

Enrolling Citizens:

A Primer on Archetypes of
Democratic Engagement with AI

By Wanheng Hu and Ranjit Singh

Enrolling Citizens: A Primer on Archetypes of Democratic Engagement with AI

by Wanheng Hu and Ranjit Singh

Contents

01 Introduction

03 The Literate Citizen

05 The Responsible Citizen

08 The Contributive Citizen

13 Conclusion

15 Acknowledgments

Introduction

While debates over the adoption of algorithmic systems in a wide range of domains have been ongoing for years,¹ generative AI systems like large language models (such as ChatGPT, Claude, and Gemini) have become an inflection point for these conversations. In response to these developments, lawmakers, regulators, academics, and technologists alike have begun to sift through technical jargon and marketing hype to take on the challenge of safeguarding citizens from potential harms from AI while maximizing their access to its benefits. A common feature of these extensive efforts is, as Michele Gilman has argued,² the importance of including citizens in various stages of AI development and governance. To do so in meaningful ways, however, is impossible without a clear vision of what citizens ideally should do. This primer takes up this imperative and asks: *What current approaches exist to ensure that citizens have meaningful involvement in the development of AI, and how do these approaches envision the role of a “good citizen”?*

In this primer, we highlight three major approaches to involving citizens in AI: AI literacy, AI governance, and participatory AI. Each of these three approaches is premised on the importance of enrolling citizens but envisions different roles for citizens to play. Practitioners’ efforts to advance AI literacy are centered on the idea that everyone needs to have sound knowledge of AI, from “kindergarten to university,” because “jobs will largely be related to AI” in the near future, and “literacy in AI and computer science will become as important as classic literacy (reading/writing).”³ Those who call for AI governance consistently highlight the critical role of public input, particularly in policymaking and regulatory processes,⁴ based on the conviction

-
- 1 Emanuel Moss et al., “Assembling Accountability: Algorithmic Impact Assessment for the Public Interest” (Data & Society Research Institute, June 29, 2021), <https://datasociety.net/library/assembling-accountability-algorithmic-impact-assessment-for-the-public-interest/>.
 - 2 Michele E. Gilman, “Democratizing AI: Principles for Meaningful Public Participation,” Data & Society Primer (New York: Data & Society Research Institute, September 27, 2023), <https://datasociety.net/library/democratizing-ai-principles-for-meaningful-public-participation/>.
 - 3 Martin Kandlhofer et al., “Artificial Intelligence and Computer Science in Education: From Kindergarten to University,” in 2016 *IEEE Frontiers in Education Conference (FIE)*, 2016, 1–9, <https://doi.org/10.1109/FIE.2016.7757570>.
 - 4 Christopher Wilson, “Public Engagement and AI: A Values Analysis of National Strategies,” *Government Information Quarterly* 39, no. 1 (2022): 101652, <https://doi.org/10.1016/j.giq.2021.101652>.

that AI’s potential harms must be responsibly managed by society as a whole. Finally, there are increasing appeals by participatory AI practitioners to involve citizens, from the very beginning, in the design of AI systems. After all, it is ordinary citizens who are often the most impacted by AI-generated decisions.⁵ We argue that these differences ultimately stem from the distinct imaginaries of “good citizens” underlying each approach. By “imaginaries,” we mean a set of collectively held visions and expectations that shape how societies and citizens orient toward emerging technological futures.⁶ These imaginaries are deeply grounded in each approach’s stance on the nature of democracy and the promise of AI.

To better understand the citizen imaginaries embedded in these approaches, we draw on the larger history of initiatives focused on public involvement in the sciences in the UK and US. The question of how to involve citizens in science and technology has been a topic of long debate over the past three decades with investments in three major science and public initiatives: Public Understanding of Science, Public Engagement with Science, and Citizen Science.⁷ Our previous work on these initiatives has shown that they relied on different imaginaries of what makes a good citizen — being **literate**, **responsible**, or **contributive**.⁸ We will focus on this lineage to analyze and elicit archetypes of normative visions of “good citizens” in the context of democratic engagement with AI. As these efforts stand today, we will show that practitioners invested in AI literacy, AI governance, and participatory AI have assumed (respectively) the imaginaries of literate, responsible, or contributive citizens.

