

Poynter.

The Craig Newmark Center for
ETHICS and
LEADERSHIP

Poynter Summit on AI, Ethics & Journalism: Putting audience and ethics first

Poynter created a ChatGPT-powered chatbot to answer questions or summarize sessions from the summit. Check it out here.



The following Poynter staff members contributed to this report: Alex Mahadevan, Kelly McBride, Tony Elkins, Jennifer Orsi, Barbara Allen

On your keyboard, press option (on a Mac) or control (on a PC) and right arrow to jump to the next page

Rapidly advancing generative artificial intelligence technology and journalism have converged during the biggest election year in history. As more newsrooms experiment with AI, the need for ethical guidelines and audience feedback have surfaced as key challenges.

The Poynter Institute brought together more than 40 newsroom leaders, technologists, editors and journalists during its Summit on AI, Ethics & Journalism to tackle both topics. For two days in June 2024, representatives from the Associated Press, the Washington Post, Gannett, the Invisible Institute, Hearst, McClatchy, Axios and Adams along with OpenAI, the Online News Association, the American Press Institute, Northwestern University and others, debated the use of generative AI and its place within the evolving ethics of journalism

The goals: Update Poynter's AI ethics guide for newsrooms with insight from journalists, editors, product managers and technologists actually using the tools. And outline principles for ethical AI product development that can be used by a publisher or newsroom to put readers first.

Data from focus groups convened through a Poynter and University of Minnesota partnership underscored discussion, while a hackathon tested attendees to devise AI tools based on audience trust and journalistic ethics.



ALEX SMYNTYNA/POYNTER

Alex Mahadevan, Director of MediaWise at Poynter, speaks at Poynter's Summit on AI, Ethics, and Journalism in June 2024.

Key takeaways:

- There is significant anxiety and distrust among audiences regarding AI in journalism, exacerbated by concerns over job security and the motives behind AI use.
- Audiences largely want to be told when AI is used in news production.
- There is a need for clear, specific disclosures about how AI is used in news production to avoid label fatigue and maintain audience trust.
- Data privacy is a significantly overlooked concern in the deployment of newsroom AI tools and should be addressed.
- Newsrooms are encouraged to experiment with AI to discover new capabilities and integrate these tools thoughtfully into their workflows.
- Continuous audience feedback and involvement in the AI development process are essential to creating relevant and trustworthy news products.
- News organizations should invest in AI literacy initiatives to help both journalists and the public understand AI's capabilities and limitations, fostering a more informed and collaborative environment.

Listening to the Audience: Insights from Poynter-University of Minnesota focus groups

Audience was the key word to emerge during the Poynter Summit on AI, Ethics and Journalism. Specifically, how to talk to readers about AI, improve their lives and solve their problems — not just those of the news industry.

Poynter partnered with Benjamin Toff, the director of the Minnesota Journalism Center and associate professor of Minnesota's Hubbard School of Journalism & Mass Communication, to run a series of focus groups to discuss AI with representative news consumers. Some key takeaways Toff found include:

- **A background context of anxiety and annoyance:** People are often anxious about AI — whether it's concern about the unknown, that it will affect their own jobs or industries, or that it will make it



ALEX SMYNTYNA/POYNTER

Benjamin Toff of the University of Minnesota presents focus group findings, commissioned by the Poynter Institute, on AI in journalism.

harder to identify trustworthy news. They are also annoyed about the explosion of AI offerings they are seeing in the media they consume.

- **Desire for disclosure:** News consumers are clear they want disclosure from journalists about how they are using AI — but there is less consensus on what that disclosure should be, when it should be used and whether it can sometimes be too much.
- **Increasing isolation:** They fear that increased use of AI in journalism will worsen societal isolation among people and will hurt the humans who produce our current news coverage.

Anxious and annoyed

Some participants felt besieged by AI offerings online.

“I’ve noticed it more on social media, like it’s there. ‘Do you want to use this AI function?’ and it’s right there. And it wasn’t there that long ago. ... It’s almost like, no, I don’t want to use it! So it’s kind of forced on you,” said a participant named Sheila.

Most participants already expressed a **distrust of the news media**, and felt the introduction of AI could make things worse.

The focus groups suggest perhaps the biggest mistake newsrooms can make is the rolling out of all things AI; Instead of sparking wonder in our audiences, are we going to annoy them?

A notable finding of the focus groups was that many participants felt certain AI use in creating journalism — especially when it came to using large language models to write

content — seemed like cheating.

“I think it’s interesting if they’re trying to pass this off as a writer, and it’s not. So then I honestly feel deceived. Because yeah, it’s not having somebody physically even proofing it,” said one focus group member.

Most participants said they wanted to know when AI was used in news reports — and disclosure is a part of many newsroom AI ethics policies.

But some said it didn’t matter for simple, “low stakes” content. Others said they wanted extensive citations, like “a scholarly paper,” whether they engaged with them or not. Yet others worried about “labeling fatigue,” with so much disclosure raising questions about the sources of their news that they might not have time to digest it all.

“People really felt strongly about the need for it, and wanting to avoid

being deceived,” said Toff, a former journalist whose academic research has often focused on news audiences and the public’s relationship with news. “But at the same time, there was not a lot of consensus around how much or precisely what the disclosure should look like.”

