



CENTER for NEWS,
TECHNOLOGY & INNOVATION.

Using global research & collaboration
to promote an informed digital society.

FOR RELEASE APRIL 9, 2024

Watermarks are Just One of Many Tools Needed for Effective Use of AI in News

A Global Cross-Industry Group of Experts Discuss Challenges
and Opportunities in an AI-Incorporated News Landscape

FOR MEDIA OR OTHER INQUIRIES:

Amy Mitchell, Executive Director

press@innovating.news

innovating.news

RECOMMENDED CITATION:

Center for News, Technology & Innovation, April 2024, “Watermarks are Just One of Many Tools Needed”

Watermarks are Just One of Many Needed Tools for Effective Use of AI in News

A Global Cross-Industry Group of Experts Discuss Challenges and Opportunities in an AI-Incorporated News Landscape

“The literature to date suggests that watermarks and disclaimers ... won’t be a silver bullet.” But they could be helpful — alongside experimentation, model transparency, collaboration and a thoughtful consideration of standards — in differentiating between harmful and helpful uses of artificial intelligence (AI).

Indeed, journalism today — the production and dissemination of fact-based news to inform the public — takes all of us: journalists to do the critical reporting, technology to enable distribution, access and information gathering, research to evaluate impact, policy to support and protect all of the above, and, importantly, the public’s involvement and interest.

So concluded a day-long discussion among leaders from around the world in journalism, technology, policy and research. It was the second convening in a series hosted by the [Center for News, Technology & Innovation \(CNTI\)](#) on enabling the benefits — while guarding against the harms — of AI in journalism.

Co-sponsored by and held at the [USC Annenberg School for Communication and Journalism](#) Washington, D.C. offices, the Feb. 15 event brought together technologists from Google, Microsoft and Scroll, journalists from the Associated Press and Semafor, academics from USC, LSE and CUNY, former members of government and researchers from UNC, Yle (the

Participants

Charlie Beckett, LSE
Anna Bulakh, Respeecher
Paul Cheung, Fmr. Center for Public Integrity
Gina Chua, Semafor
Ethan Chumley, Microsoft
Elik Eizenberg, Scroll
Deb Ensor, Internews
Maggie Farley, ICFJ
Craig Forman, NextNews Ventures
Richard Gingras, Google
Jeff Jarvis, CUNY
Tanit Koch, The New European
Amy Kovac-Ashley, Tiny News Collective
Marc Lavalley, Knight Foundation
Celeste LeCompte, Fmr. Chicago Public Media
Erin Logan, Fmr. LA Times
The Hon. Jerry McNerney, Pillsbury Winthrop Shaw Pittman LLP, Fmr. Congressman (CA)
Tove Mylläri, Yle (Finnish Broadcasting Company)
Matt Perault, UNC Center on Technology Policy
Adam Clayton Powell III, USC Election Cybersecurity Initiative
Courtney Radsch, Center for Journalism and Liberty
Aimee Rinehart, The Associated Press
Felix Simon, Oxford Internet Institute
Steve Waldman, Rebuild Local News (*moderator*)
Lynn Walsh, Trusting News

For more details, see the Appendix.

Finnish Broadcasting Company) and Oxford and civil society experts and philanthropists from a range of organizations. (See sidebar for the full list of participants.)

Under the theme of how to apply verification, authentication and transparency to an AI-incorporated news environment, participants addressed four main questions: What information does the public need to assess the role of AI in news and information? What do journalists need from technology systems and AI models? How should technology systems enable these principles, and how should government policies protect them?

The session, held under a modified [Chatham House Rule](#), continued the tone and style that began with CNTI's inaugural convening in October 2023 ("[Defining AI in News](#)"). CNTI uses research as the foundation for collaborative, solutions-oriented conversations among thoughtful leaders who don't agree on all the elements, but who all care about finding strategies to safeguard an independent news media and access to fact-based news.

Throughout the convening, participants often prefaced their remarks by describing their own optimism or pessimism about the present and future role of AI in journalism. The event's moderator, Steve Waldman, discussed this tension in his introduction: "To me the answer is we have to be both absolutely enthusiastic about embracing the many positive aspects of this [technology] and absolutely vigilant about potential risks. Both are really important." As the report explains, there are no easy answers but there are avenues to consider as AI technology advances at a blistering pace.

