26th ISOJ Synthetic Media and Journalism: Avatars, Image Creation and Manipulation

---

- Chair: Robert Quigley, professor of practice, UT Austin
- Carlos Eduardo Huertas, founder and director, CONNECTAS (Latin America)
- Claire Leibowicz, head of AI & media integrity, Partnership on AI
- Santiago Lyon, head of advocacy and education, Content Authenticity Initiative (CAI)
- Craig Silverman, national reporter, ProPublica

Summer Harlow Toward the end of the day for our final panel of ISOJ, the grand finale, absolutely. This is Synthetic Media and Journalism: Avatars, Image Creation and Manipulation. And this panel is chaired by Robert Quigley, a professor of practice, also here at UT Austin. So let's welcome them all.

Robert Quigley Sorry about that. I'm going to stand up here to do this, all right. So welcome everybody to the final panel of ISOJ 25. First of all, big, huge round of applause for Rosental. I always tell him this, and I don't just say it because he's my colleague or anything like that, this is the best journalism conference going, and I've always felt that way. Congrats to you, Rosental. Congrats to Summer, Teresa, the Knight Center staff, all the student volunteers. A whole bunch of them are my students, so congrats to them too. And this panel is about synthetic media and journalism, and it includes avatars, image creation, manipulation. I'm Robert Quigley. I'm a Professor here in the School of Journalism and Media here at UT. I teach a digital journalism class like our freshman class. I also teach copy editing, podcasting, a summer class on AI in the media, which is why I'm doing this. And I also supervise a student group of students working on AI projects. It's under Professor Christian McDonald's tutelage. He directs an endowment here that pays for it. There's so many good things happening here at the university, and I'm glad to be part of it. And one of the best parts is being here this afternoon.

In terms of this panel, this is definitely saving the best for last, so let me introduce the panelists and let them get going. We have Craig Silverman, who's an award-winning journalist, author, and one of the world's leading experts on online disinformation, fake news, and digital investigations. He's a national reporter for ProPublica, where he investigates digital platforms and online manipulation. He's also the editor of the European Journalism Center's verification handbook series. He writes with the Digital Investigations newsletter, and he previously served as media editor at BuzzFeed News, where he pioneered coverage of digital disinformation and media manipulation. We also have Carlos Eduardo Huertas, who is the director of CONNECTAS, which is a Latin American investigative journalism and collaboration platform. CONNECTAS is an independent journalism alternative, which he founded in 2012 as a Nieman Fellow at Harvard University. It has 160 members in 19 countries, and it has exposed abuses of power even in closed societies. Huertas has led the Initiative for Investigative Journalism in the Americas, which is a 10-year collaboration between his organization and the International Center for Journalists. Then we have Santiago Lyon, who's the head of advocacy and education for the Adobe-led Content Authenticity Initiative, working to combat misinformation through digital content provenance. He has more than 40 years of experience in photography, as an award-winning photojournalist, photo editor, media executive, and educator. As a photographer for Reuters and the Associated Press, he won multiple photojournalism awards for his coverage of eight wars around the globe from

1989 to 99. In 2003, he was a Nieman Fellow at Harvard before being named VP and Director of Photography for the Associated Press, a position he held until 2016. And then we have Claire Leibowicz, who's the head of AI and Media Integrity Program at the Partnership on AI, or PAI, where she has worked since the organization's inception. She has developed AI governance strategies, researches responsible AI, and informs technology policies. At PAI she has led efforts to create best practices for AI technologies impacting digital media and online information with a focus on trustworthy information and local new sustainability. She launched PAI's work on AI safety, human AI collaboration and fairness challenges, and has been a journalism fellow at Tablet Magazine and a Bellagio fellow at the Rockefeller Foundation. So obviously a power back panel. I'm going to let them get to it. So we're going to start right off with Craig Silverman, so Craig, come on up.

Craig Silverman Thank you, Robert, and good afternoon, everyone. So I'm going to be taking care of the manipulation portion of today's program. I'm going to kind of take you to the bad place, and as we move along after, things will get a little bit more hopeful, okay? But first, to kind of take a look at the range of sort of deception and manipulation that's being done using synthetic media, I want to take you to a wonderful place called Facebook that you might be familiar with. And so I was doing some reporting recently looking not necessarily for synthetic media manipulation, but kind of looking at how a 100% fake false information was sort of evolving on Facebook as they were preparing to pull back with fact-checkers in the US. And this is the kind of thing that I was sort of encountering. So this is a page with well over a million followers, and what you can see here is you know a post claiming that Elon Musk announced that he had acquired MSNBC. This is something he had sort of joked about at one point, if you remember. And you can tell it's totally legitimate because there is a photo of Elon Musk holding a piece of paper that says, contract signed with MSNB, which is obviously what a contract would look like if he had signed to buy MSNBC, right? It's very legitimate. So this is, you know, this is a page, I identified dozens of pages that were just posting 100% made up headlines. It reminded me of reporting that I had done back in like 2014, 2015, 2016. And they were posting these made up images and made up headlines all over Facebook. But the difference this time from 10 years ago is that they were easily able to generate images that were a little more convincing than just the stock image or thing you might have seen before. And this is another sample from the kind of page I came across. So if you remember, there was a Malaysian Airlines flight that just sort of disappeared off the face of the Earth a while back and was the subject of a lot of reporting and mystery, and the good news is they found the plane, apparently, and everyone is intact. There are the skeletons still sitting in the plane at the bottom of the ocean, apparently. There was a whole set of pages with just incredible archeological discoveries and mysteries solved in this kind of thing. And again, like both of these generated thousands and thousands of engagements on these posts on Facebook. And of course, you know, this might have taken somebody with a little bit of Photoshop skills at one point. Today it''s a few prompts, and you can easily generate something like this.

And that's sort of like the news realm of it, of what's out there on Facebook, but then there's what sort of is often now referred to as AI slop. And this is an image that was generated for, primarily for an audience in India, and the people who generate it, which we'll meet in a moment, really what they're trying to do is just generate images that get a very strong emotional reaction out of people, and they're sort of advising people to say, look, starving poor people do really well on Facebook. It generates a lot of sympathy, it generates a lot of comments, gets a lot of engagement. And there's a whole series of pages that are just posting nothing but really emaciated looking people trying to get folks to engage with it.

