

25th ISOJ AI in the newsrooms: What is working now and how it is helping to improve journalism

- Chair: **Trei Brundrett**, senior advisor (consultant), the **American Journalism Project's** Product & AI Studio and Startups Studio
 - **Lilian Ferreira**, general manager of strategy and metrics, **UOL, Brazil**
 - **Aimee Rinehart**, senior product manager for AI strategy, **The Associated Press**
 - **Andrew Rodriguez Calderón**, data project lead, **The Marshall Project**
 - **Zach Seward**, newsroom's editorial director of A.I. initiatives, **The New York Times**
-

Summer Harlow [00:00:00] So I want to bring out the moderator, who is Trei Brundett, who serves as a senior adviser to the American Journalism Project. He's a consultant, leading the product in the AI studio there and working closely with the startup studio. He's also co-founder and the former chief operating officer for Vox Media. Let's welcome Trei.

Trei Brundrett [00:00:30] Howdy, everybody. For those of y'all who are not from Austin, welcome to Austin and to Texas. I am happy to be back home. I've been living in New Jersey for seven years, and so I always hope that Rosental will invite me back. So thank you. Rosental. I want to take a second and I've been coming to ISOJ for many years. I don't know when the first year was, it feels like forever. And, I just want to congratulate Rosental on 25 years. That's incredible. 25 years ago, it felt like this was, you know, cutting edge. You probably had to work hard to talk some folks into even making some space and time to talk about the Internet, and journalism and the internet and bring everybody along, some of them kicking and screaming and some of us very eager and excited and he's always brought an incredible group of people together, so I'm looking forward to meeting all of y'all. It's such an international group and its practitioners and academics, and it's really special. So I just want to say, you know, thank you so much, Rosental, for making this a place for all of us to gather. So, I'm excited about this panel. And, I think, you know, there is a, I don't know if y'all have heard about this AI thing. There is a lot of noise. There's a lot of hype when, you know, at this time. And what we're going to try to do today is bring a little bit of signal, in that noise. We have an incredible group of folks who are going to give some presentations about the work that they're doing, how they are, using AI in their own newsrooms, and how they have been, you know, examples of good use of AI in newsrooms. This is something that I've been thinking a lot about and, talking to a lot of folks about, as mentioned, I've been working with the American Journalism Project as a consultant on their new product in AI studio, and we're very excited about that project, because what we're doing is, giving local newsrooms an opportunity to experiment with AI, and learn, how it might be used. And we started that by identifying, you know, opportunities, problems to solve, taking much more of a product mindset to how we might do that work. And then thinking about, you know, is AI useful here? Is it as transformative as we think that it could be? And, that's not only in how we produce journalism, but how we interact with our audiences and how we build revenue models around this new technology as well. And so there's a lot to think about, a lot to work on. There's a lot of experimentation that's happening. And I know a lot of folks are eager to learn,

you know, what's working and and you know what's not working. And one last thing I'll say about that is that I think it's incredibly important that as an industry that we are engaging with AI, that we're learning about it. I think that it reminds me a lot of the early days of the web. And why we have ISOJ is that the Internet came along and was transformative for our industry and how we served our audiences and pursued our missions. And, you know, some folks were, as I said, excited and some folks weren't. But what was important is that we learned how it was going to help us do our important work. And I think this is a moment where everybody needs to be engaging in it. There's a lot of fear and uncertainty and doubt right now. And, I think that, it's important that we learn about this new technology so that we are, well informed and, our perspectives about what the challenges are, what the dangers are, but also what the big opportunities are that we can speak from experience about what's going to be valuable, what's going to be meaningful, and what I believe is going to be another very, disruptive and transformative moment for journalism and for media overall. And so I hope, hopefully today, the presentations that you hear will be in the conversation that we have afterwards. Answering your questions will be helpful and encourage you to engage. I think this is a moment where we need, absolutely, to do that. So let's get right into it, that's enough for me. I'm going to introduce our first speaker today, panelist Lilian Ferreira from UofL in Brazil.

Lilian Ferreira [00:05:47] Hello, everybody. I'm from Brazil. Rosental, thank you for the invitation. I did a little pass with AI and I will present for you all my presentations in a video. If you could play the video, please? And then you'll tell me what you want, your ideas about it.

VIDEO [00:06:07] Hello, I'm Lilian Ferreira, general manager of strategy and data at UofL. Today I will bring you the actions using AI that we are already doing in our newsroom. UofL is the third most accessed website in Brazil, behind only Google and Facebook. Around 94 million people pass through our website every month. Even with a newsroom of around 300 people. We have many bottlenecks, and we use AI to help us in our mission to bring more and better information to the public. Our AI strategy started well before the long boom last year. We were already using proprietary machine learning with NLP to moderate comments on all articles four years ago. Since 2017, we have used the ALS algorithm for content recommendation in the 2020 elections. We sent 10,000 automatic texts, with the municipal result being the first news website in Brazil to do this. What changed is that with LLM and ChatGPT, AI was no longer restricted to developers. The entire newsroom began to have contact with this technology and use it in their daily lives. We created an area dedicated to actions in six main focuses: newsroom, public, audience, videos, data, and code. I'll talk a little more about each of them now. The newsroom immediately started using LLM to produce better titles and text. But the big leap was when we created a tool to transcribe live videos. In fact, they were the one who helped me write this text that you were listening to now. I did the text with some breaks for reasoning and changes of direction. As you can see on the side, I used the transcription tool that we developed at UL based on whisper from GPT, and that was a big leap in quality for the newsroom, because not only they would be able to transcribe recorded interviews much more easily. And then I believe that everyone here works in journalism. So you understand the difficulty of having to transcribe or to take notes while someone speaks, and then an hour of interview turns into four. So this helped a lot in the day to day work of the newsroom, because

