are tackling that as their systems are deployed across the board, and we just need to be careful how they're deployed to ensure that we don't have bias when they're used for certain uses that would be illegal under existing law.

Becca Branum: Yeah, I think that's right, and I think it's important to keep in mind all the different ways that AI can be used, and that in certain circumstances, people might not be seeking out a lack of bias; they might actually be seeking out a particular viewpoint. And so, I want to disaggregate bias as it's applied in sort of civil rights law, which I think laws apply to AI and need to apply to AI, but when we're talking about bias in a broader sense—in a sort of ideological, neutrality, or lack of neutrality type of way—I want to challenge the idea a little bit that that's always desirable because it assumes that a person is interacting with AI and expecting the AI to sort of reflect truth back to them rather than them engaging with AI in a particular way.

We know people use AI for all sorts of things: for play-acting, for role-play, for just seeing what they can get the AI to say back to them because they think it's fun to see what the guardrails will let them do. People interact with AI for all kinds of different reasons. And so, I think I want us to be careful about separating viewpoint from bias, and bias in conduct versus developers and AI systems being able to have a point of view. Again, there are certain AI systems that I use more readily than others because, in fact, I prefer the viewpoint and sort of the constitution or the guardrails that they have in place as compared to some of their competitors. I think that in the end ends up being better for everyone involved because it allows them to access the kind of information that they're looking for.

At the same time, though, that makes it all the more important for companies to be clear about what it is they're offering to the public and to their consumers. We don't have the same types of transparency rubrics, like the Santa Clara Principles, that we have in social media as applied to AI yet, and that's a pretty dangerous thing for consumers because they don't always know what it is that they're signing up for or what it is that they are getting out of an AI system. And so, companies are starting to, but need to do a lot more to make clear what their viewpoints are and what it is that you're getting when you solicit information from these systems. I think that would also go a long way to both identifying bias and making sure people can make informed choices, while not flattening the types of lawful information that are available to people.

Transparency Mandates and Compelled Speech

Jamal Magby: So Becca, what kinds of transparency requirements, if any, make sense for AI systems that generate or curate speech? Are there risks that transparency mandates could themselves create new problems?

Becca Branum: So, I'm going to separate out sort of transparency that systems and companies can take on their own, and then I'll deal with the First Amendment elephant in the room as it applies to transparency requirements. I think all companies have an obligation to be transparent to the public about when a person is interacting with AI and what it is that they're actually offering. Trickies—or sort of tricking people, deceiving people into thinking they're interacting with real people, or that they're getting information that's reliable when it isn't in fact reliable, is not a good way to design a product; it's not a good way to run a business.

I'd separate that out, though, from what the government can force AI developers to disclose to the public. When you mandate transparency, and you mandate companies say things, that implicates the First Amendment. It's not an absolute bar; there are things you can require companies to be transparent about, but it's a lot narrower of a list of things that you can force companies to say about their products than you might otherwise think, because what we have seen is sometimes transparency mandates—particularly transparency around content decisions and sort of ethical and safety judgments that companies are making—can end up being a way of controlling content, but through a transparency mandate itself. And so, I wish we could say that transparency would solve all of our problems here. I think a lot of problems can be solved with transparency and transparency norms taken on by the industry, but again, when the government gets involved, the First Amendment has a lot to say about it, and we have to be really careful about what we're mandating, lest we set these regulations up to fail.

John Coleman: We can call it transparency when it comes to companies and compelled speech more often than not when it comes from the government and it's mandated. I agree that there is just a narrow circumstance where the courts have granted compelled speech a lower judicial standard, but for the most part, compelled speech requirements, like disclosures or even disclaimers that you have to affix to AI content, will likely get some sort of heightened judicial review, meaning these kinds of regulations are going to face an uphill battle in court.

And if we're talking about labels on speech itself, if you have to include a text on an image that's been generated by AI that says, "This has been generated by AI," or something similar, studies are starting to show that this might not be an effective way of helping people decide what was actually AI-generated in the image. If it was wholly generated, you might have a case for it, but in most cases, people are using AI for all sorts of things, whether it's touch-ups on photos or moving things in the background to make it a more aesthetically pleasing image. And so, if you just have a general label, one-size-fits-all, you're not telling the viewer much about what's changed in the photo or whether it's even deceptive, which I think is oftentimes the interest that states have to regulate here.

Another issue is that I think most people interacting with chatbots today know that they're interacting with a chatbot. You have to go to a website like ChatGPT or Character.AI, and even they have disclaimers up top. But even if they didn't, the user is going to these chatbot sites, they know that they're interacting with a chatbot, so I don't know that a compelled government mandate to include a disclosure is all that necessary when most users know that they're interacting with a chatbot.

Do Existing Frameworks Suffice?

Jamal Magby: So John, do existing First Amendment frameworks adequately address the challenges posed by AI, or are we heading into territory that requires new legal approaches?

John Coleman: Let me try to answer this with a story. Several weeks ago, South Korea reported that a wolf had escaped a zoo, and authorities were asking for information or sightings about the animal. Several tips came in, and I've heard two different stories here: either someone

submitted a deepfake of a wolf walking through the downtown area of the city, or someone had posted on social media a deepfake of this wolf and authorities saw it. Either way, authorities were thrown off the case; they were searching inside of the city when the wolf was ultimately found well outside of the city limits, and the person who posted or created this deepfake now faces jail time.

