26th ISOJ Beyond the Hype: The Real Impact of AI on Newsrooms

---

- Chair and presenter: Nikita Roy, founder, Newsroom Robots Lab
- Juliana Castro Varón, senior design editor of AI initiatives, The New York Times
- Uli Köppen, chief AI officer, public broadcaster Bayerischer Rundfunk (Germany)
- Tim O'Rourke, VP of editorial innovation + AI strategy, Hearst Newspapers
- Brooke Siegel, vice president of content, Yahoo News

Summer Harlow This panel is Nikita Roy, who is the founder of Newsroom Robots Lab, an AI training and advisory firm for media organizations, so let's welcome our next panel.

Nikita Roy Well, hello, everyone. Can you hear me? OK. Perfect. Welcome, everyone. I guess this is the quintessential AI and newsroom panel that has to happen at every journalism conference. Otherwise, is it even a journalism conference? I am very excited today. We're going to keep it different. It's not just about AI tools, but really focusing on the transformation that's happening as a result of this technology, what that looks like across different ways, from local news to a public broadcaster to a global news aggregator, and then also think about how do we design future news experiences with AI.

So I'm very excited to get into all of this with you. My name is Nikita Roy. I am the founder of the Newsroom Robots Lab, where we help newsrooms implement AI, and I host the Newsroom Robots podcast, where I've had the privilege of talking to a lot of people about how they're incorporating AI and journalism. A couple of the guests here on the panel today, and I'm really excited with the perspectives that we're gonna be bringing forward in this panel. Joining us is Tim O'Rourke. He is the Vice President of Editorial Innovation and AI Initiatives at Hearst Newspapers. We also have Uli Köppen, who is the Chief AI Officer at the German public broadcaster, Bayerischer Broadcasting. We have Brooke Siegel. She is the Vice President Of Content at Yahoo News. And finally, Juliana Castro, who is the Senior Design Editor of AI Initiatives at the New York Times. And so, as you can see, a lot of diverse perspectives on AI. And we're going to start off with a few presentations from everyone about the case studies, and how AI is actually being implemented in their newsroom, and then get into a bit of a discussion to talk about what does collaboration look like between product and editorial. How do we prepare for a time when AI powered platforms like ChatGPT, Perplexity, Google AI Overviews are threatening the very relationship to our audiences by summarizing our content? What does that look like? What does that mean for us?

And before we get into all of this, I just want to set the context for why this conversation is so important. Because what we're dealing with is not just another tool, it's a shift in the interface for news. It's a shift in the infrastructure for news. It's a shift for the culture of how newsrooms are working and evolving, and I'm very excited to get into all of this. In 2022, in November, when ChatGPT was launched, that was a chatbot moment. In 2023, I think, was all about chatbots. In 2024, last year, we saw a lot of experiments, a lot of progress that newsrooms were having in regards to AI. But in 2025, things are, again, picking pace a lot with something called AI agents. This technology allows AI systems to basically go and act autonomously on behalf of a user. So if you look at something like ChatGPT's operator, you can instruct it to go and browse multiple different websites and give you the summary of news sites. What happens then if it's not AI that is interacting, if it's not a user that's interacting with the news website, but AI. And how do we prepare for that huge shift that we're going to see? And at the same time, we're seeing a lot of advancements with

multimodal AI systems. So it's not just text anymore, but AI able to interact and analyze text, images, video, audio, that's everything that powers us here in our newsrooms. And so what opportunities does that bring about for us? But what are the risks when anyone can go on Google Notebook LM and create a very human sounding podcast, and the recent update that happened with Dali in ChatGPT, you're able to create very, very realistic images. I was scrolling on my LinkedIn feed today, and somebody had actually posted about a photo that they had got from Dali. I stopped for a second because I thought it was real. It was Biden and Trump playing Lego on a table. It was so realistic that for a second, I knew it was fake because I could never imagine that happening, but that's what Dali was now capable of doing. And it's not just about verification, but if somebody who was just scrolling casually on their LinkedIn feed for a second thought something like that was true, what could happen, and the possibilities for misinformation, and how do we as news organizations play a role in building trust in this AI mediated information ecosystem.

So as we get into all of these conversations today and all of the presentations that everyone over here is going to share with you all, I want you to start just thinking about more than just these tools, but what is the news information ecosystem going to look like and how should we be leading it. This is the conversation to help us understand: how do we prepare for that upcoming shift that's happening? Because these decisions in the system is what's going to design and prepare us for the future of our industry. And so with all of that, I'm excited to get started with actually just giving us a landscape of where we are right now, and what's happening across these different news ecosystems. And starting off, I'm very excited to have Tim give us an update on how AI is being implemented at Hearst at a local landscape. So Tim, over to you.

Tim O'Rourke Hey everybody, I want to start off by saying, I'm not an AI utopian, I'm a little scared of Rosental, so I want to get that off my chest right away. So at Hearst, so I work for Hearst Newspapers, otherwise known as H&P. We're a part of Hearst Corporation, and I work on a team called the DevHub that's an editorial innovation team that works with all these many newspaper brands that you see here on this screen. They range from hyper local newsrooms with just a couple people covering large communities to major metros like the San Francisco Chronicle, Houston Chronicle. Our DevHub team there is based in San Francisco, but we're spread out across the country so that we can help all these different newsrooms with different types of innovation projects.

