Surveying Foreign Influence in AI Tools

March 4, 2026
Featuring Joseph Bodnar and Jamil Jaffer
Moderated by Leah Siskind
Introductory remarks by RADM (Ret.) Mark Montgomery

MONTGOMERY: Hey, welcome, and thanks for joining today's event hosted by the Foundation for the Defense of Democracies. I am Rear Admiral Mark Montgomery; I'm the senior director of FDD's Center on Cyber and Technology Innovation. Today's Wednesday, March 4th.

We're here today to talk about authoritarian regimes that are gaining influence over how Americans understand the world. Legacy media has left a big void in content cited by large language models, and the authoritarian governments have been eager to fill this. By optimizing their authoritarian propaganda for AI consumption, foreign influence is shaping the narratives and AI tools relied on by millions of us for our research, our education, and our everyday information.

Russian campaigns to shape AI narratives, Kremlin-aligned perspectives embedded in chatbot responses, and the rapid spread of Chinese-built AI models into the United States – all of this adds up to a serious and growing risk to American security.

On this topic, FDD is proud to present a study of propaganda citation by AI tools, a memo – a discussion here that investigates how lot – large language modules [sic] drive traffic to this propaganda, complicating the use of AI as a trustworthy research tool.

To discuss these topics today, and critically, the options available to policymakers, technologists, and the media ecosystem, I am pleased to introduce Joseph Bodnar, senior research manage – manager at the Institute for Strategic Dialogue. Joe previously held research roles at the German Marshall Fund, the Atlantic Council, and his written wide – written widely on state-backed propaganda, Russian disinformation, and foreign interference.

We also have my good friend Jamil Jaffer, the founder and executive director of the National Security Institute at George Mason's and Antonin Scalia Law School, where I'm still just a senior fellow. I can't get that distinguished fellow. Thanks, Jamil.

JAFFER: We'll work on it; we'll work on it. We – we do a battlefield promotion, if you'd like.

MONTGOMERY: There we go. Alright.

He's also an assistant professor of law and director of the National Security Law and Policy and Cyber Intelligence and National Security LLM programs at George Mason University, where my nephew's a proud graduate of the law school. Jamil's also on CCTI's advisory board. Jamil, it's great to see you here.

All right, moderating – moderating today's panel is FDD's own Leah Siskind. Leah wears dual hats at FDD. She's both director of Impact, but really most importantly, she's the AI research fellow at FDD's Center on Cyber and Technology Innovation. She has a great new memo that I recommend to each of you entitled, "A Study of Propaganda Citations by AI Tools" [sic. "AI-Amplified Narratives: Measuring Propaganda in LLM Citations"]. It's available at the entrance.

Before I turn it over to Leah and the panel, a few words about FDD. For almost 25 years, FDD has operated as a fiercely independent, non-partisan research institute exclusively focused on national security and foreign policy. As a point of pride and principle, we do not accept any foreign government funding.

For more on FDD's work on cybersecurity, technology, AI, and more, please visit FDD.org, follow us on FDD – X and Instagram, and subscribe on our YouTube channel.

Leah, the floor is yours.

SISKIND: Thank you so much, Monty. OK, let's dive right in. Today, I'm excited to discuss the findings of FDD's original research on authoritarian influence on large language models, which for – here on out, we'll just say LLMs.

We have an excellent team of researchers and a brilliant technical editor, and some of whom are in the room right now. So, I just really quickly wanted to acknowledge your help and say thank you so much. Where's Britt? Thank you. This could not have happened without you.

So, the issue we discovered is that since most legacy media outlets either – are either blocking LLM bots for training and retrieval, they've created a massive void which is being filled by state-aligned authoritarian media. For our experiment, we studied three major international conflicts – Israel/Gaza, Taiwan/China, and Ukraine/Russia. We found that when analyzing three major large language – LLMs – ChatGPT, Claude, and Gemini – that 52 percent of responses to questions about current international conflicts cited state-aligned propaganda sources.

Understanding this void is just one piece of a larger puzzle, so I've brought in two other experts to help us cover the full national security picture. I'm so pleased to be joined by Jamil and Joe to talk about the various ways adversarial actors are leveraging AI to disseminate propaganda. So, thank you so much for joining us.

OK, I wanted to start with Joe. You found that Russian influence on American LLMs manifests in a lot of different ways. Can you help recap your research for us and tell us about Russian strategy and tact – tactics for AI influence?

BODNAR: Yes. So just for some background on me, I'm mostly – I primarily track Russian influence operations. And there was prior research done that showed some covert information operations had been able to successfully influence the output of chatbots.

So, our question was, has overt propaganda been able to do the same thing? Have the RTs and the *TASS*-es of the world been able to get cited by…

JAFFER: Wait, RT is not a legitimate news source?

(LAUGHTER)

I have friends in DC who literally watch RT like it's actual…

BODNAR: Yeah.

JAFFER: …it's actual news.

BODNAR: Yeah, that happens.

JAFFER: Maybe the – the DNI maybe, even. Didn't – I didn't say it out loud.

BODNAR: So, we were interested in this question to see who besides his friends were getting served RT. And we – our approach to it was kind of varied because we were interested in a lot of different variables.

So, we looked at four chatbots – ChatGPT, Gemini, Grok, and DeepSeek – because we wanted different capabilities and different design philosophies. Then, we picked five topics related to the war in Ukraine to see if there was any particular topics where there might be data voids and some that had maybe less data voids. And we designed three layers of – this is a lot – three layers of questions – neutral, biased, and malicious – and then we asked those questions in five different languages.

So, the result of all of that – the big takeaway for me was about one in five of all – of our queries returned citations of Russian state media. And that number, the percentage of citations, increased from neutral to malicious. So, each time the level of biases increased – or bias increased, we saw more Russian propaganda.

We also saw the models behave differently. ChatGPT was just – it performed the worst in our testing, and it was also the most susceptible to being influenced by that bias. Grok was actually kind of difficult. So, it did cite Russian state media, but it also cited a lot of pro-Russia influencers on X. It would surface their posts. That didn't follow – we didn't categorize that as part of our data. More studies need to be done into that, but since they weren't directly, overtly connected to Russian state media, we didn't catalogue it.

DeepSeek, incredibly variable. Some questions returned limited Russian propaganda. Others, it was only Russian propaganda and covert IO responses that informed the entirety of DeepSeek's response. And then credit where it's due, Gemini did really well. It put warnings on a lot of things, saying that if we were looking for information on a conflict, we should consult reputable sources and not a LLM.

One thing to note, language didn't have as much of an impact as we thought. So, we asked questions in five languages, and it wasn't a variable that caused any statistical variation. There's a lot more research to be done in this space. New AI models are coming out. More tests need to be done. And AI summaries are working their way into literally everything. Google Search in particular, I think, is problematic, but that's for future talks maybe.