Key Themes & Takeaways:

1. [There is No Silver Bullet for Addressing Harms of AI in News While Still Enabling its Benefits](#)
2. [We Need to Experiment with a Number of Possible Tools](#)
3. [The Role for Policy: Consider Industry Standards as a Start, Rethink Regulatory Structures, Lead by Example](#)
4. [Successful Standards, Uses & Guardrails Require Technology Companies' Active Participation](#)
5. [Research, Research, Research](#)
6. [Best Steps for Newsrooms: Innovate with a Degree of Caution](#)

This was the second in a series of CNTI convenings on enabling the benefits while managing the harms of AI in Journalism. [Stay tuned](#) for details about CNTI's third AI convening, to be held outside the U.S.

1. There is No Silver Bullet for Addressing Harms of AI in News While Still Enabling its Benefits

Artificial intelligence will impact seemingly every industry, with the news media being no exception. This can be particularly challenging in the news and information space as the use of AI by various actors can sometimes work against informing the public and sometimes work towards it. In either case, new technological developments have added new challenges for separating the real from the fake. And while the media, technology companies and others are making attempts at identifying AI, the group was in agreement that no single solution will be completely effective. In fact, some research finds that current tools, such as labels of AI use, may actually do more harm than good.

So what is being done now to verify or authenticate AI-generated content? To date, the tool that has been furthest developed and has also received the most attention is “watermarking” — a technique where markers get [embedded](#) into a piece of content (e.g., image, audio, video, [text](#), etc.) when it is first created. Proponents say watermarks help journalists and the public identify ill-intended changes to content they find online.

Several online platforms have begun implementing their own software and/or rules for asserting the authenticity of content, many of which comply with the Coalition for Content Provenance and Authenticity ([C2PA](#)) standard. OpenAI released an [update](#) that adds watermarks to images created with its DALL-E 3 and ChatGPT software. Google’s [SynthID](#) allows users to embed digital watermarks in content that are then detectable by the software. Other companies like [Meta](#) have focused on provenance, with policies that require disclosing content as being generated with AI. Across the technology industry, companies are taking note of provenance and implementing tools to assist users.

While developing these kinds of tools clearly has benefits, the group identified several important considerations and reasons the tools alone should not be seen as a solution in and of themselves:

- **Preliminary research finds that labels related to AI in journalism can have an adverse effect on the public.** Research [examining](#) the impact of content labeling raises several caution flags. First, overly broad labels about false and manipulated information can lead users to discount accurate information, suggesting labels need to be comprehensive and explicit about which content is false. Similarly, additional [research](#) on tagging content as AI-generated or enhanced finds, “... on average that audiences perceive news labeled as AI-generated as less trustworthy, not more, even when articles themselves are not evaluated as any less accurate or unfair” and that the effects are more common among people with higher pre-existing levels of trust.

These types of content labels can also lead to an “[implied truth effect](#)” in which false information that is *not* tagged as such may be interpreted as authentic. Similar

findings exist when studying provenance. On the more hopeful side, some [findings](#) suggest that the public *can* appreciate provenance details as long as they are comprehensive — poorly-detailed provenance can lead to users discounting authentic information. To that end, researchers are exploring what specific [labels](#) to use in a given context (e.g., AI generated, manipulated, digitally altered, computer generated, etc.) and how users interpret these terms.

