



Webinar Transcript

Trust in Information in the Age of AI

October 6, 2026

 Alan Shark Associate Professor George Mason University	 Meghan Herbst Senior Research Editor WIRED Magazine
 Richard Coffin Chief of Research and Advocacy USAFacts	 Renée Cummings Data Science Professor of Practice, University of Virginia
 NATIONAL ACADEMY OF PUBLIC ADMINISTRATION	 Joel Gurin, moderator President and Founder, Center for Open Data Enterprise (CODE)

On October 6, 2026, the [National Academy of Public Administration](#) (NAPA), the [Center for Open Data Enterprise](#) (CODE), and the [Bridge Alliance](#) co-sponsored a public webinar on ***Trust in Information in the Age of AI*** with the panelists shown above. The webinar was described as follows:

This webinar will explore how AI is disrupting trust in information, how and why it can be misleading, and how safeguards, awareness, and regulation could help make it more accurate, unbiased, and useful.

The one-hour webinar recording can be viewed [here](#). CODE has also provided the following transcript, lightly edited for clarity and flow.

Alan Shark: We've got a great panel discussion today, and I'm excited that you're joining us. We're going to be talking about AI trust in the world in which we live. It is a changing narrative every moment of every day, and it's incumbent upon us to keep up on it, which is a challenge unto itself.

I'm Alan Shark, and I am a co-chair of NAPA's Standing Panel on Technology Leadership. I share that responsibility with Theresa Pardo, who is also among all of us here. Our group has done several things this last year around AI, data, and trust, and all the dimensions of that.

Now I'd like to introduce our moderator, Joel Gurin. Joel is the president of the Center for Open Data Enterprise, better known to many of us as CODE, which he founded in 2015. CODE is a nonprofit organization based in Washington, D.C, whose mission is to harness the power of open and shared data for the public good. Before starting CODE, Joel worked in journalism, government, and the nonprofit sector, including leadership positions at Consumer Reports and the Federal Communications Commission. He is also a NAPA fellow.

Joel Gurin: Alan, thank you so much.. And great thanks to all of our panelists and to everybody on this webinar. I told somebody earlier today that I was going to be moderating a webinar on trust in information in the age of AI. And they asked, Is there any? That's actually the question we are asking ourselves. We know that AI is disrupting trust and information in a number of ways, that it can be misleading, that it can be problematic, and yet in other ways it can also be very helpful and open up new kinds of information and insights. We want to talk not only about the problems today, but about how safeguards, awareness, and regulation could help make AI more accurate, unbiased, and useful. I'm looking forward to a great discussion.

I want to note that this webinar is the last of a three-part series on restoring trust in government information that has been supported by the National Academy of Public Administration and the Bridge Alliance. Those of us at CODE are very happy and honored to be able to do this series and bring some insight to these important issues. Now let me introduce the panelists in alphabetical order.

Richard Coffin is the Chief Research and Advocacy Officer at USA Facts, which is a nonpartisan civic initiative aimed at empowering Americans with facts by making government data more accessible, understandable, and usable. He was the organization's first employee in 2015, guided its initial research launch and product vision, and he oversees the organization's research, marketing, and communication efforts, and is responsible for government outreach and partnerships. Richard, welcome.

Professor Renée Cummings is an AI, data, and tech ethicist who works at the intersection of technology, power, and society. She is the first data activist in residence at the University of Virginia School of Data Science, where she was named Professor of Practice in Data Science. She's also a senior fellow at both the Brookings Institution and at the think tank All Tech is Human, and she's a member of leadership groups in data and AI at the World Economic Forum. Welcome and thank you for being here.

Meghan Herbst is a senior research editor and contributing writer at *WIRED*. She earned a master's degree in journalism from UC Berkeley and reported local news for the *Press Democrat* and *Sonoma Magazine*. Before becoming a journalist, Meghan served as a weather forecaster in the US Air Force.

