

## ISOJ 2023: Day 1

### Panel: How can journalism incorporate AI, including generative tools like ChatGPT and Bard, to improve production and distribution of news?

---

Chair: [Marc Lavallee](#), director, Technology, Product and Strategy/Journalism, **Knight Foundation**

- [Jeremy Gilbert](#), Knight Chair, Digital Media Strategy, **Northwestern University**
  - [Sam Han](#), director, AI/ML and Zeus Technology, **The Washington Post**
  - [Aimee Rinehart](#), program manager, Local News & AI, **The Associated Press**
  - [Sisi Wej](#), editor-in-chief, **The Markup**
- 

**Marc Lavallee** ...of realization, and he quoted someone that he had been talking with, who said in reference to generative AI, “this is going to change everything about how we do everything.” Now, that might feel a little hyperbolic. I think we've been hearing that about the metaverse and about blockchain and some of these other areas of emerging technology for years now. As journalists I think our backs get up a little bit whenever we hear that kind of stuff, but I think what we are experiencing and witnessing right now is there's a lot of questioning. It's this period of examination. I think we're simultaneously trying to really understand the capabilities that have emerged in the past handful of months and what we do about those and how we take advantage of those, while also really trying to make sure that we are poised for what feels like another wave upon wave of progress over the next couple of years as well. So I feel really grateful to be able to participate in this session this morning with such great panelists, and Mallery did a great job introducing, but just to put faces to names and bodies to bios, I guess, we're just going to go down the line very quickly so everybody is active here.

**Aimee Rinehart** My name is Aimee Reinhart. I work for The Associated Press on the local news AI initiative, which is Knight funded initiative. The initiative is meant to shore up interest, awareness and maybe adoption of AI tools in local newsrooms because we were seeing a widening gap between technology, and we worried that that would mean more bad things at the end for local newsrooms. I was a digital originator at the *New York Times* online. At the time there were two job positions; you could either be a producer or a night editor. Now there's like three or four departments that support all the technology that's happening at the *New York Times*. So I've been through it before. This is very similar time. The jokes are similar. The resistance is similar. And so that's how I know we're in a different place than when we were talking about Bitcoin.

**Sam Han** My name is Sam Han. As the engineer director responsible for AI at *The Washington Post*, I bring 20 years or more of experience as an AI practitioner. My AI journey started right on this campus. I had masters work here focusing on qualitative reasoning — it's a field of an A.I. Then after a few years industrial experience, I completed my PhD work at the University of Minnesota, focusing on text classification, machine learning and high-performance computing. Throughout my career, I led an AI team successfully in many different industries, including startup, e-commerce, hospitality and news media. I'm really excited about the AI potential in journalism, and I'm grateful that

you guys invited me for the panel and the symposium. I'm looking forward to having a lot of discussions on this topic throughout the symposium.

**Jeremy Gilbert** Hello, I'm Jeremy Gilbert. I'm the Knight chair for Digital Media Strategy at the Medill School of Journalism at Northwestern University. I oversee the Knight Lab, where we have a community of technologists, designers and journalists working on problems of how we can apply technology and its potential to journalism. I had the chance at *The Washington Post* to work with Sam on some natural language generation tools. We've been doing work both in the Knight Lab at Northwestern and then for me in the industry since about 2009, and we have most recently been spending a lot of time looking at what these generative models enable us to do. So I'm very excited to talk about that.

**Sisi Wei** Everyone, my name is Sisi Wei. I feel like my mics not working. I'll just take Jeremy's. OK. Hi, everyone. I'm Sisi Wei. I'm currently the editor in chief of the Markup, which is a nonprofit investigative newsroom, and we challenge technology to serve the public good. I have worked in sort of the intersection of technology and journalism since I was a college student at Northwestern, where Jeremy was a professor of mine. But my career in journalism has always been sort of related to technology in some way. I started out as a graphics editor doing data viz at *The Washington Post*. Then I went to *ProPublica* to be a news apps developer. So I've been trying to use code to serve journalists and to tell stories for a very long time. Now at *The Markup* we really think about how do we actually use those skills to then interrogate the technology that comes out and make sure that it's doing good things for people and not perpetuating harms. So that's me.

**Marc Lavallee** We were going to have one additional panelist this morning. We did actually ask ChatGPT if it would be willing to participate in this panel and it declined, and said, "as a large language model is not appropriate for me to be interviewed live on stage at a conference about journalism." So it sends its regrets.

There are so many topics around this right now and sort of pieces of the conversation. So what we're hoping to do this morning is touch a little bit on a lot of areas. Think about this as like the overture of a musical that we're going to all be living through for the next three or four acts. So we're going to sort of talk about a variety of things. There's how we're processing the moment, you know, trying to put it in context of other things that we've seen. I think what you've heard... If you if you're adding up all the numbers here, I think we have about 100 years of experience on the stage here. And also looking at who has the opportunity to innovate in the space over the next couple of years, as well as getting into a lot of the questions about what this means for the way we work, what it means for jobs, what it means for ethics, what it means for misinformation. So action packed adventure ride we've got going on.

To start us off... Aimee, to start with you and the work you've been doing in particular in the past six months or so and being kind of 411 for hundreds of organizations trying to sort of grapple with this... I'm curious about sort of how you're seeing the feelings emerge, the sentiments, the questions and so on, and sort of how you've been processing that and trying to provide answers to the field.

**Aimee Rinehart** Yeah, I personally would like to see things slow down just a tad. After GPT four was released, I saw trending on Twitter #GPTfive, and I was like, slow down. So OpenAI, you need to take a break and give us all a break, so we can catch up. I think what we get the most are questions like, what's everybody else doing? You know, everybody wants to know what everybody else is doing, what is *AP* doing... The internal guidance at

*AP* is going to be very different than what we might give to a very small local newsroom, just because *AP* has some resources. A local newsroom with three to five people might not have those resources, so it varies. But I'm really optimistic for local news to take advantage of this moment, and I think we're going to see some amazingly creative projects in the next year or two, and then it's going to get awful — we'll talk about that later.

But we, as a litmus, we wanted to do a reset for the industry, and so we developed a webinar on March 1. Usually our webinars, 150 people sign up and maybe 75 show and it's all nice. This time it was 2,500 people who signed up and it was bananas. It really got the attention of people at *AP* as well, internally, because it was like, "is this just Aimee's big idea or is this something larger?" And it turns out, (it was) something larger — also my idea. But if you go to the next slide because I forgot to grab the clicker. Oh, wait don't do that yet. We did a survey of people before the event and we wanted to ask to get a read of the room, like who's really using this? So if we do an informal survey here, how many people here have experimented with generative AI, and we're not taking names, so go ahead raise your hands loud and proud. That's it. All right. OK. That's better. OK. So the others, how many people think AI can take on repetitive tasks? OK. We have some selling to do here. OK.