- **Current labeling techniques don't differentiate between uses of AI that help inform rather than disinform.** Indeed, some content alterations are done to help inform, which current methods of labeling don't address. One participant shared an example from Zimbabwe where a newsroom used [chatbots](#) to offer information in many more local dialects than Western-trained models could provide. There are also several non-publicly facing innovative AI uses such as [fact-checking](#) radio broadcasts or combing through and synthesizing news archives (which offer particular value for local news organizations). The current research suggests that a simple label acknowledging the use of AI risks automatic public rejection of what would otherwise add value to the news product.
- **We need to clarify the intended audience for each label.** There is a lot of conversation about the benefit of provenance for the public but, as the group discussed, it may be more crucial for journalists and content creators. Marc Lavallee noted, there is “limited value in any kind of direct consumer-facing watermark or signal” and Richard Gingras said, “the value of provenance is probably more for the world of journalism than it is for the world of consumers.” Consider the example of information cataloged for [original artwork](#). Celeste LeCompte noted, “most of that information is not generally revealed to the public but is rather something that is part of an institutional framework.”
- **We need to better understand and articulate the various elements of identification: provenance, watermarking, fingerprinting and detection.** In this convening, four similar but distinct forms were discussed. First is *provenance*, which refers to a “manifest or audit trail of information” to ensure content is “always attributable.” Provenance information is imperfect, however, because it can be input incorrectly or changed by malicious actors. Second is *watermarks*, a class of identifiers embedded directly into content which is more difficult to remove and considered to be more robust than provenance. A third technique called *fingerprinting* can serve as a lookup tool, like reverse [image searches](#). Finally, there are what are termed *detection methods*. These use AI models to detect other AI models, though the development process remains challenging. They are the least robust and, as one participant asserted, more research needs to be done in this area.

Again, the conclusion is not that watermarks are bad or of no use, but that they need to be fully thought through within the vast array of AI uses and considered as just one tool among many that are needed — which leads to the next takeaway.

2. We Need to Experiment with a Number of Possible Tools

As we further develop the efficacy of tools like watermarking, participants encouraged further experimentation with additional ways to help identify and explain various uses of AI content. As one participant remarked, “Doing something is better than doing nothing while waiting for the perfect solution.”

A few ideas shared by participants:

- **SSL for AI:** SSL, or Secure Sockets Layer, was designed to be an encryption tool for e-commerce, but is now used by virtually every website as a privacy and authentication tool. As one participant stated, there’s no “values determination,” just a determination on the content’s origin. Thus, there are no false positives or false negatives. Could publishers and technologists collaborate on something similar here?
- **Accessible Incentive Structures to Adopt Standards:** Another idea was to use incentive structures for journalists and other responsible content creators to adopt certain standards and labeling techniques which could eventually become commonly understood. Search Engine Optimization (SEO), the process by which websites strive to rank higher in Google and other search engines, was offered as an example. While not a perfect corollary (with questions about [gaming algorithms](#) and the value of information that is not public facing) SEO, as originated, did offer a strong incentive and was pretty easy to adopt. How might something like that work for identifying AI content? And how could we measure its effectiveness? Getting the incentive structure right so that fact-based content gets promoted while “... inauthentic material or material of unknown provenance is lessened is really the place to focus,” suggested Lavalley. “If we get to a point where basically only bad actors are the ones not willing to use a system like this, I think that’s the threshold that we need to get to in order for it to be effective.”
- **Training the public:** Continued attention on AI literacy and education is important. As the research (cited above) shows, most of the public seems to distrust any use of AI in news content. A more nuanced understanding is important to allow journalists to use AI in ways that help serve the public. One participant shared information about how media and technology education are included in Finland’s national education plans and how news organizations there have also developed training materials for the general public. Information about how AI is used in news needs to be understandable for a non-technical audience but also allow people who would like further details to

have access to that information. As such, per Tove Mylläri, “We believe that educating people to be aware and letting them decide themselves also builds trust.” Understanding what types of training materials are most effective is crucial for increasing knowledge about how AI is used in news.

One important element for any of these tools, especially (but not limited to) those that are public facing, is communication. The publishers and others who implement these tools need to explain to their audiences and the general public what these tools do. It will take time for them to become recognized, commonly understood and utilized. The “nutrition labels” parallel provides a sense of that timeline. Nutrition labels, noted Anna Bulakh, are now generally understood and serve a valuable purpose, but that took experimentation about what kinds of information the public wanted and it took time for shoppers to become accustomed to them. In fact, it remains a [work in progress](#) and also carries with it the important consideration of who decides what goes into the label. News consumers are not yet used to “nutrition labels” for content. “Provenance is providing you with a nutritional label of content [so you can] be aware that it is AI manipulated or not AI manipulated,” Bulakh added. Provenance should be understood to mean information about the source of the material (i.e., the publisher, author, and/or editor) as well as what tools or mechanisms might be relevant to assessing the trustworthiness of the content.