And then Alan, in addition to being our MC to introduce us all, is also going to serve as a panelist in this discussion. *Dr. Alan Shark* is a former executive director of the Public Technology Institute and an associate professor at the Schar School of Policy and Government at George Mason University, where he also serves as a faculty member at the Center for Advancing Human-Machine Partnership. He is a NAPA fellow, co-chair of the Standing Panel on Technology Leadership, which is hosting this webinar, and also hosts the podcast sharkbytes.net. And he's the author of more than a dozen books on technology leadership and management.

This is really a perfect panel to have this discussion today. Welcome to all of you.

I want to begin the conversation by talking about the nature of the problem we're trying to solve. We could have a whole panel with very technical folks talking about the nature of AI hallucinations and what causes them, but we felt that would probably get us a little deeper into the weeds than most of you would want to hear. So I did what many of us do when we want a quick explanation of something: I went to Gemini, and I asked it, Can you explain concisely why AI gives inaccurate results so often? Here's what it said:

Large language models are fundamentally predictive pattern finders designed to produce plausible text rather than verified facts. Because they lack true real-world grounding, they do not "know" what is true, only what statistically fits the problem. When gaps occur in their training data, their drive to generate a complete answer often overrides accuracy, causing them to confidently invent plausible details. Even as their reasoning improves, this underlying design means they can still build logically sound arguments around entirely false facts.

Not bad for a chatbot, and a little self-reflective, I think. So without going deep into the technology, we know that there are real inherent flaws in the ways these models work. We know that the accuracy issue is almost built in. And we're also going to talk about issues of bias and ethics that are in many ways part and parcel of how AI works. We're going to begin by exploring some of the issues and then move into what could be some remedies and solutions in the latter half of the webinar.

Richard, I'm going to ask you to lead us off. Your organization, USAFacts, has conducted some recent research on how Americans are using AI and how much they do or don't trust it. And you've simultaneously looked at the question of how much people trust government information and how that level of trust in government is changing. Can you please give us an overview of what your findings tell us about trust in information in the age of AI?

Richard Coffin: Happy to. At USAFacts, our entire mission is to make government data more accessible to the public and make it easier to use. And as part of that, for several years we've been running a poll that we call the [State of the Facts](#), where we ask people what makes something fact in their mind. We've been doing this with the AP and NORC for the past seven years or so. This year we added a module about AI and had a number of interesting findings. The headline is that when we're taking data from government to the public – something that is incredibly important for democracy - every single part of that pipeline is seeing challenges right now.

At the source, people are trusting information from government less this year than they were even two years ago across every domain we asked about. The highest level of trust we saw, where people say they have very or high levels of trust in information about an issue, came on inflation. Even on that issue, only 18 percent said they had high trust in government information,

and that was down about 8 percent from last year. Then, 16 percent had high trust in information on health care, 15 percent had high trust in information on jobs and unemployment. And those were both down about 10 percentage points since our last poll two years ago. So at the source, information has become a little less trustworthy.

When you go to the intermediaries of this information, like the media and nonprofits, the organizations that got the most trust were public media, NPR and PBS. But they only got 27 percent saying they had high levels of trust. And the lowest was businesses, which got 10 percent trust from the public in terms of being an intermediary for information.

There's a big gap between what people use and what people trust. We saw, for example, the highest number of people using any source for information about government was people saying they use social media. About 42 percent said they use it to get information about government, but only 10 percent trust it. That's about a 30-point gap, which is a huge issue.

And then, to what we're talking about today, AI compounds all of this. Only 9 percent are highly confident in the information they get from AI. 34 percent think AI makes it harder to know if information from government is true. And 42 percent find it hard to verify information from AI. So, in all these different things, AI actually compounds this issue of trust in intermediaries and trust in government information.

A lot of people are still undecided, and that may be a bright spot. There's a lot of room to still get people to trust AI. But there are challenges across the pipeline as a whole.

Joel Gurin: Thank you, Richard. It's a really interesting and insightful survey. And I should mention that [USAFacts consistently ranks as one of the most trustworthy sources](#) of information in America.