the market tools until then were very good in English, but for Portuguese they were very deficient and then our tool managed to be better. Exactly using LLM. A second use that was very important for us was in live transcriptions. So we started using this tool a lot when a president, a great figure was speaking or even football coaches, we could quickly write what that person was saying. So it would generate text every two minutes when the video was processed. With that, we were able to publish text alongside the live video much faster. We have several other ideas for using the LLM newsroom, such as correction, following our newsroom patterns and diversity manuals. For example, these initiatives are in our backlog, along with several others, such as title suggestions for SEO tags, which should enter our new CLS soon. In terms of the final public, it is very difficult to develop these tools, both due to cost and accuracy, to prevent the AI from hallucinating. As UofL is a journalism website, we can't take the risk of the AI ending up saying something that is not true and compromising our credibility. We just launched a prototype, the first result for the final public, which is a chat bot to answer questions about income tax. That is really complicated to understand. You can see it in use here. We did a lot of testing. Develop the LLM a lot to be able to arrive at an AI that responds correctly and can say, I don't know. This was our biggest challenge because every time it didn't know an answer, it invented it and that was our problem. We changed the prompt a lot. We spent months working on it so that we could deliver this first MVP. Another delivery we made was Dia. Dia is here by my side. She is the presenter made with artificial intelligence at UofL. We used it on our social networks to publicize some articles. The issue right at the beginning is that we still didn't have cheap, affordable and good quality tools for us to make Dia. Then there was the voice. This has taken a huge leap forward in recent times. That robotic voice to a voice that today we almost can't determine if it's a human voice. In fact, that voice you are hearing is my AI clone voice. Did you notice? In the end, it was very expensive for us to generate this presenter using AI. We had to generate the presenter's text. We had to do lip sync. Sometimes the voice wasn't synchronized properly, then she didn't have a lot of movement, so we wanted to cover the rest of the video with images. What ended up having a lot of covering work. Instead of making our work easier. She ended up giving us even more work. And so we retired Dia for now. To generate an audience. Of course, we also use the technology of the moment LLM is text and we have a lot of text, almost 30 years of producing online content. So we started to generate new articles based on our own content. This is very important with all this movement of copying other people's texts, of asking the LLM to write and not citing a source, which is crucial for journalism. We only use the LLM in our own texts. We hired a tool at the beginning. So our idea is that when something new comes up, for example, a famous person's illness, we'd use our own database with explanations about all the illnesses. So we take our own texts and ask the LLM to rewrite them. After the text is written. There is always an editor looking and seeing if everything is correct. And then we always mention that the article was written based on data from this other article. So we managed to make text production much faster. Another use for us to gain an audience for text on big themes of interest. For example, as I said before, we used a bot to generate automatic text about elections. We are now using LLM to generate text about where to watch soccer games. As you know, Brazil is a country that loves soccer. We have several championships. We watch all European championships, and Brazilians always want to know where to watch these games as not all of them are broadcast. So we created a tool powered by LLM to help us generate these