Then let's go to the next slide so we can see how people perform. Oh, I've got the clicker, look how that works. OK. So, 60% experimented with generative AI; this room, I think not quite 68% and confident to take on repetitive tasks. This was true in a survey that we gave to 200 U.S. local news leaders took our survey and they agreed almost with similar numbers that they were confident AI could take on repetitive tasks. Once we had that number, we're like, OK, so we have people who are going to be willing to listen to what we have to say about AI being able to be helpful and not take jobs away. The other thing that because this will dovetail into something we'll talk about later is how many here are concerned about AI increasing the speed of mis and disinformation. Yeah, me too. So with this, these are the survey results: about 500 people took the survey, 63% are concerned about the increasing spread of misinformation. I used to work at First Draft, which trains journalists how to identify, verify and responsibly report on mis and disinformation, and I'm concerned, too.

So this is how newsrooms are using ChatGPT specifically. So they've confessed to us, they're generating summaries, creating outlines, keyword lists, adding hashtags, headline brainstorming and rapid construction of HTML, which I think is going to be a big boom for small newsrooms who used to have to wait for the college student to return to update the website. This is *Lost Coast Outpost*. It is a local newsroom in Eureka, California, and Hank Simms — I was delighted to encounter through another training — he is the lead technologist and the editor for it, and he likes to experiment. If you go to his website under news gizmos, he's got loads of fun experiments. This one is agendizer. He scrapes public meetings — agendas — and then he summarizes it. And what's fun about this is he's also added a component here, you could read it as a sea shanty or Scooby Doo or a pirate. You know, it's Hank's website, he gets to do what he wants to with it. So but this is a fun component, and I like to see fun added into technology whenever you can. Another experiment, as soon as DALL-E came out last July, he was out front using it. I said, "Hank, aren't you worried about putting local artists out of work?" He said, "We've never hired a local artist ever." So he's like, "This cost me a nickel to make, it took about 30 seconds to do, and we have art that gets prioritized by social media. So it's kind of a win win." His agendizer platform, on the previous slide, that costs about \$10 a month for him to run.

So ChatGPT gets most of the headlines, but OpenAI has other examples, and to me it's the functionality behind ChatGPT 3, 4, and soon to be released 5, that really excites me. Because if you go to their OpenAI's examples page, you'll just see this whole list of things that are possible. And to me I thought of right away, Parsons Structured Data. That has been one of the biggest obstacles to creating something like a universal police blotter that other newsrooms can plug into. It's really great at taking unstructured data and making sense of it. Whereas before, when *AP* pioneered in 2014 natural language generation of earnings reports... you know, that is incredibly structured data. There's no deviation, and you get your results in a very systematic, uncomplicated way. At the time it was really impressive. That's nine years ago. Now we're talking about grabbing information basically from the air and making sense of it in a sheet. So.

**Marc Lavalley** Thank you, and so as you work with all of these news organizations and you have the Hanks of the world who are just out there — the acronym from the tech community is JFDI, “just \*\*\* do it.” You’ve got people at that end of the spectrum. You have folks who are either not paying attention to this or actively avoiding trying these things out. In the past handful of months, have you sort of like landed on, at least for now, a space in between those two extremes that should be kind of the center of gravity for how small organizations in particular, or a single motivated individual like a Hank, could be doing this kind of work where they should sort of land for the moment?

**Aimee Rinehart** Well, I mean, the guidance is really to experiment with it and to see how far it can go for your newsroom, and it really depends on the size of the newsroom. Larger newsrooms... like *AP* is taking a cautious posture right now because it affects every single department of *AP*; we're an international operation and there's got to be a lot of thought put into that. Hank, his website, his rules — he gets to experiment. And so I'm looking forward to meeting more Hanks. There aren't that many Hanks out there, which is why he keeps coming up in my presentation. So if you're listening other Hanks I'm interested in the tinkers. We've heard a ton from journalism student or from comp sci students from Stanford and from University of Pennsylvania who want to help local newsrooms. And I have a lot of time for them. So also students who are scrappy, interested, we're interested in hearing more from them.

**Marc Lavalley** We'll have an open call for Hanks at the reception tonight. So so Sam, as someone who is and has been in a large newsroom for a while, I'm curious about particularly this moment at *The Washington Post* and sort of how folks are grappling with it there.

**Sam Han** Sure. *The Post* acknowledges the destructive nature of the generative AI, and its challenges and opportunities it presents. We are curious to understand what the future newsroom looks like, where AI agent can help reporters gathering information, analyzing it, providing suggestions. And we also want to explore... we are exploring, what does storytelling mean in a world where users can ask questions or having a conversation with an AI chatbot for information. What does subscription and advertising business mean in that world? While we are trying to work to answer these questions, we also want to encourage innovation using AI; such that we can position ourselves as the leader in using AI for storytelling, personalization and potentially content creation. So that's kind of the effort we are putting in. And then while we are doing all this, we want to make sure that we keep the journalistic integrity and standards in mind. And as an an engineering AI leader of the organization, I also sense the change of paradigm, where every software developer is becoming an AI developer or majority of the code is written by AI.

I saw some evidence of this change in our hackathon last month. So we had an AI hackathon open to software engineers, product managers, analyst and all. My group machine learning team did not actively participate in the hackathon. Our team provided advisory kind of support. Now, looking at the output of the projects, it's amazing. Just like software engineers, analyst product managers could generate a POC that's not imaginable. Like if I were like several years back, I would have invested like two data scientists, two data engineer or a few software engineers to accomplish it in like one quarter. I could see a rough product in like one week at the hackathon. So that's a huge paradigm shift. I think we are adept to kind of acknowledge that change, and that will prepare ourselves for that big change as well.

**Marc Lavallee** And you have some slides here to that that sort of go into some of the work that The Post has done, going back to I think this is the 2016 election cycle. And so I think as you go through this —sort of the historical lens on AI generated content and then also this this hackathon — a question to see this that I'm curious about is with these new capabilities and a much larger number of people being able to create these kinds of tools more quickly and simply than was the case in the past, are folks coming up with basically the same product ideas much faster or are they coming up with different things?

