

5 Recommendations for Ethical AI in Local News

1

Newsrooms need clear goals for adopting AI tools

Documenting the goals for the use of an AI tool prior to its adoption can help mitigate unintended consequences.

Prior to tool adoption, news organizations should understand why they want to adopt AI technology and what news goals they are trying to achieve by doing so. Any AI tool should be a complement to the work that journalists and newsrooms already do.

One concern raised repeatedly from PAI Partners was that local newsrooms may deploy tools without fully understanding how the tools work and the implications of their use. For example, a recommender system that optimizes for engagement may yield higher subscription conversions and ultimately be more profitable for local newsrooms. At the same time, this may create incentives to over-index on crime news that is easy to create and generates engagement. Further, many tools that use natural language processing, such as transcription services, work well in English but are [less precise](#) in other languages, which might create an imbalance between journalists who work in non-English language contexts and those who work in English. Therefore, newsrooms should consider the relative priorities of their goals when deploying AI tools and how satisfying one goal (like time saved or engagement) might trade off against another (such as diversity of news stories or representation).

[Documenting the goals and intended purpose](#) of a proposed tool can illuminate the distance between intent and outcomes in practice. For example, a recommender system should ideally suggest stories that include many topics and cite a variety of sources. Writing down the implicit and explicit decisions that should go into making a news recommendation can be a valuable exercise toward both enhancing transparency and mitigating bias in a newsroom.

Documentation can take the form of an ethical checklist ([like this one](#)) where newsrooms answer questions about the intended purpose of a tool, which editorial standards may be implicated in its use, how it will impact readers, and the risks of unintended uses so that proper safeguards can be put in place to minimize potential harms. Recommendations 2 and 3 offer more insight into tactics for documenting AI decisions in local newsrooms.

2

Technology must embody the standards and values of the news operation

Establishing a broader ethics code is a fundamental first step toward ensuring transparency and accountability when AI tools are incorporated into the newsroom.

Newsrooms use ethical standards and frameworks to guide their decision-making. There are a variety of ways, however, that AI tools might conflict with these newsroom standards. A [recommender system](#) built to optimize for engagement may serve more polarizing stories to audiences than a human editor would. Audience data gathered by a recommender system may also be uploaded to a central database and used by a third party for other purposes, such as data mining for marketing agencies. Additionally, third-party photo searching and tagging systems may use that same tech to support facial-recognition tools.

Most of the ethical issues raised by the use of AI in local journalism are not new ethical issues. Human decision-makers are guided by newsroom standards and ethics codes while picking stories for the homepage, even if the codes do not explicitly cover this work. Discussing how technology implicates news standards and values lends itself to larger discussions about newsroom ethics and reflections on which stories need to be told and promoted.

Mapping newsroom standards onto technology and documenting how those standards should apply to technology adoption is a critical first step toward transparency. When implementing new tools, a newsroom can refer back to these standards to ask critical questions about which standards may be jeopardized.

Some newsrooms have even adopted AI ethics guidelines to guide purpose-driven uses of technology that align with their journalistic guidelines. The AI and Automation Lab at Bayerischer Rundfunk, Bavarian Broadcasting in Germany, for example, has a [useful example](#) of AI ethics guidelines, as does the BBC, with their [Machine Learning Ethics Principles checklist](#). For newsrooms without firm ethical standards in place, the Online News Association offers a [“Build Your Own Ethics Code”](#) for news organizations, small startups, and bloggers.

3

Transparency, explainability, and accountability mechanisms must accompany the implementation of AI tools

While AI ethics can be defined in many ways, transparency, explainability, and accountability are key values that ensure the ethical deployment of AI, including in a local news context.

In both journalism and AI ethics, transparency is understood as a pathway to accountability. The news industry has elevated transparency to an [essential standard](#) in journalism, as a remedy to loss in credibility and trust, and as a means to ensure public accountability. [According to Professor Mark Deuze](#) at the University of Amsterdam, journalistic transparency can be understood as “the ways in which people both inside and external to journalism are given a chance to monitor, check, criticize and even intervene in the journalistic process.” And in the AI ethics field, [transparency](#) (which makes the elements, goals, and origin of a system clear) enables accountability, fairness, and other AI ethics goals by shining a light on [black box systems](#).

For newsrooms implementing AI, transparency is a first step to articulating clear goals for tool adoption and also for ensuring adequate oversight. Transparency and explainability can help local newsrooms understand if a system is appropriate to accomplish a particular news goal, and also benefit audiences. They can also serve as a foundation for accountability, both internally (among journalists, editors, and operations teams) and externally (to the audiences they serve) who are also implicated by AI and deserve to understand its use in news.

But what does transparency and explainability look like in practice and how can they be achieved by local newsrooms using AI tools? PAI’s [Fairness, Transparency, and Accountability](#) team has produced [extensive guidelines](#) on how to enable transparency in AI systems through documentation, which is “both an artifact (in this case, a document with details about the system, similar to a nutrition label on food) and a process (in this case, a series of steps people follow in order to create the document).” Each can provide an infrastructure for making progress towards other ethics goals. Documentation can also enable more clarity on the overall mission and purpose of a news organization, better processes for communication, and more clarity around roles and expectations through the process of discussing and documenting an AI system.