Some of the focus group participants made a similar point, Toff said. “They didn’t actually believe (newsrooms) would be disclosing, however much they had editorial guidelines insisting they do. They didn’t believe there would be any internal procedures to enforce that.”

It will be vitally important how journalists tell their audiences what they are doing with AI, said **Kelly McBride**, Poynter’s senior vice president and chair of the **Craig Newmark Center for Ethics and Leadership**. And they probably shouldn’t even use the term AI, she said, but instead more precise

descriptions of the program or technology they used and for what.

For example, she said, explain that you used an AI tool to examine thousands of satellite images of a city or region and tell the journalists what had changed over time so they could do further reporting.

“There’s just no doubt in my mind that over the next five to 10 years, AI is going to dramatically change how we do journalism and how we deliver journalism to the audience,” McBride said. “And if we don’t ... educate the audience, then for sure they are going to be very suspicious and not trust things they should trust. And possibly trust things they shouldn’t trust.”

The human factor

A number of participants expressed concern that a growing use of AI would lead to the loss of jobs for human journalists. And many were unnerved by the example Toff's team showed them of an AI-generated anchor reading the news.

"I would encourage (news organizations) to think of how they can use this as a tool to take better care of the human employees that they have. So, whether it's to, you know, use this as a tool to actually give their human employees ... the chance to do something they're not getting enough time to do ... or to grow in new and different ways," said one participant, who added that he could see management "using this tool to find ways to replace or get rid of the human employees that they have."

"If everybody is using AI, then all the news sounds the same," said one participant.



Said another focus group member: "That's my main concern globally about what we're talking about. The human element. Hopefully, that isn't

"The internet and social media and AI all drive things toward the middle — which can be a really mediocre place to be. I think about this with writing a lot. There is a lot of just uninspired, boring writing out there on the internet, and I haven't seen anything created by AI that I would consider to be a joy to read or absolutely compelling."

— Kelly McBride

taken over by artificial intelligence, or it becomes so powerful that it doesn't do a lot of these tasks, human tasks, you know? I think a lot of things have to remain human, whether it be error or perfection. The human element has to remain."

Moving forward

Toff still has more to glean from the focus group results. But the consumers' attitudes may hold some important insights for the future of financially struggling news organizations.

As AI advances, it seems highly likely to deliver news and information directly to consumers while reducing their connection to news organizations that produced the information in the first place.

“People did talk about some of the ways they could see these tools making it easier to keep up with news, but that meant keeping up with the news in ways they already were aware they weren't paying attention to who reported what,” Toff said.

Still, somewhat hopefully for journalists, several focus group members expressed great concern for the important human role in



producing good journalism.

“A number of people raised questions about the limitations of these technologies and whether there were aspects of journalism that you really shouldn't replace with a machine,”

Toff said. “Connecting the dots and uncovering information — there's a recognition there's a real need for on-the-ground human reporters in ways there is a lot of skepticism these tools could ever produce.”

ALEX SMYNTYNA/POYNTER

Panelists discuss audience attitudes and trust. Left to right: Poynter's Alex Mahadevan; Joy Mayer of Trusting News; Phoebe Connelly of the Washington Post; and Benjamin Toff of the University of Minnesota.

An overview of AI in the newsroom

Nikita Roy, International Center for Journalists Knight fellow and host of the **Newsroom Robots** podcast, and **Phoebe Connelly**, senior editor for AI strategy and innovation at the Washington Post, laid out AI projects at newsrooms and how they can inform the ethical use of the technology. Here are key takeaways from the session:

- AI tools have emerged to transition longform journalism to bullets and summaries.
- Newsrooms should prioritize “chatting” with users about their content, making it scalable and searchable. They should really think about the mechanics of how users will interact with words, from taps to swipes.
- Several newsrooms are using AI to sift transcripts of government meetings and are either training systems to write the stories or bolster



ALEX SMYNTYNA/POYNTER

International Center for Journalists Knight fellow Nikita Roy (left) and Washington Post senior editor of AI strategy and innovation Phoebe Connelly discuss AI's use in newsrooms.

their local government reporting. It isn't hypothetical.

Some newsrooms have attempted to harness AI for their journalism and business to varying degrees of success. Roy said the AI newsroom projects she sees generally fall into one of four categories:

- Content creation, which includes tools that generate headlines or social media posts
- Workflow optimization, which includes transcription and proofreading tools
- Analytics and monitoring, which includes paywall optimization and tools that can predict customer churn
- Audience-facing tools, which includes interactive chatbots and article summarizers

Journalists owe it to both themselves and their audiences to familiarize



themselves with AI tools, Roy said. Not only can AI help journalists with their own work, but understanding AI is key to keeping tech companies accountable.

“There’s so much policy decisions, so much legislation that has not been fixed,” Roy said. “This is a very malleable space that we are in with AI, and this is where we need journalists to be the people who deeply understand the technology because it’s only then that you can apply it.”

The Washington Post has taken a cautious, yet still ambitious, approach to generative AI in the newsroom, with the recent rollout of Climate Answers. The AI-powered chat interface allows readers to ask questions about climate change and get a succinct answer based on eight years of Post reporting.

