

Motivations, Goals, and Pathways for AI Literacy for Journalism

MANDI CAI* and SACHITA NISHAL*, Northwestern University, USA

Science and technology news, which includes coverage on AI technologies, serves a plurality of functions for lay audiences today, such as education, agenda-setting, and curation. Reporting on AI is particularly challenging given the lack of transparency surrounding most black-box algorithms, and the rapid pace of progress in this domain. In addition to the difficulties specific to AI reporting, journalists' ability to report newsworthy information in a way that adheres to their professional norms of objectivity, accuracy, and transparency is increasingly challenged by institutional upheavals within their profession. In this workshop paper, we make a case for the importance of AI literacy among journalists, and lay out objectives for a curriculum designed to teach journalists about AI, and how they can approach reporting on it, given their prominent role in the dissemination of scientific information to the public.

Additional Key Words and Phrases: AI literacy, journalism studies

ACM Reference Format:

Mandi Cai and Sachita Nishal. 2023. Motivations, Goals, and Pathways for AI Literacy for Journalism. In *CHI '23 Workshop on AI Literacy: Finding Common Threads between Education, Design, Policy, and Explainability, April 23–28, 2023, Hamburg, Germany*. ACM, New York, NY, USA, 6 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 MOTIVATING AI LITERACY WITHIN JOURNALISM

Journalists who report on science and technology today not only *inform* society about novel discoveries, but also *contextualize* these developments with respect to the scientific ecosystem, societal impacts, and even policy repercussions [8]. This especially applies to reporting on artificial intelligence (AI) and allied disciplines, since these technologies increasingly permeate several aspects of people's lives: from content recommendation on streaming services, to AI-powered voice assistants, to the facial recognition in our phones, and so on. In tandem with this rising sea of public-facing AI deployment, we also observe incidents of over-hyped model capabilities, data-based discrimination, biased algorithmic outputs, and misleading algorithmic outputs [7, 24, 35, 36]. Given these issues, holding AI models and their deployments accountable has emerged to be an incredibly important role of the press [12, 15].

In reality, however, hype surrounding AI and its capabilities persists in media coverage [4–6]. Kapoor and Narayanan have laid out a broader taxonomy of pitfalls in the media coverage of AI in five major American news outlets — chief among these are flawed comparisons of AI and human abilities, hyperbolic claims about AI, and neutral treatments of institutions who have vested interests in the narratives surrounding AI [26].

Interviews with science journalists highlight how data literacy and formal scientific backgrounds can help them understand, question, and contextualize dense scientific information, leading to higher quality reporting [2]. However, not all reporters who write about AI's wide-ranging impacts have such backgrounds. One case study found that journalists were often confused about the conceptual distinctions between "AI", "algorithms", and "automation", and their intuitions about AI often relied on guesswork and imagination [25]. AI literacy efforts specifically targeting

*Both authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Association for Computing Machinery.

Manuscript submitted to ACM

journalists, though not a panacea, can thus be a start towards media coverage that considers AI systems outside of the umbrella of hype and with nuanced descriptions of the underlying technology and its capabilities [14].

In this paper, we argue that every journalist, no matter their specialty, should be broadly AI literate. We propose the design of a curriculum for AI literacy, and articulate specific goals for the curriculum rooted in different parts of the news-making process. We then highlight the broader tensions and challenges within journalism that any efforts towards AI literacy would need to account for, to be practically applicable. We conclude with raising future research questions geared towards improving AI literacy for reporters, and suggest methods to design such programs.

2 GOALS OF CURRICULUM FOR JOURNALISTIC AI LITERACY

We now discuss the goals of our proposed curriculum, and elaborate how they might support specific tasks involved within news reporting. We explicitly consider how the implementation of these goals might benefit journalists across different reporting contexts: reporting formats (e.g. print, digital, TV), employment contexts (e.g. staff journalism, freelancing, general reporting, beat-specific reporting), and varied roles that journalists perform for their audiences (e.g. watch-dogging, curation, advocacy). Given these conditions, our proposed curriculum aims to:

- **Foster discussion** of five conceptual questions explored in literature on AI education, “What is AI?”, “What can AI do?”, “How does AI work?”, “How should AI be used?”, and “How do people perceive AI?”, in the context of journalists’ topics of coverage [29].
- **Support interrogation** of technical and societal aspects of AI, including datasets, areas of application, funding sources, and so on [13].
- **Offer sources** outside of the industry of commercial AI applications who can speak to the nuances and ethics of building and deploying AI, given that AI news coverage is already saturated with industry perspectives [10, 34].
- **Enable reflection** by mapping competencies of AI literacy [29] onto journalists’ role perceptions [20], to consider how they can communicate about AI for their specific audiences, and what they need to learn further.