3. The Best Role for Policy: Consider Industry Standards as a Start, Rethink Regulatory Structures, Lead by Example

Government policy and regulation can be critical in promoting and safeguarding the public good. They can also have lasting impact and as such must be thoroughly and carefully approached. Any drafted policy should consider its potential impacts both now and in the years to come. To facilitate this, the discussants laid out several insights:

- **Established standards can help build effective policy.** One part of the conversation focused on creating industry standards that, at least in some areas, can be used to inform effective legislation and regulation (recognizing these two forms of policy have [different definitions](#) which can also vary by country). Standards allow for experimentation and adjustment over time, could be developed for [different parts](#) of the system and, as past examples have shown, can then develop into effective policy. Consider, for example, the U.S. Department of Agriculture (USDA)’s development of organic food standards in the 1990s. The 1990 Organic Foods Production Act [created](#) the National Organic Program which was tasked with developing [standards](#) for organic food production and handling regulations. The USDA [defines](#) organic as a “labeling term that indicates that the food or other agricultural product has been produced through approved methods. The organic standards describe the specific requirements that must be verified by a USDA-accredited certifying agent before

products can be labeled USDA organic.” The creation of this standard, remarked one participant, took time to develop but eventually led some consumers to seek out this type of product especially because they believe it follows specific criteria. Media content could follow a similar process. Once trusted standards are implemented, the public may seek content that conforms to those protocols, which would likely incentivize many industry actors to also use the standards.

In supporting this approach, former U.S. Representative & Chair of the Congressional Artificial Intelligence Caucus, Jerry McNerney, suggested that we fund AI standards agencies and, once created, enforce those standards through law, adding that it is important to have “the involvement of a wide spectrum of stakeholders,” particularly given the unique ways journalists are incorporating AI into their work.

- **It’s time to rethink what structures of “regulation” should look like today.** Several participants agreed we need to rethink what the structure of regulation looks like in an AI-incorporated environment in which developments occur rapidly with technology that is complex. One participant rhetorically asked, “do you really understand what you’re trying to accomplish?” It doesn’t mean we should walk away from complicated issues but we should fully think about the most effective approaches.

Ethan Chumley offered that there “are close parallels and analogies” to cybersecurity perspectives “if we start to view the challenges of media authenticity and trust in images as a security problem.” Regulation will likely require more updates (1) as technologies evolve and (2) as standards are revised.

- **This all takes time!** Developing broadly adopted standards that can then develop into policy is time consuming. Bulakh suggests that “every new standard would take around 10 years to be accessible” to tool creators, distribution platforms, content creators and consumers. She points to the Coalition for Content Provenance and Authenticity ([C2PA](#)), one of the leading global standards, entering year 5, even though it has not been fully adopted. Another example is [HTML standards](#) for websites which, similarly, took many years to develop and are still evolving. We have past examples to use as models for timelines which need to be built in — let’s use them!

This is not to say that a standards-first approach is right for all areas. There may well be some aspects of managing use of AI that call for some government oversight more immediately — though we still need to be sure the regulatory structures are effective and that the policies are developed in a way to serve the public long-term.

- **One immediately available step is to lead by example.** Many nods of agreement occurred when one participant pointed out that if governments want others to adopt and utilize various standards, rules or policies, they need to do so themselves.

Governments can, in the immediate term, “start adopting these standards that are already out there in their own datasets,” said Elik Eizenberg, which would likely create momentum for the private market to follow. One of the most powerful tools governments have, Eizenberg added, is to “lead by example.” In the discussion that followed, another participant added that, conversely, many current policies relating to misinformation have built-in exceptions for politicians.

4. Successful Standards, Uses & Guardrails Require Technology Companies’ Active Participation

“Journalists [and] media houses cannot cope with these issues alone without technology companies,” remarked one participant to the nodding approval of others. Another person added, “Open standards are developed by technologists. It’s [technologists’] job ... to come together to provide access to those tools ... and make them more accessible.” Media houses can then communicate and “change consumer behavior.”