And on kind of a stranger and lighter note, you might also encounter stuff like this. You know, there are pages that simply generate images of like, you know, older people, or soldiers coming back from having served overseas, or people in front of birthday cake saying, oh, it's my birthday, I'm 99 today, nobody loves me. And it makes people engage with it and talk. And in this case, this is a little bit over the top. I mean, the guy's cake has 259 on it. I don't think people necessarily believe he's turning 259, but if you look at the numbers on it, that got 11,000 reactions, over 2,000 comments, and over 500 shares. And so whether people believe it or not, it's serving the purpose, which is to get people to engage with it. And engagement for its own purpose is obviously a good way to get your social media account to grow. Bigger following, more exposure, build it up over time. But there's actually a very clear financial motive for basically kind of all this stuff. And so people, some reporters at 404 Media, when it came to the pictures of emaciated people and some of the other slop out there, what they found was that there's a whole world of kind of YouTube videos and how-tos, telling you, hey, here's how to make money using AI-generated images on Facebook. Because Facebook has a bonus program that they've had for years, since 2021, where if your page or your account is part of it, you actually get a cash payout at the end of each month based on the engagement to your content. And so these how-to videos from a lot of these guys who are in India and other places are basically saying, here, here's how to generate a whole bunch of AI slop that works really well, that can earn you money at the end of the month. And that was also sort of related to the story that I published, which was finding dozens and dozens of pages, completely made up stuff, many with AI-generated images, which again, in the area where fact-checking is going away in the US on Facebook, suddenly their position to potentially be part of this new and actually this year expanding content monetization program on Meta. So the model there is very clear. You can generate a lot of this stuff quickly, easily, relatively compelling, gets people to react, and now you can actually get paid for it. The business model piece wasn't necessarily there before. And this is another example here from another story I worked on where we found, you know, more than $25 million worth of deceptive ads about politics and social issues that had run across Instagram, Facebook, other Meta-platforms. Let me play the video, and you can get a sense of kind of what we're dealing with here.

Video Regardless of what you think of me, this can help everyone. I'm about to put over $6,400 in your pocket that you need before Christmas. This new program is an economic incentive program that is open for every single one of you, and most people don't even know about it. So pay attention. All those who call in and enroll get sent the funds you can use in less than 24 hours. And they made it.

Craig Silverman So that sounds like Donald Trump telling you, basically, here's a free $6,400. All you have to do is fill out this form, and it's in your bank account within days. That is not actually Donald Trump. That is obviously an AI generation version of his voice. And we found just, you know, overall we found about 160,000 deceptive ads. Not all of them used AI, but a lot of them did. They used Donald Trump, Joe Biden, Eric Trump, Melania Trump. I mean, just on and on and on. In some cases, they used deep fakes where it looked like the person was talking. With that one, they didn't even really bother syncing up the audio with him talking. They just showed a bit of video of him and then the voice sounded pretty convincing. And we found these across more than 300 pages. Some of them were sort of more overtly political, but a lot of them were really leading people to very deceptive means to hand over their personal information. In some cases their health coverage was changed as a result, and they might have lost health coverage, in terms of Obamacare. In other cases, they thought they were buying a bunch of gold Trump coins, and they got signed up for an expensive monthly subscription. And the AI element of this is

really the convincing piece of it. I spoke to a guy who bought a whole bunch of Trump coins. One, he thought they were real gold coins, but they're not. But two, he felt like these were actual products endorsed by Donald Trump, and he wanted to buy them. And it's helping convince people to do things. And it is widely used in deceptive marketing, as well as the sort of deceptive content that I mentioned before. And of course, it's not just the sort of financially oriented. There have been a lot of documented cases over the last couple of years where AI generated audio, video, images, manipulated images and what have you have been used in service of, you know, political disinformation campaigns and other elements. These are just a few of the examples that have been documented from around the world. We obviously had a lot of elections last year. And so this was something that started to get into it, and you know there's always been this fear that, “Oh AI is just going to make nobody able to believe anything, right?” It doesn't necessarily go to that extent. There wasn't an example of someone being elected based on an AI manipulated video or image, but it's kind of getting baked into the information ecosystem. And there's real challenges about how we can do more to sort of call this out, and point it out to people, and help them not take action like hand over personal information, or pay for things based on it. And the dynamics here, in a way, are pretty obvious, is there's always been manipulation, there's always been disinformation. But this quote from Renée DiResta, I think, sums it up pretty quickly, is well on that, is social media makes it so easy for anyone to distribute and build an audience to get out there, which has many benefits, but also drawbacks. At the same time, when you have generative AI, and I can generate not just one image or a video, but sets of it, many, many versions to try out and to test, the ability of the cost and the time to creation goes far, far down, and it's so easy to do it right now.

But it's not just social media. So in addition to the examples from Facebook and what have you, I mean, there have been issues and examples, again, for media pointing out that in some cases, some collections, digital collections and libraries and other collections, are including AI-generated books. There were examples where there was a book about buying mushrooms and eating mushrooms and foraging for wild mushrooms that actually had some advice that could have made someone very sick or even died, and it was generated using AI. So it's starting to kind of seep into other elements of not just the obvious stuff that you might expect on social media. And on a layer above that, it's probably in everyone's pocket right now or in your bag. The phones that we use, the process that's used to capture a digital image, in its own sense, there's often a little bit of sort of guessing and processing that's happening. But at the next level is the most recent pixel phones from Google, for example, have incredible image manipulation capabilities built right into them. And this is a post here, if you look at these images, basically it's just on the phone itself saying, “Hey, insert a monster,” or “Hey, insert a car where there wasn't one before.” And the LinkedIn post is from one of the best forensic, image forensics experts, Dr. Hany Farid talking about how for many years he said, “Look, these AI generation engines, they're not perfect yet. Sometimes there's an extra finger. We can spot artifacts. There's things we can do.” And then when a colleague handed him this phone, and he realized he could insert a monster that looked like that or a car that looked like that, he said “Okay, maybe we're here.” The good news is that the blue versions, they figured out a way to potentially identify artifacts in it. But we are getting to the point where that kind of automated detection we certainly can't rely on that completely, at all. And to the naked eye, the things that we might have been spotting and using, the levels are getting to the point where it's going to be really, really difficult.

And so to complete your tour of horrors, just to summarize a few things, one, these tools are widely available, many of them that are free, produce very high quality images and video and audio, and they're getting baked into our phones and to other things we use

every day. I think a lot of people are aware of this. A lot of us know about deep fakes, about AI-generated images, but that doesn't mean all of us know how to recognize them, and it doesn't mean that we think about that every time we're consuming media, and that's something that people really hit on and exploit. And that's because the motivation of people to earn money, to fool people, those kinds of things, that's very powerful, and the development of the manipulation technologies is usually ahead of the detection, which means that platforms of all types are seeing this there. And in some cases, they are building in incentives and tools to generate, and building in ways to easily earn money, which again creates more motivation for people to use this, to think about ways to abuse it, or to push the limits. And so the bad news is there isn't one easy solution. There's no one detector or anything we can use, but the good news, and you'll hear more from other people here, is that there are initiatives being worked on in ways to kind of call this stuff out, to point it out, to build kind of metadata and other things into the ecosystem so that we can start to learn if we're interacting with potentially manipulated or AI-generated images. And that's one hopeful thing amidst this sea of AI slop that unfortunately we're all swimming in right now. So thank you.