The DevHub is the home for our editorial AI work. We also organize ourselves around three paths of work. Editorial engineering, we build the infrastructures that aren't covered by our core product and engineering team. Things like the election infrastructure, tools, templates, server maintenance, different kinds of backend processes. And then the core of what we do is in the middle there, data visualization projects. These are things like data-driven analysis. We scale great data projects from one market to our other local markets. We really try to drive kind of modern innovation, data-driven journalism across our group. And then finally, we have a set of tools and templates that we call the Click to Publish program. It allows journalists at the local level, without any coding skills, to go ahead and publish interactive journalism through this series of templates. They cover alternative formats like timelines and reader feedback loops through surveys, polls, interactive maps, interactive list experiences, so the kinds of things where reader engagement is at the core. Across all these lanes, we now, we're lucky enough to get some funding to build a small AI team. It's about four people, and we try to integrate AI into these different areas, the places that we own, but also the places that can help our journalists do more important local journalism in their communities.

This is a pretty organized-looking timeline. I can tell you it was not this organized. Two steps forward, one step back is the kind of progress that we usually see. Things could get a little messy, but we're always trying to innovate on our path to progress with our newsrooms. And Nikita mentioned in 2022, ChatGPT came out, smacked us all in the face a little bit. We got down to kind of learning about the technology, understanding how we could apply it to our local journalism network. We were able to find a couple of people in our organization, especially one engineer nearby in San Antonio who he spent his nights and weekends in coffee shops and at home learning about AI. It's not what I do on my weekends, but we were happy that he was doing it because we could draft him onto our team and really get started on this innovation work around generative AI quickly. We turned that into some internal tools and some kind of one-off projects and those did okay. So we were able to get some funding, and that led to the ability to build a small team for our local newsrooms. And then we hit what my colleague, Emilio Garcia-Ruiz in San Francisco Chronicle calls the “now what” moment, right? You got some funding, you got some people, you got all these different ideas, some of them are good, a lot of them were terrible, and they're flowing in fast. What do you do? Now what at that point? And so we kind of navigated that and figured out, okay, what's our path to really the best way that we can take advantage of this technology while respecting our standards, our ethics, and understanding that we want to innovate, but we want to innovate cautiously so we don't do damage to the brands and we don't lose trust with our audiences. And that's led us to this current period where we're enhancing and accelerating the best ideas that come out of our local newsrooms.

Just to ground us in some numbers, because this panel's about the current impact, we have about 650, 700 journalists in our network. In the past year and a half, we've been able to train about 350 of them directly on tools that we've built and other tools out there in the industry, as well as our core strategies. We've got plans in place to take that across the whole editorial group this year. We've got about 65,000 tool uses and counting. Every month use of AI tools grows. We follow the principle that a human should be in the loop, so there aren't straight publish tactics. These are things where we have human review, but they do speed up and accelerate the work that we were doing on the production side. That calculates to over 22,000 hours of human time saved. That's time that we reinvest right into our local journalism to allow our folks to do higher minded work, to allow them to get deeper into their communities, to allow them to do things that maybe they didn't have the time to do previously. And then overall, we've done about 25 big tools that we maintain and external experiences for our readers that harness AI. So we keep our pedal to metal to make sure that we're consistently learning.

We spend a good amount of time, you know, using AI to help our local reporting, not just through tools, but through bespoke analysis. Without getting into too much detail, we've done some great work on sentiment analysis, including with the governor here of the great state of Texas, Greg Abbott, and his rhetoric around the border and immigration. We've analyzed some of Trump's speeches from his inauguration to other speeches that he's given over the past few years. And then we've done a lot of work around taking big data sets and big asset sets and doing analysis on those. So our reporters have the time to pull out the learnings and then do their own work, and that's led to some really great bootstrap journalism.

We do create, maintain, and constantly update a core set of tools in our room. These are just three of them. The one on the left is we call it Producer P. It's a bad name, but it's a good product. It's an audience optimization bot, like a lot of great ones in the industry, and we've seen this really be impactful, especially for our hyper local newsrooms that don't

have audience professionals reviewing every story before they go to publication. In the middle, we have an automated quiz generator that we call MC. And the visual shows an archived content from the San Francisco Chronicle. They do an automated archived quiz that's got a human come in to check for factual accuracy, but it takes about 80% of the production work out of the building of the quiz. We also do news quizzes, sports quizzes, things that kind of engage readers at the interactive level. And then probably the one we're most proud of is Assembly, which is a public meeting monitor. Again, there's some good ones across the industry. This is ours, and we're proud of it. We're covering over 300 meetings now, and it grows every week. This tool allows whenever a public meeting that's on our scraper list gets posted, the video gets posted online, we scrape the audio out of it, create a transcript, catalog that, automatically alert the reporters or editors who cover that topic through our Slack instance, and then they can go in, download the transcript, summarize the content, and it can lead to stories. This is, again, a very hyper-local type example of how that can work. So you've got this long meeting. The San Francisco Chronicle, of course, covers like the Board of Supervisors in the city, but we don't have the resources to get to every subcommittee or in this case a public works meeting that runs three hours, 16,000 words spoken. But we had the Assembly bot set up, so it scrapes the content, produces the transcript, alerts the reporter within a few minutes of the video posting online, and that reporter can go and summarize the content. And you see this one's about 11. He picked out one that to him, knowing the city and knowing the bureaucracy, that was interesting. It wouldn't have been interesting to me, but he noticed it, right? That led to a story about how public money is being spent in the city, and if you know San Francisco, like the way that we spend money in San Francisco is a little different. So it turns into a little scoop, not a huge story, but something that kind of keeps the coverage going.