Tanit Koch spoke of the importance of technology companies in helping guard against those with bad intentions: “We need tech companies and not only because of the scale and the speed that disinformation can happen and is happening on their platforms, but simply because they have the money and expertise to match the expertise of the bad actors. We definitely and desperately need more involvement by those who feel a sense of responsibility to create open source tools to help the media industry detect what we cannot detect on our own and all of this full well knowing that the dark side may always try to be ahead of us.”

A separate point was raised about whether there was a role for technology companies, and policies guiding them, in helping users better understand certain risks or benefits associated with the use of certain technologies. “When I take a pill, there’s a warning label,” offered Paul Cheung, but these risks are not as clearly defined in the online space. Thus, consumers “have no information to assess whether this is a risk they’re willing to take.” People may use a certain technology tool or piece of software “because it’s free and fun” without knowing the risk level. And those risk levels likely vary, suggesting a model similar to the U.S. Food and Drug Administration (FDA) or Federal Aviation Administration (FAA) may be a better route than a “one-size-fits-all” approach.

Understanding these risks relates to the discussion on building trust that occurred during CNTI’s first AI convening. [As reported following that gathering](#):

Better understanding of [AI language] can also help build trust, which several participants named as critical for positive outcomes from AI use and policy development. One participant asked, “How do we generate trust around something that is complicated, complex, new and continuously changing?” Another added,

“Trust is still really important in how we integrate novel technologies and develop them and think two steps ahead. And when it comes to putting that in writing, “We need to think about what’s the framework to apportion responsibility and what responsibility lies at each level ... so that you get the trust all the way up and down, because ultimately newsrooms want to be able to trust the technology they use and the end user wants to be able to trust the output or product from the newsroom.”

5. Research, Research, Research

To date, research about labeling AI content (and how users engage with labels) is limited. We need much more data to gain a fuller understanding of the best strategies forward, as well as which strategies are likely to fall short, backfire or possibly work in some areas but not in others. To develop a deeper understanding of what, why and how certain policy and technology approaches work better than others, we need to conduct more studies, replicate findings and build theories. Approaches must also reflect geographic diversity by including researchers representing a range of local contexts and communities and by examining how people in diverse, global contexts are similar and/or unique in their interactions with AI-related news content. For example, U.S./European research that focuses solely on strategies to address internet disinformation would not serve well those places where radio is still the largest news medium. We need to provide the resources and support for this work — starting now.

Several researchers in the room noted that the existing literature does not yet fully grasp how users interact with provenance information or how the tools being developed will influence user behavior.

Felix Simon shared recent preliminary research he’d conducted in collaboration with Benjamin Toff that featured a striking finding: “For our sample of U.S. users, we found, on average, that audiences perceived news labeled as AI-generated as less trustworthy, not more, even when articles themselves were not evaluated as any less accurate or unfair.” Yet, when the authors provided users with the *sources* used to generate the article, these negative effects diminished. Further research is warranted to better understand how the public interprets labels on AI-generated articles and how to best present this information.

When it comes to watermarks in particular, “we have some directional indications about efficacy but we do have extremely limited data in this area,” offered Matt Perault. Based on the limited data and research, he presented four key research questions that need to be addressed:

- 1. Will disclaimers and/or watermarks be implemented correctly?**
- 2. Would users actually observe them?**
- 3. Would disclaimers and/or watermarks have a persuasive effect?**

4. What effects would watermark and/or disclosure requirements have on competitiveness and innovation?

Answers to these questions will assist in the development of evidence-based policy.

6. Best Steps for Newsrooms: Innovate with a Degree of Caution

Technology has revolutionized the media many times over (e.g., the printing press, radio, television, the internet, social media, streaming, etc.) with AI being the latest example of an innovation that will change how reporters gather and share information and how consumers take it in. Panelists made a number of workplace recommendations that would help journalists incorporate AI into their reporting and build a better relationship with their audiences.

- **Newsrooms need to embrace AI technology, but do it cautiously.** They must be willing to experiment with new tools to see what works for them and what helps audiences better understand the news. Offered Charlie Beckett, “The crisis in journalism at the moment is about connecting the ton of good quality stuff out there to the right people in the optimal way that doesn’t make them avoid the news.” It was further noted that quality reporting requires diligent research and thorough questions and that a similar approach should be applied to the exploration and adoption of new technology. Participants discussed a number of existing and developing technologies that could spur newsrooms to adopt the use of AI to support their work (see sidebar).