Some important background:

- It is based solely on coverage from the Post climate team — leading to a very low-to-nonexistent risk

of hallucinations, which is a term for falsehoods generated by large language models. The concept of retrieval-augmented generation — pulling answers from your own archives or database — can help newsrooms leverage generative AI without compromising journalistic ethics.

- If it doesn’t find a suitable climate article, it won’t answer. ChatGPT and other generative search chatbots will always try to give you an answer — which leads to hallucinations and other issues.
- It was in the works for six months.
- Its disclosure is a great template for other newsrooms, and offers a link to frequently asked questions and an audience feedback form.

The Post has also rolled out article summaries and has released [an internal tool called Haystacker](#), which will use AI to comb through and classify thousands of videos and images. You’ll notice that all of these AI-powered tools serve the audience

— even Haystacker will allow the Post’s visual forensics teams to find more stories for readers.

Some other AI tools mentioned by panelists and audience members:

- Quizbots, designed to engage readers with trivia about their local news. There are third-party companies providing these solutions, but some news organizations are building them in-house.
- One newsroom is building a solution to meeting transcriptions using [OpenAI’s Whisper model](#).
- Another publisher is using AI to power its podcast and create TikToks.
- A local newsroom has created an entirely AI reporter, complete with a name and persona.

“...this is where we need journalists to be the people who deeply understand the technology because it’s only then that you can apply it.”

— Nikita Roy

Ethics meeting the moment

The rise of generative AI isn't the first time journalism has grappled with ethics amidst changing technology. McBride and Poynter faculty member Tony Elkins presented a history of ethical quandaries in journalism, and guidance on how newsrooms can meet this moment. Some key takeaways:

- Journalism ethical policies are unlike any other ethical decision-making systems (including medical or legal). While many news organizations have strong policies, the industry is not licensed, and there's no governing body with formal consequences. And more importantly, journalists do a poor job explaining our values and standards to the audience. As a result, news consumers do not understand our jobs and trust has eroded over time.
- Technology is changing so quickly. The journalism industry has created

replicable ethical standards rooted in democratic values. Technology companies have a completely different set of values, and their work and their products are not rooted in seeking the truth, so they are not going to walk side by side with us. It is going to be up to journalists to distinguish our ethical standards. As AI becomes part of software updates, it is incumbent on administrators and practitioners to stay-up-to-date on AI-enhanced features.

- Our goal is to support the creation of news products with the ethics baked into the design of the product, so that we understand what the audience needs to know about our standards and our work. Any AI product must serve a consumer need and answer the audiences' questions.

As technology, particularly AI, advances at a breakneck pace, it introduces new challenges for



ALEX SMYNTYNA/POYNTER

Poynter faculty Tony Elkins outlines the history of ethical dilemmas in journalism to Summit attendees.

maintaining journalistic integrity. Unlike journalism, technology companies operate under a different set of values, prioritizing innovation and user engagement over truth and accountability.

This divergence creates a critical need for journalists to establish and uphold their ethical standards independently. The session highlighted image manipulation and AI-generated content blurring the lines of reality for the public, underscoring the urgency for the journalism industry to define and defend ethical standards in the face of these technological changes. Examples from Elkins included:

- From history, Time Magazine infamously darkened a photo of O.J. Simpson on its cover.
- During the Gulf War, a Los Angeles Times photographer was fired for merging two photos to create a new, striking photo for the newspaper.
- More recently, the Associated Press had to retract a photo of Kate

Middleton, Princess of Wales, that was likely manipulated with AI.

New tools, like OpenAI's Sora, Microsoft's Vasa-1 and Adobe Firefly will make it even easier to pollute the information ecosystem.

McBride also introduced questions we as an industry must address:

- How do we make our own AI transparent? Should we even use the term? Evolve our vocabulary?
- How do we make AI made by others transparent?
- How do we educate the public on AI's impact on perceptions of reality?
- How do we ensure that we understand the authenticity of material we are reporting on?
- How do we contribute to a healthy conversation about this?
- How do we avoid polluting the public marketplace?



ALEX SMYNTYNA/POYNTER

Kelly McBride, Poynter's senior vice president and chair of the Craig Newmark Center for Ethics and Leadership, raises key questions on AI ethics in journalism.

The Hackathon

The summit featured a hackathon, where journalists, technologists and ethicists aimed to develop AI-driven solutions that address the challenges facing modern newsrooms. Participants were tasked with creating tools and products that serve the audience, make newsroom workflows better and align with ethical standards. The hackathon served as a microcosm of the broader discussions at the summit, emphasizing the importance of integrating ethics into AI design while showcasing the creative potential of technology to transform the future of news.

Key takeaways:

- There is huge value to seeking audience input at every stage of the product development process.
- Data privacy is an important aspect of generative AI tools that is not often referenced in these



ALEX SMYNTYNA/POYNTER

Participants brainstorm AI tools for journalists with ethical considerations during the event's hackathon.

discussions. It should be.

- There are big challenges around verifying and vetting large datasets.

- Opportunity to redefine journalism's value to audiences as connector, responder, solver,

empowerer, trusted source and even collaborator. Also to reach new audiences.

- Also low-hanging fruit is to really hone focus on one key part of the demo and take an iterative approach.