So I want to talk about some of the things that I think we've done well as a local news org. Again, for us it's about the network effect a lot. We set out early. We brought Nikita in. We talked to our legal team. We talked to our product, executive leadership, newsroom leadership, and really set the expectations early on the kinds of things that we were going to do and the kind of things we were going to stay away from, and we trained consistently on that with our newsrooms. We also focused a lot of our resources on building trust in our newsrooms, on communicating with our newsrooms, on making sure that people were trained up and that people are excited about this while still preaching, we have to be careful, we have to do this conservatively. And then we treated AI as another tool in our toolkit. On my team, the DevHub, it's about pushing the boundaries a little bit in journalism, and so we wanted to do that while being respectful to our standards.

So what could we have done better? This part of the presentation is four and a half hours. Three points, so just like other types of projects, I think we got excited about some of the technology at a couple turns, and we didn't think it all the way through of how we could get to an audience. So one example, we did really awesome, comprehensive local voter guides down to the very, very hyper-local level. We used AI to help us translate those into Spanish, brought a Spanish language editor in to fact check everything, but we didn't have a great plan of reaching those Spanish language audiences in our local communities. So we could have spent that time elsewhere. We also at the beginning especially, you mentioned the chatbot era, I don't think we did a good enough job really saying like, generative AI affords all these amazing possibilities. I think that the main user experience being a chatbot infrastructure led a lot of the ideas to come down to be just that, and I think we could have differentiated more at the start. And then this is like so incredibly obvious to me now, but we, you know, at the beginning it's like, “Oh, if we build this really cool tool that harnesses this technology, you know it's going to work, and everybody's going to love

it.” Well, some tools work better for hyper-local newsrooms. Some work better for major metros that have different types of specialists on staff. So we've learned over the course of time we really have to cater our tools and our experiences to those particular internal users and make sure that it works for the ways that they do their jobs.

Finally, last slide for me, is just three small-scale projects that we're excited about right now. In the background, we're also working to bring generative AI closer to our core product, getting personalization, getting on different types of multimodal, but for three really fun, I think useful, utility-type projects we got coming. We're doing a property tax guide for our readers here in Texas, starting in Harris County, where Houston's located, and that's got a big data component, utility component, explainer component. But also we use AI to come in and help a reader kind of lay out an outline for a script, so that they can, if they need to challenge their property tax assessment in order to save a couple thousand bucks, they got a starting point if they can't afford a lawyer. So to us, that was a good application of the technology. We're also doing some good work around aerial mapping, and satellite mapping, and using AI to help us analyze that. And then finally we've got, it's a chat interface, but we kind of built it as a throughline through our food content, it's called Chowbot. And so it's built on the back of thousands of pieces of food and restaurant content that's rotated as part of a program, and we allow users to kind of go in and discover more of the human written content that we have in our local communities. The last one just published in San Antonio last week. So again, just three examples of how at a very local level we're using AI to push our newsrooms forward.

Nikita Roy Thank you, Tim. A lot of exciting innovations happening at Hearst. When we last spoke on the podcast, back in 2023, you were just getting started with Producer P, so there is so much that has happened in the last two years and excited to get into this later on. Uli, we are excited now to, over to you, to talk about how you're building the AI strategy at the Bayerischer Broadcasting.

Uli Köppen Hello everyone, it's a great honor to be here. I'm part of German Public Broadcasting. And when Nikita told us it's about the real impact of AI, I was like, “Wow, that's a lot of ground to cover.” So let's jump in.

That's the part I think we should talk more about as journalists: How does AI change the way we report about AI? And how does AI change the teams we are building reporting on AI? We specialized very early on at German public broadcasting on algorithmic accountability reporting. And we started about eight years ago, together with Der Spiegel, we did a story on the SCHUFA algorithm. This is the main credit scoring algorithm in Germany, but it was not really transparent on how it worked, but every German citizen was scored with that. And for that, we did a data donation, and we could do some reverse engineering of the algorithm and at least have a peek into how this algorithm works. Funnily enough, today I got the news that SCHUFA is publishing their score. It took eight years that this happened, but at least now they are planning to do that, which is great. We did a lot of stories on the vaccination algorithm. During the pandemic, we did a story on AI recruiting because the assessment centers that work automatically are getting more and more. And people, for the first round of the application, are sitting in front of the screen, and they are telling, theoretically, they could say what they had for breakfast for 30 seconds. And the software promises to extract personality traits for HR. We scrutinized one of those algorithms, and we worked with an actress and she did the same speech over and over again with the same tone of voice, the same mimics, and it was relatively stable, the output. But as soon as she changed her outfit, she put on a headscarf or she put on glasses, the output changed dramatically. We did a lot of technical and statistical

tests with that, and in the end, we could tell the story that for now, this algorithm is not fit for a sensible area such as recruiting. This is a story we just published recently. It's about data brokers, and they're selling our data on the internet, and the data comes from our cell phones and the apps on our phones. And we could track people from their home to their working place, and we could even track people working for the German Secret Service and for the American Secret Service, and we did that together with Wired.