Robert Quigley Okay, next we have Carlos Eduardo Huertas.

Carlos Eduardo Huertas Could you refresh, please? Could you replace the link, please? Mm-hmm. Could you check, please? Sorry, technical issues.

Hello, it's a pleasure to be here introducing this invaluable experience about the elections in Venezuela. The presidential election held on July 28 of 2024 put Venezuelan journalists to the test. It had to inform audiences following the best quality and independence standards while being in a context of extreme political polarizations. In response, two initiatives were developed to counteract the blockade imposed by pro-government forces to obstruct media outlets, spread lies, and restrict access to information. Those initiatives were known as Venezuela Vota, which became the largest streaming of an electoral event in Venezuela and La Hora de Venezuela, which intended to explain and amplify at the international level all the content produced by media outlets that were being persecuted. The limited resources of 14 digital platforms came together and set in motion different strategies to keep their word safe and sound. Artificial intelligence played an essential role in one of those strategies. I'm talking about Operation Retweet. Let me introduce, so they can explain in their own words. They are La Chama and El Pana. Chama, can you say hello to the participants of the International Symposium on Online Journalism?

Audio Chama Me da gusto estar aquí.

Carlos Eduardo Huertas Oh, okay. Well, Chama, please, could you, can you speak in English?

Audio Chama Sure, Carlos. I bet our English is better than yours.

Audio Pana On a more serious note, we really appreciate this space to share what we've done.

Carlos Eduardo Huertas Oh, she's bullying me now. Please could you explain the meaning of your names?

Audio Chama In Venezuela, Chama and Pana are female and male nouns used to show trust and closeness. This was something that our editors were after. They wanted to

leverage the trust that the media outlets were willing to place in two avatars to disseminate their content.

Audio Pana That is why we speak with a Venezuelan accent, and we look like regular chamas or panas from Venezuela.

Carlos Eduardo Huertas Okay, let's watch part of the video that introduced them to the world.

Video Hi, I'm La Chama. Hi, I'm El Pana. If you didn't know, we're the host of Operación Retweet, and we've been sharing information in Spanish about everything that's happening in Venezuela. Now we will start doing it in English as well. But before we continue, and just in case you haven't noticed, we want to tell you something. We're not real. We were created by AI, but what we share is very real, verified, and high quality content created by journalists. But why are Venezuelan journalists using AI though? As you know, persecution and repression against everyone the government considers a dissident or sees as a traitor have increased since the elections on July 28.

Carlos Eduardo Huertas And Chama, what was the response to your unusual aspect?

Audio Chama I don't mean to brag, but did you watch that video posted by Shakira a few days ago in which she hugs a younger version of herself? Well, it was generated using AI, and it had 400,000 views in 3 hours. In over a day, our first video had the same amount of views, and we are certainly not Shakira.

Audio Pana Operación Retweet has had wide coverage. That's why you can find features in top media outlets around the globe, such as The Guardian, CNN, the BBC, and El País, as well as in hundreds of publications worldwide in at least 10 different languages.

Carlos Eduardo Huertas And now that you mention it, can you explain the meaning of the Operación Retweet?

Audio Chama In Venezuela, the expression “retweet” goes beyond the former Twitter. It is used anywhere to say that something is valid and that you can share it. That is precisely what followers of Operación Retweet are doing with their content, sharing it with relatives via WhatsApp or any other social media, even in Telegram and WeChat, which the government is now trying to impose.

Carlos Eduardo Huertas And what are the advantages to using artificial intelligence to provide the oxygen needed by journalism in Venezuela?

Audio Pana Since we're avatars, we're allowed to report in a context in which uncertainty and danger increase by the minute for regular journalists. There are no laws against us.

Audio Chama It is an innovative and safe way to bypass the information blockade and the hurdles to journalistic content.

Carlos Eduardo Huertas Okay, as La Chama and El Pana say, the initiatives that came out with the Operación Retweet are an effort to provide collaborative information of quality, and that has been verified by Venezuelan journalists before, during, and after the presidential election. As part of the strategy, it was decided that the content will not be disseminated in a single channel, but that it will be shared on social media. But amid

accusations of electoral fraud and crowd doubt ensued consequently, the teams of both initiatives have explored new strategies. Since mid 2024 until now, 25 people working for the press have been detained. Twelve of them remain in prison, and ten were released by the accused of terrorism merely for doing journalism. For example, Ana Carolina Guaita was one of the first to report the knockdown of many statues of the late Hugo Chávez. She was kidnapped and was missing for over 50 hours before knowing where she had been taken to. Maduro's crackdown continued by blocking X, the main platform used by Venezuelans to read the news. This block continues today, sometimes. Then he went for WhatsApp, used by people to communicate and come together.

Video WhatsApp, llegaste al final. No te quiero más, WhatsApp. Todos, los cinco poderes están bajo ataque de WhatsApp. Yo lo voy a hacer voluntario. Hagalo también. WhatsApp, imperalísmo tecnológico. Estás atacando a Venezuela. Entregaron todas las listas. Aquí dice desinstalar. Te fuiste WhatsApp. Si te visto, no me acuerdo. Seguimos paso a paso, liberando a Venezuela del golpe de estado cibernético, fascista y criminal. (WhatsApp, you've gotten to the end. I don't want you anymore, WhatsApp. All, the five powers are under attack from WhatsApp. I'm going to do it voluntarily. You do it too. WhatsApp, technological imperialism. You're attacking Venezuela. They gave you all the lists. Here it says uninstall. Here it say uninstall, you left WhatsApp. If I see you, I won't remember. So we will go step by step, freeing Venezuela from the cybernetic, fascist and criminal coup.)

Carlos Eduardo Huertas That is why the use of artificial intelligence shields journalists from persecution and repression while upholding the right of people to be informing, exactly what journalism is about. We have put artificial intelligence at the service of collective intelligence in an unprecedented form of collaborative journalism in the region. That is where the powers come from. In two months of intense war, 528 pieces of content were shared and enacted in countries via 74 media outlets in a massive coverage. We also took advantage of technology to guarantee access to more than 25,000 electoral tally sheets, the documents that evince electoral fraud. This material was collected by the opposition, and many countries acknowledge it to be true. In Venezuela, this material is considered a crime to publish these documents. This was made possible thanks to Google's Pinpoint, which also uses artificial intelligence. Chama, what has been the most revealing aspect to this strategy?