Meghan, let's turn to you next. You have a lot of insight into what the issues are with AI and its accuracy. You recently wrote a piece for *WIRED*, one of my favorite magazines, called "[Prove Me Wrong](#)," where you compared your own knowledge as an expert fact checker with what AI chatbots can and can't do. For that article, you reviewed a lot of studies that tried to measure the accuracy of different AI models for different tasks, and you also looked into what AI said about whether or not it was accurate. Do you want to summarize for us what some of your insights were?

Meghan Herbst: Yes, absolutely. And thank you, Richard, for leading us off. I actually use USAFacts quite a lot in my job. But yes, I wrote an article for *WIRED* a few months ago after my editor asked if I might want to write something for the magazine on being a fact checker. So I dove into the research, of which there is a lot. I found about 17,000 papers on arXiv over the last decade or so dealing with AI and accuracy. And digging down into it, what the numbers show is that about 45 to 60 percent of the time, the major LLMs that people use every day, Claude, Gemini, ChatGPT, are inaccurate.

This dovetailed with what I was experiencing on my own as a fact checker. A lot of times we're looking up secondary sources and trying to go to other news articles and things like that to verify something that's been verified before. Google started doing something a few years ago called AI Overviews, where it would offer you what it thought was the accurate answer whether you wanted it or not. I was finding about 70 percent of the time it was correct, but 30 percent or more of the time I couldn't rely on it. And as a fact checker, that doesn't really save me any time in my

process, and it can lead to a lot of issues for a fact checker or even someone who's just a member of the public relying on it for accurate information.

As an unscientific but fun test, I took the fact-checking test that I was given before starting at *WIRED* and gave it to a few of the LLMs. I think I did Grok, Gemini, ChatGPT, and Claude. All of them failed the fact-checking test pretty abysmally. I don't think I would ever hire them as a fact-checker.

That's not to say that they don't have any use. They can be useful especially in locating government documents that are difficult to find. Some of the government websites can be a little arduous to get through, and their databases are a little bit difficult to navigate. So it can definitely be useful in those contexts. But on the whole I've stuck pretty close to my processes pre-AI.

Joel Gurin: Yes, AI will not put you and your fellow fact checkers out of work today. You also mentioned in your article that there was a survey of computer scientists who work on this issue, and more than half of them said that they don't think this is a problem that's going to be fixed anytime soon.

Meghan Herbst: Yes, they did say that. It is a feature that is baked into the way that these models work, that there is always going to be a chance that they will get the wrong information, and that they will compound on that wrongness by trying to offer even more wrong information, and sort of get stuck in a loop. That's something I've experienced many times myself, and I'm sure others have as well.

Joel Gurin: You also mentioned that they are not particularly self-aware.

Meghan Herbst: Right. I asked ChatGPT how often models like it hallucinated, and it told me they hallucinated 1 to 2 percent of the time. And when I went to check its source, it didn't exist. So that was an interesting answer.

Joel Gurin: Thank you. Renée, I'd like to turn to you next and get into another dimension of the problem. We've been talking so far about the fact that AI results can just be factually wrong. Your work has focused on what may be an even more insidious problem in some ways, which is the way that AI, the data it uses, and the algorithms themselves can lead to biased results that then lead to biased or unfair decisions. I know you've looked at this phenomenon across a wide range of public policy and social areas. I wonder if you could please talk about this phenomenon of algorithmic bias and the impacts it can have. And how do we move from bias to ethical AI and algorithmic justice?

Renée Cummings: Thank you so very much for that question. And let me say it's an absolute honor to be on the panel.

My background is as a criminologist and criminal psychologist. I say that I spend my time in the classroom, the courtroom, the newsroom, and the boardroom. So I've got an opportunity to explore AI from every dimension.

