DRAFT: Artificial intelligence for independent news publishers

This guide aims to provide a snapshot of the main considerations, questions and opportunities surrounding the general use and applicability of artificial intelligence (AI) technologies in journalism.

This first draft will guide conversations taking place between members of the PINF network at the Independent News Forum in October 2023, and is subject to evolve following these conversations and additional questions at the event.

This guide will help readers to:

- Better understand the definitions of AI and generative AI (genAI) technologies;
- Explore some of the possible applications of AI for journalism;
- Reflect on the risks and challenges, as well as the factors that influence the implementation of AI in newsrooms;
- Review some of the key recent resources and updates in the field of AI.

Estimated reading time: 45 minutes to 1 hour. Click the section name above or in the navigation bar on the left hand side to navigate directly to the relevant part of the guide, or to quickly go between sections.

Part 1: What is artificial intelligence and why should I care?

Artificial intelligence (AI) is a collection of ideas, technologies, and techniques that relate to a computer system's capacity to perform tasks normally requiring human intelligence[1]. It's best to think of AI not as a tool or a technology, but as a scientific discipline, like mathematics or biology.[2]

The term 'artificial intelligence' was first coined[3] in 1955, but it's only in the last few years that it has increased in popularity. In journalism specifically, a newsroom survey and report from JournalismAI at Polis, LSE from 2019[4] found that AI

[1] https://reutersinstitute.politics.ox.ac.uk/our-research/industry-led-debate-how-uk-media-cover-artificial-intelligence
[2] https://course.elementsofai.com/1
[3] History of Artificial Intelligence (Council of Europe)
[4] New Powers, New Responsibilities (2019)

technologies – such as machine learning and Natural Language Processing (NLP) – were already being used for social media scraping, automation of simple content creation, or to comb through large amounts of documents.

Then, in the summer of 2022, people started experimenting with image generator DALL-E to make all sorts of creative, strange and beautiful art, and in November 2022, OpenAI launched ChatGPT, an AI-based chatbot that can generate content by responding to prompts from users.

Not long after, Microsoft announced the new AI-powered Bing search engine, Google introduced a conversational AI tool called Bard, and ChatGPT became publicly available. The widespread accessibility of these AI tools, combined with the marketing efforts of big technology companies and the news cycle, are some of the factors contributing to the hype surrounding artificial intelligence today.

Within the discipline of AI, there are **various subsets and sub-fields** (see diagram below). Perhaps the most popular one recently is generative AI (genAI)[5], a type of artificial intelligence that creates new and varied content by running algorithms on massive data sets. The newly created content resembles human-created content, and can be in the form of text, images, audio, video, or code. These tools and their algorithms respond to inputs and questions from the end user, but genAI systems have no conceptual understanding of the contents of the data sets they learn from, or whether the outputs they provide are true or correct.

Because of this, generative AI can generate content that appears to be authentic and accurate, but may be flawed, highly inaccurate, and may amplify harmful biases and discrimination. Earlier this year for example, the Guardian published an article explaining their discovery that ChatGPT had made up URLs referring to Guardian articles that had never actually been written (these events are called hallucinations[6]).

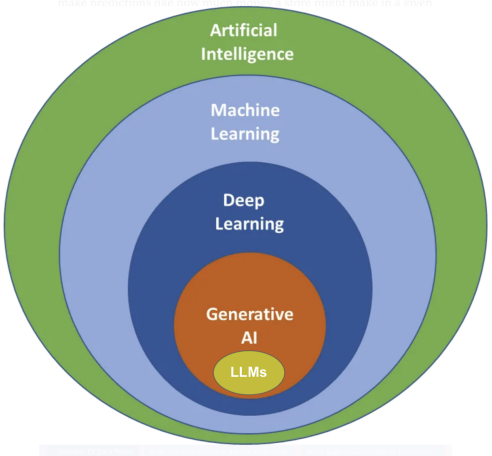

[5]The Financial Times has published a useful and clear interactive primer on genAI.
[6] "Hallucination is the term employed for the phenomenon where AI algorithms and deep learning neural networks produce outputs that are not real, do not match any data the algorithm has been trained on, or any other identifiable pattern. It cannot be explained by your programming, the input information, other factors such as incorrect data classification, inadequate training, inability to interpret questions in different languages, inability to contextualise questions." JournalismAI Generating change report (2023), glossary section).