Conversely, concerns remain about how AI models use journalists’ works. Courtney Radsch asserted, “As we see more generative AI content online, less human-created information, the value of journalism, I think, goes up. So we should be thinking about a model that can allow us to have some say over the system.” The complex pros and cons of methods such as licensing or copyright on the nature and effectiveness of the knowledge ecosystem, including that they can hinder broad distribution and reward quantity irrespective of quality, were mentioned but were covered more in depth at CNTI’s first AI Convening, as written about in the [summary report](#).

- **Journalists need to apply layers of transparency in their work, just as they expect from the people and organizations they cover.** Innovation and transparency are critical to serving local communities and fortifying the journalism industry for the digital age. But remember that “transparency has multiple meanings, each of which must be addressed.” First, journalists must offer the same level of transparency that they demand of others. At a non-technical level, that means explaining why a particular story is being told at that time. But it also means clearly explaining how AI is

used to report, create and produce content including attributions and points of reference in the output of AI-driven answer engines (e.g. ChatGPT, etc.). Research has shown that some responses by AI-driven answer engines can be biased by the articulation of the source material used to feed them. News consumers would benefit from understanding related references and attributions which can help audiences rebuild trust in the media.

- **Recognize that there is more to AI than the content being produced.** Jeff Jarvis shared his thoughts on the ABCs of disinformation, which a 2019 [report](#) by Camille François of the Berkman Klein Center for Internet & Society at Harvard University defined as manipulative actors, deceptive behavior and harmful content. Jarvis summarized, “We have to shift to the ABC framework — that is to say actors, behavior and content,” rather than a singular focus on content. While journalists may be able to identify those seeking to sew harm through diligent research and reporting, it is difficult to persuade news consumers that such people act with malicious intent.

Another concept offered was that of “fama,” a Latin word dating back to Europe before the printing press that brings together rumor and fame to form the concept of believing something to be true because it was said by someone the listener trusts. Its modern equivalence is believing in news — or conspiracy theories — because they are uttered by a trusted voice. As noted by Deb Ensor, “We don’t often think about our audiences in terms of their behaviors and how they share or trust or value or engage with their information suppliers.” Those with ill intentions then appeal to this behavior with deceptive tools, such as bots and troll farms to spread disinformation. With the advent of AI, these actors have even more tools at their disposal and journalists need partners in technology and research to keep up.

AI Innovation Sidebar

Participants noted a number of innovations publishers, editors and reporters can explore to better incorporate AI into their work, with the assistance of technologists. Some of these ideas are already in the works, while others need to be developed including, as one participant noted, broadening beyond a singular focus on large language models (LLMs) or secondary uses of LLMs. LLMs are AI models that are trained on a massive amount of text.

- **Local LLMs:** Because AI technology is Western-centric, journalists in countries that do not use the world's most popular languages are left behind. There needs to be a concerted effort to develop local language versions of large language models (LLMs), another example where industries will need to partner in order to achieve results. One participant offered an example in Zimbabwe where chatbots were trained to better interpret the various local dialects in the country. These types of localized innovations are also receiving support from users as they feel better represented. In addition to the technology companies that would build such LLMs, journalists and others have a major role to play: We need "people to create the information [through] front-line reporting and analysis that these models then ingest and generate new material," remarked Maggie Farley. This type of front-line reporting also requires "talented people," pointed out Erin Logan, who have secure employment with liveable wages and benefits which requires sustainable business models for news organizations.
- **Shared LLMs:** Aimee Rinehart shared that as a part of her Tow-Knight Fellowship on AI Studies her capstone project is a blueprint of an LLM specifically for news organizations that would "add transparency because, as journalists, we all like to know, 'who's your source?'" In this case, the source isn't the person who provided the news tip, but rather that system that supports the production of the news item. Such an LLM could do far more for journalism than simply outsource writing. A journalism-focused LLM, Rinehart added, "could resurrect the archive [and] provide a licensing opportunity for newsrooms." These opportunities are likely to be especially relevant for local news organizations, which are struggling to remain economically viable.
- **RAGs:** While LLMs are built on massive data sources, journalists often need a more narrow scope for their work. That is where Retrieval-Augmented Generation (RAG) comes in. RAGs are used to limit the scope of an LLM query, such as to particular data sets, to more efficiently and accurately pull results of the billions of data points that form the input. One participant, Gina Chua, said RAGs can be used to read documents, classify data, or even turn journalism into something more conversational (and therefore more accessible). Such AI tools can be applied at scale to rebuild local newsrooms and have the potential to "improve journalism products, which [can] then improve our engagement with communities," Chua remarked.
- **Pinpoint:** Another participant called on technology companies to develop tools to help journalists process massive amounts of raw data that can inform their reporting. One example already available is Google's [Pinpoint](#), a collaborative tool that allows reporters to upload and analyze up to 200,000 files — not just documents and emails but also images, audio, and scans of hand-written material.