There are a number of ways newsrooms can transparently explain the use of AI to their audiences. Newsrooms can use labels, for instance, to clearly identify machine-generated content, such as in reports about AI-generated media or on stories that were themselves generated by AI. (In [this story](#), for example, The New York Times marked

synthetic images with a label reading “Image Generated by AI.”) When using AI tools to report stories, news outlets could offer even greater transparency by using a sidebar, white paper, or technical repository to explain their use. [The Markup](#) and [BuzzFeed](#), for instance, have posted open-source data and code used for certain articles in [GitHub repositories](#). The New York Times runs a [blog](#) regularly posting about technical tools developed in-house. And The Wall Street Journal has posted [articles](#) explaining how some of their AI technologies drive their news coverage. Newsrooms should also use this phase of AI’s integration into newsrooms to experiment with other modes of transparency, like reporting statistics of a personalized paywall operation.

While local newsrooms may not have the resources to maintain a blog about all the uses of AI in their processes, news outlets should be transparent about the systems they use, such as listing third-party services used to personalize content, clearly explaining how a data-mining tool was used in a story to augment reporting, or adding an entry on the use of technology in a public ethics code. Any technical descriptions should be accessible to non-technical audiences so that readers can reasonably understand how AI was used.

Documentation alone, however, is not sufficient for ensuring responsible deployment of AI systems in newsrooms. The emphasis placed on ethical practices within newsrooms is contingent on [a leadership and culture](#) that values responsible journalism and AI and the integrity of news output. Business incentives may conflict with responsible use of AI tools, though transparency can also be [considered](#) as a strategic business practice. There are also limits to the levels of transparency and documentation that newsrooms can ethically provide, including legitimate reasons related to the protection of journalistic sources or audience privacy. Documentation, however, as both an artifact and process, can serve as a means to drive toward greater accountability, fairness, and explainability in local news processes and as a method of implementing each of the ethical principles outlined here.

These principles are intended to guide local newsrooms as they seek to harness the potential of AI tools and assumes that many of these organizations will not develop such tools in-house, instead relying on third-party software vendors to provide transcription services, automated story writing, or recommendation engines. Newsrooms that are developing machine learning tools themselves can refer to PAI’s [ABOUT ML Reference Document](#) for an extensive guide on documentation for machine learning tools and to the paper [“Algorithmic Transparency in the News Media”](#) for further discussion.

4

Newsroom staff need to actively supervise AI tools

Newsrooms should understand the AI tools they're using and maintain some level of control over these tools to oversee their use.

Any ethical checklist for AI tools in newsrooms should ensure a human is continually involved in maintaining a tool, outline how they are overseeing the system, and confirm that they have adequate training to do so. In many cases, tool oversight should not solely fall to an engineer or data scientist. For example, [an editor should](#) continually be involved in overseeing the stories selected by a recommender system on a news outlet's home page.

Human supervision will look different depending on the complexity of the tool and the risk of harm to both readers and a news organization's reputation. The potential harms of inaccurate transcription software are not the same as an auto writer that produces an inaccurate story. An auto writer or a recommender system will likely require more documentation than other tools. The BBC has grappled with ethical tradeoffs that emerge from the automated recommendation of content, such as how to increase engagement without compromising quality. In [this useful paper](#), which serves as a form of documentation itself, the BBC outlines their tactics for documenting content recommendation processes and dealing with this challenge.

If a tool makes use of data, documentation should include answers to questions about what data will be collected, the quality of that data, what the system is optimizing for, and the accuracy of its results. In some cases, local newsrooms may need to get additional information from third-party vendors about how their systems are built. This includes details about the data used to train the system, the choices made when building the system, any potential biases that may be embedded in those choices, and how different data, models, and systems work together. Such information can enable newsrooms to better tailor an AI tool to their use case and provide meaningful oversight.

Integrating AI tools in a newsroom also means understanding that there will be a tradeoff in how newsroom staff allocate their time and attention. Maintenance and oversight of AI means less time for staff to spend on other work. Newsrooms should be clear on the compromises they are making in building and maintaining these tools, including where attention may become scarcer because of deploying AI.

5

Distribution platforms must embed journalistic values into their AI systems

Consideration of AI and local news should include social media and other platforms for personalized news, such as news aggregators like Google News and Apple News, even if newsroom products seem separate.

While local newsroom journalists and editors may not be using AI tools widely, their work still heavily interacts with AI on the algorithmically powered platforms (like Facebook, Twitter, and Apple News) through which their content is shared. Here, we have been inspired by the work of the [News Quality Initiative](#), which “seeks to elevate news quality when algorithms rank and recommend news” and to “bring a critical eye to how platform news products are serving up local news.” Their work has made clear that “[small news outlets in particular are platform-dependent](#).” Further, the Markkula Center for Applied Ethics has launched a [News Distribution Ethics Roundtable](#) to untangle the implications of news distribution tactics that largely rely on AI techniques.

Therefore, while it is important to consider the ethical boundaries for AI in local newsrooms themselves, it is vital to the sustainability of local news that we consider AI ethics on the platforms distributing such content as well. Platforms should consider improving accountability by engaging in participatory panels, ones where journalists and platform representatives can strive to find consensus on values. By identifying these values, such panels could impact the ranking and algorithmic recommendation of news articles, as exemplified by the News Quality Initiative. Alongside other models for the participatory governance of social media like [platform democracy](#) and [citizen juries](#), this form of algorithmic governance applied to the AI-based distribution of local news might help ensure that AI tools are developed to ethically and fairly distribute local news content.