SISKIND: And just note that quite – that aligns with our research as well. ChatGPT was kind of a shit show and – and Gemini was…

BODNAR: Are we allowed to say that?

SISKIND: I think so. I hope so.

JAFFER: We're livestreaming right now, so we did. No FCC to bleep you out.

(LAUGHTER)

SISKIND: So how much of this is deliberate Russian strategy and how of it is just – how much of that do you think results in filling a void, a data void?

BODNAR: I think it's – it's both. It's certainly a deliberate Russian strategy.

So, we can look at – to quotes from people like John Mark Dougan. I hope not everybody here knows who John Mark Dougan is. He's a former Florida cop who is now notorious for running a – a group of hundreds of fake global news sites that are called “CopyCop” sites. And at a roundtable last year, he was actually complaining about how the AI models he was using to produce content for those sites had too much Western bias in it, and he said we need to train it from the Russian perspective.

So…

JAFFER: He's an American?

BODNAR: He – well, he lives in Russia now.

JAFFER: Oh, well, yeah.

BODNAR: But he was…

JAFFER: Kind of like Edward Snowden.

BODNAR: Yes. I would not be surprised if they were in the same circles.

JAFFER: There you go.

BODNAR: But CopyCop sites, they put out so much content. To my knowledge, none of them have appeared in LLM results. They're kind of just sitting there on the web. I don't know if they're being scraped and things like that.

The big player in this game is another IO called the Pravda Network, collection of dozens and dozens of sites, publishing a flood of content, seven million articles to date. And that's where the term "LLM grooming" kind of came from, originated, was studies into this because it does look like it's deliberately set up to influence the output of LLMs. And some studies have shown it's able to do so.

SISKIND: That's so interesting.

Jamil, this has been a very busy couple of weeks for Chinese AI misuse. While everyone was distracted by the Claude/DOW showdown – I know you have strong feelings, too – China was caught stealing knowledge from Claude and using ChatGPT to review one of their influence campaigns where they were trying to influence the Japanese election and attempting to silence critics of China in America by impersonating U.S. government officials.

JAFFER: Yeah.

SISKIND: This shows that despite all of their advancements, they still see American AI tools as superior. And when it comes to building their own models, they are essentially learning from American AI and then applying a Chinese – applying Chinese censorship filters on top.

So, knowing what we know about Russia's use of AI for influence operations, can you contrast this with China's approach?

JAFFER: Well, you know, what China's been doing – one, we – let's be clear, now that Anthropic's come out with a report talking about distillation, we now know pretty incontrovertibly that a lot of the Chinese tools that have made it onto the open market – the DeepSeeks, the Qwens, the Kimis – are all trained on the output of American closed-source models and – and have been distilled down.

And so, we know that they're essentially doing what they've done in every other IP area, stealing American intellectual property, repurposing it in China, and then putting it out on the market. So, we know that's part of what's going on here.

But then when you look at what they're doing with those tools, you see them, both for their own home audiences but also for foreign audiences, ensuring that when those tools are utilized, they're putting out content that aligns with the – with the Chinese state view of the world, right?

So, you see, when you ask DeepSeek or Qwen or Kimi – and they vary amongst one another – I don't want to lump them all together – with some variations – if you ask about Taiwan, they give you more state influence views than sort of the standard spin, what you might get from a normal American model, right? You get views about Falun Gong and the like. When you ask them about Tienanmen, obviously, you know, "Tienanmen what?" right?

And so, there is – it's clear they are trying to shape the – the information environment, right? Less than – it's less than undermine the models themselves, right? Which – which I don't doubt, that their narratives, right, that they're putting out through their media organizations are also being pulled in the way that your studies have shown. So that's happening already, but on top of that, when you – if you use their models – and the real problem is a lot of American companies are going deep on their models.

SISKIND: Yeah. We're going to come back to this.

JAFFER: They're implementing their models in their core operating – they're – well, they're like, "Well, look, we're not putting out information. It's fine. We're just using it to, you know, do an internal app." But the problem is if we know they're influencing the outcomes of – of news stories and research, what else are they influencing the outcome of inside your business that you don't know, right?

And so, I think American companies have to step very carefully if you're implementing DeepSeek, Qwen, or Kimi. I get it, you're saving a bunch of money, but that may come at a much larger cost down the road.

SISKIND: I want to come back to that, because it's such an important point and we need to talk about solutions and how to approach this, but first – Joe, both of our research found that chatbots tend to show confirmation bias…

BODNAR: Yeah.

SISKIND: …which means that how you phrase the prompt can – can lead to reinforcing that perspective. Can you talk about the national security implications of that?

BODNAR: This is tough, because it gives actors who want to run these info ops a shortcut. They can go to these chatbots, prompt it, and get it to produce tons of content that isn't copypasta, as I was actually just talking about with Max [Lesser] here, and feed it to users in different languages around the world. It's just a tool that speeds up these – these info ops that are becoming increasingly problematic as the world – the pace of news accelerates and people's hunger for it increases too.

SISKIND: And we – you talked briefly about the Pravda Network, which operates in over 80 countries, and they essentially launder Kremlin talking points through locally-branded sites that appear independent.

So how does this kind of distributed infrastructure make it harder for LLMs or anyone to filter propaganda?

BODNAR: Can you ask the question one more time?

SISKIND: Sure. So, we're talking about Pravda...

BODNAR: Yeah.

SISKIND: ...and how it's – it operates in so many different languages...

BODNAR: Yeah.

SISKIND: ...in 80 different countries. But because it's kind of – it's filtered often through locally-branded sites...

BODNAR: Yeah.

SISKIND: ...it's a network of sites, right? So how does this kind of distributed architecture make it harder for LLMs to identify – or any person for – to identify that this is propaganda?

BODNAR: It – I mean, if they're targeting environments where there isn't a lot of good – good news sources and they're flooding those environments with bad content, that content's going to get scraped. It's more likely to get picked up, and that's problematic.

Yeah, I...

SISKIND: Jamil, pitch in.

JAFFER: You know what – what's really interesting about this stuff?

SISKIND: Yeah?

JAFFER: You know, you might think of, do we need to bake into these tools the ability to ignore these sites, right? Remove the problem network sites...

SISKIND: Yeah.

JAFFER: ...from your – from your models. But let's say you did that, right? At least some people are going to say, "Well, that's a woke approach, right? You're gerrymandering the model," right?

SISKIND: Yeah.

JAFFER: I mean, I think that Grok would say to you, you know, the people who created Grok would say, "That's not an appropriate thing. We're not – we're not going to police that content."

SISKIND: Yeah.

JAFFER: It's part of why Grok is using all these – all these – the folks that you didn't identify, but the – the – sort of the Russian influencers on Twitter, right, or to X, right?

BODNAR: Yeah.