texts. We have data on where the games will be transmitted on a page. Then we request for LLM to generate text based on the model explaining where people can watch that game. We can generate texts of all the games that will be playing something that was unthinkable to do manually. Another of our premises is to always make things possible that were previously impossible and of course, make work easier. We invested a lot in videos. Our idea was I think everyone's idea was, you know, generating video from text. We have tested several tools on this. So the best tool that we have tested and are still testing is called Rizzle. They make videos from text and photos. It even searches images databases. This video you are seeing here was generated by a Rizzle with the photos and text we uploaded. The artificial intelligence automatically takes the main moments of the text and generates the lettering. The results were very good, with half a million organic views on YouTube and in a much simpler way in 15 minutes, anyone who has no video skills can generate these videos. Another tool that we are using is also a dubbing tool Rosc. Here we also test several tools. Don't think that the first one will be the best, nor that you will find a perfect tool that meets all your needs. We don't think the results are great, but we believe the test is always valid. By the way, this is me here that you see speaking in different languages. It's not me speaking, I used AI to speak in English and Spanish. And this comes into several ethical issues. So instead of being the person's own voice, for example, President Lula speaking in English or in Spanish with his voice, it would not be what he actually said, which is why we chose to use a more journalistic translation. In data. We already use the ALS algorithm, which to date is one of the best for recommending content, but we have a lot of user data and we didn't make good use of it. So we started to look better at this data. Clean data. Work on this data and coordinate a more user focused delivery. So our idea is to deliver the best content to the person at the right time, in the right place. We already tried to do something similar a while ago before the LLMs appeared, and we ended up being blocked. Exactly. Because all technology available about time couldn't understand context. Nowadays, on the UofL home page, which I'm showing here, there's this entire part which is made for delivery according to the entity that the user browses the most. But we got stuck a lot with each combination. Another very important point is that the line today is done in English, so all the correlation is done between the words in English. So when I use an alarm in Portuguese, it takes the words in Portuguese, translates them into English, generates the correlations in English, returns and translates that into Portuguese, which is not always as accurate as possible. A very easy example is that in Portuguese we have gender, the el presentador and the presentadora. In English it is not. In English there are a lot of references in it and in Portuguese it can also be a feminine or masculine it. And this completely changes the meaning of the sentence. So we want to use this model name Sabia and our new system, which is an LLM in Portuguese. So language is another barrier that we face in using the LLM. Perhaps the dances and codes are the most obvious and expected, but no less important. We just developed the hub AI, which has its symbol on the side. It is an internal tool into which we plug both the LLM and various data sources. So to research several different data, the editorial team will be able to use it, but it is also what we use to generate the tool for the final public on income tax and it is also what we are using now to generate new codes. So before we had a huge queue and many bottlenecks in the development team and using LLM, we can provide much more with codes being created much faster. Just using the free GPT chat, we were able to generate these calculators here that you can see by my side. And our idea is to improve them, adopt them to

our websites and code. With this, we are able to develop many more calculators, games, and quizzes to greatly increase both the delivery to the final public and the data part because of these games. These quizzes, we are able to have user data and we also manages to increase the audience. So it ends up bringing together many of the parts that I mentioned earlier that are our main focuses. We are in a time of change and knowledge of this new technology that is certainly already changing how we produce and consume content, and that will still evolve a lot. For anyone interested, these are the eyes I use to generate this video. There was still a lot of manual work that it is a path of no return.

Trei Brundrett [00:17:40] All right. Well, we've had one AI presentation already today. We're on, we're on theme. All right the next presentation is from Amy Reinhardt, who is at the AP.

Amy Reinhardt [00:17:59] Hi, everyone. Congratulations, Rosental it's great to be here for the 25th anniversary celebration. I hope my presentation, answers a little bit about, is the question I always get is what is everybody else doing? So, at AP, we did our first survey on AI in the fall of 2021 with almost 200 U.S local news leaders and this report that we're that we published on Monday is different and that we invited participation from news people globally and the focus is on what's everybody doing; what's on everybody's mind generative AI. So this report started over dinner the night before the first Nordic AI summit in 2023. I was talking with Nick Decapolis and Natalie Helbig about how different our survey results from 2021 would be now. And Nick said, let's test that out. So the survey was available for the first three weeks of December in 2023. 24 questions total, with some questions that dig a little bit deeper into more questions. And we received almost 300 responses back. The responses were largely from North America, with representative representation in Europe, Asia, Africa and South America. 35% were editors, 20% were executives, 58% identify as men. And on April 9th, we released the second report on AI in the news, Generative AI in Journalism The Evolution of News, Work and Ethics in a Generative Information Ecosystem. My AP colleague Ernest Kang and I are co authors with Nick Decapolis. He literally wrote the first book, on AI in news called Automating the News. And Natalie Helbig, a law professor at the University of Amsterdam who was a leading voice on guardrails on AI in news, with the European Union. Assistant professor, in News technology at the University of Amsterdam Hannah Schols and a researcher, Charlotte Lee. So we had lots of people digging into this data. So the top level findings are that many of the explorations of AI are happening in content production. Likely considered the safer area and unknown area, few newsrooms are going beyond existing workflows to see how this technology is likely to inform newsgathering, product production, and distribution. Respondents indicated that they are knowledgeable about generative AI, with more than 80% saying they knew about generative AI, and almost 74% indicated that they or their news organization had already used generative AI in some capacity. So many new rules are emerging to grapple with the changes introduced by generative AI, including for leadership, editorial, product, legal and engineering. One person put they needed an AI intern. I was like, please don't outsource them. So almost half of respondents indicated that tasks or workflows have already changed because of generative AI. New work is created, in devising, effective prompts and in editing outputs. Perceived efficiency gains are variable, and additional research we think is needed to evaluate any real performance gains across a range of common tasks. So overall, these findings underscore the need for training