Criminology is one of the first places that we woke up to the question of machine bias. The US criminal justice system has been using algorithms for more than three decades now. In 2016, [ProPublica published a report](#) that looked at machine bias in sentencing. And that report spoke about these zombie predictions that were being deployed, and the ways in which the criminal

justice system was outsourcing our critical decision-making to risk assessment tools and their algorithms. We know that algorithms assign, they reject, they decide, they recommend, they assume.

Data is the lifeblood of this technology, and all we have spoken about so far on the conversation really are questions around data. When we think about algorithmic decision-making, and when we think about algorithmic bias, we're thinking about systemic challenges based on the history of the data that we are using. Where did we get this data from? How did we gather this data? What were the research questions that we asked?

I always say that data does not have the luxury of historic amnesia, and data also has a very, very long memory. So when we think about algorithms, and we think about the power of algorithms, we think about the fact that data has been the world's greatest geopolitical game changer pretty much in the history of the world.

Hence the reason we have some of the most powerful nations in the world now fighting over this question of data, AI, and the future of this technology. When we think about algorithms and we think about the kinds of challenges we see, we know they're linked to the kind of data that's being used. Questions around data, questions around design, and questions around decision-making compile to amplify bias.

When we have repeated bias based on data and based on these decisions, we're now seeing something called data trauma. These algorithms are really undermining legacies or changing trajectories in real time. We need to think about questions around redress or the ability to contest these algorithms or correct your data, or the inability to correct your data, and sometimes the lack of due process in the ways in which these systems are being built. There are challenges around hiring, around sentencing, around loans, around creditworthiness. How do we even define our questions around risks?

Algorithmic bias is the basis of some of the biggest questions we're seeing at the moment, questions around the governance of this technology, questions around the safety of this technology. When you have a technology like AI, beyond the fact that it is a tool, AI is a power structure. It's a system that's now being used to decide how the world moves forward and who moves forward and at what cost. When we think about questions from data centers to questions around data ownership or questions around data governance or data sovereignty, these are all critical questions about how we are building algorithms, how we're deploying algorithms, and the kinds of decisions algorithms are making or not making.

We're sort of on a data precipice at the moment, when we think about the future of data and technology and the kinds of long-term impacts algorithms have on individuals, community, and society. Bias can range from predictive analytics and policing to the ways in which police are deploying algorithms now on the streets, or the way we think about our healthcare, our insurance, or education, or the ways in which we're just thinking, How do we coexist with this technology as we try to build a future?

Joel Gurin: Thank you so much, Renée. That's a deep analysis of how these issues are embedded in the entire information pipeline. It's not like you can just tweak an algorithm and fix it. It's the data itself, how it's collected and used, and the power structure around data and its analysis.

Alan, we're now going to segue into talking about what can be done. You've studied technology and policy for decades, with a special focus on AI in recent years. I'm curious what your thoughts are and what kinds of policies could address these issues we've been talking about here. What could the AI companies themselves do to make their results more trustworthy, more unbiased, and also to help users know what's accurate and what's unbiased versus what may be biased or false?

Alan Shark: There are so many ways to answer that, but let me just say that I'm a little bit more optimistic than my fellow panelists here. I find far fewer hallucinations year by year, and I've been a heavy AI user the last four years. It's getting far more accurate. I think it is going out of its way to try and at least show us sources. Now, I have to check sources like everybody here, because there are times when it will make up a source, and that's still very, very troublesome. But it still helps me get to where I want to go a lot faster than if I didn't have that tool. It's a lot faster than Google search of the past.

The thing that worries me is not just the accuracy of the data, but the accuracy of the summation of data. Because sometimes AI will take certain pieces of data, put it together and give a very interesting argument that looks very defensible. And it's based on clay. It's not defensible.

What that means for us, especially those of us in the university setting, is that we've got to teach our students not to be over reliant on AI. We've got to set the alarm bells and give them good examples of what to watch out for - how you can use it as a useful tool, but not overuse it. And to make sure they know that what you're about to see may look right - it may look better than anything you might write yourself - and it could be 100 percent wrong.

