

Poynter.

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).

(Your newsroom here) guidance and policies on using AI in our work.
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.

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 [REDACTED], who is supported by the following AI committee members: [REDACTED]. 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.

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]

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.

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.

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.

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