Audio Chama This initiative has drawn attention to the situation of Venezuelan journalists, the climate of insecurity and the risks deriving from their job in a place with restricted liberties and persecution.

Carlos Eduardo Huertas The collaborative effort has garnered positive attention among audiences. They have extensively commented on the content despite restrictions to freedom of speech. They are thankful to see that journalism in Venezuela is not defeated, and that it has been resourceful and not to carry on reporting.

Audio Pana This is an unmatched effort. As a result, a total 235 professionals based in 31 cities in Venezuela and 15 abroad have multiplied the reach of the content, raising awareness of the situation in Venezuela.

Carlos Eduardo Huertas Well, after two months, the social listening tool Brand24 revealed that our reach was about 41 million followers. That means we multiply our estimate by seven. Another striking data, the tool has a metric called the ABE, an estimate of how much a media campaign will cost. The amount was estimated at $2,780,000. What do you think about this data?

Audio Chama This amount of almost three million can't be equated to the value of the collaborative work of over 100 colleagues who are still committed to reporting information of quality in a context of repression. The challenge faced by this alliance is to continue working regardless of these adverse conditions for democracy.

Audio Pana Correct, Chama. That is why we ask you to help us disseminate our content using the hashtags, #VenezuelaVota and #LaHoradeVenezuela.

Carlos Eduardo Huertas The situation in Venezuela, as everybody knows, is getting worse, and the country's journalism needs our help. As you can see, each new follower is a citizen that we pull out of the quicksands of disinformation that is exposed to a more diverse offer of media outlets to stay informed. Feel free to reach out if you are interested in contributing to keeping valuable journalism alive in Venezuela. We are open to new synergies that lead to more and better journalism, and thank you for listening.

Robert Quigley Thank you, Carlos. Thank you Carlos. So you've heard the bad uses of AI and the good uses of Ai. Santiago is going to start talking about how we can navigate some of this. Santiago Lyon.

Santiago Lyon Hello everybody, it's great to be back here at the ISOJ. 45 years ago, Rosental walked into a bar in Madrid and was flipping through the classified ads section looking for a place to live. And a man at the bar looked over and said, you know, you'll never find anything in that newspaper. And Rosental said, why? And he said, because I'm a journalist too. And I have an apartment that I would be willing to rent for you when I take a year's sabbatical. Sure enough, Rosental rented my father's apartment. And a year later, when it came time to hand over the keys, I was there as a 10 or 11-year-old kid and thus began my long friendship with dear Rosental. Thank you. So here we are 45 years later, and a lot has changed. I went on as you heard in the introduction to have a very exciting, and dangerous, and amazing, intense career in journalism as a war photographer, as a media executive. And I feel very fortunate to be working on this content authenticity initiative now, which in many ways is a logical extension of my life's work in journalism.

So this initiative was kicked off by Adobe way back in late 2019, really in response to what we saw coming around the corner, which technology companies have research labs and you really get to see what's coming. And the research scientists at Adobe could see AI was coming, could see generative AI was going, and we decided that we needed to do something to make the world a safer place, to make the internet a safer place. And so everything that you're about to see pretty much is open source, which is the spirit of this content authenticity initiative. Images like this have been making the rounds, this one purporting to show flooding from North Carolina last year. We saw earlier ample examples of the kind of internet slop that's out there caused by AI. And it's not hard for images like this to go viral. They make an emotional connection or whatever the connection is. So as we've been working over the last five years, the community that we're working with is growing and growing and growing, and we're actually approaching 5,000 members now. And you can see here, it's really a who's who of, in this case, major media companies, both news agencies, as well as consumer facing news organizations, social media companies, hardware, software, advertising, auditing, and we really think it's the weight, size, and heft of this community that's going to make this initiative ultimately useful.

The initiative is focused on this concept of provenance. Provenance is a term that we might be familiar with from the art world, the provenance of a painting, who owned it,

where was it, et cetera. In this particular case, we're talking about the basic facts about the origins of a piece of digital content. Where did it come from? How might it have been edited along its journey? And then sharing some or all of that information with the viewer so that they can get a better sense of what it is they're looking at, and as a result, make a better informed decision. So provenance really is about proving what things are, as opposed to detecting what is false. Detection technology, while certainly useful on an ad hoc basis, has a number of problems. It's not scalable. It's not particularly accurate. It's invariably an arms race with bad actors staying one step ahead of the latest software. But provenance is a little more durable. And so the journey thus far has been interesting. We've been working with creators around attribution. We've been with news media companies around trust and bolstering the existing trust model that's at the core of all journalism. We've been working with generative AI companies around transparency. Corporations are now becoming interested as it relates to brand reputation, and consumer protection, insurance companies, auditing, law enforcement. How is a court of law nowadays supposed to know that digital evidence is what it appears to be? And a whole host of other examples that you can see here. And really what this tells us is that here we are in 2025, and there's pretty much zero empirical evidence about the origins of anything that we consume online. And so we pass judgment on things as to whether to trust them or not based on rumor, hearsay, gossip, trend, any number of things that sometimes serve us well, but have been increasingly confused and disrupted by the phenomenon of generative AI, which has been around for the last couple of years. And generative AI has really cast a shadow of doubt over everything. It's created what some people refer to as the liar's dividend. In other words, if I can lie about something, the payment I get for lying about it is that I call everything into question. And for that reason, provenance is likely to become table stakes or fundamental to just about everything digital going forward.

So with that in mind, I want to introduce this concept of content credentials. Content credentials come from an underlying standards organization that I'll get to in a second that sits within the Linux Foundation, but content credentials are really a digital nutrition label for online content. So in the same way you would go to a supermarket and pick up a food product and quickly see what's in it, here you can see what's in a piece of content by using content credentials, applicable to all sorts of digital formats, as you can see there. So, content credentials, then, in order for them to work, we have to work throughout the digital supply chain. The content credentials themselves, the technology, comes from the standards organization called the C2PA that I'll describe in a minute. But in order to get this to work, we're working in a number of areas starting with capture. So here we're working directly with the hardware manufacturers, cameras, smart phones, and also generative AI engines to establish the origin of the file upon creation. Already you can buy cameras and smart phones that come out of the box with this technology embedded into it, and you as the user can choose to turn it on or off depending on what the default setting is. And it allows you to establish the origin of the file. Next we move on to editing and here we're busy getting this technology into editing programs. Adobe programs of course, Photoshop, Lightroom, coming to Premiere for video soon but because this is an open source initiative, we're also working with Adobe's direct business competitors like Capture One or Photo Mechanic. This allows you to track whatever changes are made to the file in the editing process. I crop an image, I tone it, I darken it, I lighten it. It allows you to go back and review all of those things. So it creates a digital edit history. The next area we're looking at is publishing, and here we're working with publishers of all kinds as well as the content management system environment in order to maintain the integrity of the content credentials so that they can become visible to the viewer. So the reason we're doing that is because then the viewer can engage with the content credential represented by that little CR icon and see some or all of the edit history. So this is the symbol, if you like, that we're

working to make as ubiquitous and as widely understood as possible so that when people see it, they'll recognize that it indicates that there's underlying information about the origins or the provenance of that particular file, and if they're interested, they can click into it and see what's going on.