JAFFER: That's a problem.

BODNAR: Yeah, and – and when we talk about Pravda, those aren't hard sites to identify, and we don't have to get rid of them completely.

SISKIND: Yeah.

BODNAR: All these Pravda sites are catalogued. A lot of researchers know what they are. AI companies know what they are. You – you don't have to blacklist them completely. You can down-rank them...

SISKIND: Yeah, exactly.

BODNAR: ...in your results. You can label them as state-affiliated media, so end-users have that context.

SISKIND: Yes.

BODNAR: There's ways to do this that isn't censorship.

SISKIND: Yes. I would never advocate for blocking these sites. People need to be able to read and understand...

BODNAR: Yeah.

SISKIND: ...you know, alternative points of view. But they should know what they're reading.

BODNAR: Exactly.

SISKIND: So, I'll come back to this.

Jamil, so, it's beyond influence operations. We know that China is also, you know, actively stealing our intellectual property and using this knowledge to improve their own models. What does this say about the state of the AI race between the U.S. and China? I think it's like...

JAFFER: Yeah.

SISKIND: ...we're all following this, but I think it's good to kind of level-set on where are we right now?

JAFFER: Well, you know, there is sort of a popular theory of the case that China's winning. They've got all these – the models are coming in cheaper. They're coming in better or faster, or whatever the thing is. And a lot of – as a result, a lot of American companies, as we talked about earlier, are starting to implement those models into their ecosystems, thinking, "Well, it's just cheaper, faster, better, or whatever. Who cares if it's Chinese? I've got all sort of Chinese equipment in my operations, so what's the big deal? One more thing at the core of my AI."

The problem, of course, with that is that we don't know how that's going to play out in the long run, right? It's one thing to say I've got – and we've learned this lesson. We learned this lesson from Huawei and ZTE, British Telecom. But Huawei, at the core of its network, even as we told them not to do it, they're now having to rip it out. A lot of American local telecoms.

But you can't rip out AI models from your infrastructure, right? You can't rip out the decisions you made based on that knowledge. And remember, the decisions will have to be modified by a slight amount, right? You – your thinking on a given issue has to only move by a little bit to make huge differences down the road in your business; could be imperceptible in undetectable amounts of change, right? Something as obvious as no Tiananmen, right? No – you know, no pro-Taiwan messaging.

We just don't know, and this idea that American companies are blithely adopting models, you know, is a problem. But going back to your earlier question, the core of your question was where we are in the race. I think what it demonstrates is we're still winning.

SISKIND: Yeah.

JAFFER: Right? We don't need to necessarily make America great again on AI. We're great on AI. We're winning the AI race. The problem is we can't allow it to be – we can't allow the Chinese to get ahead. And so, there's all these debates about, you know, what chips we sell to China, right? How much, how little. And these are – these are important debates to be had, and there are reasonable views on all sides of this.

But the reality is if you give China high-end American capabilities, right – whether that's alignment of the – alignment and – and access to our core IP on our models, or the highest end of American chips, or even maybe the midrange of American chips – that's going to accelerate their ability. And a lot of people say, "Well, look, you know, they're already going to build them, whatever," right? But that two, three years delay, that might matter a lot, and that's why it's been so controversial and such hotly debated issue. There are important economic things to consider, as well, but at the end of the day, we are winning the AI race, and China's borrowing our models because they know we're winning.

SISKIND: Yeah, that's a very helpful recap.

Joe, our research found something counterintuitive, which was that Google actually surfaces less propaganda over Gemini...

BODNAR: Yeah.

SISKIND: ...than LLMs do. Oh, I'm sorry. Sorry, let me rephrase this.

My research found that when we compared searching for certain topics through LLMs or through standard Google Search...

BODNAR: Yeah.

SISKIND: ...that Google would surface less propaganda because – for – for Google...

JAFFER: As compared to Gemini?

SISKIND: Yes.

JAFFER: Oh.

SISKIND: Because for Google, the link is the product. It doesn't have to – it can rank the *New York Times* highly even if it can't read it. But an LLM needs to actually be able to read the content in order to summarize and cite it.

BODNAR: Yeah.

SISKIND: So, which means that paywall journalism is effectively invisible for LLMs.

BODNAR: Yeah.

SISKIND: State media, which is free and accessible, fills that void.

BODNAR: Yeah.

SISKIND: So, the question becomes – so you start – we started to talk about this before.

Should AI companies be required to rank or weigh media outlets by credibility? And if so, who decides what "credible" means? And Jamil, you're welcome to weigh in on this, as well.

BODNAR: Yeah. Yeah, I mean, who decides what's credible is a very difficult thing to say. I don't want to wade into that. What I would say, just to repeat an earlier point I made, is that when you know some content, some domains are tied to state sources, and in the past have pushed false and misleading content, downrank them. Try to surface what is credible. Again, we can have a debate about what that is.

But I think to your point about Google being a little bit – Google Search being a little bit better at this, they've had a lot of time to wrestle with this problem. They've understood what that the links are the threat vector. Like, that's the threat. I don't know if these AI companies have come to terms with that yet. And we've seen, like, Google has worked to get rid of data voids. They know how search engine optimization has been exploited by bad actors. I don't know if AI companies have really gotten their head around this, and that's why our reports are hopefully pushing them to – to do something.

JAFFER: Yes, I've got a few thoughts. One, I think that – I think we all know what's credible and what's not, right? And you ought to apply a common-sense rule here, right? If it's Russian state media, right, or a Russian state affiliate, or a Chinese state affiliate or Iranian state affiliate or North Korean state affiliated, let's not – let's be American companies and think through that for a second, right? And you don't need to be woke, or you don't need to be – you know, you just apply a common-sense rule, and I don't think that's – I don't believe that's censorship. I think that is perfectly appropriate. It's something we have always done, we've always thought through. I don't see it as censorship.

SISKIND: It's the opposite of woke.

JAFFER: What's that?

SISKIND: It's the opposite of woke.

JAFFER: It's the – right. Well, I mean, but there are those who would say, "Well, you know, you've really got to have – you've got to give the American people full access to whatever they want so they can figure out for themselves and not shade it at all, and not up- or down-rank, because – but you know, we've always understood our media companies to be able to help curate content for the American public in ways that are pro-democracy and pro-freedom.

And let's be real: Iran, Russia, China, and North Korea are not pro-democracy or pro-freedom, and that's OK to say that. It's not problematic to say that. It's common sense, and it's not insane.

But here's the other piece of it: I actually think that the markets can help solve – markets can help solve this problem.

So, to date, you've have market pressures on *New York Times* and *Wall Street Journal* and all those other media outlets to paywall their sites. But now, if the driver, right, is – are the LLMs, you almost want to be the LLM-trained data. They've been fighting and suing all these LLM companies. But the real truth is that you want to be that training set because you want to be cited. You want the click-through to your stuff so you can then – you can then monetize that. And so, if you're not doing that, you're actually undermining your own economic rationale.