Next Steps & Additional Information

This is the second in a series of convenings CNTI will host on enabling the benefits of AI while also guarding against the harms. It is just one of many initiatives within the larger work CNTI, an independent global policy research center, does to encourage independent, sustainable media, maintain an open internet and foster informed public policy conversations. Please visit our website: www.innovating.news to see more of CNTI's work and sign up to receive updates, and, as always, please [contact us](#) with questions and ideas.

Finally, please see the Appendix below for numerous valuable resources shared by participants of this event as well as other acknowledgements.

Appendix

List of Participants

- **Charlie Beckett**, Professor/Founding Director, Polis, LSE (*CNTI Advisory Committee*)
- **Anna Bulakh**, Head of Ethics & Partnerships, Respeecher (*CNTI Advisory Committee*)
- **Paul Cheung**, CEO, Center for Public Integrity (Fmr., now Sr. Advisor, Hacks/Hackers)
- **Gina Chua**, Executive Editor, Semafor
- **Ethan Chumley**, Senior Cybersecurity Strategist, Microsoft
- **Elik Eizenberg**, Co-Founder, Scroll
- **Deb Ensor**, Senior VP of Technical Leadership, Internews
- **Maggie Farley**, Senior Director of Innovation and Knight Fellowships, ICFJ
- **Craig Forman**, Managing General Partner, NextNews Ventures (*CNTI Executive Chair*)
- **Richard Gingras**, Global VP of News, Google (*CNTI Board*)
- **Jeff Jarvis**, Director of the Tow-Knight Center for Entrepreneurial Journalism & The Leonard Tow Professor of Journalism Innovation, CUNY
- **Tanit Koch**, Journalist/Co-Owner, The New European (*CNTI Advisory Committee*)
- **Amy Kovac-Ashley**, Executive Director, Tiny News Collective (*CNTI Advisory Committee*)
- **Marc Lavalley**, Director of Technology Product and Strategy/Journalism, Knight Foundation
- **Celeste LeCompte**, Fmr. Chief Audience Officer, Chicago Public Media
- **Erin Logan**, Fmr. Reporter, LA Times
- **The Hon. Jerry McNerney**, Senior Policy Advisor, Pillsbury Winthrop Shaw Pittman LLP, Fmr. Congressman (CA)
- **Tove Mylläri**, AI Innovation Lead, Yle (The Finnish Broadcasting Company)
- **Matt Perault**, Director, UNC Center on Technology Policy
- **Adam Clayton Powell III**, Executive Director, USC Election Cybersecurity Initiative
- **Courtney Radsch**, Director, Center for Journalism and Liberty at the Open Markets Institute
- **Aimee Rinehart**, Local News & AI Program Manager, The Associated Press
- **Felix Simon**, Researcher, Oxford Internet Institute (*CNTI Advisory Committee*)
- **Steve Waldman**, President, Rebuild Local News
- **Lynn Walsh**, Assistant Director, Trusting News

About the Convening

CNTI's cross-industry convenings espouse evidence-based, thoughtful and challenging conversations about the issue at hand, with the goal of building trust and ongoing relationships along with some agreed-upon approaches to policy. To that end, this convening adhered to a slightly amended [Chatham House Rule](#):

1. Individuals are invited as leading thinkers from important parts of our digital news environment and as critical voices to finding feasible solutions. For the purposes of transparency, CNTI **publicly lists all attendees and affiliations present**. Any reporting

on the event, including CNTI's reports summarizing key takeaways and next steps, can share information (including *unattributed* quotes) but cannot explicitly or implicitly identify who said what without prior approval from the individual.