When it comes to AI, Adobe, like a lot of companies, has built generative AI products. Adobe's product is called Firefly. It's text to image, text to video, very similar to OpenAI's DALL-E 3 and Zora, or Midjourney, or Stability, or Bria, or any of those other companies. And if you were to come across this image, you would click on the little CR icon and the nutrition label opens up. The good news is that everything that comes out of Firefly as well as OpenAI's new tools gets a content credential attached automatically upon download, and the rest of the AI ecosystem seems inclined to go the same direction. And so we're working with them to make this an industry standard. Now all of this technology comes from the Coalition for Content Provenance and Authenticity or the C2PA. The C2PA is a standards organization. Those companies there are on the steering committee. As a standards organization, it sits within the Linux Foundation, which means that it's unencumbered by IP, nobody owns this technology. It's available for anybody and everybody to use. It's also about to become an ISO standard next month, which will give it that stamp of approval from the International Standards Organization. And currently there are about 500 technology companies with representatives hard at work in the C2PA to make this the gold standard, the most robust provenance standard out there, and that's clearly the case at this point. Adobe's role is to chair the technical working group and run hard on this, but we do that with the other members of the steering committee and the other 500 or so technology companies. So when we think about these two things then, we have the C2PA, which is the blueprint upon which everything is based.

And then we have the Content Authenticity Initiative, which is tasked with accelerating implementation of the C2PA standard across a wide variety of industries, journalism, one of them, but many others as well. And so we do that in a number of different ways. We build open source tools. We spend a lot of time and effort and money at Adobe building open source tools that sit on GitHub for developers to implement. We run a very vibrant Discord or chat channel for those developers to get their questions answered. We also do a lot of community events and public speaking. My job is really traveling around the world, talking about this stuff, growing membership, working with lawmakers, policy makers, so they have a good understanding of the opportunity here. And then the third thing we do, obviously, is productize this into Adobe products, but then other members of the initiative are also productizing it. And we think that over the next year or two, we're going to see this notion of provenance as a service become available for enterprises, companies, large, small, in order to help accelerate implementation. Last year was pretty good. We got a lot of stuff done. This is just a kind of an overview. OpenAI, Meta, Google, TikTok, LinkedIn is starting to deploy this. Meta and Amazon again. The U.S. Department of Defense is starting to put content credentials on their repository of official military images. YouTube is starting to implement it, and our friends at Time Magazine gave us a little pat on the back as well. So really good progress, you know that combination of capture, edit, publish, the progress that we've made in all of those areas is very encouraging and we're really starting to see this take off.

Of course the technology is only going to get us so far, and so it's very important to remember the two other elements of this that we're also focused on. One is the area of education. And by that, we mean media literacy in the classroom, and so we've developed free curricular materials to help educators around the world educate young people. We're also talking about societal education, and as we get a critical mass of implementation, this

notion of public service announcements. And then we're talking about consumer education. Dear reader of the Daily Bugle, here's what this little thing means, et cetera. And then the other element that's also important is the policy element. We're seeing around the word legislation making its way through parliaments and other processes that is likely mandating transparency when it comes to AI and possibly including transparency when it comes to other file formats. And so we think it's very important that those lawmakers and policymakers are well-informed, so we spend a lot of time in those centers of power talking to all of those folks.

In the news media industry, progress has been somewhat slow. I would say the standout organization that's made the most progress with this is AFP, the French news agency, who are working together with their preferred camera provider, Nikon, and a French water marketing company called IMATAG, to really get this onto their images that they distribute around the world. And the rest of the news media has now gathered under the auspices of the IPTC, the International Press and Telecommunications Council, which is a 55-year-old or so organization aimed at standardizing metadata, and they've become the watering hole for journalism organizations. And we're starting to see a little bit of progress there, which is good. So it is in many ways a glimmer of hope or a lighthouse as we advance in these somewhat turbulent times. And if there's a call to action here, it's really to join us. Membership is free. We welcome individuals, companies. By joining us you get into the community, you get information, you get invited to events, there's information sharing, best practices, etc. And really we take great strength in numbers and think that there is a real benefit for individuals and organizations to consider taking a look at our work because really success is dependent on this notion of ubiquity. A world where everything has content credentials and the content that doesn't will be questioned because of the absence of underlying information. And to be clear, we're not trying to be the arbiters of truth here. Far from it. What we're trying to do is provide additional information so that consumers of news and really anything else online can make better informed decisions about what to trust, which will help guide them as they make their way through this perilous digital reality that we live in. Thank you.

Robert Quigley Thank you, Santiago. Claire is going to wrap it with us, for us with some more strategies on navigating this. So Claire Leibowicz.

Claire Leibowicz Thank you. So I have the privilege of going after two journalists, a former journalist and I guess technology adjacent person now if he would describe himself as such. And I'm going to take us back a few years to 2019, when I found myself, well before the trend where we see everyone on using DALL-E's latest model, I don't know if some of you have seen, to turn themselves into anime this week. And I had brought together, through my work at the Partnership on AI, which is a multi-stakeholder non-profit, an unlikely cohort to think about the future of media that would be AI infused. And we had a computer vision scientist from a company in our partnership speaking with a radio journalist and a human rights advocate talking with a platform policy maker, and all we're thinking about what we then described as the emergent threat of AI-generated media to public discourse, to civic life, to information quality, but also the opportunities, which we saw Carlos talk about in terms of storytelling, and privacy preservation, and how we balance those two things.