So, I think there's an opportunity here to get an economic deal done when these companies start realizing it's actually their materials being undermined by – you know, the Pravda's and the RTs of the world. And then I think the third piece of this is, the companies actually have an incentive to get the right data out there, right?

And at the end of the day, if your AI capabilities aren't trusted, safe, and secure – people don't – if people started getting a lot of Russian content and they don't want that – so one, it has to require consumers who are educated and understand what Russian content, Chinese content looks like.

But if the American people – and TikTok, I guess, is a counterpoint to this – but as the American people start to understand that that's what they're getting, they're not going to want that. So, they're going to – they, I think, will, over time, if we're able to educate the American public, will – and the Western public – will start downgrading and using the tools that are less cautious and less cognizant of these problems, right…

SISKIND: Yes.

JAFFER: …and so, I think there's a – there's incentives for investors and innovators to do the right thing here too, and to get the right kind of content, the right kind of training data, and the right kind of output data, so we get real valid, accurate content coming out. And we've seen that with hallucinations.

SISKIND: Yes.

JAFFER: This is the moderate version of hallucinations. They're not hallucinations, because they're intentionally put in messaging, but it's still messaging.

SISKIND: Absolutely. I – I feel that both the media and the AI companies have like a deep vested interest in collaborating with each other.

Oh, I – I lost my place.

Jamil, DeepSeek became a cultural moment in the early 2025 when we learned that the Chinese could produce a competitive model that is significantly cheaper for consumers.

JAFFER: A stolen, as it turns out…

SISKIND: Yes.

JAFFER: …competitive model.

SISKIND: Yes. Can you talk about how they censored controversial topics? You've – you've kind of…

JAFFER: Yes…

SISKIND LEAH: …hit on this a bit already, but…

JAFFER: Yes. Well, we see – we've seen a lot of – we've seen a lot of reporting about this, right? So, what we know is they've gerrymandered the inner workings of DeepSeek so that when you ask questions about topics that they don't want to talk about, it's primarily – to be fair to them, it's primarily for their domestic audiences, right?

But it turns out, when you start implementing them here and started using them here – and people started adopting DeepSeek because it's cheap and so fast – you started getting that same bad content.

And so, we saw it, right? Taiwan stories came out the Chinese side, right? And again, the various models, they're different, right? DeepSeek isn't the same as Qwen, isn't the same as Kimi, but they all have a layer, because, unlike in the United States – albeit, you know, some debates today about whether that's more or less true in any given scenario, there's – we had this sort of notional separation between, you know, church and state, between private and public.

In China, that's not a thing. And so, as they build these models – if you're an American, you got to know, if I'm using DeepSeek, I'm getting Chinese state propaganda at some level, and that's what you saw. You saw – you saw Falun Gong downgraded. And there's other – there's even more pernicious stuff that we should talk about that's – that's not for messaging only, that's about the output you get from – as they're coding stuff for you, if you use DeepSeek to code…

SISKIND: Yes.

JAFFER: …but that's a larger problem. But you got problems on Falun Gong, you got problems on Taiwan, you got problems on Tiananmen Square, and those are the obvious ones right. But if it's happening with the obvious ones, you could almost be guaranteed it's happening with the non-obvious stuff too. It's just more subtle. It's more cautious.

The Chinese are not stupid. They know what they're doing. This is a plan and executed operation, and the Russians are really good at this. China's getting better every day. And if we start embedding these models and start use them aggressively, as you see in the uptake, this – this cultural moment as you called it, you know, you do it at your peril if you're an American company.

SISKIND: So, yes – and it – besides DeepSeek, now that we have Kimi and Qwen and like so many other ones, what should Americans understand about when – what they're handing over when they use a Chinese built AI model…

JAFFER: Yes.

SISKIND: …like that is trained on PRC data, and their censorship laws…

JAFFER: Yes.

SISKIND: … like what are the data risks?

JAFFER: So, a couple of things. One, people say, "Well, look, I can just run DeepSeek locally, and it's fine. I'll just put it on my computer, nothing'll get out." Well, A, are you sure? Right? Can you guarantee nothing's getting out? Right? Have you – are you moderating the backend of your network to make sure nothing's leaving?

I think that's a open question, right? But number two, right, if you're getting – if you're getting content that's been gerrymandered, right – let's say you're now using it the way you would use Claude Code, to create applications. Those applications may be gerrymandered in ways that you just don't understand.

And it turns out that at least a couple of reports, CrowdStrike being the most prominent one, has determined that in fact, when DeepSeek thought it was building something for a Falun Gong practitioner or a Falun Gong group, they gave more vulnerable code.

Think about if that's happening at the core of your company. You've invested hundreds of millions of dollars in the company. You think, "I'm just going to use DeepSeek on the cheap, save a few – save a few bucks on the tokens, right? Claude's too expensive." Whatever it might be, right. "I don't want to use Cursor, so I'm going to use DeepSeek. It'll just be faster, easier." That's a problem, right, because it could be burying that – those vulnerabilities in your code, and you don't even notice. And if that's what's happening, how will you know when that code is exploited? You may not ever know.

SISKIND: Absolutely. So, I think that's very scary. And I don't think we talk nearly enough about that.

So, let's talk about some solutions. We are at an inflection point, I believe, where AI is becoming as foundational as the internet itself…

JAFFER: Whoa.

SISKIND: Yeah. And in my research, I say that the onus is on AI companies, because they need quality journalism, aswe discussed, to create quality AI.

If they continue to – if the media companies continue to withhold their information, the AI models need better – they need news streams of data, they need quality data to train on. If they continue in this direction, like the outputs are going to be lower quality.

They need it for training and for retrieval. One thing that – you know, when we're talking about the ability, or the lack of ability, from the models to detect propaganda, one thing that stood out from our study was that for all the different conflicts that we studied, it would – I noticed from ChatGPT, to be honest, would cite *al Mayadeen*, which is a Hezbollah publication, which is designated as a supporter of terrorism, and there would be questions of Taiwan about maritime issues.

Like, it was – it was so random, but the fact that the – the models themselves could not distinguish between what is propaganda and what is a reputable source – was pretty shocking.

JAFFER: Yes.

SISKIND: So, I…

JAFFER: Clearly, common sense could help you figure that out.

SISKIND: Exactly, exactly. So, I believe that AI companies have a responsibility – I personally believe that AI companies have a responsibility to label state aligned propaganda.

BODNAR: Yeah.