**Sam Han** I think it's different and fast. I think both. I see a variety of different ideas. Like, one of them was in recipe kind of area: a question AI could answer would be how can you change certain ingredients if you are vegan on the regular kind of menu? And if you live in the UK with a different weight, can you translate from pounds to kilograms or grams? Changing that, you know; all that is possible. And like a data scientist might not (think about) but a product person was out there to come up with those ideas. So it was very interesting, very interesting.

**Marc Lavallee** That's great, and so I'm curious about the pre-history of this antique robot here.

**Sam Han** Haha. So the Heliograf project, we put this project in production about seven years ago, and Jeremy was the product newsroom owner and I was the technical owner of the project. So the goal of Heliograf was trying to generate news in real time. As we get the real time data, can we generate articles automatically and cover everything? For example, in Olympics, whenever the major medal decisions are available, we'd like to have coverage right away for all the sports there. And similarly, in the election, we want to cover every election — Congressional, Senate, State Governor and all — in real time as we get the election results update from AP. So that was the main goal, which a journalist would not want to do, right? They can spend more effort in more creative writing, and this is more like a raw task that a machine can do for them.

So this is kind of the architecture, listening into the real time data from AP for elections and IOC for Olympics, and observing what events would be worth reporting and generate the article and distribute to multiple channels automatically. So here's an example for Olympics. So it provided total medal counts update automatically. Another example you'll see is as soon as the Women's Golf Medal is awarded, we wrote an article right away. We did this for all events. Here's an example we used for the election. So for this governor's race; before the poll closes, we had information about that race. And then whatever you see on the bottom, the highlighted portion is where the machine generated content and the rest was prepared by journalist. As the results come in during the election night, we updated that first part automatically, and then as the race is called, we put the winners there and finalized the article. So that's how we implemented the product. And we had a

few data scientist, a few software engineers, a few newsroom folks, and all worked together for at least a quarter to prepare this product.

So this is the hackathon POC we had last month, and the goal of this particular project was we want to provide conversational interface to users about election, about particular race, about particular candidates. So here's Sarah who is about 24 years old. She recently moved from California to Pennsylvania. She's not really into politics but wants to know more about elections in California and Pennsylvania. There's a demo. So she's exploring the California state election site, and she is looking through and she's interested in Senate candidate Alex Padilla. This is the existing interface we have. Now with this chat bot, you will see a different kind of option, and as she scrolls down near that U.S. Senate race, there's a chat bot interface and it suggests a few questions, but she can enter/ask questions. She wants to know more about Alex Padilla, and our chat bot pulls up information about her. Then she asks for information about Senate race in general. This has nothing to do with a particular race, but why do we have two Senate races in the same state at the same time? The answer is coming from ChatGPT. It's not our data. It's the understanding of election system by the large language model, and it provides answers like that. Then she wants to know more about the Pennsylvania governor race, and she can ask about it, and we pulled that information from our database, but the interface shows the results as a conversational kind of format, and it links to the particular race site and it shows all the information. Now, the user asks questions about soccer game, which is out of context, and the model recognized that it is out of context and saying we don't have an answer. And you see that response, and then the next question is about the election. It just asks who is the winner in upcoming election? And the model does not have enough context, and says "I cannot answer that question," which was also an important component of this POC.

So as you can see that this POC we can deliver the product within a week with a few software engineers. The reason why they could do this is because ChatGPT provided a few automations. So, for example, given a question, understanding the intent has been a very difficult task. That's one of the reasons why we couldn't open up that Heliograf with a free form questions or contents. But now with ChatGPT, it can categorize whether the question is about election in general or is it looking for a particular election result or asking questions about candidate, and it can route the question to different category. So in addition to that, once ChatGPT determines it's about the particular election result or a candidate, it can formulate queries automatically. We provide our schema to the system and with that same question, ChatGPT generates query for us so that we don't have to code that particular query. So using that we can surface that result and feedback to the ChatGPT and ChatGPT generates the answers for the user. So then all generic election information like why we have two Senate races in the same year in the stat, ChatGPT has enough information, knows enough about the election and it can surface that answer right off the bat. So as you can see, it shows great potential of using technology for the newsroom and the election.

**Marc Lavallee** Which is fascinating, especially the idea that they can come together in a week. One thing I want to I want to tease out of this to make sure I have it right is if you go to the regular ChatGPT website and you ask these questions, it will either not answer it or it will fabricate a potential answer, right? Hallucinate a senator in California, which may become reality soon. What you're doing in that routing that you described is actually providing the data for an authoritative real time answer coming from *The Washington Post* that is then crafted into natural language, right? So it's kind of the best of both worlds approach when that chat bot is actually something that you're building using these

technologies as opposed to it just being something that's out there on the on the open internet, right?

**Sam Han** Exactly. That's what we did. I think we took the best part from both worlds.

**Marc Lavalley** Yeah. Which I think anyone who's had an Alexa or Google Home in their house for many years, the idea that they will actually understand what it is that you're trying to do, can't come soon enough. I think to a lot of these things... Jeremy, I'm curious from your perspective around having been there and actively engaged in these prior waves of innovation around smart speakers and the shift to mobile and many other things that litter our past and nightmares. How do you put this moment into context alongside a lot of those other waves of progress and innovation for journalism?

**Jeremy Gilbert** I think that... I think that I'm going to switch microphones with Sisi. I think that we fall into the same trap every time, and we are absolutely in the same trap right now, which is that what we have to compare this technology with is the way we have been doing things before. So the temptation is to say, can this thing — a large language model like GPT — do what we did by hand or what we did with simpler computers or what we did with more rules-based AI systems. And almost every time we go through a technology change like this, it turns out the answer is not can we do the thing we did before the same or better, but rather what are the different things we can do.

So if we think back to when we first encountered smartphones, and especially for those of you involved in journalism education in this room, how many people either thought or were under pressure to create backpack or mobile journalists who could do everything that another kind of journalist in broadcast or in audio broadcast or in text-based journalism do just from a phone. It turns out that smartphones were as revolutionary as we thought they were going to be in that moment, but the answer wasn't that everyone who had a smartphone would do every kind of journalism every time they left their office to do reporting. But rather that there were different ways we would eventually use smartphones, social media, recording, etc. that didn't necessarily mean everyone was a multimedia journalist, but it did change the way you were a journalist.