The hackathon led to six imagined technologies, which ranged from apps to websites to software. All the theoretical inventions sought to help people, answer questions and improve the quality of life for news audiences. While the exercise was theoretical, one group is actually taking steps to try to pursue and get funding for its idea, an AI-powered community calendar.

As the working groups conceptualized their visions, they identified plenty of ethical considerations. Here's what some of them came up with, and what they learned through this exercise.



ALEX SMYNTYNA/POYNTER

PolitiFact editor-in-chief Katie Sanders presents her group's AI tool, Vote Buddy, for tracking precincts and candidates during local elections.

Vote Buddy

PolitiFact editor-in-chief Katie Sanders helped conceptualize a tool that would serve as a guide to local elections.

Vote Buddy was meant to be a local news product, which required detailed information about precincts and candidates and their positions. Seemingly endless details stacked up as her team considered the experiment, she said, which called for more and more journalistic firepower.

Her team noted almost immediately that “the ethical concerns were abundant.”

They started by asking hard questions about use and users. Sanders said it was important to understand exactly what the team wanted to create, consider the problems it would solve for users, and make sure there was an actual need; and if audience members/

users would be comfortable with the means by which the AI tool provided the information.

“As we started to tease out what this service could be, we also realized how much human manpower would be needed to pull it off and maintain it,” she said. “The experience showed me that your product is only as good as the amount of time and energy that you set aside for the project.”

Just because it's an AI product, she said, doesn't mean it won't eat up resources, especially when it comes to testing and rooting out any and all inaccuracies.

“Hallucinations around something as serious as someone's vote are just unacceptable,” she said. “I felt better about having been through the experience, roleplaying what it would take.”

Living Story

Mitesh Vashee, Houston Landing's chief product and technology officer, said that many journalists are simply afraid of AI, which creates a barrier to journalists learning how to use it at all — especially ethically.

He said it's helpful for journalists to start their journey toward ethical AI use by playing around with AI tools and discovering practical uses for it in their day-to-day work. That way, "It's not just this big, vague, nebulous idea," he said, "but it's a real-world application that helps me in my day. What's the doorway that we can open into this world?"

His group conceptualized Living Story, a "public-facing widget that appears at the article level, which allows readers to interact with the story by asking questions."

Vashee said that journalists' fear that AI would replace them has been front and center in many of his conversations.

"We've made it clear at Houston Landing that we won't publish a single word that's generated by AI — it's all journalism," he said. "It's written by our journalists, edited by our editors, etc. ...That being said, the editorial process can get more efficient."

He said that as newsrooms look to implement new technology to help with efficiency, more work needs to be done to define roles.

"What is truly a journalist's job? What is an editor's job? And what is a technology job? I don't know what that full answer looks like today, but that's what we will be working through."



ALEX SMYNTYNA/POYNTER

Houston Landing's chief product and technology officer, Mitesh Vashee, Mitesh is discussing Living Story with Teresa Frontado, co-founder of the 51st.



ALEX SMYNTYNA/POYNTER

Erica Perel from the University of North Carolina participated in a group that envisioned an AI tool, The Family Plan, for better understanding loved ones' news consumption.

The Family Plan

One hackathon group identified less with workaday journalism and more with theoretical issues adjacent to journalism.

“(Our group was) mostly educators and people in the journalism space, more so than current working journalists,” said Erica Perel, director of the Center for Innovation and Sustainability in Local Media at the University of North Carolina. “The product we came up with dealt with bias, trust and polarization.”

The Family Plan was a concept that helped people understand what news media their loved ones were consuming, and suggested ways to talk about disparate viewpoints without judgment or persuasion.

Their biggest ethical concerns centered on privacy and data security.

“How would we communicate these privacy and security concerns? How would we build consent and transparency into the product from the very beginning?” she said. “And, how could we not wait until the end to be like, ‘Oh yeah, this could be harmful to people. Let’s figure out how to mitigate that.’”

CityLens

The hackathon team behind CityLens envisioned it as a free, browser-based tool that would use interactive technology to help users learn about and act on their local environment.

Smartphone cameras would capture a local image and then users could enter questions or concerns, which theoretically would lead them to useful information, including, “how to report a problem to the right entity, whether a public project is in the works at that location, and what journalists have already reported,” according to the team’s slides.

It would also offer an email template for reporting concerns like dangerous intersections, unsanitary restaurants, code violations, malfunctioning traffic devices, etc.

“I really liked the audience focus,” said Darla Cameron, interim chief product officer at The Texas Tribune.

“The framing of the whole event was, how do these tools impact our audiences? That is something that we haven’t thought enough about, frankly.”

Cameron said for their group, the ethical concerns involved boundaries and the role of journalists.

She said that several of the groups grappled with questions about the lines between journalistic creation of data and the tech companies’ collection of personal data.

“How can journalism build systems that customize information for our audiences without crossing that line?” she asked, noting that there was also a concern about journalists being too involved. “By making a tool that people can use to potentially interface with city government ... are we injecting ourselves as a middleman where we don’t have to be?”



ALEX SMYNTYNA/POYNTER

Hackathon participants review dozens of sticky notes filled with ideas for AI product concepts before breaking into groups.

Omni

Omni is “a personalized news platform that delivers the most relevant and engaging content tailored to your preferences and lifestyle,” according to the presentation of the group that created it.