SISKIND: This is not impossible. Companies have done this before. YouTube did this at one point. And I think you can avoid like the domestic controversy, like, "is this being woke?" issues by, you know, focusing on foreign propaganda. I…

JAFFER: Although that's even generated controversy in the whole woke moment, right? If we label foreign propaganda, there are those in the United States, those who want to make America great, who say, "Well, you're gerrymandering the dynamic." So, it's not unanimous.

SISKIND: Definitely. Like, no matter what it'll create controversy…

BODNAR: Yes.

SISKIND: …but I think people have the right to know what they're reading.

BODNAR: Certainly. I – I was thinking of the dust up – I mean when Musk took over X – or then Twitter – he got rid of the state media labels…

JAFFER: Right.

BODNAR: …but before he did that, he actually added state media labels to the BBC and the NPR and there's just…

SISKIND: Oh.

BODNAR: …there's a difference between these two things, where one has editorial independence, but funding, and the others do not.

JAFFER: And it's OK to say that, because we live in a free and democratic society, and we shouldn't shy away from being OK with treating Russian content different than American or British content. That's OK.

BODNAR: Yes, 100 percent.

SISKIND: That's why we said "state-aligned" in our research. Because, yes, we're not talking about NPR. We're not talking about BBC, or NPR when it received government funding. Because the – yeah, these organizations still have editorial independence even if they take money from the government.

So, the – all of the outlets that we focused on were – they're operating in an environment where there is no press freedom, and their editorial line parrots the government's take on things.

BODNAR: From my really nerdy perch, I think it could be a little bit more challenging than just the commonsense approach. There is the commonsense approach, but there's also new domains being created every day by these actors. There's old domains that are then being filtered through mirror domains.

So, it's going to take – I think the onus is on AI companies. But they're going to have to collaborate with researchers and civil society organizations to create that repository that gets continually updated because threat actors are continually evolving. We can't just have one list and drop it there.

SISKIND: For Russia, specifically, how can we protect American AI tools from Russian influence?

BODNAR: I think I've talked about it a bit here. I mean, we've talked about labeling. I just talked about creating a repository. They can do things like testing against these vulnerabilities and then making the results of those tests public. So, we know how safe these models are, and if we should be on the lookout for Russian sources slipping through. And I think that's something that should be done on a continual basis.

SISKIND: There is some debate on whether the LLM grooming is effective or not. What's your take on it?

BODNAR: It's so hard to know. I mean, we don't have enough data…

SISKIND: And just to redefine that that's flooding the zone with tons of – tons of media.

BODNAR: Yes. It's putting enough content out there that it gets vacuumed into the training data and influences what chatbots say. And it's so hard to know how effective they are, because we don't know what sources they're ingesting. We don't know how they weigh those sources. We just don't have enough insight into the way these models operate.

All we can do – we're not totally blind, because we can read them, and we can get our own results, and we can sit on panels, and we can talk about them. But we just don't have enough data to know if an LLM has been groomed.

SISKIND: Interesting. Jamil, what steps should the United States take to prevent Chinese models from continuing to profit from intellectual property theft?

JAFFER: Well, we've got to take action against the state itself. We know that the – that this intellectual property theft that's happening at massive scale and has for decades has really funded the Chinese economy. It's grown the Chinese economy massively and it's run by the nation state, right?

When they hacked in American systems, American companies, stole intellectual property and brought it over to China and repurposed it, that was a state-run plan.

The same thing is happening with AI models. It may be through their private sector companies, but again, in China, there's no real difference between state and private. Every Chinese private – quote, unquote, "private" company has a CCB – CCP Party apparatchik at the top, right, even a committee.

And TikTok-ByteDance, no example – no exception, by the way. Even with the current, sort of, segregation of U.S. TikTok, it's not real, let's be honest. And so, the only way that we're going to prevent the Chinese from stealing American intellectual property and repurposing it for economic purposes and against our national security is to press our state up against theirs and put pressure on them.

So, tariffs have been one tool the president – this current president has used to effectuate that. Joe Biden had different tools – President Biden had different tools.

At the end of the day, though, none of these have worked enough because we haven't been tough enough. The reality is if we're going to fight fire with fire we have to punch the Chinese in the teeth in public, regularly, until they – until they start owning up to what they're doing and stop doing it.

And the more we think that we can negotiate with them and cut a deal before they really take ownership, the more likely it is that they will continue to do it because they get away with it over, and over, and over again.

And that's been a bipartisan problem. It was – Trump one was very tough on China, right? Trump two on China is a lot more – it looks tough on tariffs, but, in reality, we're selling them a ton of chips. We're trying to do a deal – the president says repeatedly over and over again how much he loves Xi Jinping, how much he thinks he can do a deal, he's going to Beijing.

That demonstrates weakness. The president likes being strong and tough. Every time he says, "I love Xi Jinping, I'm going to go do a deal with him," the more Xi Jinping – by the way, same thing with Vladimir Putin in Ukraine – they see weakness, they see blood in the water and they will wait him out and they will ultimately extract a better deal for themselves that's bad for America, bad for the president, his own view of himself, and ultimately, not the right way to negotiate these things.

SISKIND: That's super interesting. So, the signaling – the willingness to negotiate is actually a weakness?

JAFFER: That's right. And in fact, you see our National Security Strategy, it says, "Well, it's fine for trying to have their space. We're not trying to dominate them. We're not trying to win this thing."

The reality is, why shouldn't we be trying to win against China? We should be trying to win, right? "Making America Great Again" is not just about doing it at home. It's about a broadened presence. He's showing that in the Middle East. He showed to Venezuela. Why isn't he showing with China?

Why is he willing to have their space and we have ours too? No. We should want to win against China. That's OK. We should just say it out loud.

SISKIND: Yeah. That – yeah, that's very interesting.

Joe, one of the more alarming findings in our research – in my research – was that when users would click through to verify the answers the outputs in AI, which is like what good responsible AI use looks like. We're all taught that this is how you prevent – or how you discover hallucinations if you click – you ensure that you don't – that you're not having them by checking the sources to be safe.

BODNAR: Right.

SISKIND: So – but, in my research, I found that, the models would produce like pretty comprehensive answers, but then when you would click on the citations, you would get a much more extreme perspective than the LLM itself. So, the LLMs were kind of moderating themselves. But obviously, the citations were not.

How does that dynamically, like, map on to what you've seen in Russia's disinformation strategy?

BODNAR: We actually didn't go down that rabbit hole. So, I'm interested to hear your take on it, because it seems like you did follow that rabbit hole.

SISKIND: I mean – yeah. I can say from our research, like we – we studied, like I mentioned before, three different conflicts, and there were certain actors who really dominated citations.

Al Jazeera was far and away the leader on a range of issues, but, of course, on Israel-Gaza coverage. Even in neutral questions that we asked about the conflict, they were – I think they were in about – in over 70 percent of the questions that we would – neutral questions that we would pose, you're going to get an *Al Jazeera* citation.

BODNAR: Wow.