I guess what I would say right now is we don't necessarily know how we are going to use these tools. We absolutely need to be experimenting, but that every time we try to say, "Is GPT a better reporter than I am? Is GPT a better writer than I am?" It's a false comparison that just doesn't actually really matter. As we start to explore what these tools and systems enable us to do, we're actually going to have much more sophisticated conversations soon, as these things become normal because we will start to change the way we behave; that's much more important.

So when Sam and I were working on Heliograf, we tried very hard not to say, "How can we make Heliograf write a story just like a *Washington Post* journalist?" Because that wasn't nearly as important. We had great *Washington Post* journalists who were doing human writing. On the other hand, *Washington Post* journalists didn't have the time to cover all the stories we wanted them to cover. It would be impractical every 30 seconds to update the vote count in 500 stories. So we were able to tell different kinds of stories in different ways that freed up the humans to do the stories they should do. And that's what we should be talking about as we look at GPT. Where does it enable us to do more of the things that we cannot do, serve more of the needs of our audience that we cannot serve, rather than is this going to replace the journalists in the local newsroom that Aimee is talking to? I think

sometimes we fall into this trap of asking the wrong question instead of trying to explore: what does the technology enable us to do?

**Marc Lavalley** There's a great headline in *The Onion* years ago when they launched their mobile site. They said, "*The Onion* just got smaller and harder to read," which if you come at it from a newspaper mentality is exactly how we thought about it, right? How do we jam these stories down to mobile? You know, but it did over time, we learned how to actually take advantage of the form factor and the literal mobility of people being able to stay informed on the go. It has recrafted our industry. So to your point, I think that there's that that opportunity as well.

Sisi, from your perspective in leading a news organization, seeing that people can take advantage of these technologies in certain ways today, but that, you know, I think there's also a perception that it's going to get continuously easier over the next year or so. Like, how do you think about budgeting time and advocating for experimentation versus the having to feed the beast of coverage and things like that?

**Sisi Wei** Yeah, that's a great question, and we do have some benefits in that because *The Markup* doesn't cover day to day tech news, there's a little bit less pressure. But at the same time, we're thinking like, OK, what can we actually interrogate about something like ChatGPT? Is there something to be found in the topic itself? And then, similarly, how can we use it to investigate other things, right? And what could it actually enable us to do that maybe was not possible or just resource wise to heavy to do something like that before?

I think for us... I mean, overall, I would say it's sort of this classic journalism advice, which is to always stay curious, right? To basically encourage people in the course of their work as journalists to not limit themselves to any sort of boxes that we've been told about and to see in every instance in which you're making a decision, is there something here that you that you could do if you were thinking about the biggest version of your story? Right. And is there something that whether it's image generation (or) large language models can actually help you do that you just never thought was possible? The point on this to is we've gone through these types of changes within journalism many times. Jeremy talked about the phone. I feel like if you think about when TV journalism first got started, everyone was like, is it going to eat everything right? Is all other journalists going to be invaluable? That obviously hasn't turned out to be the case. In my own personal experience too, when data journalism and data analysis became something that the industry really embraced and started investing in... There's a lot of things that sort of happened where everyone was trying to figure out, is all journalism going to have to be data heavy in some way? Does everybody have to learn how to code? That was a conversation we had many times over. While I still support journalists definitely learning these skills, the shift here is always it's not like we then use data to do the journalism we did before. To your point, we're now just discovering completely new things, enabling journalists to prove things that they could never prove and that they had to rely on experts for. So what could we do? And I guess I'm towing a line here of trying not to give too many things away. But I will say, I think in the realm of capacity, could you use generative AI to allow you to create massive amounts of ways at the same time to interrogate another piece of technology? That's the types of things that sort of journalistically we're thinking about. But then when it comes to writing articles, we like to joke a lot that our articles are artisanal baked and written. But again, it's because we don't do the day-to-day coverage that would make a ton of sense for *The Washington Post* in every locale. That kind of stuff. Yeah.



**Marc Lavallee** Super helpful. Thank you. So one thing I want to add to the conversation here, we've been talking mostly about words so far, and that's a lot of the recent attention, but before ChatGPT arrived on the scene, there were a couple of months last year it was all about these image generation tools like DALL-E and Midjourney and so on. Jeremy, I think you have some slides around images and audio, too. All formats are up for grabs right now.

**Jeremy Gilbert** Absolutely. So if we put those slides back up. Anyway, these are two examples of projects that we're doing in the Knight Lab that combine work. So this first one comes especially from Zach Wise, who is a professor in the Knight Lab working with some student teams; what his team set out to do is not unlike what *Lost Coast Outpost* is doing, which is to say, can we use large language models and have them understand the content of a story? So we rather than writing our own stories, we're picking ones that were already published. But this from *Grist*, a startup that covers climate, and we first used GPT3 to try to understand, what is the content of this particular article. So we put it into OpenAI's playground prompt. It's very similar to ChatGPT, it works with the same underlying technology. We prompted it by saying, "OK, write us a sentence that visually describes a picture that would illustrate this particular article." So then it comes up with a sentence: "a picture of a dark blue ocean with a melting iceberg in the background illustrating..." So what we have done is we've actually shaped the particular prompt that's generated over time by a series of experiments that said, well, what is going to get us a better visual output? Then we take that prompt from OpenAI and we put it in Stable Diffusion. Stable Diffusion is a system just like the large language models around text that say, if I have a word, what is the next most likely word to appear. A diffusion model just says I have a pixel. I'm trying to make an iceberg is the next pixel I add to it more or less like an iceberg. So you can see 16 different images generated.

What we have been doing in addition to just throwing it into Diffusion, we've actually been training Diffusion. So in addition to the large language model that we inherited, we said here are examples of journalistically editorial oriented images. So can we help the large language model actually output the kinds of images that we want? We've looked at this for a number of different stories, a fossil flower that's trapped in amber or whether protein bars are actually good and what happens when the FAA has software disruptions. What you can see — especially if I had shown you the early images — is if you take the untrained AI tuned large language model, then you get back things that are visual but not necessarily journalistic in nature. So I think a big part of the work that we're trying to do here is not just can we turn out editorial images, but can we as journalists train these systems to work in the way that we work? To uphold the values that we value?

There are a lot of questions that we're asking ourselves, most of which initially are in the realm of can we. So can we turn a text article into AI art? Sure. Can we train an AI to make editorial art? Well, that's actually a little bit more difficult. It's a little more complicated. Can we have that AI art actually be representative of the subjects that we're covering? Well, that's the most difficult still. But we also have to ask these questions like, should we do these things? If we let AI make art for journalistic stories, what are the ground rules? What are the guardrails that we need to have in place so that we know, especially if humans are less involved — see what happened with *CNET* and calculating compound interest — can we ensure that the outputs that we make, that we put out to the public don't harm our credibility?