The following subsections delve into these goals, and elaborate upon their inspiration, considerations, and examples.

2.1 Foster Discussion of Conceptual Questions Around AI

To foster conceptual discussions about AI, we borrow the five overarching questions presented by Long and Magerko (2020): “What is AI?”, “What can AI do?”, “How does AI work?”, “How should AI be used?”, and “How do people perceive AI?”, given their suitability for AI education for non-technical learners [29]. Our curriculum would also echo that computational literacy is not required to achieve basic AI literacy. Many reporting positions do not explicitly require data, statistical, or computational literacy, and presenting these skills as pre-requisites to AI literacy may discourage reporters from learning about the subject.

Practically, journalists either report on selected “beats” that represent a topical or thematic specialization [30], or conduct general assignment reporting on broader topics. We argue that fostering discussion about the overarching questions about AI, with respect to case-studies within specific beats, can provide journalists a more grounded, concrete understanding of AI concepts. To this end, databases of public-facing algorithmic deployments [16, 19] can be categorized and discussed, based on specific beats or general interest topics. For example, with regards to the criminal justice beat, journalists could be prompted to consider the overarching questions with examples of data-driven risk assessment algorithms and the problematic input of criminal record data into these algorithms [31, 37]. In this case, a potential answer for “What is AI?” is “a technology used to predict human behavior, often used by law enforcement”, which

would be abstracted further as more examples are considered. For “What can AI do?”, an answer could elaborate on the ways that risk assessment algorithms accurately or inaccurately predict recidivism. The questions “How should AI be used?” and “How do people perceive AI?” offer opportunities to contend with harmful predictions made by risk assessment algorithms that target individuals from under-represented minorities.

Conceptual clarity about AI can support news discovery in the long-term, by enabling journalists to quickly identify examples of controversial, impactful, or novel applications of AI that can be “newsworthy” to their audiences [23].

2.2 Empower Journalists to Interrogate AI Systems

Should they choose to pursue a more technically-focused investigation, journalists must be capable of asking concrete questions about AI models, their training datasets, and the institutions responsible for building these algorithms [13]. Our proposed curriculum would present existing methods employed by journalists to this end, including reverse-engineering algorithms and code inspection (which would require some level of computational literacy). Simpler approaches like algorithmic auditing by varying inputs and inspecting the outputs can also be taught [13, 15]. Though a highly technical code review would require some computational literacy, understanding the general structure of a model, artifacts or potential biases in training data, the role of the designers, and who the algorithm impacts, does not. Such simple but important information can especially help staff journalists build partnerships with technologists within their newsrooms, to conduct more technically-substantiated reporting [33, 39].

Journalists can start from some of the overarching questions discussed in the first subsection, reformatting them as “What can this algorithm do?”, “How does this algorithm work?”, “How should this algorithm be used?”, “Was this algorithm audited?”, and “How do people perceive this algorithm?” to consider a specific deployment of an AI system, as a guide for who to interview as a source and what to center the conversations around. A lack of transparency in the answers to any of these questions should be explicitly noted in news coverage, in order to raise awareness about the “black-box” nature of some AI systems.

2.3 Offer Sources Outside Industry

In the last decade, mainstream news coverage on AI in the UK and the US has primarily centered consumer-facing AI and industry sources [10]. As journalists further develop a story, leaning heavily on industry sources who have a vested interest in the success of an AI product does not make for a complete, nuanced public image of AI.

Within the proposed curriculum, we seek to gather and present a preliminary list of sources, consisting of academic researchers, AI ethicists, politicians and lawyers interested in AI regulation, and individuals from the general public who have reason to believe they have been affected by AI technologies — all of whom are willing to speak to journalists about a critical evaluation of AI and its impact on people across socio-economic demographics. We acknowledge that this initial list will likely be biased towards featuring people who have the time and energy to speak with journalists, and it requires journalists to slowly build relationships with harder-to-reach populations. Still, it can provide a good starting point for journalists to cover AI from nuanced, broader, and less biased perspectives.

2.4 Enable Reflection by Mapping AI Competencies to Journalists’ Role Perceptions

When presented with technical or complex information about AI, journalists may ask how investing time in learning about AI tangibly supports their ability to develop newsworthy stories for their editors and audiences. To address this question, our proposed curriculum maps the competencies attached to each of the five overarching questions distilled by Long and Magerko (2020) to common journalistic role perceptions [29]. Journalists’ perceptions of their roles vary

by personal, social, cultural, and professional factors [17]. Normative role perceptions of journalists have included the gatekeeper, the advocate, the curator, the watchdog, the public representative, and so on [17, 20].