For doing such stories, we need interdisciplinary teams, and I guess this will be something we'll be covering today. We have to invite programmers into our newsrooms, and we have to invite product people, being able to do such stories. And this is how AI affects right now the way how we build teams and the way how we tell stories. And I'm always lobbying for more teams covering algorithmic accountability stories because I think it's so important to raise AI literacy with the stories we are doing.

Another very important part for us, and well, an area where we can do much more with AI is personalization. I'm from a public service broadcaster. We are not interested in click baiting or leading people into tunnels. We're interested in getting our stories out and reaching as many people as we can also with the help of AI. What we are trying to do is we try to find public service ways of personalization, and this is a product we launched one year ago. It's called Regional Update, and it's a personalized audio news briefing. It works like that. You type in your zip code, or you're getting located, and then an algorithm is gathering the news for you that happened around this place. And you can customize it, you can say, I want to have news that are not older than 24 hours, and I have two minutes to listen, and then you can get back into a lean back position, and you're getting alerted as soon as something happens in the area of your interest. What we're doing is we are taking the linear audio stream, we are segmenting it with an algorithm, and we're tagging it with a geolocation. And with that we have the infrastructure that allows us to offer a personalized audio news briefing. And before people had to accept the regionalization we offered for them. Like we said, “Okay, you are located in Franconia, which is part from Bavaria, so you're getting the Franconian news update. But now people can say, “Okay I'm located in between three different regions, I'm interested in all three of them. So please update me there.”

What's interesting for us is not only the product. We are very excited about the product, and we're getting very good user testing. But we're also interested in building up infrastructure behind the product. And this little cloud here for us, where we are storing the news bits with the right metadata is for us the insurance that we can personalize down the road how we want. Because I'm quite sure that right now we don't know how we want to personalize in two years. And we can use it platform agnostic, we can use it for our own platforms ,but we can also use it for smart speakers or whatever we want to use it for.

Workflows, you just mentioned agents Nikita, and I guess that's the next big thing how our workflows will work in a different way. Already now our workflows are AI supported. This is an AI assistant we rolled out in our newsroom last fall, and this should help our journalists to build different versions out of stories because we want our journalists to focus on investigations, and analysis, and field reporting, but we don't want our energy wasted in doing different versions out of a story, and sometimes we are doing really a lot of versions. So this AI assistant helps you to write different tones of voice for different radio stations, and it speeds up the workflow in our newsroom already now. And we are already experimenting with integrating agents. And agents is just a fancy word for stacked AI workflows. You can talk to your agent and you can say, “Give me the voiced update for our listeners on a certain radio program.” And then the agent breaks it down in different steps,

it gets out the right bits of news, it summarizes it for you, and then it can even use synthetic voices to voice it for you. So I'm sure this kind of workflow will change how we work in newsrooms. We're also using AI for community management a lot. We started very early on to use large language models for sifting through a lot of user comments, and we built an alert system for our newsroom to be able to get out the right comments that keep up the conversation and to be there on time before the user is already gone. This is very important for us. And this is very new, we are launching that right now. It's called Comment Digest, and it's a summary of all the comments that happens behind articles so that we are hoping to get a better entry point for people into this kind of conversation. You're getting the topics people are talking about, and you can jump right to the comments that are covering those topics.

AI quality, very important point, where we try to also do our part. We just published an AI that controls AI output, because we all know the phenomenon of hallucinations. And our newsroom was really laughing at me when I was saying to them, “Okay, you have to check the output of our AI agent”. And they were like, “Okay I can type faster, give me something better.” And this is part of the solution we are testing right now, and it seems to work quite well. It's an AI that tests AI output and that alerts the newsroom to the parts they should check. So this speeds up the process, and it seems to work quite well. It's open source. You can take it, and steal it, and make it part of your tool chain. No, of your Tool Chain, that's the word. It's behind the QR code, and newsrooms are already using it.

This is my favorite part, and I have to cut it short because I still have one minute. This is where I have put a lot of time in last year to make our organization AI ready. And as I've told you, we are coming out of a very beautiful bottom-up movement. We had a lot of data teams. We had an AI team already five years ago, and we are coming out of a movement where people were already working together from all the different parts of the organization. But when ChatGPT arrived, we had to combine it with a top-down movement, because we have to make overarching decisions. For that, we built up an AI board, and this AI board is responsible for prioritization. We own the ethics workflow, and we are responsible for strategy. And we are working on an AI network because German Public Broadcasting is a complicated organization. We have nine different broadcasters, we have 20,000 journalists, and we want to make sure that the experts from all over the organization can come together and collaborate. And for that, we have built up the internal network, and we have done an external network that should bridge the gap between journalism, academia, and also industry. We are meeting regularly. You can sign up for our newsletter if you're interested. It's aiformedia.network. We are also doing English-speaking events. We're doing hackathons, and we make sure that we open source our code. Thank you.

Nikita Roy Thanks, Uli. Very exciting about the whole infrastructure you've built for your newsroom, and excited to get into that. Now we have Brooke, who's going to talk about the work that's being done at Yahoo News.