Similarly, we are looking at what's happening with audio. So this is a project that the Knight Lab is working on with Michigan Public Radio and the *Associated Press* — supported by

the Knight Foundation. Thank you. The whole point behind Minutes is that Michigan Public Radio just discovered that I think about 70 communities around the state regularly publish council meetings on YouTube. Just put the whole thing out there. But that alone isn't actually journalism, it's just access to information. So we need to take those YouTube streams and we need to turn them into something journalistic. For example, this is Williamston; it's a community in Michigan, and here is the council meeting minutes that they published. We're using a different tool from OpenAI called Whisper, and the whole idea is can we improve on the speed and the accuracy of transcription, because that's a really important thing for journalists.

Here's the council meeting. I'm not going to make you sit through it. It would be painful, but eventually YouTube will generate a transcript. The transcript is OK; It's a good benchmark to compare it to, but it doesn't happen instantaneously. It happens when, according to time and according to cost, YouTube has time to make it, and that's not necessarily when a journalist should be doing their work when YouTube happens to make the transcript available. So this is one of these classic comparisons where smaller is better. The error rate for even YouTube when measured against human transcription is 2.961. So what we're sort of saying is, there are three errors per paragraph roughly, and then Whisper itself has different levels of sophistication. So the base model, if you don't do anything, has a much lower error rate, but higher than if we train it in different ways. So again, the whole idea here is can the lab figure out what is the rate of error that we can live with, especially if we compare — on the left, you see what a human transcription has done with on the right what we're doing out of Whisper. So we're constantly looking at it and trying to evaluate are we getting better at transcribing in the kinds of ways that a journalist would need?

Then we have the ability to tune these models, and this is, I think, maybe the most important thing for us to think about as people who create tools and explore these tools. We always can turn the knobs in different ways to make the models work better or worse at different factors. So if we take the CPU computationally, if we increase the power of the CPU, it's going to cost more, but everything is going to work faster. If we increase the GPU, it actually works much faster, but it costs a lot more and we can play with other things like the size of the model. So how many examples has the system seen before we ask it to do a particular task? The more examples it has, the slower it's going to transcribe, but the better the output will be. So what we're doing in all these cases, we're trying to dial in to what is the right mix where the tools become more usable in a very specific, very journalistic function.

**Marc Lavalley** Thank you. I think one of the things on this gets into I think a lot of the work that you've been doing as well is you can do some of these things off the shelf. Right? You can use ChatGPT on the web for free or you can pay them \$20 — things like that. At the complete other end of the spectrum, you can essentially build these large language models from scratch. *Bloomberg* has announced BloombergGPT because they have more money than anyone and they can do that. Then there's this whole set of gradations in between as well. Whisper is something you can run on your computer for free; you can also use it in the cloud and so on. So I'm curious about from your perspective both today, you can map the capacity of a news organization to what's available to them — free, low cost, high cost — but then what are things that you are looking for to exist over the next year to take more of these bespoke capabilities and make them more widely available to a larger set of organizations?

**Aimee Rinehart** Well, I think we're in a good moment for local newsrooms to take advantage of it, especially the tinkerers — the people who are somewhat fearless and are willing to break a few things to just experiment. The experiment we did with Michigan Radio, Whisper blew everything else out of the water in terms of accuracy on transcription. I feel like that has never been more accessible to local newsrooms, which is why I'm really excited for them... the 3-to-5-person newsroom, the ones who are really struggling sharing basic information in the community that comes down to city hall meetings like Jeremy showed you. It also includes things like police blotter items, school lunch menus, TV listings, I mean, that's what keeps people subscribing. It's often not the journalism, unfortunately. So we need to build in more systems to help with the flow of information. Frankly, it's a bargain right now. There is an arms race between these six companies that are dealing with it and they're lowering the prices.

So I think again, for the next year or two, local newsrooms stand to really come out the winners on this because they're able to experiment. I work for *AP*. It's got 3,000 staff members. We're global. We're going to take our time to make sure that everything we do is sound journalistically and supports our members and customers. That said, local newsrooms, again, like Hank, it's up to him to decide. So I think the real experimentation is going to be out in the field, and probably not from the big newsrooms unless you are lucky enough to have Sam Han on staff. But unfortunately, Sam Han is a unicorn in most newsrooms. There's not a lot of newsrooms that can afford that, and Sam probably has lots of other options, too, frankly.

**Marc Lavalley** I do imagine you're one of the most popular people at *The Washington Post* these days.

I think on this on the cost front, yes, these things are inexpensive now — you and anyone jump in here... this is a little speculative. But, you know, years ago, I used to be able to take a Lyft or an Uber basically across the country for like seven bucks because it was subsidized by VCs, and now that is no longer the case. When we think about building systems on top of these and routines, is it more of these teaser rates that we should be concerned about or is there something that's going to make the costs continue to go down over time? So is this the right time to hop on because we'll just be able to do more, faster and cheaper with these systems over the next couple of years?

**Aimee Rinehart** Well, Sisi and I live in New York, right? So they warned us. They said if you take Uber now, you're going to put the taxis out of business... I do think they do suppress it at the beginning and right now it's a good time to experiment because it will be cheap. But I do think there'll be a dependency that comes on and prices will increase, but Sam also knows.

**Sam Han** Yes. So I think right now it's more like nuclear arms race, right? So everyone just wants to build the largest, the best, regardless of the cost. But behind the scenes there's also an effort to make this more economical. How do you make training cheaper such that we don't spend \$500,000 to have one large model, right? But if the competition really works out and we have many players offering that option, then I think we're in a good place. However, if a few or handful of big techs just survive this competition and those are the only options we have, then we might have to think differently. We might have tech tie in. Right now, we work with the news aggregators; they have great control over the news consumption, and they have that model also and have full control and we are really tied to that particular model or technology. Then there will be challenges to the journalism that we want to have. So with that, I think we might have to look at options of using open source

language models or having fine-tuned local versions of it just in case we are getting into that kind of situation.