initiatives and for more fine grained evaluations to measure actual shifts in productivity. Almost half of respondents indicated that tasks or workflows have already changed because of generative AI. AI models take on roles as collaborators and are used as a sounding board to bounce ideas off of. As an editor said, that catches things that may have been missed so much like a self-service checkout system in a supermarket. Respondents indicate that new work is created for them, when they use these systems, primarily in terms of having to edit or proofread the outputs of the AI to ensure it is acceptable. And that is part of the big question mark of, is this really going to save us time if we have to go back and reread everything anyway? So, let's see here. Oh, there we go. So there is an unmet opportunity to design new interfaces to support journalistic work with generative AI, in particular to enable the human oversight needed for the efficient and competent checking and verification of outputs. Journalists will need well-designed editing interfaces in order to effectively use generative AI for various tasks. And let's face it, we're all here because of a user interface called ChatGPT. So that really, made people aware of the capabilities of AI, in a very felt way and it's going to have to translate something similar into the newsroom. So respondents are also open to getting help from generative AI for tasks related to analyzing, getting or processing data and information which are perhaps not, coincidentally, also the kinds of work activities that respondents rated as boring, repetitive, and tedious. So respondents indicated concerns about human oversight, accuracy and bias. And the industry is grappling with how to balance the benefits of generative AI with the need for ethical journalism practices, including the banning and limiting of use for particular use cases, such as the generation of entire pieces of published content. Overall, editors, managers and executives, and technologists were the roles that respondents thought should be more responsible for ensuring effective and ethical uses of generative AI. So while many organizations are developing or following guidelines for the ethical use of generative AI, there is a call for clearer, more concrete guidelines, training and enforcement to navigate the ethical landscape effectively. On top of guidelines, there is recognition that additional training is needed to support responsible use. Other strategies that might also improve responsible use of generative AI like a robust procurement of tools that include AI and automation, as well as internal testing and auditing, but those were rarely mentioned, though we want to underscore it here that we think it's a really important thing to review in your newsroom before adding anything into it as a tool. I also appreciate that 20% of the respondents had their arms folded firmly across their chest to suggest that responsible use of AI is to not use AI at all. That is certainly a familiar strategy for digital disruptions. So respondents expressed a degree of uncertainty about whether tech companies should be allowed to train models on news content, with some emphasizing the negative commercial impacts and others advocating to advance the accuracy and reliability of models which could benefit society. So respondents were asked if they thought other companies should be allowed to train on their AI models on news organizations, digital reporting and information. And the biggest group of respondents, almost 54%, seemed torn as they responded maybe to that question. So no clear answer there from the news industry. While recognizing the potential for revenue generation and advancements in AI tools, these respondents emphasized the need for careful consideration of copyright issues, transparency, and accountability to protect intellectual property and journalistic integrity. Additionally, respondents indicated a need for transparency in how their data is used and whether proper attribution is given to the original creators. So almost 33% opposed the idea of allowing other

companies to train models on their digitized information and expressed concerns about copyright infringement, unauthorized use of proprietary content, and the potential negative impacts on the competitiveness and sustainability of news organizations. And then there's this 14% group, who advocated for allowing companies to train on their digitized archives, arguing that such collaboration could advance the field by improving AI accuracy and reliability, benefiting the news industry and society. So collaboration is seen as vital for producing accurate, fact checked content while adhering to professional and ethical standards. So some takeaways for what this research might mean in the coming year. Usage policies, such as guidelines need to be made more concrete to better steer toward responsible use around specific tasks and cases. Tools themselves could be evaluated more rigorously to ensure alignment with journalistic expectations and norms. Human oversight, responsible experimentation, and the creation of dedicated support and learning structures were called for by respondents. And there are many claims around efficiency in AI tools, but no concrete evidence. We think additional research is needed around which tasks and use cases actually benefit, and efficiency and performance gains and quality output for a range of tasks. Design and prototyping might be used to explore more powerful interfaces to support human oversight. In the newsroom, and generating, of editing and generating outputs. So, they, we also are calling for exploration of genuinely new experiences rather than just optimizing the same old workflows. And I think that is really important for all of us to get out of what we've been doing and get into something that's wholly different. So new programs, new training programs are needed not only in prompt writing, but also in responsible use and adherence to usage guidelines, and then thinking systematically about how to evaluate and refine workflows or how to develop something entirely new. So, AP received grant funding from the McGovern Foundation to train journalists on how to report on AI across every beat. When we conducted research on AI in the US in 2021, when we asked news leaders if they might be interested in having their reporters trained on how to cover AI, several said AI is not something that is present in their coverage area. So this training is meant to inform reports on every beat about how algorithms are determining bank loans, decisions in social workers offices, schools, policing and hospitals. So if you live in an area that has one or more of those agencies, AI is absolutely in your coverage area. So we'll be giving journalists the vocabulary and information needed to start asking questions about how these institutions are making decisions and if they are being assisted by computer programs. We've already held our first of four flagship webinars. On March 28th, an AP reporter on the democracy team provided the training on AI and the election. She offered this slide as how to stay on top of deep fakes. A few tells she mentioned were, the Pope's eyeglasses, half of a chain of his necklaces missing, and his right hand is kind of mangled. And to also check to see if other newsrooms are reporting something similar. So, people like examples. And I'm just going to share a few here because the two people who I'm going to mention are actually in the room. These are people who inspire me by the creative uses of AI. And that is Paris Brown, who's in the second row here. She leads a three person newsroom at the Baltimore Times, and she leverages ChatGPT as an editorial assistant. She uses Dall-E as an image creator and uses Zing AI to offer audio versions of articles to improve accessibility. Jose Nieves, editor of the Spanish language El Toque is a Cuban journalist based in South Florida. He has a forex rate in real time, using natural language processing to read social media posts that buy and sell foreign currencies. He built a small GPT to write headlines

on social media posts, and El Toque also offers a price monitoring system that scrapes data from online stores. Unless his audience know the price of coffee, eggs and other essentials. So this moment, I think, is for small newsrooms and the curious. I look forward to learning more and sharing during our panel. Thank you.