Adriana Lacy, an award-winning journalist and founder of an eponymous consulting firm, explained that the group started with some nerves about its tech savvy.

However, members quickly found their footing — and ethical concerns. It became obvious that for Omni to work, its inventors would have to contend with the ethical issues surrounding personal data collection, she said.

“Our goal was figuring out how can we take information ... and turn it into various modes of communication, whether that’s a



ALEX SMYNTYNA/POYNTER

Journalist Adriana Lacy explains Omni, an AI tool offering a personalized news platform tailored to user preferences.

podcast for people who like to listen to things, a video for people who like to watch video, a story for people who prefer to read,” Lacy said. “Basically, compiling information into something that’s super personalized.”

Much of the information they would need to gather was essentially first-party data.

“We had some conversations about how we could ethically get readers to opt into this amount of data collection and we could be compliant in that area,” Lacy said. “We also discussed how we could safely and securely store so much data.”

Their other big ethical concern was figuring out how they could integrate the journalistic process into the project.

“So much of our idea was taking reporters’ writing, video and audio and turning that into a quick push

alert, a social media video, a podcast, an audio alert for your Alexa or Google Home — anywhere you choose to be updated,” she said. “The question remains: How can we apply our journalistic ethics and process into all these different types of media?”

Calindrical

One team is even looking to launch a real product based on its session at Poynter.

Dean Miller, managing editor of LeadStories.com, said his team of four focused on “the community-building magic of granular local newsroom-based calendars.”

He said their idea, Cylindrical, would bring real value to busy families and much-needed time to newsrooms, so the group has bought specific URLs and is working on paperwork to make the idea a reality.

“Our goal is a near-zero interface,” he said. “Think Mom driving (her) son to soccer, calling or texting to ask when (her) daughter’s drumline show is tonight, and where, and getting the info immediately and sending the info to Grandma and Dad.”

Miller said the group proposes to use AI to both collect event information

and to “assiduously” reach out to organizers to verify.

He said Poynter’s focus on AI ethics was helpful and necessary.

“(The) hackathon process was an early and quick way to surface bad assumptions,” Miller said. “We were spurred to focus our thinking on privacy protection, data security, user power and how to stave off the predations of Silicon Valley’s incumbents.”



ALEX SMYNTYNA/POYNTER

Poynter’s Summit on AI, Ethics and Journalism brought together news editors, product leaders and academics to discuss a wide range of issues related to AI and journalism.

Principles of ethical AI product development

Throughout the hackathon, teams met regularly with Poynter experts to discuss ethical hurdles in building their AI tools. Data privacy was a glaring issue, as was accuracy and hallucinations. Based on a day of conversations and rapid product ideation, Poynter developed a list of nine principles of ethical AI product development.

These principles are as close to universal to any newsroom as possible, but are not mandates by any means. For example, you probably won't find a third-party AI company that adheres to perfect journalistic ethics — and will be willing to sign a pledge to do so.

But, we hope these principles will guide a development process that puts audience trust and service first. Remember, you are trying to solve your readers' problems using artificial intelligence, not your own.

1. Transparency

- **Open development process:** Be transparent about the development process of AI tools, including the goals, methodologies and potential limitations of the technology.
- **Stakeholder involvement:** Involve a broad range of stakeholders, including ethicists, technologists, journalists and audience representatives, in the AI development process.
- **Clear disclosures:** Always provide clear, detailed disclosures about how AI is used in content creation. This includes specifying the role of AI in generating, editing or curating content. ([See ethics guidelines.](#))
- **Audience engagement:** Involve the audience in understanding AI processes through accessible explanations and regular updates on AI use. ([See ethics guidelines.](#))

2. Ethical standards and policies

- **Comprehensive guidelines:** Develop and enforce comprehensive ethical guidelines for AI use in journalism, covering all aspects from content creation to audience interaction. ([See ethics guidelines.](#))
- **Procurement agreements:** Create a contract — or build within your contracting agreements, ethical principles you expect third-party organizations to abide by while working with your newsroom. This may not necessarily be enforceable, but should attempt to align your ethical AI principles with those of the companies from which you are procuring tools and systems.
- **Regular reviews:** Conduct regular reviews of ethical guidelines to ensure they remain relevant and effective in the face of evolving AI technologies.

3. Accountability

- **Defined responsibilities:** Establish clear accountability mechanisms for AI-generated content. Identify who is responsible for overseeing AI processes and addressing any issues that arise.
- **Corrections policies:** Implement robust — public — processes for correcting errors or addressing misuse of AI tools, ensuring swift and transparent corrections.

4. Fairness and bias mitigation

- **Bias audits:** Regularly audit AI systems for biases and take proactive steps to mitigate any that have been identified. This includes diversifying training data and implementing checks and balances. Further, data bias should be a core fundamental feature in regular newsroom AI training.



ALEX SMYNTYNA/POYNTER

Bradenton Herald Senior Editor Ryan Callihan takes notes during Poynter's AI summit.

- **Inclusive design:** Ensure that AI tools are designed to be inclusive and consider the diverse experiences and perspectives of different communities. AI committees and teams developing AI tools should be as diverse as the newsroom — and preferably, reflect the demographics of the audience to be served by tool.