As journalists exercise their specific roles and attend to the informational needs and expectations of their audiences (e.g. laypersons, policy experts, science enthusiasts, etc.), it could be helpful to consider specific competencies that help them build stories more effectively. For instance, watchdog journalists (e.g. at ProPublica) could benefit from competencies like the ability to recognize the human architects of AI systems, specific problems that are challenging for AI systems, and signs of AI failure [36]. These would enable monitoring of situations in which AI has been problematically deployed by institutions, and subsequently interviewing system developers and designers. For journalists who act as curators of exciting information (e.g. at ArsTechnica), developing a greater imagination about futuristic AI applications and a greater intuition for understanding inter-disciplinary work within AI would be quite useful.

3 OBSTACLES TO BUILDING JOURNALISTIC AI LITERACY

Interventions such as curricula and training programs are enacted within a broader socio-technical system involving labor considerations, power dynamics, resource availability, and so on. In this section, we consider the tensions within the present institutions of journalism and how they may hinder AI literacy efforts.

Several institutional disruptions in the news industry have intensified journalistic work over the last few decades, including a massive shift towards digital modes of publication and distribution, the collapse of subscription-based profit models, dwindling staff jobs, general reporters and nonspecialists taking on science reporting, and a proliferation of non-journalistic science information online [1, 2, 9, 18, 20]. When reporting on AI, journalists must further contend with the breakneck pace of scientific publication, while addressing misinformation from unreliable non-specialists and social media [11]. Combined, these factors have led to a lack of institutional support and time availability for journalists, which can detract from their ability to fact-check and verify information — a process already marked by pragmatic compromises and a high reliance on expert sources [3, 38]. This can also lead journalists to view the acquisition of technical training and data literacy as prohibitively time-consuming [2].

Press releases from institutions and corporations have traditionally helped journalists to rapidly detect newsworthy information about novel technical developments and applications [22]. However, over-reliance on press releases can lead to exaggerated reporting [41], and AI-related coverage is susceptible to this as well [28]. It is thus not enough for journalists to simply understand technical terms and processes surrounding AI: they must also understand how to question expert claims and interrogate demonstrations of algorithmic applications. Increasingly, there are calls for the cautious release of AI artifacts from within the scientific community [27, 40]. However, nuanced communication from scientists cannot substitute for well-informed coverage from journalists, which reaches a different audience.

4 THE PATH FORWARD AND FUTURE WORK

Motivated by the proliferation of AI-hype in the news media, this paper puts forward a set of objectives that an AI literacy initiative for journalists would seek to fulfill. We draw from extant scholarship on AI literacy and journalism studies to give examples of some tangible activities that could help achieve these goals, while taking into account the high variability of reporting contexts, employment conditions, and domain expertise among reporters. As we set out to achieve these goals, we must also take into account the lack of time and institutional support for reporters, and exaggerated reporting that can ensue from the widespread reliance of PR material. Due to these variations and limitations, AI literacy for journalism can never be one-size-fits-all, and we call for future work that leverages frameworks of participatory design and value-sensitive design to build effective curricula, tooling, and educational material to support

this endeavor [21, 32]. Ultimately, to ensure that the public understands the benefits, risks, and impacts of AI in their lives, we need to make sure that their informants, educators, and advocates, i.e. journalists, do so too.