Diagram adapted from a Medium post by Amol Wagh

A term frequently employed in relation to artificial intelligence is **machine learning (ML)**. At a fundamental level, machine learning can be understood as the use of data to answer questions. More formally, it refers to the use of algorithms that learn patterns from data and are able to perform tasks without being explicitly programmed to do so. A defining feature of machine learning systems is that they improve their performance with experience and data – meaning, they learn.

Deep learning (DL), a sub-field of machine learning, has networks capable of learning unsupervised from data that is unstructured or unlabelled. Also known as "deep neural learning or deep neural network", it is one of the most advanced contemporary applications of AI, powering a broad range of image, voice and text recognition tools.[7]

Finally, **Large Language Models (LLMs)**[8] are also powered by artificial intelligence that has been trained on vast quantities of text data to produce human-like responses to dialogue or other natural language inputs. In order to produce these natural language responses, LLMs make use of deep learning models. GPT-3 and GPT-4 (underpinning the ChatGPT interface) are examples of large language models.

There is a lot of **hype** surrounding AI, as well as many legitimate **concerns** (particularly in relation to its misuse, harmful applications or negative impact on journalists), which we'll look at in more detail in the following sections. However, AI has the capability to support existing newsgathering and news production techniques (as well as experimentation) and free up journalists' time. But this can only be possible in the context of a responsible use of artificial intelligence, and one that aligns with a news organisation's principles, mission and strategic goals.

Navigate directly to the resource section.

Part 2: Potential applications of AI in news and journalism

Hundreds of newsrooms and journalism support organisations around the world have already been experimenting with artificial intelligence one way or another. The latest report from the JournalismAI initiative, published in September 2023, surveyed 105 news and media organisations from 46 different countries about their use of AI and associated technologies.

Of those surveyed, 28% were newspapers, followed by publishing groups (20%) and broadcasters (16%). More than 75% of respondents said they currently use AI in at least one of the areas of newsgathering, production and distribution, while more than half of those surveyed said that "increasing efficiency and productivity to free up journalists for more creative work" were the main drivers for integrating AI in their newsroom and processes.

---

[7] JournalismAI Generating change report (2023), glossary section.
[8] This article by Timothy B Lee and Sean Trott explains LLMs with 'a minimum of math and jargon'

The majority of respondents had also already experimented with generative AI specifically, to write code, generate images or create summaries of articles and news stories.

When asked about the main areas of work they expected AI to influence, the newsroom leaders surveyed identified four main areas:

- Fact-checking and disinformation analysis*;
- Content personalisation and automation;
- Text summarisation and generation;
- Using chatbots to conduct preliminary interviews and gauge public sentiment on issues.

Some of the specific ways in which AI is currently being employed by organisations to gather, produce and distribute their news highlighted in the JournalismAI report include:

- Automation of interview transcription (using tools such as Otter.ai);
- Discovery of trending topics and news events by sifting through large amounts of data, using Crowdtangle, Google Trends or Dataminr;
- Collecting, verifying and cross-checking facts and claims;
- Writing headlines and article summaries;
- Proofreading and editing, using tools such as Grammarly;
- Writing code;
- Matching content with audiences at scale, or tailoring specific content to niche groups, by developing personalisation and recommendation systems (for example, like The Times has been doing with JAMES, its AI-powered tool[9])
- Adapting written text to other formats and mediums, such as audio, using voice-to-text;
- Scheduling and optimising social media content, using tools like Echobox and SocialFlow, and AI-driven SEO tools.

**It's interesting to note that respondents identified fact-checking as being one of the main areas where AI technologies can be of help to journalists, while simultaneously we have seen many cases of generative AI specifically contributing to misinformation and generating false or non-existent content. We can assume that respondents were likely referring to the potential positive impact of other AI technologies and softwares, and not to the use of generative AI for fact-checking.*

Using generative AI to write content

Local News Now, a Virginia-based media company, and Teyit, a Turkish independent fact-checking organisation, are two examples of news organisations experimenting with AI for content generation in newsletters.