Trei Brundrett [00:30:29] Thank you. Amy. All right. Next up is Andrew Rodriguez Calderón from the Marshall Project, which I'm proud is also one of our collaborators at AJP on our project on AI Studio.

Andrew Rodriguez Calderón [00:30:44] Okay. Hi, everyone. Good afternoon. Can I get some energy? Good afternoon. Thank you, you all just had food. I'm expecting some energy. So it's a privilege to be up here with you all. I debated where to start this presentation. I decided to start it personally, and I'm really happy that I did. Taking a page from Wendy's keynote. So about me, I'm a child of Colombian immigrants to the US who have been here for over 30 years. I'm queer. And I'm also the son of a formerly incarcerated mom. I mention all of this because I think that as we talk about AI and we focus on the future, it can be easy to sometimes forget the past and the lessons that we can learn from the past. I myself, I'm a member of various identity groups. I'm sure many of you are as well, who have been the subject of exploitation, extraction, or exclusion by both the media industry as well as the tech industry. When we adopt new technologies, we bring all of that history with us. And I think that it's really important to develop practices as we use this new technology that acknowledges that past, at the Marshall Project, which is a nonprofit newsroom that covers the criminal justice system. We're very well aware of extraction, exploitation and exclusion as it's woven into the DNA of our criminal justice system. So when we were lucky enough to receive a grant from AJP, to continue our experiments in AI, we asked ourselves, how can we use this technology, in a way that doesn't perpetuate some of those ills? And, so first, you know, for us to be able to do that, we had to think about, well, what does AI do? And what is it? And so one of the ways in which we think about AI is that it's all about scale. It can amplify and augment the not so good and definitely the worst parts of us. I'm sure I don't have to give you many examples of that. It also in that way amplifies biases, inequities, and power dynamics, not just because of the nature of the training data that is used, which is often pulled from the internet, which we all know can be kind of a cesspool, but also because of the way that it's sometimes implemented and the promises that are made by tech companies about what this technology can do, you know, visa, vis-a-vis, recruitment tools, etc. And so if you keep all that context in mind and you ask yourself, how can I improve journalism, you might be hard pressed to imagine how that can happen. I'm actually really grateful that I'm going third, because I think we've just gotten some wonderful examples of how I can be used to improve journalism. Personally, in my role at the Marshall Project as a data project lead, my practice sits at the intersection of collaborative design and computation or machine learning. And so my contribution to this panel is to think about how collaborative design can help us use AI in the service of the people who are affected by the issues that we're covering? So before I jump into that practice, I'll just say to kind of get it out there, I oversimplified what I believe to be the potential uses of AI and to these kinds of reduced points, also as a provocation to stimulate some of your questions. But I believe that I can offer the people that we serve stories and news products to, ways of empowering them by creating alternative models for accountability, and can

also allow us to address fundamental information gaps at scale. And I'm going to give you an example of how we've done that, at the Marshall Project. But first, I'll talk to you a bit about collaborative design. And so essentially what collaborative design purports is to present a schema for you to engage with people, to get them involved in the design and development of whatever it is that you're building. There are many different versions of this. Some people call it design thinking. The reason that I like to use the term collaborative design is because woven into the DNA of collaborative design is this notion that anything that you create doesn't exist in a vacuum, and that you have to keep in mind the history, the context, the socio technical background of whatever it is that you're, working on, as well as the people that you're working on. But if I had to simplify it, I would say that collaborative design is about listening. It's about reflecting, then designing, prototyping, and then sharing back with the people who you're designing with. And so as an example, I'm going to talk to you about book bans in US prisons, which interestingly, this use case I'm going to give you is from the early days of ChatGPT. Now it seems kind of basic, but hopefully it'll still be informative. So last year we focused on the issue of book bans in prisons or prison censorship, in part, spurred by all of the censorship that was happening in the education system. People who are incarcerated have long suffered from difficulty accessing information. Inside of the prisons and jails that they're housed in and sometimes information that's critical to their reentry process. This reporting started with us wanting to build a tool where we requested banned book lists from every single Department of Corrections across the country, as well as policies. We believed when we started reporting that it would be really useful for people to be to see what books were being banned, and that it would also inform decisions that family members, as well as prison educators, prison librarians and the formerly incarcerated were making about what books to send in, but as we spoke with people and we listened to them, we realized that we were completely mistaken. Yes, these lists would be useful, but they were more of a historical document. People who actually were interacting with folks who are trying to get books into prisons or retain them, told us that the system is so fragmented, even when you look at a single state, because every agency is different with different correctional officers who have different relationships to every person who's incarcerated, that you couldn't possibly use a single list to make a decision about what would get in and what would not get in, because sometimes getting something and involve knowing who was working in the in the mail room on a given day. And so we then asked, so should we publish these lists? People still wanted them, so we published them, but we also asked them, what could we do that would be more helpful to you? And what people said to us is, can you get all of the policies? Because what we really need is to know how the Department of Corrections is making decisions about what books are allowed in, in part so that we can take action when we know that somebody has not been allowed to retain or to receive a book that they're supposed to. So we reflected on that, and we thought about the different ways that we could make those policies accessible, but also actionable to folks, because if anyone's had to read any kind of administrative policy, you know that it often just looks like a bunch of gibberish and it's very hard to make sense of it. So we designed and we started thinking about this information that we were getting from folks who were close to the system, and we came up with a prototype of some prompts that allowed us to pass in the policies that we were receiving and then output summaries. One thing I'll say is that we, you know, we're conscientious about some of the things that Amy and Lillian have spoken about, which is we don't want to completely