Brooke Siegel There we go. Hi, I'm Brooke Siegel. I'm the VP of Content at Yahoo News. So what does that mean? Many of you might know Yahoo News, we are an aggregator. We aggregate from tons of high quality publishers in 12 international markets, many of whom I'm sure are in this room. In my role overseeing content, I'm touching all of the various distribution points where that content lives, so that's curation. Choosing which stories should appear to our audience. It's the original content that we do. So I think one of the ways that we really differentiate ourselves as an aggregator is we have original reporting along with all of our premium publishers. There's the visuals that help identify Yahoo and our brand out in the universe. We have commerce. So it's a very big

ecosystem. How big is it? We are the number one news and information site in the US. We reach 190 million users a month. That scale is always staggering to me. It's one of the things that really excited me about taking this role, which I started this past summer. So there's Yahoo News, entertainment. There's games. You might be familiar with finance, sports, mail, search. We are really a portal to the internet, and one of the ways in which we think about Yahoo is as a guide to the entire internet. There is so much content out there. How do we help our audiences find both the information that they want to know and the information that they need to know? And balancing those two things are core to the principles of what we're doing. So how big is 199 million? We reached nine out of 10 Americans. For context, only eight out of ten people put ketchup on their burger. I don't know what the rest of folks are putting on there.

So how we're thinking about AI at Yahoo News. Really, it's about how we are leveraging AI to bring more quality to our readers. We're using AI to tackle scale. We have enormous scale. The strategy is staying with the humans. And the AI, the way that we're incorporating it is to leverage it to bring more quality, premium content to our audience, and also integrating it into our newsroom, this was echoed with previous presenters, so that editors can focus on the creative, high-impact work that they're doing and passing off some of the work to AI that shouldn't necessarily be taking up as much of their time.

So I wanted to walk through just a few ways in which we are incorporating AI, starting with the hot button topic of personalization. So we have tens of thousands of articles that are coming into our pipeline every single day. How do I find the right story to serve us to our audience? AI is helping us tackle that. And we have since rolled out a new algorithm in the summer to our app. And we've seen 165, I'm looking at, 65, I'm nodding, percent increase in time spent per user just by improving the personalization features. So that's not just at the “What is the clickiest content to serve to a reader?” It's based on what you've seen before, what are you showing interest in. It's also incorporating editors into training that personalization algorithm so that we're IDing the right publishers to tell the right stories. It's about discovery, so it could be that you're in this echo chamber where you're really just seeing what is the most popular content. What our algorithm is aiming to do is personalize it to you so that you are discovering content that is unique to your specific interests. And because of the vast amount of content that we have in our network, we're able to really find that fine-tuned approach to personalization.

But it's not just about giving readers what they want, it's also about giving readers what they need to know. So at the top of our app we have what we call our “top stories” feature. So again, we're using AI to comb through thousands of stories, identify when there's a trending news topic, and then surfacing that to editors, and this is always the key, right? It's keeping editors in the loop. So editors are then reviewing that cluster of stories and choosing the right articles, the right publishers, the right angles, so that when our audience is going into that top story cluster, they're really getting an array of perspectives on a single topic that was curated for them by an editor. But we wouldn't have been able to get to that place. We couldn't have the editor sifting through those thousands of stories. So it's really about working hand in hand.

This is one of my favorite features. You've seen it on the internet. It's the key takeaways. What differentiates it on Yahoo is we are only pulling the summary from the article itself, so it's really like a movie trailer for the article. It is not giving you everything you need to know. It's teasing what is important in that story that you've chosen to read. And we find that when we do the key take aways, people engage with the story even more. So, rather than find that the summary is enough to satisfy the appetite, people are like, “Okay, this is what

I'm interested in. I'm now going to go deeper." They can use the bullets to identify where in the story they want to lean in. And those are just some of the ways in which we're starting, excuse me, those are some of ways in which we're just starting to explore how we are leveraging AI to create a premium quality experience for our audience, and I'm really excited about all to come. Thank you.

Nikita Roy Thanks, Brooke. It's very exciting to see just how much personalization is really helping people get the news they want, and the impact that's happening in the news discovery process. And finally, to end us off with the presentations, Juliana from the New York Times is going to talk about how should we be designing AI experiences?

Juliana Castro Varón Just me saying, not a chatbot. Thank you for having me. Today I'm going to start, if you'll excuse me, talking about myself and my career. I promise we'll get to AI very soon. That's me. And now, here's a list of things I have done before getting this job. I ran an independent book publisher called Cita Press. I also wrote a book about beauty. I also work for a few art museums, mainly contemporary art museums. Actually one here in Austin. My cartoons, they're all kind of quirky and imperfect, like that one right there. I have an MFA actually from here, Go Longhorns. I also, that did not help me. I still got lost at the Texas Union today. So I know what you're thinking, that looks like the profile of an unemployed person. But in fact, I work for a newspaper. I am part of the AI initiatives team at the New York Times. We're a tiny, relatively tiny, small team composed of some editors and developers. And me, I'm the only designer. And unlike many product development teams at the Times, we sit in the newsroom. That means we help people, we help journalists, tell stories. But I begin today talking about my background in art and publishing because I am interested in and in many ways devoted to words and language. And because I'm a designer and that's my background, I am very often worried, and it's my job to worry about how they're packaged, especially now that the majority of our interactions with words and language happen through screens.