**Marc Lavallee** And there's some opportunity here as well, because the story about that particular school board on that particular day may be important to one community, but the pattern is something that exists in communities across the U.S. So through some kind of collective effort, we can potentially build some of these systems for ourselves. I think, you know, everyone in this room has lived through one of the biggest lessons of the past decade of what happens when you build your audience and your relationships on somebody else's platform. Then one rando buys it and does a bunch of stuff to it. So I think being in control of our own destiny is both an opportunity, but also an obligation for this to work for us.

So this is a lot about how we can take advantage of these technologies and build our way to a better place that is more comprehensive coverage and things like that. I'm curious also, there are the people who can take advantage of these tools — reporters, editors, basically everybody in a news organization — and as with prior waves of technology, there's a lot of concern and consternation about what does this mean for jobs. I don't think UBI is coming fast enough to take care of us, so I'm curious Sisi from your perspective in both leading a newsroom and then also reporting on this space, how are you thinking about and possibly fielding questions from staffers about what this what this means for their jobs and the way that they work.

**Sisi Wei** Yeah. I have two thoughts on this one. Well, first, a preamble that luckily none of my staffers have come to ask me if generative AI is going to take their jobs, but it's probably because our method is always using technology to investigate other technology. So that's that. But I will say this, which is every single job that I have had in journalism for the last, I think, 10 to 15 years has been a job that didn't exist 10 to 15 years before. Right. So the question is always like there will be rote tasks that sort of Sam has talked about a lot that humans don't want to do anyway, that we have seen every advancement in technology sort of like take care of a little bit.

You can think about this. I think a very good example from ten years ago, one of the first things I ever worked on is sort of every month we all report on the Bureau of Labor Statistics that come out. We just all do it. When I started at *The Washington Post*, it was a thing that you manually check their website, you refresh it until the number shows up, and then you put the number into your article, and you publish it. I even remember back then we were building graphics that would update monthly as well to allow you to see what happened over time, and I was like, "Can we just like we'll write some code so that my computer checks the website and refreshes it until the number shows up and puts it in the graphic?" We did that ten years ago. I feel like it's a micro example of how hard it is to predict whether jobs are actually going to be taken and how that is balanced with the jobs that are going to be created based on the new things we're going to do.

I think until we sort of see more examples of why we should be alarmed, I'm not actually that scared on the jobs front. I think new jobs are going to be created and new opportunities are going to show up much faster than like all of these things that we had like manual labor-intensive tasks being replaced. It's not like we are hiring people to do that over and over again in journalism specifically.

The other example that I'll share as well is also from ten years ago, which is when I was a student, Jeremy was a professor at Northwestern — the first time. One of the things that I

had the privilege of sort of working on only a tiny bit tangentially was something called narrative science. This was an effort to write templated articles automatically ten years ago... more than ten years ago. The use case was: there's tons of sports coverage out there, but journalists are not covering every single little league game that our kids are playing in, and parents would love to read articles about their kids little league games. So, could we create some sort of template such that depending on the difference in the score, it would write the article slightly differently and then publish it hundreds of times, every time there's a game. That was possible ten years ago. It did not become an insanely popular thing that everybody decided to use. It is now ubiquitous. ChatGPT can probably do it faster, but it's a use case that we tried out because the technology. I think this is important, which is the large language models or even the visual models are doing something just much faster than we used to. It's about that investment. But we tried a lot of these ideas because they were literally possible many, many years ago and they didn't catch on. So now it's about when it's faster, when you can experiment more, what can we actually dig out of it? Because experimentation doesn't cost us a year of 50 full time staff trying something. That's why I think going back to jobs: the more creative we are and the more we try to do different things, the more naturally those jobs will start to create themselves.

**Jeremy Gilbert** I mean, I think we also — and Sisi is absolutely right — have to acknowledge how much the way we tell stories has always been bounded by technology. I mean, when we were talking earlier this morning about *NBC News*, what is possible in broadcast television shaped how we tell the stories; what kinds of stories we tell, how long those stories are, how many characters are in each of the stories. I think we have to acknowledge that, especially if we're talking about text, we are still allowing ourselves to be bound by what made sense when we were transmitting over telegraphs. So we said, “Oh, we've got to put the least important information at the bottom.” Everybody gets the same version of the story; we have a lead and then a nut graph. This was all really because of what the technology enabled us to do. So I think we have the opportunity to say instead of, can we do it, what should we be doing? What kinds of stories do people actually want? What kinds of signals do we, especially those of us who are publishers, have in first party data that can help us know what our audience wants and serve that using these new technologies, instead of saying let's be shaped in terms of what we do because of what the technology allows.

**Aimee Rinehart** But that's a great example... so *AP* pioneered the telegraph storytelling, and you're welcome for the inverted pyramid. But yeah, that's a great way of them looking at the technology back then being like we better get it all upfront because we may lose the connection. So it really is adaptive in that way, and I think it's... yay *AP*.

**Marc Lavalley** I think this is all speaking to a massive opportunity here, which is not how many reporters are there in the U.S. and how many stories can they crank out a day and that is the capacity of news, or how many pages can be printed by 5 a.m.. It's how many stories can be told that are not told today or are or under covered, whether that's a story that is framed for an audience that is that is arguably rarely, if ever, had proper coverage, or topics that we've just never thought of as being, if not pure journalism, journalistic in the delivery. So we think about that that that sheer volume of opportunity that we may be able to make progress on today. I think Sisi to your point about job creation, the idea that there's a potentially economic model for those types of things that are the things that people maybe care about more in their daily routines. There's a tremendous amount of opportunity there. It is going to require potentially among people who have been in the industry some amount of unlearning of certain ways of doing things and certain maybe rote things that they that they take joy in. It's like I love writing that headline that way and things

like that. But Jeremy, from your perspective, I'm curious about, as you are continuously kind of training and working with the next generation of people going into industry, how is this changing the way that you think about preparing folks for a career? And are there other things that you're weighting more or less as a as a result?

**Jeremy Gilbert** Absolutely. For those of you in the academy, there are such delightful conversations like, can't we just ban ChatGPT under the honor code? And I said, "I suppose you can do whatever you want to do. However, are you going to ban spellcheck?" What are we going to do when every Google doc has a gen AI component? We can no more ban these things realistically than we can ban our students from using smartphones. It's just not going to happen. So we have to think about what are the assignments that we are genuinely worried that students would complete with GPT and should we have ever made those assignments. If a large language model with almost no useful information can do the same job as our undergraduate, then we're probably not training our undergraduates particularly well.