I think AI companies have gotten the message. They have to, since there's so much competition. We'll see how long that lasts; I believe there's going to be a lot of consolidation in the not too distant future. But right now, they're hearing us, and they are trying to make improvements. I'm not an apologist for them, but I'm one who finds more good than not, as long as we have our eyes open, and that we don't just take any output for granted.

Joel Gurin: Thank you, Alan. We are trying to balance what the risks of AI are versus what the benefits can be, which the whole country is talking about at the moment on a lot of different levels.

I'd like to shift and talk about the role of government in all this from a specific perspective. I do want to acknowledge that there's a tremendous, very important debate right now about whether AI could pose an existential risk to humanity in different ways. But we're not really talking about that in this webinar. We're talking here just about the role of AI in trustworthy information and unbiased analysis and unbiased guidance.

So from that perspective, I would like to talk about best practices might be for government information in the age of AI. Richard, I'm thinking back to what you found about how people are losing trust in both the information and in the AI that brings it to them. And Meghan, one problem that we've talked about before this webinar is that government agencies can also be fooled by AI if they use it to produce their own documents. You've said that as a fact checker, you need to be very careful: You can't just go to a federal source any more and assume that it's completely accurate, because they may have used AI to produce what they publish. I wonder if you could give us an example and what your advice would be to government agencies to avoid putting out misleading information themselves.

Meghan Herbst: Sure. This is something that I have come into contact with over the past couple of years, especially as a fact checker. It used to be easier to trust all of the information, or at least most of it, that was put out in official reports by the government. It's been more difficult lately. Last year there was a report released by the Department of Health and Human Services that was found to have been produced largely by AI and have sources that were hallucinated and didn't exist.

This is part and parcel with what we're seeing on the fact-checking team when we're reaching out to the government to verify claims and stories. We often get met with indifference or a lack of answers, whereas before we could count on receiving that verification. When we have both unreliable documents and the department's not willing to speak, it becomes a much more difficult environment to fact check and trust information coming from the government. We just have to get multiple sources, more than we did before. So that has definitely been a challenge. I would say that as far as the government is capable, utilizing more of their own brainpower and a little less of the LLMs would be helpful.

Joel Gurin: Fair enough. Richard, I also wanted to talk about some of the work that USAFacts is doing along these lines. You all are working on the improvement of government information in the context of AI in a couple of ways. You're campaigning to make federal data more AI ready. And you're also using a process to improve federal data yourselves at USAFacts and make it available through what are now called MCP servers, which are a very effective technical way to make it easy for people to use the data. I wonder if you could talk about those efforts, why you all have decided they are so important, and where you hope they lead.

Richard Coffin: Happy to. To what Meghan said, we have found similarly that the brand of government is struggling. The data, though, is still pretty great. And to what Renée said about the power of data: I mentioned the downward slide in trust we found in that poll that we put out. But the one thing we haven't seen change over the past 10 years is people's high level of trust in data writ large. People really have not lost trust in it. That creates both a lot of opportunities and some danger if the data is bad.

We're doing our best to make the best of government data more accessible to the public, and we have a two-pronged approach.

Our first thing is to fill the gaps between what government will do by itself, what these AI companies will do by themselves, and the public need to consume information from government through AI tools. As a philanthropically driven company, we can actually experiment and create things like MCP servers and ChatGPT plugins and Claude connectors. We can actually see if we can bring government data to people through these tools and if it helps.

Then, as the second thing, we are taking those findings from what works and what doesn't back to government and working with them to make the data more AI ready. If this super-authoritative source of information is still going to matter to people, can it get through people's natural ways of asking questions through LLMs and come out accurately on the other side? We spend a lot of time trying to make sure that the context people get is the context you would get from an expert who understands the data and not just something that AI decided is the answer.

For example, there are multiple statistics about the number of jobs in the country that you can find on the BLS website, but only one of them is the one that we actually talk about every

month. Can you give people the right answer or explain why it's different from other numbers? That's what we're trying to do.