2. CNTI does request the use of photo and video at convenings. Videography is intended to help with the summary report. Any public use of video clips with dialogue by CNTI or its co-hosts requires the explicit, advance consent of the subject.
3. To maintain focus on the discussion at hand, we ask that there be no external posting during the event itself.

To prepare, we asked that participants review CNTI's Issue Primers on [AI in Journalism](#), [Algorithmic Transparency](#) and [Journalistic Relevance](#) as well as the report from CNTI's [first convening event](#).

Resources Shared by Participants

Participants at our convening event shared a number of helpful resources. Many of these resources are aimed at assisting local newsrooms. We present them in alphabetical order by organization/sponsor below.

Several news organizations were mentioned for their use of AI in content creation. One that received recognition was the [Baltimore Times](#) for its efforts to better connect with their audience through the use of AI.

Participant Aimee Rinehart shared a blueprint for her [CUNY AI Innovation](#) project. This project aims to create a journalism-specific LLM AI model for journalists and newsrooms to use.

A novel radio fact-checking algorithm in Africa, [Dubawa Audio Platform](#), was discussed to show how countering mis- and disinformation can be done in non-Internet-based contexts. The platform was initiated by a Friend of CNTI, the Centre for Journalism Innovation and Development ([CJID](#)). The Dubawa project received [support](#) from a Google News Initiative grant.

Information was shared about the Finnish [government](#) and academic community's campaign on AI literacy, [Elements of AI](#). This project aims to raise awareness about the opportunities and risks of AI among people who are strangers to computer science, so they can decide for themselves what's beneficial and where they want their government to invest. Free educational [material](#) also exists for children. Curious readers may also learn about related research [here](#).

The topic of transparency was discussed and a 2022 [report](#) by Courtney Radsch for the Global Internet Forum to Counter Terrorism (GIFCT) provides an important overview of transparency across various industries.

A number of participants shared information about technological tools:

- Google's [Pinpoint](#) project helps journalists and researchers explore and analyze large collections of documents. Users can search through hundreds of thousands of documents, images, emails, hand-written notes and audio files for specific words or phrases, locations, organizations and/or names.
- Google's [SynthID](#) is a tool to embed digital watermarks in content to assist users with knowing the authenticity and origin of content.
- Microsoft's PhotoDNA creates a unique identifier for photographs using its system. This tool is used by organizations around the world — it has also assisted in the detection, disruption and reporting of child exploitation images.

Information was shared about [Schibsted](#), a Norwegian media and brand network, organizing the development of a Norwegian large language model (LLM). It will serve as a local alternative to other general LLMs.

The Tow Center for Digital Journalism at Columbia University released a recent [report](#) by Felix Simon titled “Artificial Intelligence in the News: How AI Retools, Rationalizes, and Reshapes Journalism and the Public Arena.”

A participant shared an innovative use of AI in [Zimbabwe](#) in which the AI model has been trained using local dialects. The chatbot is more representative of the users in that region when compared to other general AI language models.

We appreciate all of our participants for sharing these resources with CNTI.

About CNTI

The [Center for News, Technology & Innovation](#) (CNTI), an independent global policy research center, seeks to encourage independent, sustainable media, maintain an open internet and foster informed public policy conversations. CNTI’s cross-industry convenings espouse evidence-based, thoughtful but challenging conversations about the issue at hand, with an eye toward feasible steps forward.

The Center for News, Technology & Innovation is a project of the Foundation for Technology, News & Public Affairs.

Acknowledgements

CNTI sincerely thanks the participants of this convening for their time and insights, and we are grateful to the University of Southern California’s Annenberg Capital [Campus](#), the co-sponsor and host of this AI convening. Special thanks to Adam Clayton Powell III and Judy Kang for their support, and to Steve Waldman for moderating such a productive discussion.

CNTI is generously supported by [Craig Newmark Philanthropies](#), [John D. and Catherine T. MacArthur Foundation](#), [John S. and James L. Knight Foundation](#), [the Lenfest Institute for Journalism](#) and [Google](#).