5. Data privacy and security

- **Data protection:** Adhere to strict data privacy standards to protect audience information. This includes secure data storage, handling and clear consent mechanisms for data collection. Expand your organization's data privacy policies to control for AI use.

- **Ethical data use:** Use audience data ethically, ensuring it is collected, stored and used in ways that respect user privacy and consent.

6. Audience service and the public good

- **Audience-centric design:** Develop AI tools that prioritize the needs and concerns of the audience, ensuring that AI serves to enhance the public good and journalistic integrity.

- **Community engagement:** Engage with communities to understand their needs and perspectives, and integrate their feedback into AI product development.

7. Human oversight

- **Human-AI collaboration:** Ensure that AI tools complement rather than replace human judgment and creativity. Maintain a significant level of human oversight in all AI processes.

- **Training and education:** Provide ongoing training and support for journalists and staff to effectively use and oversee AI tools.

8. Educational outreach

- **AI literacy programs:** Implement educational programs to improve AI literacy among both journalists and the public, fostering a better understanding of AI's role and impact in journalism.

- **Transparent communication:** Maintain open channels of communication with the audience about AI practices, fostering a culture of transparency and trust.

9. Sustainability

- **Long-term impact assessment:** Evaluate the long-term impacts of AI tools on journalism and society, ensuring that AI practices contribute to sustainable and ethical journalism.

- **Iterative improvement:** Continuously improve AI tools and practices based on feedback, audits, and new developments in the field of AI and ethics.

Next steps in Poynter's AI ethics work

The first Poynter Summit on AI, Ethics and Journalism and its two days of discussions and hackathon yielded:

- An update to Poynter's AI editorial guidelines starter kit for newsrooms (see appendix);
- Principles of ethical product development for technologists and product managers in any newsroom;
- Ideas for six ethics- and audience-centered AI products;
- New data on audience feelings about AI;
- Recommendations for AI literacy programs, specific AI disclosures and takeaways that will help participating organizations —



ALEX SMYNTYNA/POYNTER

Attendees gather in the Great Hall of the Poynter Institute for the Summit on AI, Ethics, and Journalism.

and any using this report — to experiment ethically and effectively with AI in their newsroom.

Poynter set out to accomplish the above, and begin regular AI ethics discussions that can hone editorial guidelines as technology advances. We aim to convene another summit next year that will bring in more U.S. organizations and international newsrooms. The agenda will include more open discussions and panels, per participant feedback, and will lead to updates to Poynter AI ethics guide, new audience research and another opportunity for newsrooms to refocus AI experimentation around audience needs.



Appendix

The Poynter AI Ethics Guide for Newsrooms

Based on feedback from the Summit, Poynter has updated its ethics guide and template for any newsroom looking to use generative AI in its work.

Please send us your edited version so we can see how you are proceeding. (ethics@poynter.org).



Template for a basic newsroom generative AI policy

Every news organization needs an artificial intelligence policy. The following is designed to be a starting point for newsrooms as they explore various uses for AI in their work.

*This policy was developed by [Kelly McBride](#), [Alex Mahadevan](#) and [Tony Elkins](#). It will be updated often. It was last updated on **Aug. 8, 2024** (updates listed at end)*

Please send us your edited version so we can see how you are proceeding.

(ethics@poynter.org).

_____ guidance and policies on using AI in our work.

(Your newsroom here)

Last updated: _____

Generative artificial intelligence is the use of large language models to create something new, such as text, images, graphics and interactive media. Although generative AI has the potential to improve newsgathering, it also has the potential to harm journalists' credibility and our unique relationship with our audience.

As we proceed, the following five core values will guide our work. These principles apply explicitly to the newsroom and throughout other non-news departments including advertising, events, marketing and development.

Transparency - both internal and external.

Externally, as we use AI in our journalism, we will document and describe the tools with specificity. When AI tools influence audience-facing content, we will tell the audience in ways that both disclose and educate news consumers. We will work with editors and designers to create disclosures that are precise in language without being onerous to our audience. This may be a short tagline, a caption or credit, or for something more substantial, an editor's note. When appropriate, we will include the prompts that are fed into the model to generate the material.

Our transparency works on multiple levels. Internally, it facilitates conversation and creativity. It will be clear to our peers and our bosses whenever we are using generative AI. This will facilitate collective learning and help us create applicable, transitory policies as the technologies evolve.



Poynter Summit on AI, Ethics & Journalism: Putting audience and ethics first

Externally, communication and disclosure ideally create opportunities to get feedback from the audience, as well as educate consumers. As journalists, part of our job is to empower the audience with news literacy skills. AI literacy – understanding how generative AI works, what benefits it brings to the information ecosystem and how to avoid AI-generated misinformation – is a subset of news literacy.

Accuracy and human verification - All information generated by AI requires human verification. Everything we publish will live up to our standards of verification. Increasingly in all of our work, it is important to be explicit about how we know facts are facts. This will be particularly important when using AI. For example, an editor should review prompts, and any other inputs used to generate a story or other material. And, everything should be replicable.

Audience service - Our work in AI should be guided by what will be useful to our audience as we serve them. We have made a promise to our audience to provide them with information that [fill in reference to newsroom’s basic mission statement].