REFERENCES

- [1] Stuart Allan. 2011. Introduction: Science journalism in a digital age. *Journalism* 12, 7 (Oct. 2011), 771–777. <https://doi.org/10.1177/1464884911412688>
- [2] Josh Anderson and Anthony Dudo. 2023. A View From the Trenches: Interviews With Journalists About Reporting Science News. *Science Communication* (Jan. 2023), 107554702211491. <https://doi.org/10.1177/10755470221149156>
- [3] Aviv Barnoy and Zvi Reich. 2019. The When, Why, How and So-What of Verifications. *Journalism Studies* 20, 16 (Dec. 2019), 2312–2330. <https://doi.org/10.1080/1461670X.2019.1593881> Publisher: Routledge _eprint: <https://doi.org/10.1080/1461670X.2019.1593881>
- [4] Emily Bender and Chirag Shah. 2022. All-knowing machines are a fantasy. <https://iai.tv/articles/all-knowing-machines-are-a-fantasy-auid-2334>
- [5] Emily M. Bender. 2022. ChatGPT hype in the Washington Post. <https://medium.com/@emilymenonbender/chatgpt-hype-in-the-washington-post-c4e1355ed31b>
- [6] Emily M. Bender. 2022. On NYT Magazine on AI: Resist the Urge to be Impressed. <https://medium.com/@emilymenonbender/on-nyt-magazine-on-ai-resist-the-urge-to-be-impressed-3d92fd9a0edd>
- [7] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? . In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Virtual Event Canada, 610–623. <https://doi.org/10.1145/3442188.3445922>
- [8] Deborah Blum. 2021. Science journalism grows up. *Science* 372, 6540 (April 2021). <https://doi.org/10.1126/science.abj0434>
- [9] Joshua A. Braun and Jessica L. Eklund. 2019. Fake News, Real Money: Ad Tech Platforms, Profit-Driven Hoaxes, and the Business of Journalism. *Digital Journalism* 7, 1 (Jan. 2019), 1–21. <https://doi.org/10.1080/21670811.2018.1556314>
- [10] J Scott Brennen, Philip N Howard, and Rasmus Kleis Nielsen. 2018. An Industry-Led Debate: How UK Media Cover Artificial Intelligence. (2018).
- [11] Michael Brüggemann, Ines Lörcher, and Stefanie Walter. 2020. Post-normal science communication: exploring the blurring boundaries of science and journalism. *Journal of Science Communication* 19, 3 (June 2020). <https://doi.org/10.22323/2.19030202>
- [12] Madalina Busuioc. 2021. Accountable Artificial Intelligence: Holding Algorithms to Account. *Public Administration Review* 81, 5 (2021). <https://doi.org/10.1111/puar.13293>
- [13] Tania Cerquitelli, Daniele Quercia, and Frank Pasquale (Eds.). 2017. *Transparent Data Mining for Big and Small Data*. Studies in Big Data, Vol. 32. Springer International Publishing, Cham. <https://doi.org/10.1007/978-3-319-54024-5>
- [14] Mark Deuze and Charlie Beckett. 2022. Imagination, Algorithms and News: Developing AI Literacy for Journalism. *Digital Journalism* 10, 10 (Nov. 2022), 1913–1918. <https://doi.org/10.1080/21670811.2022.2119152>
- [15] Nicholas Diakopoulos. 2015. Algorithmic Accountability. *Digital Journalism* 3, 3 (May 2015). <https://doi.org/10.1080/21670811.2014.976411>
- [16] Nicholas Diakopoulos, Daniel Trielli, and Grace Lee. 2021. Towards Understanding and Supporting Journalistic Practices Using Semi-Automated News Discovery Tools. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (Oct. 2021), 1–30. <https://doi.org/10.1145/3479550>
- [17] Wolfgang Donsbach. 2012. Journalists' Role Perception. In *The International Encyclopedia of Communication*, Wolfgang Donsbach (Ed.). John Wiley & Sons, Ltd, Chichester, UK, wbiecj010.pub2. <https://doi.org/10.1002/9781405186407.wbiecj010.pub2>
- [18] Sharon Dunwoody. 2021. Science journalism: Prospects in the digital age. In *Routledge Handbook of Public Communication of Science and Technology* (3 ed.). Routledge.
- [19] David Freeman Engstrom, Daniel E. Ho, Catherine M. Sharkey, and Mariano-Florentino Cuéllar. 2020. *Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies*. Technical Report. Administrative Conference of the United States. <https://www.ssrn.com/abstract=3551505>
- [20] Declan Fahy and Matthew Nisbet. 2011. The science journalist online: Shifting roles and emerging practices. *Journalism* 12, 7 (Oct. 2011). <https://doi.org/10.1177/1464884911412697>
- [21] Batya Friedman, Peter Kahn, and Alan Borning. 2002. Value sensitive design: Theory and methods. *University of Washington Technical Report 2* (2002), 12. Publisher: Citeseer.
- [22] Oscar H. Gandy. 1980. Information in health: subsidised news. *Media, Culture & Society* 2, 2 (April 1980). <https://doi.org/10.1177/016344378000200201>
- [23] Tony Harcup and Deirdre O'Neill. 2017. What is News?: News values revisited (again). *Journalism Studies* 18, 12 (Dec. 2017), 1470–1488. <https://doi.org/10.1080/1461670X.2016.1150193>
- [24] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Yejin Bang, Andrea Madotto, and Pascale Fung. 2022. Survey of Hallucination in Natural Language Generation. *Comput. Surveys* (Nov. 2022). <https://doi.org/10.1145/3571730>
- [25] Bronwyn Jones, Rhianne Jones, and Ewa Luger. 2022. AI 'Everywhere and Nowhere': Addressing the AI Intelligibility Problem in Public Service Journalism. *Digital Journalism* 10, 10 (Nov. 2022), 1731–1755. <https://doi.org/10.1080/21670811.2022.2145328>
- [26] Sayash Kapoor and Arvind Narayanan. 2022. Eighteen pitfalls to beware of in AI journalism. <https://aisnakeoil.substack.com/p/eighteen-pitfalls-to-beware-of-in>
- [27] Percy Liang, Rishi Bommasani, Kathleen Creel, and Rob Reich. 2022. The Time Is Now to Develop Community Norms for the Release of Foundation Models. <https://hai.stanford.edu/news/time-now-develop-community-norms-release-foundation-models>