To answer your question about whether existing law is sufficient to cover an instance like this, well, whether or not AI was used to intentionally throw the authorities off the trail is probably covered by a law like obstruction of justice if this happened here in the United States. So, regardless of whether AI was used, a tip provided intentionally to police to throw them off the case could potentially be actionable here.

Now, I don't know all the facts in this particular case, and I certainly don't want to condemn this guy if the facts are that he provided a tip or that a different person provided the tip having seen it on social media. So, not to assign any guilt here, only to say we do have obstruction of justice here in most states, if not all, and perhaps providing a false tip intentionally could get you into hot water.

On the other hand, posting on social media is a form of protected speech, and if that's the concern—that a person engaging on something newsworthy, an animal having escaped a zoo, and authorities are now going after him because either he provided false information as a form of a joke or otherwise, or because he was just providing AI and it was illegal to do so—that's probably a little bit more questionable here in the United States. So, long story short, I think existing law is going to cover many of the harms that people are concerned about now. The better approach, rather than coming up with a large regulatory scheme to deal with the harms that AI can cause, let's take a step back and let's look at existing law first. And if we find any gaps in the law, let's identify them and let's address them specifically, rather than sort of the broad approach that I've seen happen in other states and at the federal level.

Becca Branum: Yeah, I'll just quickly chime in to say that I think that that's right, and I think it's important to keep in mind all the different ways that AI can be used, and that in certain circumstances, people might not be seeking out a lack of bias; they might actually be seeking out a particular viewpoint. And so, I want to disaggregate bias as it's applied in sort of civil rights law, which I think laws apply to AI and need to apply to AI, but when we're talking about bias in a broader sense—in a sort of ideological, neutrality, or lack of neutrality type of way—I want to challenge the idea a little bit that that's always desirable because it assumes that a person is interacting with AI and expecting the AI to sort of reflect truth back to them rather than them engaging with AI in a particular way.

We know people use AI for all sorts of things: for play-acting, for role-play, for just seeing what they can get the AI to say back to them because they think it's fun to see what the guardrails will let them do. People interact with AI for all kinds of different reasons. And so, I think I want us to be careful about separating viewpoint from bias, and bias in conduct versus developers and AI systems being able to have a point of view. Again, there are certain AI systems that I use more readily than others because, in fact, I prefer the viewpoint and sort of the constitution or the

guardrails that they have in place as compared to some of their competitors. I think that in the end ends up being better for everyone involved because it allows them to access the kind of information that they're looking for.

At the same time, though, that makes it all the more important for companies to be clear about what it is they're offering to the public and to their consumers. We don't have the same types of transparency rubrics, like the Santa Clara Principles, that we have in social media as applied to AI yet, and that's a pretty dangerous thing for consumers because they don't always know what it is that they're signing up for or what it is that they are getting out of an AI system. And so, companies are starting to, but need to do a lot more to make clear what their viewpoints are and what it is that you're getting when you solicit information from these systems. I think that would also go a long way to both identifying bias and making sure people can make informed choices, while not flattening the types of lawful information that are available to people.

Magic Wand Principles

Jamal Magby: For our last question, I want you both to get out your magic wands, if you have them available, because I would love you guys to tell me: if you could shape one principle or safeguard that guides AI policy over the next few years, what would it be and why? And John, because you're our guest, I'm going to start with you, and then Becca, we'd love to hear from you.

John Coleman: That we have an affirmation that people have a right to share AI content, they have the right to receive AI content, that they have a clear right to speak anonymously when it comes to AI-assisted content, and that people will be free from compelled speech when they use AI. I think if we got affirmations on all of those issues, we're going to be in a much better spot in society, and we're going to be on the leading edge of AI and individual rights.

Becca Branum: And I think for me, what I hope we can keep in mind moving forward, and the principle that I hope will anchor everything, is that the government cannot and shouldn't be in a position of deciding for us the types of ideas, perspectives, or information that are worth our time. That is the core of the First Amendment, and I think we should hold on to that pretty jealously and not hand it over either to the government or to companies as a monolith to decide what it is that we should be thinking and should be doing with this new technology. Whether we like it or not, and some people are not the biggest fans of AI, AI is being incorporated into a lot of products, which means it's really important for us to protect our ability to use those tools and that technology in a way that's consistent with our First Amendment values, and make sure that we can use it to communicate, learn, and think the ways that we want.

Conclusion

Jamal Magby: That's fantastic. Well, John and Becca, thank you so much for being here today. I really, really appreciate your time and your thoughtful candor throughout this conversation.

Becca Branum: Thank you.

John Coleman: Thank you.

Jamal Magby: Thank you for listening to Tech Talks, presented by the Center for Democracy & Technology. I've been your host, Jamal Magby. Tech Talks is edited by Jacob Kaufman and produced by Drew Courtney. Check out more of CDT's work by visiting us online at cdt.org/techtalk and on various social media @CenDemTech. That's C-E-N-D-E-M-T-E-C-H. Thanks for talking tech.

[Theme Music Fades Out]

Speaker Specifications

- **Jamal Magby:** Host / Moderator.
- **Becca Branum:** Guest; Deputy Director of the Free Expression Project at the Center for Democracy & Technology (CDT).
- **John Coleman:** Guest; Legislative Counsel on AI and Free Expression for the Foundation for Individual Rights and Expression (FIRE).

Accessibility Note

This transcript includes thematic headers, speaker tags, and sequential timestamps to remain fully compliant with ADA accessibility guidelines for individuals utilizing assistive screen readers, or those who are deaf or hard of hearing.