[9] https://www.journalism.co.uk/news/the-times-employs-an-ai-powered-digital-butler-james-to-serve-personalised-news/s2/a739273/

Teyit's weekly newsletter provides a selection of verified claims and events to its audience (by generating summaries of claims previously verified by a person), while Local News Now's hyperlocal website ARLnow introduced an AI-generated morning newsletter that features an AI-written introduction and summaries of journalist-written stories. In an interview with Nieman Lab, the founder of Local News Now explained that he sees the AI-generated newsletter as a complementary product that doesn't take away valuable time from reporters: "AI is not a replacement for original local reporting. It's a way to take what has already been reported and repackage it so as to reach more readers."

The Marshall Project, a non-profit covering the US criminal justice system, has enlisted AI to write policy summaries: since the previous year, the newsroom had worked on requesting and analysing the policies of the state corrections department, which establish what publications and materials incarcerated people are allowed to read. After thoughtful consideration, the team settled on a four-step 'human-in-the-loop'[10] approach that used ChatGPT to produce summaries of the policies. Among the team's takeaways: "A machine-human hybrid approach opens up new reporting possibilities without compromising on editorial integrity, and helps already strapped newsrooms overcome resource constraints while allowing reporters, designers, and product teams to prioritise resource decisions by revealing what is must-have now, what is a need-to-add later, and what shouldn't be greenlit at all."

In the UK, one example at a local level is the newspaper NottinghamshireLive, which announced[11] to its readers in August that it would start experimenting with AI-written bullet point summaries at the top of articles, noting that these will always be reviewed by an editor prior to publication (as well as including a disclaimer explaining they were written with the help of AI).

The industry is still in a phase of exploration and it is still too early to show whether these experiments are benefitting newsrooms. More time and more research are needed to study the impact (positive or negative) of generative AI on editorial content and strategies, as well as on newsroom processes.

Collaborating to develop AI-powered solutions

In 2022, German regional public broadcasters Bayerischer Rundfunk and MDR built an NLP-based tool to analyse published content, as well as the comments of audiences, in order to identify feedback for the newsroom and find potentially underreported topics.

The project, called 'What's there, what's missing', was developed collaboratively as part of the JournalismAI Fellowship, which brings together international teams of journalists, developers, product managers and others to explore innovative solutions

[10] "In machine learning, human-in-the-loop (HITL) is used in the sense of humans aiding the computer in making the correct decisions in building a model." (Source: Wikipedia)
[11] https://www.nottinghampost.com/news/news-opinion/editors-letter-nottinghamshire-live-launches-8645265

to improve journalism through the use of AI technologies. This year for example, the newsrooms that are part of the 2023 Fellowship cohort[12] are working on projects including a tool for tailoring existing content to niche audiences, and a tracker that collects and monitors when members of parliament mention issues in which they also have a financial stake, among other ideas.

Collaboration between newsrooms of different kinds, sizes and locations could prove an effective way for news organisations to experiment with AI technologies. It can enable resource and expertise-sharing (particularly if, for example, some of the participants are more advanced in their use of AI), and provide an environment for testing the applicability of AI-based tools or processes that is not limited to a single's newsroom's content, workflows and audience.

Of the newsrooms surveyed in 2023 by JournalismAI, the majority were in favour of more collaboration not only between newsrooms, but also between newsrooms, academic institutions and other media organisations globally, due to its potential to "lessen the disparity" between small and large newsrooms, as well as between organisations based in the Global South versus the Global North.

Experimenting with AI in US local newsrooms

Local News AI is a two-year initiative from the Associated Press, helping local newsrooms identify and adopt artificial intelligence-based solutions. Among other activities and resources (included later in this guide), they are working with five local newsrooms from areas including Michigan and Puerto Rico, to develop and implement projects using artificial intelligence and generative AI.