If government data is really accessible through these tools, it will help the public understand information in our democracy, which is really important. It's valuable to members of Congress who can actually use it to legislate, and it's valuable to government officials who can use it to govern.

Joel Gurin: Thank you, Richard. Alan, you've written recently about the new European government rules on watermarking, which look like they might be adopted in the US as well. I wonder if you could talk about what those rules are, what they're intended to do, and if they could have a global benefit in making AI more trustworthy.

Alan Shark: The quick answer is maybe. You know, Europe has always been a group that likes to regulate. They're the ones who had the first Data Act. They're the ones who had the first AI Act. And now, as of August of this year, they came up with this digital watermark.

I think it's interesting. The reason they passed that was the concern over deepfakes and a lot of harmful things that are being spread around the internet. And their feeling was if we could find a way to recognize this early on and alert the public, we're doing two things. One, we're reestablishing trust by at least making the effort. And two, we're actually flagging these people who are violating basic rules of decency and the like.

The question is, Will this occur in the United States? It already is. I've queried my best friends, ChatGPT, Claude, you know, the gang, and my brain trust sometimes, and all of them are saying that one way or the other, they're probably going to be doing it. Some of our U.S. companies, the big six, have already said they are using watermarking or will be using it. They're trying to make sure that they're compliant with their customer base in Europe and elsewhere across the globe.

The problem is that as soon as one company comes up with something that may be a solution, you have people coming up with ideas about how to thwart it. Just like license plate readers, you can go online and find these so-called clever ways to disguise your license plates so that they're not recognized. I think it's going to be a battle of technology. But I think the companies that we deal with are going to need to embrace it.

The fear of embarrassment can be a deterrent. Let's assume you're a professor and you are using AI in ways that probably are not ethical in terms of plagiarism or whatever. Just the fear alone that there are now detection tools that are better than before might be able to check you.

There are a lot of nuances that have to be worked out. How much of whatever they're writing is AI? And how do you draw the line? Is it grammar? Is it rewriting a sentence? Or is it totally just taking what AI gives you? But I think it's a step in the right direction. We need to know if something's deepfake or not.

Joel Gurin: Thank you, Alan. Renée, I'd like to come back to you now and return to this deep issue of bias, both in data and in algorithms, and the need for a more ethical framework for data, algorithmic analysis, and the conclusions we draw. Are there safeguards that can help protect the public from biased data or biased results from AI? Anything ranging from policies that the companies might adopt to government regulation to prospects for litigation. What do you see as some of the possible correctives to the risk of bias in these systems?

Renée Cummings: Thank you so very much. Of course, it's the question of data governance, which is so critical. Data is about real-time decision-making. It's about insight, oversight, foresight, and hindsight, right? Those are the critical questions that we try to answer with data. We know the biggest questions around data are those of fairness, accuracy, accountability, and trustworthiness. These are the things that we're trying to figure out.

We were speaking about the many AI tools that we've been using and looking at questions around sources and citations. The other challenge is the kind of information that you're getting and the limitation of the information. It's beyond the sources: It's the diversity of data that the models are being trained on and the diversity of information that you're actually receiving.

You may access the technology and you may get a particular response, but whether or not that response is the most dynamic or the most diverse or the most broad, those are the challenges. These technologies have historically been trained on Western data, so you may not be getting the kind of intellectual diversity that you may get if you choose different and more distinctive and dynamic approaches. That's something that's very critical.

When we're thinking about AI and the use of this technology, we always come back to data. And when we think about data, we know there are regulations and rules as to how data needs to be used. The GDPR in 2016, the General Data Protection Regulation coming out of the EU, set the standard as to how we're supposed to be thinking about data, using data, processing data, and building with data. Large language models have now presented an extraordinary challenge to that.

In the U.S., some states have been doing exceptionally well when it comes to the ways in which data should be used and not be used - questions around data bias and algorithmic bias. In California recently, there was a battery of legislation around surveillance and algorithmic bias and how we need to use this technology. New York City has long deployed legislation that deals with hiring tools and how they should be used and the questions around bias there. Vermont and Colorado and Texas have also been thinking about algorithms and questions around data.