SISKIND: Which was very interesting – and by the way, a lot of people that – a lot of this comes down to media literacy and understanding. Like, just because you have a slick website does not mean that this is the kind of liberal Western journalism that we're used to.

JAFFER: Well, we know in the case of *Al Jazeera*, it's not.

SISKIND: Absolutely. It's run by the royal family, so it is – they have no press freedom. It's completely aligned with the government. But because they don't have a paywall and it's totally open – they're doing a lot of other things. In fact, they're far and away the, like, the best I've seen ever at optimizing their content for AI.

But I bring this back because – to our study – because when we would ask specific factual questions with numbers, like, "What percent of Gaza at this moment in time is being controlled by the Israeli military?" The LLM would balance different sources, and it would come back with 40 percent.

And then you would – it would cite Al Jazeera in the same line, and that article, the headline would be, "Israeli military controls 70 percent of Gaza." And then they would include another one by a Hamas-linked NGO that would say, "Israel controls 80 percent." So, if you're – if you're trying to be an educated reader and read your sources along with the material, it's going to give you a very different impression.

BODNAR: Yeah.

JAFFER: That's a hard problem to solve for, right? Because you want – the LLM is ultimately getting you a right-ish answer, right? As it tries to moderate the various sources it's looking at, but then it's giving you sources that have a lot worse and not quality answers. How do you solve for that? To your point, maybe it's labeling, maybe it's the like.

What I wouldn't want to see is a requirement from the federal government, right, or state governments or worse the EU, right, that you have to label things and the like. I think there are strong market incentives, to your point and to what Joe said, for these companies to do – to do this smartly and in an effective way. We haven't fully seen it yet. A lot of people are saying, "Well, now we need to regulate because they're not doing it today."

I think that's the wrong answer. I think – or that we need to give equal time. Right? It's such a funny thing because equal time used to be a problem for conservatives. Now it's – now it's like the popular thing amongst conservatives. I don't really get it. But Ronald Reagan had it right, which is the market can figure it out. It will – this will work out over time and the LLMs are seeing that result playing out. They're going to figure this out. We don't need a sprint to regulation to solve these problems. What we need is more people like you guys talking about these things and doing the studies you're doing.

BODNAR: I'm a little more skeptical about market forces here, I'm going to be honest. Because we can look at social media as an example. They did initially react to this and then they've rolled it all back and they're all doing OK. Like, state media labels are gone. Crowd tangle for research into Meta products is gone. And those platforms are still – they still have users going to them every day…

JAFFER: But to me that's an education problem…

BODNAR: …the market hasn't solved them.

JAFFER: Right? To me that's – so that's – you're right. The market hasn't gotten to see optimal results. You could have two – you have two views of that. It could be that the market doesn't have sufficient information for consumers to make good choices. Right? That we're lacking information, that's not a good market. Right? Or you could – and your answer could be how do I solve for that? Allow more information to flow in the market, educate consumers, have them be smarter on how they consume. Or you could say the government should step in and regulate, right, and solve for this.

My view is that latter approach, the government regulating, is actually going to cause more losses and cause more lack of confidence in what you get. Particularly in America, we're skeptical of government intervention information. Rather than allowing more information flow in the market, educate consumers about what is real and what's not real, and have them – have them put the pressure on the market for the market to change.

SISKIND: Can I ask a question to you, Joe? Because you brought up the EU and we haven't really talked about Europe at all. And your – your study obviously focused on five different languages, had a more European focus.

BODNAR: Yeah.

SISKIND: And the EU has been pushing AI companies towards adopting the disinformation code and flagging sanctioned media and chatbot outputs. The US is obviously taking a very different approach. Are there things that Europe is doing on specific problem that we should be borrowing?

BODNAR: I don't know if they're doing enough on chatbot responses. That was part of our recommendations, was that they do need to clarify their position on how chatbots handle sanctioned media, because I don't know if – if that's clear. Those laws were designed for social media platforms. And now that these chatbots are on the rise and their user bases are getting them into the realm of being a very large online platform, which would put them under regulatory pressure or obligations, they need to clarify what that looks like.

SISKIND: Interesting.

JAFFER: One other thing to just say about regulation, right, is that regulation also tends to lock in the larger players that are already in the market. So, if you want innovation and you want people to come into the market and be more innovative, the more you regulate, the less likely you are to have those starts who can come in and afford the cost of regulation to do that.

Now, the reality is in the LLM world, it's already very costly to be a startup player. You need huge amounts of investment in data and the like. But the more you regulate, the less innovation you're going to get. Proof – proof's in the pudding. Look at innovation in Europe, they regulate every day, every single thing, right? They regulate first, innovate second. We innovate first, regulate second. That's the right approach. I'd go America than in Europe.

SISKIND: Well, also, AI agents are going to disrupt all of this anyways. And for social media labeling, who knows, maybe in 10 years we'll have like 20 new social media companies, and we won't be reliant on the monopoly that we have now.

JAFFER: We'll all be reading Moltbot.

SISKIND: Yeah, exactly. Or our bots will be reading Moltbot.

JAFFER: Exactly.

SISKIND: Jamil, how should we – we talked about the risks of American companies using Chinese models and – or building tools on top of Chinese models. How should we encourage transparency around the use of Chinese models in America?

JAFFER: Well, look, I do think it is – it would be appropriate for the U.S. government, particularly for critical infrastructure providers and the like, to say, "We want to know when you're using Chinese models. Government contractors, we want to know if those are embedded in your capabilities that you're supplying to us, right? Or that you're supplying to – or that you're using for critical infrastructure operations that go to America's national and economic security."

And so, I think that's a reasonable form of transparency. If you're using foreign capabilities in the core of your operations, you should be able – you should be required to tell the public and the government about that. I think it's different when it's, you know, if you're using American capabilities, I don't think you should necessarily have to disclose that.

But if you're using foreign capabilities, I think that's perfectly appropriate. We've done it in a – a variety of other contexts. There's nothing unusual about that. And to be candid, we barred the use of certain foreign capabilities – Huawei, ZT – the most prominent recent example is Kaspersky before it. I don't even think that's necessarily a bridge too far. We're not there today. And actually, the adoption's been significant.

SISKIND: Well, DeepSeek on federal devices.

JAFFER: Yeah, on federal devices. But the real problem is – the real pernicious problem is in American companies and critical infrastructure providers. And the fact that having watched the Huawei ZT thing go down, having seen – having heard the debate over TikTok – albeit, you know, all bollocks up in the last year or so – the reality is that American companies should know better. American boards of directors should know better.

And maybe this is partly educating company executives and boards of directors. These are bad tradeoffs they're making for the slightly cheaper tokens, the slightly – slightly cheaper inference. You're making huge mistakes that could go at the heart of the value of your company. And by the way, the core intellectual property that you're now pouring into these models, some of which may be going back to Beijing.