The projects include automating summaries of public meetings, automating the transcription of recorded videos and summarising transcripts to create the initial framework of an article, and sorting of news tips and coverage pitches from the public and automatically populating them into a coverage planner.

Additional reading:

- *David Caswell (formerly of BBC News Labs), a consultant, builder and researcher focused on AI in news, has published a report examining the ways in which newsroom leaders at larger organisations are thinking of the practical and strategic implications of artificial intelligence. He provides other examples of projects using genAI in the section called 'Practical deployment of generative AI in newsrooms';*
- *This article dates back to February, but Nicholas Diakopolous, Professor in Communication Studies and Computer Science at Northwestern University in the US, wrote an initial roundup of the possible uses of generative AI in news production (including summarising texts, translation and document analysis), and it can be interesting to compare how some of these uses are being applied today.*

[12] https://blogs.lse.ac.uk/polis/2023/06/01/meet-the-2023-journalismai-fellowship-cohort/

Part 3: Risks, challenges and considerations about the use of AI in journalism

For any positive opportunities it could provide, artificial intelligence also poses a great deal of risks and challenges for publishers and journalism, spanning editorial quality and accuracy, ethics, technical competency and sustainability, to name a few.

Limitations of generative AI in particular include, for example, a lack of accuracy and an increased potential for contributing to misinformation; the exploitation of underpaid workers to train AI systems; the replication of harmful biases[13] such as racism and sexism; and issues of plagiarism or copyright infringement.

In both JournalismAI surveys, newsrooms have emphasised the importance of the human-in-the-loop approach, deeming human intervention as being crucial in mitigating harmful risks, algorithmic biases and inaccuracies, especially in matters of creating and fact-checking content, and providing context to news events and issues.

The pressure of updating skill sets, processes and newsroom roles

The emergence of AI and its rapid evolution demands at least a basic level of AI literacy if news organisations are to experiment with it, or be able to make a sound judgement on whether it can help or hinder their work.

Arguably, some of the skills required to better understand this discipline – sharp judgement and the ability to analyse and interrogate multiple sources of information – are a core part of journalists' arsenal and already employed in traditional newsgathering and reporting (whether they're covering the technology and AI beat, or not). However, obtaining and maintaining at least a basic updated level of AI literacy, and an understanding of its potential applications, does require time, training and resources, which are often difficult to free or to come by in thinly stretched and underfunded local newsrooms.

AI tools require more oversight than traditional software

Partnership on AI (PAI), a US non-profit collaboration between academic, civil society, industry, and media organisations, recently published a draft guide[14] for AI procurement and use. It divides the potential AI tools newsrooms can use into 5 categories according to their purpose (such as lead generation or data analysis), highlighting some key differences between AI tools and other tools used by newsrooms, and sets out principles and a step-by-step guide for AI adoption. Perhaps crucially, it emphasises why AI tools should be carefully utilised, and always in conjunction with human assessment and oversight:

---

[13] https://www.rollingstone.com/culture/culture-features/women-warnings-ai-danger-risk-before-chatgpt-1234804367/
[14] https://partnershiponai.org/pai-seeks-public-comment-on-the-ai-procurement-guidebook-for-newsrooms/

"Several features differentiate AI tools from other technologies, like traditional software; for example, they are sometimes unexplainable, can lack context to draw appropriate and/or accurate conclusions, and can return unsafe outputs. Traditional software relies on a rules based system where the outputs are the same every time. AI tools are iterative and often make decisions without explicit programming. Unlike with traditional software, we don't always have insight into how AI systems arrive at their conclusions or the factors involved. AI tools therefore require an additional layer of oversight that might not have been previously necessary with traditional software that is 'plug and play' and produces the same results using the same processes everytime", the guide states.

Developing transparent guidelines for a responsible use of AI

As noted earlier in this guide, newsrooms have been using AI-powered tools like data scrapers or analytics dashboards for years to do their jobs, without particular disclosures. However, when generative AI became widely available and newsrooms started experimenting with content generation, it quickly became apparent that a set of principles or guidelines would be necessary: both to ensure that experiments wouldn't go awry, and also to inform the wider public, which now has the same level of access to generative AI tools and interfaces that journalists and newsrooms do.