But I want to get back to this question about data governance, which is so critical to how we think about data and how we use data. How do we audit the kinds of decisions the data continues to make? We have seen some major challenges in healthcare, in child protection, in child welfare, in housing, in hiring. We have seen data do spectacular things, and we have seen data do some very dangerous things. And as a critical data scientist and a professor of data science, I know that data is not only solving some of the world's greatest challenges, but data is also opening some significant wounds.

I always say that there is no raw data. All data has been cooked in some way, form, or fashion. Every data set that we've used has been used somewhere before, has been used to make a decision before. It's critical to bring a very sophisticated understanding to data. It's more than just the numbers. There's this perception of objectivity and accuracy: it's math, it's science, it's data. You know that the data is always right. But I can tell you from being a criminologist and dealing with policing data for a very long time - we always say that you can interrogate data long enough, and it will confess to pretty much anything. And that's what data does.

So when we think about data, we've got to bring a very critical thinking lens. We've got to understand the social structure, the political structure of data. We've also got to think about this inability sometimes to correct the data, to seek redress, to contest these decisions that are

being made. And the work that brought me into this space of data and AI, which is due process, and the questions of justice, fairness, and equity.

I love data. I spend all my day thinking about data, how to make data better and more accurate, and how to make better decisions. I'm also committed to public interest technology. How do we use data in the interest of the public? How do we build tools that empower the public? And that's why I say it's more than data and AI literacy. It's about building agency. It's about empowering people to understand that this is just not simply data - that data is power. That's why we're seeing this global grab for AI.

This question of data is so much bigger than sources and citations and all the great things that we can do with it. We truly have to understand the decision-making power of data and the kinds of ways in which data has now changed the world. It has reimagined just about every business model, from an entrepreneurship perspective and a competitive advantage perspective.

We're seeing some great things. Alan, when it comes to AI, I am very optimistic. I am committed to accelerating intelligence, and that is the work that I do, the building of large language models and working with companies to ensure that the inherent risks are as limited as possible. But there are people who are being impacted by data in real time, whose lives are being changed, who are unable to contest this very strong and powerful system, and we've got to balance the two.

Data is the future, and we need more and more of it. But we've also got to understand this really extraordinary power data has to decide on access, opportunities, futures, and legacies as well.

Joel Gurin: Thank you so much, Renée. It's interesting. When we launched the Center for Open Data Enterprise back in 2015, people were already talking a lot about the power of data. What's happened with AI is we just added an accelerant to all of that, because AI can be used to scrape and pull in so much data, because it can analyze it so quickly, and because it can make data available to so many people rightly, wrongly, or in some confused way. Like anything powerful, it has to be treated with care and respect.

I also want to flag that you mentioned a number of states that are taking action, and I think that's really important to underline. We tend to think so much about federal regulation and the role of the federal government in data, AI, and everything related. But the ability to take action at a local level is really worth noting, because it's an empowering concept for people who want to take this on.

I see we have some great questions in chat. From Matt Rumsey, here's a question primarily for Meghan and Richard (and I'm paraphrasing a bit). You all are in the business of getting good information out to the public. Trust in information has been trending down for years. It's harder than ever to sift thoughtfully produced content out of all the slop. How do you grapple with the ways that AI is supercharging the sheer volume of content available online on any given question or topic?

Richard Coffin: Thanks for the question, Matt. There are two things I want to throw out there. First, I want to be super clear about what we see in our numbers with AI. As I said earlier, there are more people that distrust it than trust it, and more people that are scared about the future of it than are excited about the future of it. However, the majority is actually neutral or undecided. I think the people that have strong opinions have been the loudest so far, but I actually think there's a huge opportunity right now to sway a bunch of people into the good column for AI. If

we can make it easier for them to access good data and authoritative sources, I think it'll make it easier for people to both get good information and trust the tools that they're using.