So among other things, one of the things we're pushing really hard on and it is not going over terribly smoothly with our faculty is that a reporter is not automatically a writer. These large language models are much better at generating things that look like writing than they are at doing meaningful reporting right now. For example, one of the things we need to train young journalists to do, people starting their careers, is to think about how can we ask questions that lead us in interesting directions and then use large language models to tell those stories in all kinds of different ways. The reporting is probably the thing that has the most value. The large language models can support the reporting, but they don't as easily mimic the reporting as they mimic the output of news stories. So we're spending a lot of time trying to ask ourselves questions about what does it really mean to be a journalist and what does it really mean to train someone to be a journalist. Our third and fourth graders still learn arithmetic, even though they are very aware that there are calculators. So I'm not saying we don't need to teach writing at all, but if we have assignments that ChatGPT does just as well as our students, those assignments are no good.

**Sisi Wei** So I will just say because I adjunct taught journalism students for a while, and ever since ChatGPT came out there's been a homework assignment that I really want to happen, but I don't teach students anymore. Which is that one of the most important skills in journalism in a newsroom today I think is critical thinking. The ability to really... it doesn't matter sort of like what it is, maybe it's something a source said (or) maybe it's how you're writing something. It doesn't really matter the topic, but that you can think critically on it. When ChatGPT came out, I was like, you know, I think the hardest thing for me as a teacher is that I can't use my time to teach critical thinking in very specific ways because there's a limit to what I can generate as the homework assignment. So here's what I mean, which is I think it's great to teach journalism students how to edit, even though their job is not going to be editor out the gate. But through the act of editing, they really interrogate every sentence, and it's a little different than fact checking, but includes fact checking. I've always just thought that journalism professors could use ChatGPT to generate articles on different topics that each individual student is interested in and then assign them to edit that piece of work into an actual good piece of journalism. This is where like 2% of it being wrong would be great because I want you to catch that. Or like how it's phrased is not quite right, and because you're interested in this topic and can do more reporting on it, you can change it. To look at it from that angle... Like I would never write 50 versions of an article for my students to check, but now you can actually do something like this.

**Marc Lavalley** That's great. So we're going to take questions in a second here. For folks in the room, if you want to line up at the at the microphones, and then folks on online, just put them in the in the chat, and it will magically appear on this on the iPad here. As folks are lining up, I guess one quick last question for you. Small, tiny topic: misinformation. So some folks are talking about this as potentially like the text apocalypse. What do we do?

**Aimee Rinehart** I mean, fight fire with fire. We have to have tools that identify artificial intelligence. That would be great, and Santiago Lyon is in the room to talk about this later, but to have some kind of footprint on images that carry over and don't get snipped out from social media platforms. It's a big problem we had. *AP* fact-check team was very busy two weeks ago with the Pope and the Donald Trump images, and those images came up fast and furious and got circulated online. Luckily for right now, just right now, the hands are a little wonky, the ears are funky. So if you take a very close look at it and you're not on one of your tiny smartphones, then you might be able to see that it is generative AI. But I think within the next year that will pretty much erase because it's evolving so quickly. So we do need more tools, better tools, and those tools have to keep up with the technology they're trying to screen for.

**Marc Lavalley** That's great. Yeah, so provenance is a big piece of it and Santiago is doing a session on that at some point. So we have one mic here, right? OK. That looks like either a real or fake Bill Adair.

**Bill Adair** It's the real Bill Adair. Thank you guys so much. This has been a great conversation. I think you've raised two really important points. I'd love to hear you elaborate on them a little bit more. One is, Jeremy, you raised the need to re-imagine form and it's something that some of us have been talking about. I'm the founder of *PolitiFact*. We said when we started this, "Hey, let's re-imagine the story form," and it hasn't happened so much. I wonder what newsrooms and educators can do to reimagine the journalistic forms with this amazing new tool. The other thing that restricts that, that you all were just talking about is culture... is the culture of what we do that inhibits that? Because we've always done it the same way, as you were saying with the telegraph. So how can we get outside the old ways of thinking?

**Jeremy Gilbert** I believe very, very strongly in human centered design. I think part of our challenge is that even human centered design research we have done around journalism and journalistic story forms is too often accepting that these are the options — we can tell a newspaper style story, we can tell a TV style story, we can tell a radio style story — instead of going out and trying to look at how news consumers, especially younger news consumers, consume information and what are the ways that we can serve those people. So, for example, even if we keep the inverted pyramid story form, which I'm not trying to attack *AP* specifically on, we can, for example, play with the linguistic complexity so we can make stories for English as a second language readers. We can make stories for a fourth grader, a 12th grader, and anything in between. We can suss out based on reading speed, for example, which version to serve someone. So we can start to say maybe the story looks to any individual the same as we would have told it, But using a generative AI, we can explore how complicated it should be or what kind of viewpoint we should take. So I think these things are possible, but they start by really studying what people want instead of saying this is what we want to give them.

**Marc Lavalley** So we've got about five minutes and five people in line, so use your time sparingly.

**Jim Anderson** Hi, my name is Jim Anderson, and I'm from Mexico. I'm co-director of our NGO called Factfile. My question is should media outlets make it visible every time they use artificial intelligence in their content? Is there a standard of when they use it... when they use this explanation for their readers? And the last question: Is there any currently discussion about at the editorial level about this and these warnings when you are using artificial intelligence?

**Sam Han** I can answer. So in our newsroom, one topic that we are really working on together and we have a guideline saying if anything is created by AI, we want to have a human in the loop, make sure that we review it before publish. Then when we publish, we want to make sure that we let readers know that they are generated by AI. So we have policies like that to enforce.

**Marc Lavallee** I think there are few efforts on this front. One I'll mention very quickly is the partnership on AI, which involves about 100 plus organizations, recently published a synthetic media framework, for creators and publishers to try to sort of start to lay out some of those guidelines — not just for journalism, but really for all media as well.

**Sisi Wie** Um, I'll add something else, too, because I think when this came out... I think it was a couple of days ago. It's not in journalism, but I thought it was a great similar aspect. In New York City, it was announced that sometime in the near future, if you were using AI to filter job applicants in the process and that a real human being would not actually see the applications until the filtered group that you had to disclose that in your job posting. So even thinking about how do we use AI in very different ways, making that disclosed. I think transparency is going to be key.