outsource that journalistic process to a machine. So we first went in and we chose each section of the policies that was relevant to book bans, so that we weren't just passing everything in and we had a bit of control. We also fact checked the summaries that were output afterwards. And so then we prototyped and this is what the summaries looked like after we were done and we built them into a tool supposed to be a video, basically, like the list is on the bottom, you can see the different books. And then if you look at the top which says read Policy summary, you can see the summary of the policy that we got, as well as a link to the actual policy, which we put on a document cloud so that we could retain it. And then we went back to the same people that we had been speaking to from the beginning. Mind you, they were already aware of the work that we were doing. So they were very generous. We got to talk to even more people because they were saying, hey, the Marshall Project is doing this. It's really great. By the end, we were being invited to conferences with really radical librarians. It was really cool. And they told us, yeah, that's cool. Like, that is what we wanted. Thank you. But actually what we really need is to be able to compare the policies. So can you go back to the drawing board, take these policies and make it so that if I put the Alabama policy next to the New York policy and every other state, I can see if my state has an appeals process, because there are some systems that don't have appeals processes. And then I can go to the Department of Corrections and say, why does every other state have an appeals process and you don't? So we went back to ChatGPT. We did some prompt engineering, and we managed to come out with a series of headers that we could apply to every single policy summaries, again, all being fact checked. And if there was ever a section that we didn't have information for because the policy didn't contain it, it would say this policy does not have information about a review in an approval system, for example. And that would create a space for accountability, because knowledge gaps can sometimes be sites for action. And so just to recap, the news and tech industries are historically exclusionary, extractive and exploitative. It's important to keep that in mind. I think collaborative design encodes a different set of values and allows us to engage differently with folks and AI and design and journalism together as tools in that design process can produce empowering stories and products, can create alternate pathways for accountability that are more ground up and can also address fundamental information gaps at scale. My main takeaway is that I don't think AI is the end all, be all. It is a tool, but it is not neutral. And so I think that it's important that we have a framework that informs the way that we use AI as a tool and how it perpetuates, motivates our journalism. Thank you. Sorry. One more thing. If you're interested to know more about this use case, you can type this into Google and you'll see a full write up with examples of prompts. If you have any questions about anything else that I've said, please write to me. I'm happy to share books and texts about this design process.

Trei Brundrett [00:41:12] Thank you Andrew. That was awesome. Not only because of the work, but also the gifs that you included. It's always important to have good gifs in your presentation. All right, next up is Zach Seward from the New York Times.

Zach Seward [00:41:36] Cool. All right, we'll go back. Thanks everyone. First, Lillian, Amy and Andrew's presentations really blew me away. So thank you for all of those great insights. Once upon a time, not so long ago, there was ChatGPT and people had a lot of fun with it. You could use it to write poems, formal emails, or a pitch for a Bravo reality show about the bats that live