And fast forward to now, and I'm going to skip some of these examples, but I'll maybe point out something slightly different than my colleagues have pointed out. This is just from last week, if you went to Google and searched about what was going on with SpaceX, you might see this magazine promoted Science Magazine as one of the top searches, and

there were all of these kind of AI-generated images and falsified articles relying upon AI, showing that the companies, even though they are putting forward efforts to kind of de-prioritize this material, are still succumbing to kind of that confluence of generation and recommendation cheapness that Craig and Renee had spoken about. Here's just a few more that based on YouTube policies should not be on YouTube or should at least be labeled as having had AI components, but they were not. These are out of Spain. And even last week, this is an imagery, but many of you may have seen that someone in Italy claims to have put together the first fully AI-generated newspaper. Some may see this as a positive way of getting out information. Others may be more negative. And this is particularly intriguing. About six months ago or around the election, someone put out this deep bake emblazoned with the PBS News logo, prompting PBS to have to actually respond on its Twitter and say “this was not from us,” showing that it's also brand reputation, as Santiago was saying, that could be implicated by the presence of this technology. So clearly, we are interested not only in how we understand news as records of reality today, how we make and consume news, how do consumers interpret these signals and disclosures alongside content, and even the legitimacy of our trusted news institutions in this moment, prompting many policymakers to respond. And I'd say the past two years, really, there's been a lot of government activity, whether that's at NIST in the United States, or even just last week, or the last few weeks, unlikely bedfellows, China and Spain, put out policies mandating AI labeling of content. So there's a lot of legislation around the transparency protocols that Santiago was broaching.

But the question that's animated my work, where we work with all these different institutions, was how can AI policy, and that could be norms, standards, actual regulation, support media, both media institutions and artifacts as records of reality, which is what we, I think, really care about. And what should newsrooms do, which is obviously a main point for this audience, but also what responsibility to other institutions have to the news? What does it mean for image generators to bake in these signals that Santiago was talking about so that people trust news downstream? And that's something we think about, this pipeline of institutions that are all integral to news being vibrant, and accurate, and trusted. And so two years ago, after lots of this convening, and this was in 2023, we put out our own technology policy, in essence, which delineated how to develop, create, and distribute synthetic media responsibly, or at least how to start. It couldn't be completely comprehensive. And we got 18 institutions to initially join, ranging from the BBC, and the CBC, and media, to Adobe, and OpenAI, to civil society organizations, and even Bumble, who is not a newsroom but cares a lot about image manipulation as a threat to trust, just as newsrooms might in terms of their business being so contingent upon trust. It's also in Spanish and French. I encourage you to take a look. And what was important to us was that we'd be able to articulate that even if we have a value like transparency, what does that mean for a newsroom to enact transparency, but what does that mean further upstream for some of these technology companies to do so as well? And this is a ton of text, but it's just extracted from the policy to give you a glimpse into what we were kind of articulating for newsrooms to do. We would talk about prompt adjustments, when you realize certain harmful categories of synthetic content. Some of the C2PA reference came a year before many of those platforms actually joined on to C2PA, which we think hopefully catalyzed their involvement, and other norms around labeling and how they go about doing that. And what was meaningful was that all of these institutions agreed to actually work on in-depth case reports to show how they enact this guidance. Not only creating a body of almost case law for all these different instances of positive and negative uses of synthetic media, but also showing how institutions actually deploy this in the real world, which is very different than just a policy static on paper. And I'll go through, this is just a whole mosaic of those different cases, and those have informed a lot of the policy making

happening in the United States. So a big goal of ours is how do we build out a body of evidence? Including the news media's integral to AI policymaking that will then get funneled into the government activity that's happening in the AI realm.

So two cases I really want to zoom in on because they show how there's not really agreement about what it means to do right by your audience in the A.I. age. So the CBC and the BBC took completely different approaches. And I'd say Carlos' work kind of matches the BBC's approach, which was that you could use deepfake technology to tell a more truthful, honest, authentic, privacy-preserving story, and do that in an ethical way that services audiences. Of course, the BBC, a different context than Venezuela. The CBC said, actually, we don't want to deploy deepfake technology to tell a story, and that was a story about romantic fraud, where they wanted to preserve the identity of the source. They felt their audiences were not ready for that, even with adequate disclosure, whether it's the avatar saying, I'm not real. Or you'll see what the BBC did. So the CBC writes 10 pages, if you want to read the whole thing, but they have this really intriguing quote saying that even with public labels conveying that content is synthetic, they didn't feel there was sufficiently broad public understanding to prevent them from being misled or misinterpreting it as the CBC kind of misleading them by using AI technology. The BBC, on the other hand, talked about how anonymization is a thing we've been doing forever, and this is just merely another mechanism for doing that better. And they did two things to label it, and a lot of this legislation says “label AI content,” but not a lot it tells you how to actually do that. So many of our partners come to us to ask about how we do that. They had an on-screen caption saying faces have been digitally altered. But they also have this audio narration at the beginning, showing that people were consuming this in a different medium, and they even explain their logic for why they are bringing this to the public. They say that the deepfake technology allows us to capture these emotional stories while maintaining the anonymity that Alcoholics Anonymous requires. So they're almost justifying it in their disclosure and providing that explanation for audiences. And in the piece, the commissioning editor talks about there's a real authenticity afforded by the technology. And this is kind of this kind of paradoxical idea that through this technology, you can actually be more truthful or more authentic with your audiences.

And this question of how we communicate manipulations, whether it's through the content credential symbol, or what actually gets people to care, or what do we owe audiences when we use this technology or report on it, is one that's going to become increasingly prevalent in your organizations likely, but also as policy activity mandates it. And I want to walk through just a few tidbits from the cases that talk about how to do this or things that complicate this picture of how we provide transparency, which is an intuitively appealing thing, but sometimes can backfire or not be ideal. So Microsoft, in their case, talks a lot about the use of content credentials on LinkedIn, so giving people image context on LinkedIn. And they did some user research that said very subtle changes in language, saying it was certified by them, or verified by, or signed by, significantly impacted how audiences understood what that disclosure was conveying. Meta's case talks about this wild kerfuffle, where they originally labeled content as made with AI, and many people perceived that as being inherently punitive. The idea being that AI equals false or nefarious, even if it's just editing a pixel out of the background. So they changed the language to soften it to AI info, and they're still working through how they convey AI manipulations without implying that, you know, editing out a blemish on the face is something nefarious or false for a creator to create. Code for Africa, oh, we're going to skip this. Code for Africa talked a lot about their adamance that we have to label all types of edited content, whether it's a blemish on the face or something much more significant like an avatar or AI slop. They have this really profound quote which says, and I'll skip to the

bottom part, “Our assumption that even if, if we use non-disclosure, even a benign usage rather than a malicious usage, if we don't disclose that, that has the potential to erode the trust relationship between organizations and their audiences.” So they feel we better over-label because we're just going to diminish trust if people interpret AI as being false or misleading, more like kind of the CBC's approach to people's skepticism. And Google's case is quite interesting because they're a distributor who has a slightly different challenge than a newsroom. But they talk about this idea of the implied authenticity effect, that if you don't label everything, people are going to assume that anything that's unlabeled is inherently authentic. And again, authentic doesn't mean truthful, which is why journalism is so important, but users don't have that kind of complexity in mind when they're thinking about AI. So we have some other work on this, but I think a few insights to remember about how we provide transparency that complicate our sense of the transparency policymaking. People are not perceiving just AI generated as a merely technical and neutral signal. They perceive it as punitive, they do see it often as commenting on its truthfulness. And that is something that requires media literacy, but also the platforms and other conduits to this information to be very privy to and responsible for. Obviously, most edits we see on AI, even though we've seen some here, are not materially significant. They are those minor edits, and what does that mean for how we have these regimes of labeling that may trigger these systems and labels even when they're really minor? I'm going to skip that one. Social media platforms, interestingly, have been willing to over-label, so there's been this inclination of we're better off giving people more context. So what that means for how people then interpret news and interpret media that's unlabeled is an interesting question that we haven't seen a ton of evidence for.