So, that's why. Every time that a knowledge organization decides that the best way to repackage information and knowledge is through a chatbot, I am curious, and in some cases, I'm appalled. Now, I get it. Chatbots are useful. They have kind of this back and forth quality to them. They're clean, and given how the internet looks right now, it would make so much sense that we would want, that we could crave an interaction that is not only kind of natural language based, but that does not look like this. It is normal that readers would want that, but I have a number of problems with the fact that most of our AI-based interactions would look like this, whether it is a search bar with a summary or a chatbot. I'm not saying that's never a good fit. I'm just saying in the best of cases, the context is buried, and that's when it's present at all. But again, it makes sense.

Let's think of an example. Let's say you are, and let's kind of get outside of journalism, let's say you're an airline. Let's you're always on time. And you're perfect, and you never have any problems. But people still come to your website with questions about policies, and in the very rare occurrence that you're late, they want to know how to fix that, and it would make sense to put a chatbot instead of having many call representatives. It makes sense, and I get why people would do it. However, before large language models, chatbots were deterministic. That meant that for every question there was one possible answer given by a set of rules and a logic. Most recent AI chatbot, the ones that have us all here, are probabilistic. Sorry, I got excited. That means that for every input, they can have many outputs. And that's what makes them so exciting, but also that's why it makes them so dangerous. There's not a single answer for every question, but rather they take a guess. They take a statistical guess. In a way, chabots can make too much sense, like when Air

Canada lost against the passenger because a chatbot made up an answer that was not consistent with airline policy. And I want, I'm willing to guess, that the answer was not, “you can fly in a horse next time,” but something that made sense, maybe too much sense. Oh, I have an animation there, and I forgot. Here you go.

So again, in some cases, the chatbot may be a good fit. But if you're in a decision-making capacity, I would ask myself once or twice whether or not that's the best fit for your audience. So now, talking about audience, I want to talk about the different types of users that my team and I build products for. On one side we have, a slow clicker, on one side we have readers, and on the other side we have journalists. We build products for reporters of the New York Times. And reporters and readers of the New York times have in common that they're curious people. But they come to us maybe slightly different. Reporters may come to us to make sure that they're asking the right questions, but readers want answers. What they both want, however, is context, a way to make sense of the facts. So we've established that one of the things that makes AI so powerful and so dangerous is that it can understand meaning. But most often the value of AI happens when we are able to not generate anything new, but rather create structure out of massive datasets.

So then, what are the best cases of AI? Thank you for asking. I'm going to talk about some examples. Here are my three favorite. Semantic search, that is searching by meaning and not by keywords; multimedia search; and structured research can assist journalists using AI. And here are a few examples that we've used. Last November, my colleagues obtained videos from a group that was spreading misinformation about American elections. The problem was that these recordings, videos, and just like Zoom videos, totaled more than 500 hours. And when they received these recordings from a source, there weren't 500 hours until election night, so they needed help. They came to our team, and our team helped them find passages, and trends, and highlights. And of course before publishing, they went back to the original recordings to make sure that they were not missing any important context. That's the human oversight and review. In fact, these journalists who have covered election interference for many years were uniquely positioned to make sense of the truth. They knew whom and what to ask for, and how to make sense and how to interpret those findings for readers.

Our team believes that AI alone is pretty much worthless, at least in our profession, but over and over again, we've learned that expertise paired with AI is a very powerful combination. When the travel desk prepared to celebrate the 20th anniversary of 52 Places to Go, they came to us for help to dig through the archives. And we found that, for example, sustainable travel had been a thing from the very beginning but people, writers were using very different language throughout to talk about that same trend.

And finally I want to end today making sure that I acknowledge that we do have a bot. The Connections Bot has a feature that analyzes the most common wrong answers for each day puzzle. And they use AI to, again, guess what could be a theme that players were thinking of when they make that attempt. And yes, even those are reviewed by the games team before publication. So in acknowledgement of that one, great bot, here are a couple facts about me to close. I'm okay with bots, but only sometimes. And I'm very happy to be back in Texas for a few days. Thank you so much.

Nikita Roy Thank you, Juliana. So, no bots, no chatbots. And I think what you raised was a really important point that generative AI is a huge design challenge as well that we are facing right now. It is a revolution of the entire user experience of news and one of the things that we have to be thinking about going forward. But with all of these projects that

we had discussed over here in every single newsroom, one of the key things is how do you allocate for resources across all of these interesting ideas that are coming up in the newsroom? And how do you say “yes”? Which are the ones you should say “yes” to? And which are the ones you should say “no” to? How do you decide that? So I want to ask you, Tim, you're working in the local news ecosystem. You have a small, mighty team that's powering all of Hearst newspapers’ AI innovation. So how are you approaching what you say “yes” to, and how do you allocate those resources.

Tim O’Rourke So, we really strongly value collaboration, so we don't make these decisions, you know, off by ourselves in a small room somewhere. But we really look at, number one, because we are a local news network, you know, can whatever project we're going to do, can it actually get down to the hyper local level, to the neighborhood level, to something that gives us unique value in our communities. And also, for us, we overvalue if something is scaleable. So is the research we do, the work that we do to build something, is that something that we could take to other markets, and it could be just as successful? So those two criteria. And then we also try to, as I said in my presentation, weigh tools that help our journalists, with new experiences for our readers, as well as building processes that help accelerate our journalism. So balancing all that, and then letting a lot of people weigh in, and apply a formula, and go from there.