SISKIND: Yeah. I think reading a couple weeks ago – or just learning that Airbnb is using Qwen was kind of a surprise moment for me to – to understand that.

For both of you, what does successful collaboration between policymakers and technologists look like to you and the media?

BODNAR: I'll let you kick that one off.

JAFFER: Well, look, I – I tend to believe that technology and policy can go hand in hand. There doesn't need to be this sort of divide of the valley over here and DC over here. There's a lot of common ground that we have and a lot of common shared values. A lot of people tend to think that there's this real massive gulf of values between us, and there really isn't. I think part of the problem is that we've sort of forgotten – we've sort of lost our way in thinking about what makes America special and different. Right?

What makes America special and different is the opportunity we give people to come here, become part of the American vision and dream, and integrate into our society and grow the society writ large. In recent years, we've forgotten that sort of core of America and we've allowed that to infect all the way we think about the way that companies collaborate with the government, right? We see it playing out in this Anthropic/DOW thing, but it's got – it goes back to back in the old days when Google and Maven happened. Right?

We're just seeing that thing re-crop its head up even as we're in a more sort of collaborative environment with industry and government than we've ever been. So, I think there's a lot of room for collaboration, I think there's a lot of room for cooperation. I don't think we need heavy-handed government regulation to make that happen. I think the government and industry can and should work together effectively.

And the one thing I – I really hate about a lot of our big – our big companies, they often will say, "Well, you know, now that we've grown to a certain size, we're really global in nature." But the real truth about all of these global corporations is that you never call the Chinese MSS when they have a problem. They never call the Iranian MOIS, right? They call NSA. They call FBI. They call DHS, and there's a reason for that. It's because they're American companies first.

There's a reason why they're still headquartered here and not in Moscow and not in – not in Bermuda and not elsewhere, because they benefit from being US headquartered and American. They benefit from a worker base, they benefit from innovation, they benefit from all of those things – tax policy and the like.

That means that they should have some responsibility too, and the market should force that responsibility upon them. And we've allowed – we've allowed this sort of "We're a global corporation" thing to go for too long.

SISKIND: Interesting. Joe?

BODNAR: I'm going to leave the policy questions to you two. I'm a OSINT nerd. So, if you have any questions that get into that data, I'm happy to…

SISKIND: Well, that's fine. Well – so before we turn actually to the audience Q&A, I do have a final question. We've spent a lot of time talking about sophisticated state actors and complex technical vulnerabilities, but the average person – for the average person using ChatGPT to research a news story, they have no idea that all of this is happening.

What is the single most important thing you would want an everyday AI user to understand after today?

BODNAR: If that's a tough question, and I think it gets to your point of – I just want to say verify your sources. To make sure the information that it's giving you – go check – go check those citations, make sure it's actually being – what's being represented is – in your chat is actually what was written by the reputable source that it's citing.

SISKIND: But I think we have to go much further than that because it's – checking your sources is not – it – we're at – like…

JAFFER: They've got to be reputable.

SISKIND: Yeah, where we're – I think you need to incorporate some media literacy with it because right now, we're just driving. Like, imagine, you know, K-12 systems are rolling out these tools and they're all trying to figure out how to incorporate this in the learning environment. And we're tell – like, it's like you – when you're teaching children, you have to cite the source, and we're driving people to these different sources and saying, "OK, this is – you know, this slick site is clearly some legitimate source."

But I think people need to also understand – they have to learn, just like you do, AI – AI trainings and how to use LLMs. You need to understand more of – you have to learn about the media ecosystem as well.

BODNAR: 100 percent.

JAFFER: And I think part of this goes back to, you know, I – as a political conservative, you know, we've always been concerned about sort of what we perceive, rightly or wrongly, as a – as a liberal bias in the media, right?

But the real truth is that as much as you might not like the liberal bias – the quote-unquote, "liberal bias in the media," you certainly don't want Russian media taking that place or Chinese media taking that place. And unfortunately, I think that conservatives have – we, conservatives, and my part of the party have lost our way.

We've gone – we've gotten so upset with American mainstream media that we're willing to fight the battle on behalf of foreign state actors who are doing things even worse and saying, "Well, we need to open everything up." That's crazy.

It is better to have American media with whatever biases it has and bring in the more balanced view of those rather than say, "Well, we just need to have open scenario and have everyone come in and – and demanding equal time for all of the people." That's crazy and wrong-headed and not where conservatives should be.

SISKIND: Well, as a liberal, I think it's a great point to end on. Fully agree.

(LAUGHTER)

Well, thank you so much. I guess we can move to Q&A now, if folks have questions?

ANAND: Hi. Thank you, guys, so much for being here today. My name is Shivane Anand. I'm with the Washington Institute for Near East Policy.

So, my big question, more geared towards you, Jamil, but there's this notion of the digital Silk Road and the Chinese effort to sell their cheap – cheaper AI models to countries that don't have an abundance of options and – and maybe don't have the ability to choose.

What do you say to them as a threat that Chinese AI presents? And what is the angle that the United States and its foreign policy and its foreign tools can play to sort of combat that in an effort to win this great power competition?

JAFFER: Yeah, that's – it's a great question. What I'd say to them is how has – how has collaborating with the Chinese for the last two decades worked for you, right? If you're a country in Africa, or a country anywhere else in the world – Southeast Asia, anywhere else – and you've been working with the Chinese, how has it worked for you on minerals? How has it worked for you on development? They throw a bunch of money at you; they bring in cheap labor. Is it your labor that's being used? No. They're giving you loans. How's that going for you? Debt diplomacy, right? We know how it's played out. Expect the same thing with AI.

They're do – they're trying to addict you to their capabilities, like they did with Huawei and ZTE, and they're using you to collect and they're using you to shape and using you to put you in their camp. If you want to be a player, you want to help your – you want to help your country, you want to gain the leg up, relying on the Chinese who have ulterior motives is not the right play.

And that – and those countries might say, “Well, the US has ulterior motives as well. The – Europe – Europe has ulterior motives. This is – this is – everyone's got a problem, and so we're just going to roll our own.” Maybe so, but if you're going to pick one, you've got to decide who's the better actor in this domain? You've tried the Chinese for 20 years. How's it gone?

SISKIND: I – if I can add a comment to that, some of the responsibility there also – or, like, it behooves the American AI companies also to invest more in open-source models.

JAFFER: I agree.

SISKIND: They have all said that they are going to invest more in that, and in the AI action plan by the federal government over the summer, they said they were going to encourage a lot more open-source development like…

JAFFER: It hasn't happened.

SISKIND: No, it hasn't happened. And so, we don't want the world to be – to be divided into – to, “You're going with China, you're going to the US,” but that we need to create more options. We need to make ourselves more competitive for the developing world.