Joel Gurin: So hopefully, good information can drive out bad. Thank you. Meghan, thoughts on that same question: Given the volume of information, how do you find what's worthwhile?

Meghan Herbst: I'll often try to work in two modes, one where I'm utilizing some LLMs and one where I am not, when I am just using a standard search engine. I do notice that there's a huge proliferation of information that looks like it's just been churned out of a content mill, and I'm sure others have experienced that too. I would say that what's most important is to be very militant about media literacy and try to find the closest to primary sources that you can. We all know there are some platforms and websites that are more trustworthy than others, like encyclopedias and some news sources, but the most trustworthy are the best primary sources that we can get, including first person sources, research, and data itself. So just try to get down to the kernel as often as you can.

Joel Gurin: Yes, and insisting that AI supply its sources.

We have time for one more question from the chat. This question from Rob Alexander is for the university faculty on this call, so Alan and Renée, that would be the two of you. Rob says he would love to see what you're putting in front of your public administration. Are there guidelines that you propose for your students about how to use AI and data responsibly?

Alan Shark: I encourage the use of AI, but at the same time, I'm very strict about how it's used when they do papers. They've got to signify what prompt they used and how they used it. We do examples of policy summations. And we also try to make everybody become better fact checkers in a world of slop. We are overwhelmed: We're now into zettabytes of data every single year. I mean, we've never had this much stuff. It's getting more difficult to navigate what's real and what's not. We've got to get back to the discipline of fact-checking ourselves because we can't always rely on other people.

Joel Gurin: Thank you. Rob has just posted to clarify his question and Renée, I'll ask you to answer now from this new perspective. Rob, I now understand that you're not asking about AI use in the classroom, but its use when your students become professional public administrators. What do you teach them about how they should use these tools in their future professional life?

Renée Cummings: My students are data scientists, so I lecture and I train data scientists into being ethical. They're actually building the large language models that we're using.

But I also do a lot of training for public administrators and government decision makers. And I say the same thing over and over. Bring a requisite and sophisticated level of due diligence to anything that you are doing with data. Check, check, double check, triple check.

Understand that you have a requisite level of duty of care if you're using a large language model or any kind of data tool or data decision making system. Understand their liabilities, and understand that you have a duty to warn, and you also have a duty to protect. Warn about the limitations, and protect people's privacy, protect people's civil and human rights, and of course, ensure that these algorithms are not deploying any kind of bias. Always think of the question of due process, given the power of data decision-making. Always think about equity, inclusion, and fairness and accuracy. Ensure your models can always be audited. That is so critical.

And also remember to bring an internal sort of ethical vigilance to everything that you do, because data is a changing world. It's a fast world, and you don't ever want to make a decision with data that you're going to regret. Remember the speed, the scale, and the kinds of ways in which data and algorithms have the ability in the blink of an eye to deny access, deny opportunity, and undermine legacies. So think critically.

Joel Gurin: Thank you very much. Those are very important best practices.

Now we're going to end with a lightning round. I'm going to go around clockwise and ask you each, What is one thing that you hope people listening to this will take away as a really important point? Just one or two brief sentences.

Alan Shark: I think we have to become fact-checkers ourselves. We have to empower people and train them and help them and guide them in learning how to be able to discern what's before them, because it's overwhelming.

Richard Coffin: The pipeline of information from government to the people is facing big challenges. AI has the chance to amplify those challenges, but also a phenomenal opportunity to reset and make incredible fixes to them, too.

Meghan Herbst: I would just say we should remember that LLMs are statistical inference systems. They lack their own intentionality. We are the brain that we're bringing to those systems and we need to remain vigilant about what we are taking away from them.

Renée Cummings: It's about agency. Don't give up your agency to any technology. Empower yourself about how you're using this technology. And of course, always empower others as we all take this tech journey into the future.

Joel Gurin: What a fantastic panel. Thank you all so much for being here to share your insights and thanks to everybody on the webinar who joined. Have a great rest of your day.