Nikita Roy And you speak about collaboration, so collaboration between product and editorial. How has that been evolving in order to support all of these AI initiatives?

Tim O’Rourke Yeah, for us, we're a team that kind of sits at, it's journalist developers, and we sit at this nexus of product, and engineering, and the technical, and then, of course, our newsrooms, and our first priority is to serve our newsrooms and help them innovate. So when we make these decisions, it's about bringing in different types of thinking and people from different departments and making sure that strategically, we're getting as aligned as we possibly can be while knowing that we're going to be taking some risk as we innovate.

Nikita Roy Uli, I'd love to hear from you. You’re spearheading the entire AI, and you spoke about, you want teams to be part of that collaboration. And there was a question here from the audience which was asking, how do you get buy-in for AI tools with skeptical journalists? Has there ever been an example where you got a resounding “yes” from the newsroom when you were building a tool, and how do you approach that?

Uli Köppen An interesting example is the regional update I've been showing because in the beginning we were coming to newsrooms and saying, “Okay we need metadata.” And metadata is not very sexy for journalists, like you have to type them in, and you have to make sure it works. But we came with a prototype, and they could see with the prototype that their linear audio was used in a different way, and that we can potentially reach younger audiences with that. And then we got the buy-in, and at first, they really had to type in the metadata, and then we also had to promise you don't have to do that forever. We have an AI doing that for you, but we need you as a trainer for the algorithm. So we needed this buy- in, and we got it by prototypes. And also, those prototypes can help you for the management, because if I get there and say, “I need a hell of a lot of data to build a new infrastructure for metadata.” They're like, “Well, is that really important?” But if you come with such a prototype, and you can show this is what we can do out of this metadata, then it's much easier.

Nikita Roy The importance of showing them, instead of just telling, being the key over there. Brooke, I'd love to hear from you as well, that collaboration between editorial and product, how do you bring about all of these teams to support the innovation?

Brooke Siegel First off, I'll say anything that has editors not doing tagging projects. I'm like, we're all for. So yes to the metadata. I think the shift that's been happening, I would say, to my mind in the past year, is there used to be people who specifically worked at the nexus between product and editorial. They were the liaison. They spoke both languages, diplomat, translator. I think how that's shifting now is democratizing it. So editors need to understand how to work with product and technology. It is going to be part of our jobs. It's going to be a part of the skillset, and we need to learn that language. So when we have the various projects that come up, rather than have one team that focuses on it that's always the go-to, that's also going to have just one perspective, we want the diversity of perspectives of the whole newsroom. So each project has different point people on edit with different backgrounds, whether it's data visualization, or editors, or our copy chief, standards, or all of them working closely with our product and engineers to develop those deep relationships so that it's happening across the newsroom, not just at one specific point.

Nikita Roy And Juliana, just from you as well, because the New York Times, you have your five-person AI initiatives team that is going and approaching the editorial team and building tools for them. How does that collaboration work?

Juliana Castro Varón First of all, we never build anything that people are not asking for. We experiment with many tools. What we ship to readers goes through many more filters than what we ship for journalists, because it's the journalist's job to check that things are right, but we don't put the burden of verification on the readers. So usually we pitch internally, and then we each do a little bit of research as to whether or not what we are pitching internally actually would be useful. We ask people what is taking a lot of time, and if they could have a magic thing that fixes something, what is it that it would fix? And then we start from there. We start from actual use cases. Sometimes we build small prototypes, kind of like almost custom for one project. And then we see that structure of research works for similar kind of workflows in different desks or something. But we work very directly with journalists on stories. And then if we realize that kind of like that pattern of work works for many more journalists, then we try to scale it. But we start small.

Nikita Roy Start small and scale from there. And I think one key question with all of the audience facing products that we were talking about, Juliana is against the chatbots, but at the end of the day, how do we position our messaging around being transparent around the use of AI that's happening either internally or audience facing? We have one of those as the audience question, so I'd love to pass that along. Uli, I'd like to hear from you from the regional update, how you were putting that messaging across.

Uli Köppen Yeah, I'm coming out of data journalism, and there we have a great tradition in telling how we are getting to our stories. We always publish a methodology page with our investigations, and we just kept doing that also for the product part. So if you're interested, you can dive in, always. We always make plain how we're using technology, and we have a rule in, well, telling people how we are using AI. We are trying not to overdo it. Like, for example, if we use AI like Photoshop, or like a grammar help, we are not declaring that because it's just, I guess, destroying the kind of discussion we should have because people are thinking about AI when they shouldn't. They should think about AI when AI can be confused with reality, for example, in the news. We don't want to use it there. Or they

should think AI when a picture or a text was completely created with AI. Then we have a disclaimer on the first level, like in the audio, or on the picture, then it's important to do that, but not overdo it.

Nikita Roy Interesting. How is it at Hearst, Tim?