In the following sections, we juxtapose each approach to public involvement in AI with a corresponding science and public initiative that shares similar imaginaries. We illustrate how practitioners’ imaginaries of citizens are embedded in specific perceptions of AI’s promise and democratic ideals. In short, *we describe how practitioners imagine the space for citizen action*. However, we would also like to note that their approaches are not mutually exclusive. A citizen can be literate, responsible, and contributive at the same time. One approach is not better than the other, and each has its potential drawbacks. They will inevitably need to work together to achieve a more robust development of AI in the public interest.

5 Fernando Delgado et al., “The Participatory Turn in AI Design: Theoretical Foundations and the Current State of Practice,” in *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, 1–23, EAAMO ’23 (New York, NY: Association for Computing Machinery, 2023), <https://doi.org/10.1145/3617694.3623261>.

6 Sheila Jasanoff and Sang-Hyun Kim, *Dreamscapes of Modernity: Sociotechnical Imaginaries and the Fabrication of Power* (University of Chicago Press, 2015).

7 While there is much internal diversity of approaches to accomplish these initiatives, we will focus on the more traditional, institutionalized approach that represents each initiative. Our discussions of these science and public initiatives will primarily draw on cases from the United Kingdom and the United States.

8 Wanheng Hu, “Imagining the Model Citizen: A Comparison between Public Understanding of Science, Public Engagement in Science, and Citizen Science,” *Public Understanding of Science* online first (2024): 1–16, <https://doi.org/10.1177/09636625241227081>.

The Literate Citizen: Public Understanding of Science and AI Literacy

*“AI potentially becomes one of the important technology skills in the twenty-first century.... AI literacy means having the essential abilities that people need to live, learn, and work in our digital world through AI-driven technologies”*⁹

In 1985, the Royal Society of the UK published a report titled *The Public Understanding of Science*, also known as the *Bodmer Report*. Recognizing that science was then facing dwindling public support and funding, it urged various stakeholders — including schools, scientists, media, industries, and the Parliament — to take action to improve the public’s understanding of science, which was deemed “a major element in promoting national prosperity, in raising the quality of public and private decision-making and in enriching the life of the individual.”¹⁰ The report became iconic for the Public Understanding of Science initiative, which took off in the UK during the 1980s and 1990s. It was characterized by major scholarly and practical efforts to measure the public’s “scientific literacy” along with their attitudes toward science through surveys. Furthermore, significant investments were made to popularize science through means such as science museums and popular science magazines.

For proponents of this initiative, science was an enterprise exclusive to scientific experts and overwhelmingly beneficial for society and individuals. They presumed an electoral democracy that served as the conduit for scientific funding and autonomy through scientifically literate,

9 Davy Tsz Kit Ng et al., “Conceptualizing AI Literacy: An Exploratory Review,” *Computers and Education: Artificial Intelligence* 2 (January 1, 2021): 2, <https://doi.org/10.1016/j.caeai.2021.100041>.

10 Report of a Royal Society, ad hoc Group, and The Royal Society, “The Public Understanding of Science” (London: The Royal Society, 1985), <https://royalsociety.org/topics-policy/publications/1985/public-understanding-science/>.

and thus supportive, citizens who would vote for science. Literacy was a means for citizens to not only better access the benefits of scientific knowledge and technology but also prepare for democratic participation in science-based decision-making.¹¹

Literate citizens should have sufficient knowledge and understanding of science/AI, be familiar with it to identify ordinary ways of using it in their everyday life, and offer their support to ongoing developments in these fields.

More recently, interventions focused on AI literacy have adopted a similar imaginary of a literate citizen; they pay significant attention to surveying the public's competence in and perceptions of AI. Their broader aim is to cultivate AI-literate "future citizens" through education at various levels from kindergarten to K-12 school and college.¹² Practitioners invested in AI literacy imagine that citizens should know, use, and evaluate AI technology to their advantage, especially in a future job market where mastering AI tools will become an essential skill. They tend to limit AI design to the domain expertise of technologists, while citizens are primarily considered users who are afforded the many benefits of using AI tools.

Despite these similarities with the Public Understanding of Science initiatives, there are also differences, which we interpret more as a matter of emphasis: (1) AI literacy advocates emphasize that citizens must go beyond supporting and start adapting to ongoing developments in AI that will ultimately transform society; (2) although they view AI as a generally positive technology, they also acknowledge its potential risks and ethical challenges, which are made more visible compared to the Public Understanding of Science initiatives; and finally, (3) the ideal of democracy in AI literacy interventions is imagined less as a means to protect AI, and more as an egalitarian end in which all citizens are entitled to the benefits of AI and an education in AI skills.