And what journalism knows well, and I actually think journalism should bring to the platform environment, is that simply disclosing the presence of AI is an imprecise form of transparency for audiences. You need that rich context about where something came from, how it was made, beyond just whether or not a certain technology was used. And to explain if AI is being used, why is that the case? And I just want to complicate a little bit of the kind of policy momentum around standardization of visual signals for these labels. When Santiago presented the nutrition label, it seems very intuitive that you're just providing context. But what I've found really animating for the field and I'd love for journalism to be involved in, is there a normative decision about what gets put in this nutrition label? In what order? What gets bolded? And those decisions, as the LinkedIn anecdote shows, are actually really integral to how people grasp what they mean, and how they are consuming the content that we produce. So I think standardizing this is very important, but also being sensitive to what we're choosing to put and how we're conveying that across the board.

So my concluding point is I've been very heartened by the fact that these two questions are animating many of our partners, whether it’s a platform, shockingly, or it’s a newsroom. What do we owe our audiences? And what are they ready for? And what's interesting is that's inconclusive. And you know your audience best, so I think we still want to embolden institutions to ask that question themselves. But as you can see, we also need to share much more about how audiences are adapting, what our legislation means for people's interpretation of labels and transparency, and also make sure that lots of the technology companies who are producing these models and kind of enforcing a lot of these policies, are doing what they say they will do and actually holding them to account. So thank you.

Robert Quigley Thank you, Claire. So a lot of great information there. Claire mentioned that a new model of DALL-E came out while you were all getting ready to come to this conference. It's changing really, really fast. You've heard the word “trust” a lot. At one

point, we had a word cloud years ago at ISOJ to see what people were talking about, I think, on Twitter at the time. But I think trust would be a big one during this conference, and so I guess my question is, generally, for you all, how optimistic are you that we can win this battle of trust in this age where it's getting better and better to be untrustworthy? Or am I looking at that the wrong way? And then Carlos, for you, have you been able to measure how people reacted to your synthetic reporters? Has there been a feeling that that's still trustworthy information since you played out as verified journalistic stuff, or is there still an AI problem there of people not feeling trustworthy? And whoever wants to start first, feel free.

Carlos Eduardo Huertas I'm going to start. Actually in Operación Retweet, it was very clear to the end that the avatars are not journalists. They do not pretend to be journalists. They are to pretend to be a chama and a pana, like two young people talking about the product that made by professional journalists, and that helps. But it's amazing how it's a very borderline because some people are so confusing. Even, you know, for example, in a panel or in a conference like this one, after the conference, some people approach me and say, oh, the guys are in Zoom? These kinds of questions, you say, "Really, you don't understand that they are not real?" And even if you are very clear about what he's doing, I believe the audiences could be confusing. But this is the new world that we are living now.

Claire Leibowicz Can I just give an example? So several years ago, we did user research. We did a diary study for the academics in the room that we presented at a computer science conference called CHI, and we talked to lots of Americans around the country, across the political divide, and asked them what they thought of Twitter's manipulated media label, which was in 2020, when lots of Trump artifacts were being labeled with manipulated media. And we talked to one Republican in the middle of the country and he said, "Oh, I thought that label was Twitter telling me the media was manipulating me." Thank you for laughing. Implying that there was a complete, that's not what they intended, and people bring their own, they were ostensibly providing transparency. They're doing the right thing, they're responding to this skepticism and giving people more context, and yet someone brings their own way of interpreting the world to that. And to me, that's both a humbling and daunting, indicator of why this is so challenging to help induce trust when people are coming to this from many different perspectives

Santiago Lyon I think it's also important to recognize that we're at the very beginning of what's gonna be a long journey. AI is not going away. It's here to stay. And what I see when I travel around the place, I often ask audiences, so how many people in the room, and I can ask it here as well, how many people in the the room have actually experimented with AI, whether it's large language models or generative AI? Maybe a quick show of hands? And it's not very many for a room like this. And I think it's really important that we not necessarily embrace this, but that we at the very least become familiar with it. Because over time AI is going to become normalized. Already people are starting to use it to save time, there are some efficiencies. But like any tool, it can be used for myriad purposes. Like a kitchen knife can be used to slice tomatoes or can be used to kill somebody. So this notion of kind of taking control of the narrative. And as journalists, I think there's a tremendous opportunity for journalism to explain how all of these technologies work in a way that I sometimes see, but I wish I saw more of, so that readers and the consumers of news media can really understand what's going on. And then, you know, when it comes to the implementation of labeling, which is what I'm focused on, as I mentioned earlier, I think education is a huge component of this. And I'm curious to see how governments around the world modify their education policies so that media literacy becomes much more

robust, because it's a very irregular landscape right now. In a few countries it's taught well, but in the vast majority of countries, media literacy is not taught at all, or if it is taught, it's not taught well. And that's a tremendous opportunity, I think, for governments and countries to really help their populations understand what's going on here, not to mention the societal education and the consumer education.

Robert Quigley I got a question specifically for you, Craig, from Josh B. These are on this iPad, these are questions from you all. Why do you think platforms aren't doing much about these fake AI ads? You'd think that ad managers could easily check ads and take them down quickly.