Nicholas Diakopolous and Hannes Cools, a postdoctoral researcher at the AI, Media, and Democracy Lab at the University of Amsterdam, have been collecting the guidelines published by a variety of news organisations, and analysing patterns and similarities. The principles they reviewed emphasise oversight and "meaningful human involvement and supervision in the use of AI", and "reject the idea of replacing journalists with machines and highlight the importance of the decision-making role of humans when using generative AI tools."

"Artificial intelligence can help the work of journalists by refining raw data in the same way as software like Excel does, and as support for writing articles, like using online thesaurus or proofreaders and spell-checkers. AIs are tools, but in no way sources of information", states Swiss outlet Heidi.news in its brief guidelines[15].

Taking collective action to question AI and shape how it's used in the public interest

In an article recently published on Nieman Lab, Mike Ananny, associate professor of communication and journalism at the University of Southern California Annenberg School, and Jake Karr, deputy director of New York University's Technology Law and Policy Clinic, make an interesting case that being able to control the language of AI equates to maintaining press freedom.

The authors argue that the conversation surrounding generative AI has carried a tone of "inevitability about the technology and its ability to replace, improve, or swallow the news", and that journalists have participated in this conversation in a

[15] https://www.heidi.news/cyber/la-redaction-de-heidi-news-prend-position-sur-l-usage-des-intelligences-artificielles

"largely defensive, reactive posture", focused mostly on potential job and revenue losses in the industry and on figuring out how genAI could be used in a way that meets journalistic standards of accuracy and rigour.

Instead, they challenge journalists to draw inspiration from the likes of the Writers and Screen Actor Guilds in the US[16] to "find their collective voice on genAI" and use it to critique, reject and re-shape genAI systems that don't meet the standards of journalistic work.

This sentiment was echoed in the JournalismAI survey, in which respondents voiced concerns about the fact that it is technology companies driving AI innovation, and the concentration of power in these companies that are profit-driven and lack transparency could "further commercialise journalism, boosting poor quality and polarising content, and leading to a further decline in public trust in journalism."

In what could be an interesting starting point, in the midst of the European Union's negotiations on an AI Act, 26 organisations representing thousands of creative professionals around the world (including news, magazine and publishing) have released a set of Global principles for artificial intelligence (AI). These guidelines aim to provide an ethical and accountable framework for the development, use and regulation of AI systems, and signatories include various European and country-specific publishers' associations.

Part 4: AI in the news – key updates and resources

Advances in artificial intelligence – both generally and within journalism – are moving at a rapid pace. Day by day, new publishers are deciding how (and if) they will incorporate generative AI in their work, publishing guidelines or statements, and experimenting with genAI. Technology companies are continuing to advance their research and development of AI-powered products, and other industries as well as governments are grappling with the use and potential impact of AI.

Below are some key updates that have taken place in recent months in the field of AI, as well as some resources to explore if you are considering the potential use of AI and genAI in your organisation, and some people to follow who are regularly thinking, talking or writing about updates in this field.

Artificial intelligence in the news

BBC has announced its generative AI guidelines (October 2023): The BBC has stated that it will announce a number of projects that explore how genAI could play a role in "both what we make and how we work" across a range of fields, including journalism research and production, content discovery and archive, and personalised experiences. They are also attempting to safeguard the public interest, by preventing web crawlers from using BBC content to train AI models.

[16] https://www.poynter.org/business-work/2023/artificial-intelligence-writers-guild-unions-journalism-jobs/

The UK government will hold its first AI Safety Summit and side events: On 1 and 2 November, the government will bring together key countries, as well as leading technology organisations, academia and civil society to inform rapid national and international action in the development of AI. To complement the summit, a series of events will take place under the AI Fringe, including keynotes and panels from industry partners such as Ada Lovelace Institute and The Alan Turing Institute.