Tim O'Rourke Yeah, I think we, again, coming from a data journalism background, we kind of do overdo it a little bit on the methodology, and we want to use, like the trust panel that was up here was talking about, we want to use plain language to describe how we're doing our work, where the AI comes in, where the humans are looped in, and really kind of try to explain that in a real way where real people, not just people who work in newsrooms or work in technology, can understand that. And then through our work with our legal team, I think, we definitely put in more disclaimers like where within the body of the tools or the experiences so that everybody's on the same page. And so we're really trying to be overly transparent at this point, understanding that it's sensitive, and that I think it's new not just for our readers, but for our journals who are using the technology as well.

Nikita Roy There's another question here from Jeff about who is responsible when an AI system produces any false or defamatory content, the AI developers, the newspaper, the human editors? How do we navigate that when we are creating audience-facing products? And I think this is maybe the chatbot era, and I think maybe Brooke, if you have any insight on that?

Brooke Siegel I mean I think to start, and this goes back to your previous question a bit, it's about baby steps, which I think is what really we're all talking about. It's like this is new technology, and we need to be introducing it to our audience in such a way that they trust it, they trust how we're using it, and that we are leveraging it to create more quality journalism. So one of ways that we're using AI in the background is sifting through tens of thousands of stories, and then finding stories that we find, the AI, that might not be up to our quality standards because it's trained by editors on what our standards and practices are, and those stories get floated up to editors, who are then going through it as a human to see if it should be published, or if it's something that we need to review deeper. So I think using AI to create more quality is what is going to improve trust and get both reporters, and newsrooms, and audiences that are engaging with the content more comfortable.

Nikita Roy I think that also comes back to the whole question about design. And I want to ask you, Juliana, how do we design for trust then when we are building out these audience-facing products? We are building chatbots, or personalization, or new forms of interfaces that we probably haven't even explored yet. How do we start to think about how do we meet audience and share that with them?

Juliana Castro Varón I think it's not too different that one is being done in terms of being transparent with the user, having like the pieces I talked have either a methodology, or its own piece explaining what we did, and how we did it. We don't publish AI generated copy, and so when we stand by a piece that we wrote, each journalist or the people who are part of it just stand by it. I think it's mainly about making sure the reader knows how this story happened. If it's about one story, or if it's about kind of like the whole archive or something, being transparent about that, and that we also have public principles that speak to how we use it for reporting, and then how we use it for the readers. If we use it in any capacity at the newsroom, that's disclosed. And I think there's, like, it's really, there are a ton of discussions as to how big is the disclaimer, and so on, that are more design specific. But I

think the baseline is that whenever it feels meaningful, and people would not expect that the copy, either be generated or automated or something like that, people want to know. We also, in order to be transparent, we constantly are having research done with audiences, readers and non-readers of the Times, to understand kind of what people want to read and what they're general feeling is toward it. So I think, yeah, trying to be transparent, but also trying to understand what people are into and want.

Nikita Roy Yeah, exactly. Go and ask your audience if they're OK with it. Do they want it? Are they comfortable with it? And I think AI means different things to everybody, so also explaining what that is. And we're all talking about transparency, and how we are using AI in our newsroom. But Uli, I think also I want to get into, you spoke a lot about algorithmic accountability. And one of the things is also how do we educate our communities about AI, the spread of misinformation that is possible without alienating those communities. And that's another audience question from Misty. And I want to get into that, because we can go into all of the different ways of how AI is able to create images, and then suddenly, that trust is even lost from your platform. How are you thinking about that?

Uli Köppen Well, we do have a great fact checking team, which has a more and more important job, and we are trying to combine their skills, with the data skills, and the deep fake detection skills. So I guess in getting more technical knowledge into fact checking teams is one of the reactions we are having and getting more people into fact-checking. But of course, well, we are on lost territory because it's much too much out there. We can't fact check everything. We have to focus on certain stories, and there we need our good journalistic stomach feeling what we should focus on. But I guess that's our most important reaction. And of course again, investigative teams working around AI. And interdisciplinary teams, looking at algorithms, really being able to scrutinize those algorithms is largely important. Because when we ask questions to firms, they always say, “We can't tell you. It's proprietary.” And then you need those skills, and you have to dive in and really check and audit those algorithms, if possible. It's not always possible, but if it is, then it's crucial to do so.

Nikita Roy I think that's where a key win for journalism is the rise of open source AI that's happening. We don't always have to go and rely with the big proprietary open AI models. There is a huge movement going on and hopefully that is going to accelerate over the next few years that we'll see. But tying into that same thing a bit, I want to pass a question along to you, Tim, as well. There's an audience question over here that’s from Paulo. A TV station's website is using AI to set the news agenda, research the news, write the content, which will then be presented in video by an avatar. That's all possible. The only human action is now to approve whether a subject will be aired or not, and the goal of this is to increase advertising inventory with minimal work. Do you believe this could compromise the quality of journalism? And is this the type of extreme automation that is a path of no return?

Tim O’Rourke Well, that's a big question. I will say there's parts of that that I think could be very valuable to an audience. It's a little far for where our company would go right now. I have read about some great uses of avatars, where going on live TV and reporting on something could put a reporter in danger, and so they use an avatar for protecting the reporter and the future reporting that they could do, and I think that would be a great example of when you would do that. But for our local news network, we would want our editors, you know, making the key news judgment decisions. We want to gracefully kind of build in personalization, but we're being very careful in the application and anything that faces the public.