underneath Congress Avenue Bridge. It started to seem almost like what generative AI couldn't do, right? This whole presentation for me at the last minute? I'm kidding, but it wasn't that bad. If you go to Po and look for an auto presentation, you can try it yourself. Maybe create an inspiring pop rock ballad about Rosental with the song Generation Apsuno? The one I'm not kidding about. You can find me later and I'll, I'll play it for you. You know, the whole thing. It can be a blast. But a year later, it seems clear that introducing the world to generative AI through parlor tricks like that created some distorted impressions of what the technology is good for, or at least what it's best at. How many of you have stared at a powerful LLM tool like this one? The cursor in an empty text box, blinking menacingly at you as you wonder. You know, good question. What can you do to help me today? If you've done that, you're not alone. AI apps, though popular, suffer from lots of users who try one of these pilot parlor tricks. Right? The first page of my novel for me? Sure. And then just never return. Just 14% of monthly users of AI apps are daily users, according to one analysis, which compares to 51% for popular apps like games and social media platforms. I'd like to think that's due in part to a fundamental misunderstanding. I mean, generation, it's right there in the name, so I get confused. But generative AI's most powerful use, as Andrew just showed us, is not in creating entirely new text or images. It's creating structure out of messy data that already exists. This talk is building on a point that I made briefly at the end of my last public talk, which just happened to be here in Austin at South by Southwest last month, where I shared some examples of journalists using generative AI to create structure out of unstructured prose, one of which was Andrew's inspiring project on banned books. Another was a custom GPT developed by Jay Mark that parses complex government audit reports in the Philippines to help expose corruption there. And another was a product recommendation site generated by intelligently mining the links in the archives of the newsletter "Why is this interesting?" The common element in each of these generative AI projects is that they aren't really generating something new. They're creating summaries, extracting information, and structuring data in a more usable form. So today, I'd like to just build on that argument with a few other examples. On March 11th, New York City's Committee on General Welfare held a nearly six hour meeting, which, as is its wont, at which several important issues were discussed things like funding cuts for homeless shelters and rental assistance programs, backlogs and applications for food stamps, the city's response to migrants seeking asylum, and so on. But if you go to the city council's website to find a record of the meeting, you wouldn't know any of that. You'd be staring at a page like this one, wondering if this is really how the nation's largest city operates in 2024. And it is. Vikram Oberoi, an engineer in New York with an interest in local government, thought there had to be a better way. Now the City Council does publish transcripts and video recordings of its meetings. If you can find them in this interface. And so Oberoi wrote detailed prompts for GPT four to break up the transcript into timestamped question and answer pairs he called chapters. And give each of those chapters a helpful title, and then summarize the salient points of the discussion. That's one of the prompts he used of many to do that. A totally complex process, but all done using a LLM. And having done all that, Oberoi had all the structured data he needed to create a far more accessible version of the same exact meeting, navigable by chapter, with helpful summaries of yes to the chapters and the video synced to the transcript. Not just for that six hour meeting of the Committee on General Welfare, but every City Council meeting since he started this project a few months ago, an incredible contribution to civic engagement and local journalism by someone who doesn't even consider himself a

journalist. Give it some thought, and you'll find plenty of ways to use this approach to your advantage as a journalist. A colleague of mine at The Times, covering immigration, was frustrated by New York City's habit of releasing important statistics ad hoc in the middle of City Hall press conferences. They weren't published anywhere else. The mayor would just say it in the course of a Q&A with reporters. So what my colleague did was give the transcripts to a LLM and asked it to spit out all the statistics. And so all of a sudden, for each and every one of the mayor's press conferences, he had a kind of Harper's index of the press conference, which allowed him to much more quickly generate and find the statistics that mattered most to him. Now, this one, you know, as opposed to Oberoi's example that I was talking about earlier, is using an extremely simple prompt that anyone here could employ without a lick of code. And yes, you should fact check these data if you plan to publish them. But in my experience, labs are very good at this kind of task and far less likely to invent information than when generating text whole cloth, which is yet another reason why creating structure is the best land use case. You know, real life is messy. Journalism at its best helps people make sense of that mess. Using AI to give structure to messy data is therefore a pure form of journalism. The neural networks underlying large language models can seem like enigmas. What you're looking at on screen is not only what you might see after a wild night on Sixth Street. It's a visualization of training a neural network which, even amid mysterious calculations over vast arrays of data, produces gorgeous fractals like this one AI hues toward structure and it loves structure, it's not random. One final example. I've been talking mostly about creating structure out of text, but there's just as much potential in creating structure from images. A lot of you are traveling here from out of town at your employer's expense, I assume, which means lots of receipts like this one and a painful data entry process when you get home. Well, next time, try feeding your receipts to your favorite multimodal LLM and asking for the data you need. And there you'll have it. The most beautiful kind of structure out of chaos. Thank you. If you want, the slides or all the references in this talk, you can, just use that QR code right there or go to my website. And I would love to hear from you too at the contact info on screen. Thanks.

Trei Brundrett [00:49:41] All right. It's time for discussion. Question and answers. We have a few that have come in, so if you're on Zoom or Slack, send us more. We have an awesome set of panelists. And I'll just start filling in. If y'all don't. And y'all have better questions than me. All right. Why don't we start? This is a question from Julietta via Slack. What are some of the privacy issues stemming from AI and Andrew, I know, y'all thought a little bit about this at Marshall Project. So I'll start with you and then other folks join in.

Andrew Rodriguez Calderón [00:50:17] Sure. Is this working? Okay, great. Yeah, we thought so. For example, for the project that I just showed you all. We didn't have to worry about privacy because the policies were already publicly available. And it's interesting because I actually think that there's tremendous potential. And, you know, Zach, you were just alluding to this as well. Like, there's tremendous potential to take publicly available information that people really need, but they probably don't know is there. Or even if they do, it's hard to parse and make it actionable for folks. But we also have some internal experiments that we're running. For example, we have a magazine that goes out to, I think now over like 6 or 700, carceral facilities in the US as well as in Canada and Mexico. And there are surveys that we put in there. And so

people send us responses and we're trying to figure out the best way to use that information. We do use it already with what we call reader to reader. But if we were to use it for journalism, how do we protect the privacy of these folks who are already in a very, very vulnerable situation? And part of it is, like isolating the model so that we have a local version that we're using that's on devices that we own rather than exposing it to any kind of third party like OpenAI, Google Cloud, etc.