AI literacy builds on a positive recognition of AI's utilities and potential and emphasizes the critical importance for citizens to develop AI skills to remain economically competitive. What it obscures, however, is the agency of the citizens in the development of AI. With a focus on adapting to AI's transformative potential, it tends to pay less attention to the critical interventions citizens can make in shaping how it operates and its possible future(s).

11 Hu, "Imagining the Model Citizen."

12 Kandlhofer et al., "Artificial Intelligence and Computer Science in Education"; Ng et al., "Conceptualizing AI Literacy."

The Responsible Citizen: Public Engagement with Science and AI Governance

“Responsible engagement with AI by administrators requires public engagement and ‘a prior debate with society,’ because ethical challenges ‘permeate all layers of application of this technology.’”¹³

After over a decade of interventions and reflections, the Public Understanding of Science initiative in the UK evolved into Public Engagement with Science, shifting the focus from educating the public to engaging the public in governance. In 2000, the UK House of Lords published a symbolic document that marked this shift — the *Science and Society* report. The report addressed concerns over public unease with rapid scientific developments, including genetically modified organisms, as well as the deep anxiety among scientists about public mistrust in science. Acknowledging the uncertainties and risks in scientific advancements, it called for “changing the culture of policy-making so that it becomes normal to bring science and the public into a dialogue about new developments at an early stage.”¹⁴ The forms of engaging the public in policymaking about science included citizens’ juries, consensus conferences, and standing consultative panels.

Advocates for this shift have recognized lay people’s credible insights about science and the risks brought by scientific advances, envisioning an ideal of responsible citizens who would actively undertake the duty to manage the power and uncertainties of emerging science and technology alongside scientific experts. They presupposed a deliberative democracy through which technoscientific research and its applications are governed responsibly.¹⁵

13 Wilson, “Public Engagement and AI,” 1.

14 House of Lords Select Committee on Science and Technology, “Science and Society” (London: The Stationery Office, 2000).

15 Hu, “Imagining the Model Citizen.”

AI has sparked ongoing debates over the core ethical principles and values that guide its governance;¹⁶ these principles are often used as starting points for regulations.¹⁷ In governing the profound impact of AI, all citizens are potential stakeholders. There is increasing emphasis on public engagement in the policymaking processes for AI governance, which is often proposed as a solution to ensure that AI follows collectively agreed-upon values.¹⁸ For example, the various US federal agencies such as the National Telecommunications and Information Administration,¹⁹ the Federal Trade Commission,²⁰ and the White House Office of Science and Technology Policy²¹ have all requested public comments on diverse AI accountability policies in recent years. The imaginary of citizens in AI governance broadly aligns with the Public Engagement with Science initiative — both frame engagement as a matter of responsibility.

Responsible citizens should respond to the emerging consequences of science/AI in their everyday lives and be deeply committed to and involved in the collective decision-making processes to govern it as stakeholders.

Proponents of public engagement with AI governance put citizens' concerns and perspectives up front in AI's design and deployment, with a precautionary stance toward AI's potential harms. Indeed, the articulation of the rights of citizens and ethical guidelines for AI often mirror its risks in terms of safety, discrimination, privacy violations, opacity, and lack of due process.²² These efforts align with deliberative democracy as the process through which cit-

16 Anna Jobin, Marcello Lenca, and Effy Vayena, "The Global Landscape of AI Ethics Guidelines," *Nature Machine Intelligence* 1, no. 9 (September 2019): 389–99, <https://doi.org/10.1038/s42256-019-0088-2>.

17 Ben Wagner, "Ethics as an Escape from Regulation.: From 'Ethics-Washing' to Ethics-Shopping?," in *BEING PROFILED: COGITAS ERGO SUM: 10 Years of Profiling the European Citizen*, ed. Emre Bayamlioğlu et al. (Amsterdam University Press, 2018), 84–89, <https://doi.org/10.2307/j.ctvhrd092.18>; Thilo Hagendorff, "The Ethics of AI Ethics: An Evaluation of Guidelines," *Minds and Machines* 30, no. 1 (March 1, 2020): 99–120, <https://doi.org/10.1007/s11023-020-09517-8>; Anaïs Ressaygue and Rowena Rodrigues, "AI Ethics Should Not Remain Toothless! A Call to Bring Back the Teeth of Ethics," *Big Data & Society* 7, no. 2 (July 1, 2020): 1–5, <https://doi.org/10.1177/2053951720942541>.