Privacy and security - Our relationship with our audience is rooted in trust and respect. To that end, as we utilize AI to customize content or develop products that alter the way content is delivered, we will protect our audience’s data in accordance with our newsroom’s privacy policies. Most privacy policies forbid entering sensitive or identifying information about users, sources or even our own staff into any generative AI tools.

As technology advances and opportunities to customize content for users arise, we will be explicit about how your data is collected – in accordance with our organization’s privacy policy – and how it was used to personalize your experience.

Therefore, we will disclose any editorial content that has been created and distributed based on that personalization.

Exploration - With the four previous principles as our foundation, we will embrace exploration and experimentation. We should strive to invest in newsroom training – internal or external – so every staff member is knowledgeable in generative AI tools.



Logistics

The point person/team on generative AI in our newsroom is _____, who is supported by the following AI committee members: _____. Coordinate all use of AI here. This team will also be the source of frequent interim guidance distributed throughout our organization.

The team will seek input from a variety of roles, particularly those who are directly reporting the news.

You should expect to hear at least monthly communication from this team with updates on what we are doing and guidance on what activities are generally approved.

In addition, members of this team will:

- Monitor your content management systems, photo editing software and business software for updates that may include AI tools. Because software changes quickly and AI is being added to nearly every technology product, it's important to delegate appropriate team members to stay knowledgeable of updates.
- Write clear guidance about how you will or will not use AI in content-generation.
- Edit and finalize our AI policy and ensure that it is both internally available and where appropriate, publicly available (with our other standards and ethics guidelines).
- Seek input from our audience, through surveys, focus groups and other feedback mechanisms.
- Manage all disclosures about partnerships, grant funding or licensing from AI companies.
- Understand our privacy policies and explain how they apply to AI and other product development. This includes regularly consulting with editors, lawyers or other privacy experts that influence newsroom policies.
- Innovate ways to communicate with the audience to both educate them and gather data about their needs and concerns.
- Outline a clear process on how the policy will be updated, as specific as the number of meetings per month of your committee, who is on the committee, etc.



Poynter Summit on AI, Ethics & Journalism: Putting audience and ethics first

Current editorial use:

All uses of AI should start with journalism-centered intentions and cleared by the appointed AI group. Human verification and supervision is key. Here's the form you should use:

[Create a form to collect information, suggested questions]

How do you want to use AI?

What is the journalistic purpose of this work?

How can you gather knowledge on audience needs and attitudes about your intended use?

How should the audience's needs and attitudes inform your AI use?

How will you fact-check the results?

Will any material be published?

Which journalists will be responsible for overseeing this work and reporting out the results?

Which editors or managers will oversee the work?

What are the risks, the bad things that might happen? (ie. hallucinations, copyright or legal issues, data privacy violation) What safety nets can you devise to intervene before negative outcomes harm your newsroom's reputation?

What are the privacy implications of this use, and how will we protect user data? Have we run this by the experts who create and maintain our privacy policy?

Tools to use, tools to avoid

Include list of approved tools

What tools to avoid

Here are the generative AI sources we encourage you to use:

[List any LLMs that you have a formal relationship with, or that you are comfortable encouraging your colleagues to use.]

Please avoid the following services:

Entering our content: [Your policy will either approve entering your content into a specific LLM that you have a partnership with, or forbid entering your content into any LLMs]



Poynter Summit on AI, Ethics & Journalism: Putting audience and ethics first

Editorial use:

Generative AI is generally permitted for the following purposes (but please still fill out the form, so we know what's going on.)

Research - It's fine to ask a publicly available large language model to research a topic. However, you'll want to independently verify every fact. So be wary. It is fairly common for AI to "hallucinate" information, including facts, biographical information, and even newspaper citations.

Headline experimentation - Asking AI to generate headlines is a form of research. The same caveats apply. Also, be sure to put enough facts into the prompt that the headline is based on our journalism and not other reporting.

Summary paragraphs - **Do [or do not]** use AI to generate article summaries that appear at the top of our work. Our policy (below) is that we do not enter our content into any large language models. [Or state which LLM are permissible.]

Searching and assembling data - You are permitted to use AI to search for information, mine public databases or assemble and calculate statistics that would be useful to our audience. Any data analysis should be checked by an editor

Visuals - **Do/Do not use** _____ service to create illustrations for publication. All illustrations must contain the following credit: _____.

Do not use AI to manipulate photos unless they are for illustration purposes and clearly defined. Visual journalists need to be aware of software updates to photo processing tools to ensure AI-enhancement is not/is being used according to our policies. Do not publish any reader submitted content without first verifying its authenticity.

Fact-checking

- Use of AI alone is not sufficient for independent fact-checking. Facts should be checked against multiple authoritative sources that have been created, edited or curated by human beings. A single source is generally not sufficient; information should be checked against multiple sources.



Poynter Summit on AI, Ethics & Journalism: Putting audience and ethics first

Social media use

Use of verbatim GPT content **is/is not** permitted on our social channels.

- If verbatim GPT content is allowed, it should be edited for accuracy and voice by a human
- Audience teams should do regular content audits to ensure social copy/posts meet ethical guidelines

Privacy and security

- No personal information from your staff or your customers should be entered into programs
- None of our intellectual property should be entered into a program, outside of what is outlined above.
- Staff working with AI tools should have a clear understanding of our organization's privacy policy.
- When using generative AI to customize content for audience subsets, AI disclosure should include how user data was used to do so.
- Wherever possible newsrooms should run the LLMs behind generative AI tools locally, meaning on hardware that you own, via tools like AnythingLLM or GPT4All.