**Marc Lavallee** It's fascinating. Next question.

**Tricia Crimmins** Hi, my name is Tricia Crimmins. I'm coming from *The Daily Dot* and also Columbia University's Graduate School of Journalism. My question is off of Aimee and Jeremy's points about imaging. So I know that in my understanding of DALL-E, it pulls from all other images on the internet, and then I also know maybe a month or two months ago, there was the thing where people were putting their photos into like an AI generated thing and they were able to see themselves in different time periods, as a Viking, as a conquistador, like all these things. So, Aimee, I know that you said that Hank said that they've never hired a local artist, but it seems like inevitably the imaging is going to be drawing from other art. I know Janelle said earlier, plagiarism is plagiarism is plagiarism. So how do we square that?

**Aimee Rinehart** I mean, we'll find out, right? Getty is suing Stable Diffusion, and we will see. It will probably take a couple of years to go through the court. This is usually tried on the topic of copyright, so we'll see.

**Jeremy Gilbert** The only thing I would add is if I were talented enough to have my own artistic style, I would be thinking about training a model to at least help me do my work. So I think there is a range of degrees. Am I totally basing off of other people's work without attribution? Am I building off of my own work? I think we'll see. There are court cases that will decide some of the larger questions, but I think we'll also see artists... just like we see writers now using AI to improve their own output. We'll see artists doing the same thing.

**Sisi Wei** I'll just add one last thing, which is that the copyright office is actually paying attention now, and I think it was yesterday or the day before they announced that they



were going to go on these listening sessions to talk to people about how they use it, including artists. Copyright is such a key and live issue that at least at *The Markup*, we are not going to publish any sort of Midjourney/DALL-E created images because of that sort of unclear ethical line of is this an amalgamation of thousands of people's images all together. That said, on the flip side, when it comes to like looking for things... like it's always when the thing becomes the news. Like when you're looking at does it create biased imagery for some reason, then that would be the exception.

**Neil Chase** Neil Chase from *CalMatters*. Aimee, thank you for bringing up Hank — every question has to mention Hank. In addition to being a very cool technologist, he has created a level of journalism in a community of maybe 60,000 people with zone stories and everything that's better than most people's journalism they have in their own communities. So the technology is cool, but for a good output.

For all the other newsrooms in the world that don't have a Hank, where's the line between the tools that we will build and the function you want to have in a newsroom? If I'm running a newsroom of one or two or ten people, what am I expecting tools to provide for me, maybe that the *AP* provides or that a software vendor provides? Versus what are the people in... Jeremy, you started to get at this with the reporting skill on the street. Where's the line between what I need to do in my newsroom and what the tools I'm going to have in the next whatever five years are going to do for me.

**Aimee Rinehart** I think because it's such an exciting time, there's a race to be like we need AI everywhere and I'm not sure where. That's kind of not the great way to adopt new tools. I think you need to see where your newsroom pain points are and I and see what could be alleviated. Sometimes if you don't know where the pain points are, ask people what they really hate every day about their job. People will gripe endlessly, but if you ask them for new ideas they will be like "Eh, I don't know." So you'll get more answers with the pain points. To think about what is out there, I think, again, open-source code — the person standing right behind you is working on a lot of that with our team and with lots of others. I think if you are adept with some form of technology, you might want to start at least learning prompts to ChatGPT, to learn how to code for your website. Hank also wrote his website; it's a bespoke website, so it's a very flexible website.

**Neil Chase** Yeah. Thanks.

**Serdar Tumgoren** Hey, everybody. Sedar Tumgoren with Stanford University and *Big Local News*. Aimee, you know me too well. I'm going to ask a code question really for all of you. So actually, your friend and colleague Ernest, who some of our students are working with for a local news AI project, he just sent them a heads up about a lawsuit — that I actually hadn't really been that aware of — which is this class action lawsuit against GitHub, Microsoft, OpenAI about a particular tool called Copilot, which I was very excited to have some of our students start using. But it's given me a little bit of a pause about should we be using it to do the code generation part with the issue being that a lot of these swaths of code that it's generating are taken whole cloth out of open-source projects and have very specific licenses for giving credit back to the creators. So I guess I've been kind of thinking about well should I use it as an assistant to help them explain code to themselves or debug code, but maybe nudge them away for the moment at least about using the actual generated code? Am I being too uptight? Are you all discussing this stuff in your respective institutions? And if so, where is your thinking leading?

**Aimee Rinehart** Yeah, I was alarmed to see that because I think open source is a really great way to level the playing field for technology and technology adoption. But we talked about this at length before we started the project with KSAT in San Antonio because we were worried that maybe a student would borrow too much code from a repository and then it would be it would be somehow plagiarism. So we warned against that. And to me, I feel like ChatGPT and other technologies are great for checking code and checking your work in general, so that maybe is how I would pursue it, especially because you're at a pretty substantial organization.

**Marc Lavalley** I want to... That's a great topic; we should talk about that more. I want to get the last question in very quickly under the wire.

**Audience Member** Oh, I feel like this is a really tough last question then, because I'm really curious about how you all feel about the potential for AI to exacerbate inequality. What I mean by that is looking at some of the brands that are up on stage today — even *The Markup*, whether it's tech and technological awareness and being ahead of the game in that sense — is there a bottom up approach that doesn't wait for innovation to trickle down from the *AP* and from *The Washington Post*, and is there a way to distribute the economic and resource benefits that will come from generative AI across the news industry. So that the small publishers, the ethnic publishers, (and) the community publishers have as much of a role as all of you do in kind of building this future.

**Marc Lavalley** You're speaking straight to my heart right now. That's one of the biggest things we need to figure out. Personally, I think there's a huge opportunity here in this in this moment, partly driven by the fact that a single motivated individual can be doing a lot more, and the fact that there are these efforts — like what Serdar is doing with *Big Local News* — that are trying to create systems, build them once that allow them to work for thousands of small news organizations. So I think the collective action that we can will as an industry with that in mind, I think is a huge opportunity. In contrast to a lot of the innovation that has occurred over the past decade in visual journalism and video journalism and things like that, that inherently require teams of people in large organizations in order to be able to accomplish. Any last takers on that before Rosental gets the hook.

**Jeremy Gilbert** I mean, I think especially for people in this room, if you're going to explore the use of automation like this, we need to be helping and looking at the needs of organizations like you're talking about. If we only start with *The Washington Post* will only help *The Washington Post*, we need to start looking much more broadly.

**Sam Han** So similarly Heliograf and other projects we did in the past required really machine learning specialist or data scientists to work on now, does not require that. Any software developer can find a topic or a project and use a technology to implement certain things. So I think it's so much better now.

**Marc Lavalley** With that, thank you all so much for this, and thank you to the audience for the great questions.