18 Wilson, "Public Engagement and AI."

19 NTIA, "AI Accountability Policy Request for Comment," National Telecommunications and Information Administration, April 11, 2023, <https://www.ntia.gov/issues/artificial-intelligence/stakeholder-engagement/request-for-comments>.

20 FTC, "Commercial Surveillance and Data Security Rulemaking," Federal Trade Commission, August 5, 2022, <https://www.ftc.gov/legal-library/browse/federal-register-notices/commercial-surveillance-data-security-rulemaking>.

21 OSTP, "Request for Information: National Priorities for Artificial Intelligence," The White House Office of Science and Technology Policy (OSTP), May 23, 2023, <https://www.whitehouse.gov/wp-content/uploads/2023/05/OSTP-Request-for-Information-National-Priorities-for-Artificial-Intelligence.pdf>.

22 See, for example, how the AI Bill of Rights is framed in the United States: Office of Science and Technology Policy, "Blueprint for an AI Bill of Rights | OSTP," October 2022, <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>.

izens as stakeholders decide whether and how to engage with AI based on their preferences on issues ranging from informed consent for data collection to using AI for automated decision-making.

The AI governance approach calls for institutional checks and balances grounded in ethical principles and regulatory guidelines that are maintained by active engagement of citizens, civil society, regulatory bodies, and technology companies. It emphasizes the rights and obligations of a citizen in the face of AI development and deployment. Despite the proliferation of such guidelines and principles in recent years, effective measures to facilitate substantive public discussion and engagement in AI governance remain underexplored and underdeveloped.²³ This gap results from an unclear understanding of the knowledge and capacities citizens need to engage with AI and shape its design, as well as ongoing cultural debates over what constitutes the “social values” that should guide AI development.

23 Gilman, “Democratizing AI.”

The Contributive Citizen: Citizen Science and Participatory AI

“Calls for participation in AI design argue that participation can enable AI systems to better reflect the values, preferences, and needs of users and other impacted stakeholders, or more broadly, that participation will empower stakeholders in shaping the design of AI systems.”²⁴

Concurrent with the development of Public Engagement with Science in the UK at the turn of the twentieth century, Citizen Science projects emerged in the US when the Cornell Lab of Ornithology started to recruit the general public to collect data on birds.²⁵ First described by Rick Bonney in his 1996 article “Citizen Science: A Lab Tradition,” Citizen Science is generally defined as “public participation in scientific research, in particular, with members of the public partnering with professional scientists to collectively gather, submit, or analyze large quantities of data.”²⁶ These projects become popular because citizen participation not only improves the feasibility and cost-effectiveness of large-scale scientific research but also provides an avenue for science education. Typically serving as volunteers, citizens may participate in any aspect of the scientific research pipeline, from formulating research questions to designing methodology, collecting and interpreting data, and developing applications.

24 Delgado et al., “The Participatory Turn in AI Design,” 1.

25 Gwen Ottinger, “Reconstructing or Reproducing?: Scientific Authority and Models of Change in Two Traditions of Citizen Science,” in *The Routledge Handbook of the Political Economy of Science* (Routledge, 2017), 351 identifies two distinct traditions of CS: “social movement-based citizen science” (initiated by activists) and “scientific authority-driven citizen science” (organized by scientific institutions). We focus on the latter in this primer.

26 Rick Bonney et al., “Can Citizen Science Enhance Public Understanding of Science?,” *Public Understanding of Science* 25, no. 1 (January 1, 2016): 2–16, <https://doi.org/10.1177/0963662515607406>.

In contrast to the optimistic outlook in *Public Understanding of Science* or the reflective attitude in *Public Engagement with Science*, practitioners involved in Citizen Science projects tend to take a neutral stance on science: it is a tool that can be improved through citizen participation, especially for large-scale issues. These projects operationalize participatory democracy in the execution of science, which provides a useful approach to generate scientific knowledge by valuing lay labor, expertise, and intelligence.²⁷

Contributive citizens should be involved in the process of creating scientific knowledge and building AI and rely on their own standpoint, perspective, and available expertise to shape the collective future(s) of science/AI and articulate its potential implications.