Ongoing training

Regular training on AI tools and experiments will be available and at times even mandatory. This training will be delivered by a combination of members from the internal committee and outside experts.

Creating custom GPTs

All custom GPTs must be approved by the AI oversight group. Know that the systems you develop with ChatGPT's custom GPT program will not be private. Any custom GPT code should be publicly available.

Use reliable sources to train custom GPTs. One of the best ways to create solid and useful output is to limit and control the sources a custom GPT draws on to material we can vouch for. In many cases this will mean limiting our custom GPTs to our own material.



Poynter Summit on AI, Ethics & Journalism: Putting audience and ethics first

Commitment to audience AI literacy

Along with this AI policy, we have developed an AI literacy [page/article](#) to help our audience understand how and why we're using generative AI. This material will be regularly updated to reflect our most current experimentation. As our language evolves, we will be better able to describe specific AI applications and tools. On the page, we link to resources, articles and other materials to/this article will:

- Help our audience understand the basics of generative AI.
- Explain why newsrooms use AI in their work.
- Build a more robust vocabulary for describing AI.
- Avoid AI-generated misinformation.
- Use chatbots responsibly to seek out factual information.
- And create responsibly using new generative AI tools.

Guidelines for product teams and technologists

Our product/technology team is committed to understanding and staying up to date with all tools, software or companies we use or partner with. We will:

- Vet third-party vendors and their usage policies before testing any AI product.
- Make sure any product we use adheres to our own data and privacy policies.
- Perform comprehensive testing on all software and tools for reliability and accuracy before using them for any consumer-facing content.
- Ensure all software settings are correct, and in accordance with our policies, before using any LLM.
- Keep up-to-date on the latest software updates for products we use.
- Provide best-practices, documentation or training for new tools to internal users.



Poynter Summit on AI, Ethics & Journalism: Putting audience and ethics first

Updates

March 27

Paragraph starting “Use reliable sources...” was added under “Custom GPTs”

Aug 8

Added in fifth value of privacy and security

Enhanced transparency values

Expanded the role of the committee members to include expertise on privacy

Added in material to encourage newsroom embrace AI literacy as part of news literacy responsibilities

Added guidelines for product teams and technologists



Speakers

- Alex Mahadevan**, Poynter
- Benjamin Toff**, University of Minnesota
- Burt Herman**, Hacks/Hackers
- Jay Dixit**, Open AI
- Joy Mayer**, Trusting News
- Kelly McBride**, Poynter
- Nikita Roy**, International Center for Journalists
- Paul Cheung**, Hacks/Hackers
- Phoebe Connelly**, The Washington Post
- Tony Elkins**, Poynter

Participants

- Adam Rose**, Starling Lab for Data Integrity
- Adriana Lacy**, Adriana Lacy Consulting
- Aimee Rinehart**, Associated Press
- Alissa Ambrose**, STAT/Boston Globe Media
- Annemarie Dooling**, Gannett
- April McCullum**, Vermont Public
- Ashton Marra**, 100 Days in Appalachia; West Virginia University
- Conan Gallaty**, Tampa Bay Times
- Darla Cameron**, Texas Tribune
- Dean Miller**, Lead Stories
- Elite Truong**, American Press Institute
- Enock Nyariki**, Poynter
- Erica Beshears Perel**, Center for Innovation and Sustainability in Local Media, UNC
- Ida Harris**, Black Enterprise
- Jay Rey**, Tampa Bay Newspapers
- Jennifer Orsi**, Poynter
- Jennifer 8. Lee**, Plympton and Writing Atlas

- Jeremy Gilbert**, Northwestern University
- Jessi Navarro**, Poynter
- Joe Hamilton**, St. Pete Catalyst
- Kathryn Varn**, Axios Tampa Bay
- Katie Sanders**, PolitiFact
- Lindsay Claiborn**, VERIFY/TEGNA
- Lloyd Armbrust**, OwnLocal
- Meghan Ashford-Grooms**, The Washington Post
- Mel Grau**, Poynter
- Mike Sunnucks**, Adams Publishing Group
- Mitesh Vashee**, Houston Landing
- Neil Brown**, Poynter
- Niketa Patel**, Craig Newmark Graduate School of Journalism at CUNY
- Peter Baniak**, McClatchy
- Rodney Gibbs**, National Trust for Local News
- Ryan Callihan**, Bradenton Herald
- Ryan Serpico**, Hearst
- S. Whitney Holmes**, The New Yorker
- Sarah Vassello**, Center for Innovation

- and Sustainability in Local Media, UNC
- Sean Marcus**, Poynter
- Shannan Bowen**, North Carolina Local News Workshop
- Teresa Frontado**, The 51st
- trina reynolds-tyler**, Invisible Institute
- Yiqing Shao**, Boston Globe Media Partners

Poynter.

The Craig Newmark Center for
ETHICS and
LEADERSHIP

 **IFCN**®

 MediaWise

POLITIFACT

Tampa Bay Times