Craig Silverman Yeah, so I mean, there's a couple of pieces on it. One is always a factor with digital platforms is scale. So the amount of ads running at any given moment, whether it's on Meta, whether it is on Google's platforms, is just astounding. And so there is no ability for ongoing human review of everything. And so they are obviously training models to kind of spot ads and try to say, “Does this one look OK? Does this one need to be held back to human review? Does this need to rejected?” And I think, you know, one issue is they are not investing as much in that oversight element and review element as they could be. And I say that not just my own opinion, but having spoken to a lot of people who have worked on those teams at these places, you know, it's easy to spend more money at these companies on the things that are revenue generating, and actually doing more to stop ads from running prevents you from earning more money. So the incentive is not the best in that scenario. And I think that's a challenge. You know the other part with these ads is that it is genuinely a challenge to detect any type of manipulation or detect any kind of deep fake. There is no one system you can run a video or an image through that is going to hit you with 100% accuracy, and we probably never will have that. So to be fair to them, that is difficult, but what isn't difficult is that a lot of times it is big public figures like Trump or celebrities or business leaders that are in these ads that are deceiving people, and the ability to do basic facial recognition of well-known people like that, of which you have massive data sets of images to train on, that I think there's a bit of a less of an excuse for. I mean, one positive thing is Meta did recently announce they're rolling out more facial recognition for well- known celebrities. They're running tests in parts of Europe and other places. But I think, there has been a lack of investment. I think they like the scale, they like making it so it's very easy for you to run ads, and they don't want to put friction that prevents people from spending money on the platforms. And that being said, there are some technical challenges, but I think if they put the full weight of their engineering teams, they could certainly do more.

Robert Quigley And here's one from Kerry Cochran. This is for you, Santiago. Why should the credential nutrition label be trusted? Can't AI generate this?

Santiago Lyon Well, the technology behind the content credential is made up of three things. It's made up number one of what we call cryptographically secure metadata or asset hashing. It's technology not dissimilar from what you use for online banking. So it creates a unique alphanumeric code that links the identity of the content creator to the content, using what's called a signing certificate, where you essentially buy a piece of digital ID that you then link to your content. That's one thing. The other thing that we use is invisible watermarking so that even if the content credential, even if metadata gets stripped off, the invisible water mark is generally resistant to that. And then the third thing we use is fingerprinting technology, where you store some of the information about the original file on the cloud, and you're able to compare it with the original image at hand. Those three things are stronger than the sum of their parts. In other words, in isolation, they all have

vulnerabilities, but when you combine them, it's quite robust. And so, the interactivity of the content credential is key here. It's easy to fake the content-credential logo, but if it's not interactive and doesn't take you to information about who the original signer is, then it's clearly not very useful. So a lot of this has to do with not only implementing the technology but educating people about how to interact with the content credential. And so we spend a lot of time with focus groups and design teams trying to figure out what information is valuable because it's not a binary proposition here where everything is either AI or it's not. It can get quite complex blended content, and you can have provenance information that can be quite voluminous. And so then the challenge is how do you display that in a way that's useful to the consumer or to the viewer? And one of the bigger issues we're thinking about is really what constitutes a material change to an image. If I decide to saturate or make a little more intense a sunset picture that I make on my smartphone because that's the way I saw it and the sensor on the camera was limited and didn't show it that way, that's a fairly innocuous change. If I put my head on someone else's shoulders, that's a major change. So what's the degree there, and how do you communicate those degrees to the viewer? So there's obviously work to be done on all of this, but the core technology is pretty robust. Really now it's a question of adoption, socialization, education, and especially as it relates to policy and standardization.

Robert Quigley So another question from Raul Flores says, do you think there'll be a time when AI becomes government regulated? If so, how far away do you think we are from that regulation?

Claire Leibowicz They're all looking at me. So also, what do you mean, I think it's always interesting to ask what you mean by AI regulation, right? Because AI policy touches so many other domains. It touches copyright, it touches intellectual property, it touches even in the context of like other parts of the law with privacy and all of those realms. So I think there is kind of precedent from so many other areas that's just confronting AI now. And the courts almost can't keep up with the way in which, let's say, let's use copyright. That's going through the courts, thank you, New York Times, but that's going to go through the court for a while. And there is precedent, but that, I would argue, is AI policy making that will really dictate the behavior of companies if they go in a certain direction. And it already is affecting how I think they conceptualize the world and navigate whether they're covering their, you know what, or they're trying to figure out how to move that forward. In terms of in the US, and this is a global question, Europe has been ahead, the EU AI Act is about to come to fruition, it's being implemented. The DSA, you could argue, is AI adjacent. There are so many realms. So we do see it in the US, the executive order under the Biden administration is being stripped away, but that was largely symbolic as well. At the state level, there are lots of different AI policies happening. So I think we do see it, and I think what's fascinating is that AI is touching so many fields, realms of life, that it's just going to kind of confront existing legislation. And that's a question that those of us in AI policy ask, when do you need an entirely new regime for thinking about information or fairness or bias that warrants a new piece of legislation versus we can rely upon existing norms? I saw the big zero, so it's telling me to be quiet.

Robert Quigley I was begging for one. I was like come on, one. So because I want to ask one last question, so there's several students in this room who are hoping to go into journalism. Like how would you like position yourself if you were back in that spot? Like what are some of the things they should be doing?

Santiago Lyon I think they should be learning how to use the technology to their advantage. I mean, I did a little experiment the other day. It's a truism in journalism that

you get the answers to the questions you ask. And you know, stupid questions get stupid answers. Good questions get good answers. And so I asked ChatGPT to come up with 10 tough questions for Vladimir Putin for a theoretical interview, and I was really impressed with the way the questions that in this case, ChatGPT formulated, because they were cutting, they were accurate, they were precise, they were difficult. And I think just something as simple as that can help journalists do their jobs. And I think there has to be a willingness to use these tools in a way that improves journalism, and I think there are many ways to do that. It's not a substitute for journalism. It's not a substitute for a good human, you know, journalistic practice, but I think it can certainly enhance it and make it better.

Claire Leibowicz Can I offer, not as a journalist, but having some degree of technical literacy beyond use that allows you to call BS on the part of some of the tech companies, I think would be enormously, not only empowering for your own sense of your ability to ask these tough questions and determine what is the right question to ask of a technology company or technologist, if you are interested in that, but also just how to be someone in forming the world in a moment where technology is touching everything. I see this in my career and I'm sure in journalism just taking an intro computer science class would make you even better positioned to interrogate what's going on in the 21st century.

Robert Quigley It was an honor being on this panel with amazing people who are thinking through these issues, and it was an honor being here at ISOJ and wrapping up the International Symposium on Online Journalism, for Friday anyway, so thank you.

Summer Harlow Thank you so much, Robert, Craig, Carlos Eduardo, Santiago, and Claire. This was a wonderful ending to a wonderful conference. Now, if everybody will just bear with me for a couple of seconds. Oh, y'all can leave. You don't have to stay, sorry. Thank you again. First, just a really, yes, another round of applause. First, just a quick announcement for those of you who are attending the Coloquio mañana. Si van a participar en el Coloquio mañana, empieza a las ocho por la mañana, por el desayuno, no a las nueve, la bienvenida y todo.