Recruiting ordinary citizens, especially users and stakeholders, to partake in AI design has become increasingly prevalent within the AI community,²⁸ giving rise to participatory AI projects that share Citizen Science’s imaginary of contributive citizens. The forms of citizen participation include commenting on design ideas through questionnaires and interviews, group discussions with the project team, and collaborative prototyping and decision-making.²⁹ By enrolling citizens as active participants with their own values, preferences, and needs, participatory AI projects imagine participation as a way to shift power to citizens and make substantive contributions to building ethical AI systems. For participatory AI practitioners, democracy is a means to calibrate AI to ensure that it serves impacted communities more effectively and ethically.

27 Hu, “Imagining the Model Citizen.”

28 Jennifer Lee et al., “Power and Technology: Who Gets to Make the Decisions?,” *Interactions* 28, no. 1 (December 23, 2020): 38–46, <https://doi.org/10.1145/3442420>; Abeba Birhane et al., “Power to the People? Opportunities and Challenges for Participatory AI,” 2022, <https://doi.org/10.1145/3551624.3555290>; Delgado et al., “The Participatory Turn In AI Design”; Lara Groves et al., “Going Public: The Role of Public Participation Approaches in Commercial AI Labs,” in *2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT ’23: the 2023 ACM Conference on Fairness, Accountability, and Transparency, Chicago IL USA: ACM, 2023)*, 1162–73, <https://doi.org/10.1145/3593013.3594071>.

29 Delgado et al., “The Participatory Turn In AI Design.”

Interventions focused on contributive citizens raise questions such as: (1) what is the nature of the contribution that citizens are making? And (2) do processes employed to work with citizens ensure their meaningful participation? AI safety practitioners have reflected on these questions extensively when analyzing the challenge of aligning AI with human values and interests — commonly referred to as the “**alignment problem**.”³⁰ Participatory methods have been frequently proposed as possible resources to resolve the alignment problem, but they raise challenges around whose values and interests AI should be aligned with, and how practitioners should elicit those parameters from contributive citizens. In this context, participation is typically deemed central to the search for a set of rules to be encoded and refined to guide the behavior of already-built AI models.³¹ It is organized either by inviting a small number of people from “differently situated” groups historically excluded from social power,³² or through crowdsourced input surveyed from a larger set of people to inform the reinforcement learning stage of model development.³³

Participation is, in effect, instrumentalized in two ways: (1) affording a degree of political legitimacy to AI systems that have already been built, and (2) refining technical safety measures by translating values into code. This type of participation is not designed to grant citizens the agency to contest, reject, alter, or assent to the development or deployment of AI. In short, the nature of the contribution that citizens make is directly related to the processes employed to solicit their contribution.³⁴

30 Iason Gabriel, “Artificial Intelligence, Values, and Alignment,” *Minds and Machines* 30, no. 3 (September 1, 2020): 411–37, <https://doi.org/10.1007/s11023-020-09539-2>; Jiaming Ji et al., “AI Alignment: A Comprehensive Survey” (arXiv, February 26, 2024), <https://doi.org/10.48550/arXiv.2310.19852>.

31 Tyna Eloundou and Teddy Lee, “Democratic Inputs to AI Grant Program: Lessons Learned and Implementation Plans,” *OpenAI* (blog), January 16, 2024, <https://openai.com/blog/democratic-inputs-to-ai-grant-program-update>.

32 Stevie Bergman et al., “STELA: A Community-Centred Approach to Norm Elicitation for AI Alignment,” *Scientific Reports* 14, no. 1 (March 19, 2024): 6616, <https://doi.org/10.1038/s41598-024-56648-4>.

33 Anthropic, “Collective Constitutional AI: Aligning a Language Model with Public Input,” *Anthropic Blog* (blog), October 17, 2023, <https://www.anthropic.com/news/collective-constitutional-ai-aligning-a-language-model-with-public-input>.

34 See, Sherry R Arnstein, “A Ladder of Citizen Participation,” *Journal of the American Planning Association* 85, no. 1 (2019): 12, <https://doi.org/10.1080/01944363.2018.1559388>; Gilman, “Democratizing AI” for a more detailed articulation of these dynamics.

The participatory AI approach aims to open up the black box of AI and empower citizens as stakeholders to shape AI design from the beginning. It strives to incorporate the standpoint and situations of impacted communities as an intrinsic part of building AI systems. This approach makes a moral case for valuing the perspective of those who are closest to the dysfunction of such systems, thus positioning them with authoritative and unique expertise on how they work (or do not work) in everyday life. It has moved the engagement of ordinary citizens one step further from principally governing ready-made AI systems to directly shaping AI systems-in-the-making.