- [28] Author Zachary C. Lipton. 2019. OpenAI Trains Language Model, Mass Hysteria Ensues. <https://www.approximatelycorrect.com/2019/02/17/openai-trains-language-model-mass-hysteria-ensues/>
- [29] Duri Long and Brian Magerko. 2020. What is AI Literacy? Competencies and Design Considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–16. <https://doi.org/10.1145/3313831.3376727>
- [30] Melanie Magin and Peter Maurer. 2019. Beat Journalism and Reporting. In *Oxford Research Encyclopedia of Communication*. Oxford University Press. <https://doi.org/10.1093/acrefore/9780190228613.013.905>
- [31] Lauren Kirchner Surya Julia Angwin Mattu, Jeff Larson. 2016. Machine Bias. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- [32] Michael J. Muller and Sarah Kuhn. 1993. Participatory design. *Commun. ACM* 36, 6 (June 1993), 24–28. <https://doi.org/10.1145/153571.255960> 479 citations (Crossref) [2023-02-15] QID: Q114298314.
- [33] Alfred Ng and Maddy Varner. 2021. Nonprofit Websites Are Riddled With Ad Trackers. <https://themarkup.org/blacklight/2021/10/21/nonprofit-websites-are-riddled-with-ad-trackers>
- [34] Dennis Nguyen and Erik Hekman. 2022. The news framing of artificial intelligence: a critical exploration of how media discourses make sense of automation. *AI & SOCIETY* (June 2022). <https://doi.org/10.1007/s00146-022-01511-1>
- [35] Safiya Umoja Noble. 2018. *Algorithms of oppression: how search engines reinforce racism*. New York University Press, New York.
- [36] Inioluwa Deborah Raji, I. Elizabeth Kumar, Aaron Horowitz, and Andrew Selbst. 2022. The Fallacy of AI Functionality. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Seoul Republic of Korea, 959–972. <https://doi.org/10.1145/3531146.3533158>
- [37] Alex Rosenblat, Kate Wikelius, <!--danah -->boyd, Seeta Peea Gangadharan, and Corrine Yu. 2014. Data & Civil Rights: Criminal Justice Primer. *SSRN Electronic Journal* (2014). <https://doi.org/10.2139/ssrn.2542262>
- [38] Ivor Shapiro, Colette Brin, Isabelle Bédard-Brûlé, and Kasia Mychajlowycz. 2013. Verification as a Strategic Ritual. *Journalism Practice* 7, 6 (Dec. 2013), 657–673. <https://doi.org/10.1080/17512786.2013.765638> Publisher: Routledge_eprint: <https://doi.org/10.1080/17512786.2013.765638>.
- [39] Craig Silverman, Ruth Talbot, Jeff Kao, and Anna Klühspies. 2022. How Google’s Ad Business Funds Disinformation Around the World. *ProPublica* (Oct. 2022). <https://www.propublica.org/article/google-alphabet-ads-fund-disinformation-covid-elections>
- [40] Irene Solaiman. 2023. The Gradient of Generative AI Release: Methods and Considerations. <http://arxiv.org/abs/2302.04844>
- [41] Petroc Sumner, Solveiga Vivian-Griffiths, Jacky Boivin, Andrew Williams, Lewis Bott, Rachel Adams, Christos A. Venetis, Leanne Whelan, Bethan Hughes, and Christopher D. Chambers. 2016. Exaggerations and Caveats in Press Releases and Health-Related Science News. *PLOS ONE* 11, 12 (Dec. 2016), e0168217. <https://doi.org/10.1371/journal.pone.0168217>