JAFFER: Well – and look – and the other piece of this, to your question, what can the US government do about it, right? I do think we made a mistake by sort of taking all of America's soft power off the table, right? And I think that was a huge mistake.

There – no doubt there were – there were problems of those programs, there were – there were programs that had gone wildly off track. That's something that needs to be rolled back and – and – and corrected for. That doesn't mean you throw the baby out with the bath water.

And unfortunately, the moves we've made on American soft power have dramatically undermined our influence in these very countries where we need to be winning hearts, minds, and AI adoption for our American companies.

We want to collaborate with technology companies. Winning AI adoption for our companies is something that the American government should be actively engaged in. And taking our soft power off the table, taking USAID out of the game, that's a mistake.

Again, doesn't – not to make excuses for all of the problems of those programs – there's – a lot need to be fixed – but throwing the baby out with the bath water is not the right answer either.

ANAND: Thank you.

WILLLIAMSON: Hi. I'm Lauren Williamson, I'm with the Alexander Hamilton Society. Phenomenal conversation here.

And so, my question is whose responsibility is it to educate consumers on what they'll come into contact with on social media or with ChatGPT or Claude or DeepSeek? And kind of who or what do you think has enough credibility at this point to be taken seriously by the consumer? For this scope we can just say the US, but of course that applies globally. And what kind of pushback do you expect from both companies like OpenAI and state governments as well?

So, I think kind of the question at the heart of this is who gets to decide the narrative that consumers accept or create the tools, I think, that consumers use to weigh how credible AI platforms that they're using are? And so, I think that could look different for each demographic of the population, but I'm wondering what your thoughts are on that. Thank you.

BODNAR: I – I have an icky DC phrase in my…

(LAUGHTER)

…in – in my throat, that is: “whole of society approach.” I mean, these companies have responsibilities for saying what their vulnerabilities are. Researchers have responsibilities for continuing to test that and make sure they're being honest with us. Journalists have a responsibilities to educate people. Schools, like you said. It's – everyone has a little bit to give here, I think.

SISKIND: I don't think that's icky. I think that's totally right.

And one of the things that we've done in FDD is kind of studied countries who are more used to being inundated with disinformation and how they're – how they're educating their population around it.

So, it's everything from tools to, as we mentioned, education in the classroom. Teaching kids from a young age, not just, these are the tools you're going to use, but these are the risks associated with them.

BODNAR: Yeah.

SISKIND: So, I completely agree with you. It has to come from everywhere.

JAFFER: Yeah. And it's going to start early. I mean, K-12, kids are using these tools today, they've got to learn how to be skeptical and how to – how to sort of push back with the AI.

I guess I shouldn't say this publicly. But I share a – I share ChatGPT account with my son, so I can watch all the questions he's asking, he can see all what I'm asking, and it is awesome. It is wild to watch. Because I see him – I see him argue and ask questions. And it's clear that, like, somebody – I don't know if it's the public school system in Arlington, or us, or whatever, but somebody has taught him how to argue with the AI and get to the truth. And it's awesome.

And if it can happen to him and he's 16, it can happen every kid in America. But you have – you as a parent and as a family, and as educators, you cannot be scared of these tools. You cannot take them off the table.

If you're in a law school like I am, you can't be like, “Oh, we'll just – we'll just ignore that whole AI revolution.” You've got to adopt. You've got to embrace. And you've got to teach your students, your children, K-12, all the way through to my parents – you know what I mean – my mom. You've got to teach them how to ask the questions the right way to get to the truth.

SISKIND: Yes.

JAFFER: And that critical thinking is the thing – ultimately, everyone's always, “Oh, you know, AI is going to take over the world and they're going to replace all these jobs.” The reality is that, if we teach our kids and our families, and our – and our adults right, this is actually an enhanced human creativity.

It's going to dramatically uplevel people. It's going to be a leveler, not an underminer. And I think there's a real opportunity here, but critical thinking and pressing, and educating yourself.

By the way, AI – think about all the – all the education that the internet put on our fingertips. AI gives you easy colloquial access to all that information. You can retrain people in jobs. You can – you can gain a huge amount of learning and uplevel people that didn't have access to good schools, that didn't have access to college education.

This is a dramatic game changer for the average American, including – the average person around the globe – including people who don't have access to a lot of these things. I think there's a huge opportunity here.

SISKIND: That's great. And I know – I know you're a real China expert, but I didn't know you were kind of modeling that kind of surveillance at home with your own…

(LAUGHTER)

JAFFER: There you go. Yup.

SISKIND: Do we have…

JAFFER: It is – it is like the CCP in the house.

SISKIND: Do we have time for one more question, perhaps, or we're good? OK.

BROOKS: Hi, Mary Brooks. I'm also a professional researcher.

So, my question for the researchers, as we are trying to figure out how to integrate new technology into our workflows without completely giving over control to the LLMs – kind of one of the following two questions.

One, are there any tricks or tips that you have found to be really, really effective, that kind of turbocharge your work without, again, giving over too much control to the LLM? Or do you have a fun anecdote example from your research that was, like a woah, aha – like, amazing moment where you really found kind of like a shocking thing that you wanted to share?

SISKIND: I can say, as I mentioned, *Al Jazeera* was just masterful at optimizing their content for AI. And so, the next step in my research journey here is to look at Common Crawl and the presence of different outlets in training data. I think that's the next step to analyze.

JAFFER: The new SEO?

SISKIND: Yeah, yeah.

BODNAR: Yeah. Are you talking about, like, future research or in my own research how I use AI?

BROOKS: In your own research.

BODNAR: Simple things. I don't want it to think for me, but if I have a task that could be easily automated, like, for example, I'm looking at an information operation, I'm pulling dozens and dozens of posts, I want to know what the key word is in each of those, or what country was mentioned in each of those, throw in ChatGPT and – or Claude and see what it says. It just makes it faster. I don't feel like I lose anything from that process.

JAFFER: My favorite technique is, I'll take – I'll take output from one LLM and put it into another and say, “Critique this. This is what – this is what Gemini told me, critique it, give me your own view. Is this right? Is this wrong?” And I'll go back and forth. I really want to build an app that does that for me in real-time. So, I've been – I'm going to have my son work on that thing and get that done. That way we all can buy it.

BODNAR: I will also say super technical things. It's been helpful for me, because I don't come from a technical background. So, if I want to look at like the page source of the domain, or something like that, I can plug it in and ask questions about that too, which will give me a good lead to then go ask an expert to verify it for me.

SISKIND: That's really helpful.

Well, thank you so much, everyone. Thank you to our panelists. We really appreciate your time and for sharing your expertise.

(APPLAUSE)

SISKIND: Oh, forgot to mention. For more information on FDD and the latest analysis on these issues, we encourage you to visit FDD.org, and we hope to see you again.

END