Trei Brundett [00:51:29] Anybody else want to join in? All right. That's a good answer. All right. Jesus via Zoom asked is anyone scared about the potential dangers from AI, specifically regarding misinformation like deep fakes? Who's scared of this? Amy, I know that you're scared, but I'm going to ask. I know that, you know, I've thought a little bit about this and as you've done the survey, you know how people are thinking about misinformation, deep fakes.

Amy Reinhardt [00:51:58] Yeah, I know everybody's worried. And we have a lot of elections this year, not just in the US. I do think we will see generative AI images, but there's still friction points to creating those accounts for now. And so I don't think we're going to get flooded. And I also was really inspired by people's OSINT skills on the Kate Middleton stuff. People were very activated to find that out. I hope they can be activated for something more substantive, around the election. You know, I'm more concerned about the trust and safety teams being reduced at platforms. I'm also concerned about funding around mis and disinformation researchers being cut. So we'll know less than we did, probably, in 2016 about this election. So those are the big concerns. And then you can just, you know, like everything else this year, you just drop AI on top of it.

Trei Brundett [00:52:53] Anybody else want to speak on misinformation?

Zach Seward [00:52:56] I mean, I think it's a huge concern and something we're looking at really closely. I can imagine a time not far from now when we are going to want our entire newsroom equipped with a set of tools to make sure that the person on the other line on it, what you know, we used to think of as an innocuous phone interview is actually the person that they say they are. You know, if your mother says she loves you, check it out. Well, now it's like if your mother says she loves you, make sure it's actually your mother.

Andrew Rodriguez Calderón [00:53:28] If I may, just throw in one small point as well. Like in the previous, panel discussion, there's a lot of conversation about crowdsourcing. And like with the collaborative design, I was talking about this notion of bringing folks into the journalistic process, I think can in some ways the anecdotal as well, though obviously not a wholesale solution. But if you get more people involved in journalism, you get them thinking and engaging with this way of consuming information. And I think that can potentially provide some pathways for, addressing misinformation and disinformation.

Lilian Ferreira [00:54:01] The main question and main factor here. I would try to explain in English a joke that I read the days before in Portuguese, to people that are watching a video with deep fake and one call to another, hey, did you see that the president says that? And the

other says No, this is not real, can't just see he has six fingers. And then the the one who trusts that call to another. Friends, do you know President Lula has six fingers? So I think that we need people to not trust that because they will trust that people has six fingers.

Trei Brundrett [00:54:41] All right. Six fingers. It's true. What do we. Here's another question. Maya via Zoom. What do you think eyes roll is? And the lack of trust in the industry. Or on the flip side, how can it be used to build more trust? Let's take that one. This is an interesting one, actually. There's been some really interesting research recently, on AI and trust.

Andrew Rodriguez Calderón [00:55:09] I can jump in just to say that I think that I have this kind of notion and I'll throw it out for you all that I think news organizations need to be a bit more like tech companies, in the sense that they should be creating products that allow people to engage with the world more effectively. And I think that I can be a tool in your arsenal to do that, and in that way can, potentially create trust by becoming a place where people go to make sense of things. And I don't just mean information. For example, like at the Marshall Project in Cleveland, along with Signal Cleveland, we just created a judge guide to help people navigate the overwhelmingly complex nature of judicial elections there. But that's true basically everywhere. Like my husband is from France and he is very engaged civically, and even he's having trouble making sense of U.S. elections. And we often talk about journalism being a democratic institution. So why don't we create products that help people navigate the complexity of democracy? And I think if we do that, then maybe we'll get some more trust.

Trei Brundrett [00:56:13] Zach, at the Times, how y'all been thinking about trust and bringing AI into the newsroom since I think that's your title in the newsroom. Somehow.

Zach Seward [00:56:25] I'm just the AI.

Trei Brundrett [00:56:28] Zach is AI in the newsroom.

Zach Seward [00:56:29] One thing that's really been striking. We have this great audience insights group that does all sorts of surveys and a variety of topics, including, of course, people's perceptions of AI and how, you know, different labels on different features or exploring are perceived. And the unfortunate thing for me is finding it's all over the place. Everybody has their own, you know, conception of what is meant by AI and brings that baggage to bear when the term is used. For one, like we're sort of blurring the line, even in this panel between traditional machine learning and generative AI. And that's a really important distinction. But, I mean, I totally get why it's unreasonable to expect a reader just looking at the word artificial intelligence to make that distinction. And so, you know, the conclusion we draw from that is, as always, transparency is going to be incredibly important. And it's not sufficient to simply say this. This was generated by AI because that's not, helpful explanation to most people. And certainly if it's something that we're creating and putting in front of readers, it's on us to be clear about the process through which that was done and what the potential risks and drawbacks are and how we addressed those risks and so on. You know, transparency is an age old, you know, value in journalism. And I think it becomes heightened, in this case, I guess the optimistic view would be

we do that well as an industry that, should, you know, result in higher trust, in, in journalism. But we'll see.

Trei Brundrett [00:58:19] All right. My time takers have indicated that we are done, which, first of all, thank you all for joining us. And, and thank you to everybody who sent in questions. I know more, heeded my call and sent in a lot more than we were able to answer. And finally, thank you to all the awesome panelists for sharing all their work.