Ongoing negotiations over the European Union's AI Act (September 2023): Back in June, Members of the European Parliament adopted a negotiating position on the EU's first AI Act, and discussions are currently underway with individual member states to define the law's final form before it is passed. In September, Algorithm Watch and 118 other organisations signed a statement calling for the closing of a loophole in the Act, which relates to the classification of AI systems deemed 'high risk' (such as those used to monitor students, assess consumers' creditworthiness, evaluate job-seekers, etc). The statement argues that unless the loophole is closed, it could leave it up to those individuals or companies developing or deploying AI systems to decide themselves if the systems they created are 'high risk'.

Developing organisational policies for generative AI (October 2023): Also this past June, Nesta and Newspeak House launched the Civic AI Observatory, an initiative to support civic organisations plan and adapt to the rapidly evolving field of Generative AI, including providing guidelines, case studies and helping to form communities of practice for specific industries. The Observatory will hold an in-person unconference on 20 October, and it has recently published a round-up of policies and guidelines on using and procuring generative AI tools from various institutions (including newsrooms and the government).

European AI & Society Fund (ongoing): This fund was set up in 2020 to ensure that civil society can be an active and effective partner in shaping European policies so that AI serves the needs of people and society. Fourteen foundations support this work, which focuses on funding, learning and collaboration to build capacity in civil society organisations to shape AI policy and advocacy. In summer 2023, the Fund conducted a landscape review of the philanthropic institutions that are funding AI work in Europe.

People to follow to stay updated about AI developments

For practical resources and interesting reads:

- Charlie Beckett, founder of Polis at LSE, and Tshepo Tshabalala, manager of JournalismAI;
- Mattia Peretti, freelance AI consultant (currently with Internews);
- Aimee Rinehart, senior product manager for AI strategy, Associated Press;
- Nicholas Diakopolous, Professor in Communication Studies and Computer Science at Northwestern University, who edits the Generative AI Newsroom publication on Medium;
- David Caswell, AI consultant;

- Chris Moran, head of editorial innovation, the Guardian.

For coverage of AI policy and technology updates:

- Melissa Heikkilä, senior AI reporter for MIT Technology Review, also writes The Algorithm newsletter;
- Madhumita Murgia, artificial intelligence editor at the Financial Times;
- Karen Hao, AI reporter and contributing writer at The Atlantic;
- Timnit Gebru, founder of the Distributed AI Research Institute (DAIR) and former co-lead of the Ethical AI Research Team at Google .

Six resources to explore further

- JournalismAI is a global community of practice that provides research, training and resources – every year they organise an Academy for Small Newsrooms, a Fellowship programme, and the JournalismAI Festival, which are all free to access. Their online Slack community recently opened up to the public, and for those just beginning to explore the topic, their Starter Pack and Discovery newsletter course can be a useful starting point;
- Partnership on AI (PAI) has developed and maintains a public database of AI tools and how they are being used by newsrooms, as well as the level of technical expertise needed to use them. The list currently features 75 examples, including the Press Association's RADAR and Full Fact's fact-checking software, and tools that some newsrooms have likely already been using (such as Chartbeat for analytics, or Parse.ly for audience engagement and data visualisation);
- After publishing a study of US newsrooms' use of AI in 2022, The Associated Press' Local News AI initiative has developed a free self-guided course for newsrooms on how to build an AI strategy across newsgathering, production, distribution and business operations. To accompany it, there's a short scorecard designed to assess a news organisation's readiness to start using AI;
- Other self-guided courses on AI and genAI basics include:
    - Elements of AI (developed by the University of Helsinki);
    - Google News Initiative's Introduction to machine learning;
    - How to use ChatGPT and other generative AI tools in your newsroom, from the Knight Center for Journalism in the Americas.
- Joe Amditis of the Center for Cooperative Media in the US has published a comprehensive handbook for small publishers interested in experimenting with (and making the most of) generative AI prompts;
- Newsletters focusing on AI developments include The Happy Journalist, written by Jacqui Merrington, and Develop AI by Paul McNally, founder of a company with the same name, based in South Africa that reports on AI, trains journalists to code with AI and builds tools for media businesses.