While it affords a process to shift power to citizens in determining how systems are designed and operationalized, the participatory AI approach may entrench a static ideal for “ethical technology” — namely, as long as the technology itself integrates stakeholder views and passes their evaluation, it is good to go. Practitioners of participatory AI have recognized this challenge and framed it as “participation washing.”³⁵ This view may ignore the social context of technology application, which is dynamic with possible new ethical implications that emerge over time. Besides being contributive participants in AI design, citizens must also be reflective users, critical monitors, and committed regulators — or, in short, literate and responsible — to ensure their autonomy and dignity in living with AI.

35 Mona Sloane et al., “Participation Is Not a Design Fix for Machine Learning,” in *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, EAAMO '22 (New York, NY, USA: Association for Computing Machinery, 2022), 1–6, <https://doi.org/10.1145/3551624.3555285>.

Discursive Features	AI Literacy	AI Governance	Participatory AI
Which science and public initiative best resonates with this approach?	Public Understanding of Science	Public Engagement with Science	Citizen Science
What is the imaginary of good citizens?	Literate citizen	Responsible citizen	Contributive citizen
Who are the knowledge makers?	Building AI is the domain of engineers and computer scientists.	Variously situated stakeholders — engineers, computer, scientists, regulators, and civil society — need to be involved in regulating AI	Building AI must involve impacted stakeholders, whose input is essential to complement technical know-how.
Who is a good citizen?	Literate citizens should know, use, and evaluate AI technology to their advantage.	Responsible citizens should be aware of their stakeholder status in AI and seek to play a role in its regulation.	Contributive citizens should employ their standpoint and available expertise in building AI.
What is the role of democracy?	Democracy adapts to AI development and promotes citizens' capacity for AI skills inclusively.	Democracy is the process through which citizens collectively decide when and how to govern AI based on their standpoints.	Democracy serves as a means to calibrate AI tools with the values, preferences, and needs of impacted communities.
Potential limits	Practitioners may focus solely on adapting to technological shifts in AI as a proxy for societal change.	Practitioners may tokenize public discussion and engagement in AI governance.	Practitioners may instrumentalize participation to perpetuate exclusion instead.

Conclusion: Enrolling Citizens and AI in the Public Interest

This primer reviews three imaginaries of “good citizens” — literate, responsible, and contributive — that permeate contemporary discussions on AI literacy, AI governance, and participatory AI. These approaches advocate for involving the public in the deployment, regulation, and design of AI, respectively. In unpacking the divergent visions of citizens’ ideal roles that underlie these approaches, the primer clarifies different conceptions of public interest in the context of AI development: (1) as an egalitarian vision of AI’s benefits for AI literacy practitioners; (2) as citizen rights in contending with AI risks for AI governance initiatives; and finally, (3) as community values to calibrate AI in participatory AI. Each conception is intertwined with specific presumptions about the promise of AI and the nature of democracy in the process of enrolling citizens.

The challenge of what citizens should do — and how they should exercise their agency — is central to the imaginary of “good citizens.” These imaginaries are grounded in particular expectations of a commitment to significant work from citizens and visions of desirable futures. Proponents of AI literacy treat this commitment as a matter of building personal competency; in contrast, practitioners in AI governance and participatory AI frame it as ethical responsibility. For the latter, simply being a member of society shaped by AI is enough justification for citizens to have a role in its design and regulation, and literacy is merely a means to that end, not a required qualification. However, while the AI governance approach largely focuses on responding to already developed systems or setting conditions for their operation, the participatory AI approach often focuses on how to empower citizens in the process of building them.

While it is largely impossible to come up with a universal standard for building AI in the public interest, and all approaches will remain local and situated, this primer invites a critical

reflection on the underlying assumptions about technology, democracy, and citizenship that ground how we think about the ethics and role of public(s) in large-scale sociotechnical change. Each approach has its own drawbacks, and the uncritical adoption or dominance of any single “ethical” approach — which might benefit certain communities in their respective contexts — might be “unethical” for other communities situated in different contexts. What is expected from us as citizens and how we should exercise our agency in the face of AI-based systems are not easy questions to answer, but they remain essential considerations for all ongoing efforts. We hope that readers use this primer as a resource to: (1) parse through the various discussions on diverse approaches to involve them as citizens in AI; and (2) identify the expectations from them across these approaches in shaping our collective AI-based future(s). After all, the future is not given; it is made.

Acknowledgments

We express our heartfelt appreciation to Michele Gilman and Stephen Hilgartner for providing feedback on an earlier draft of this primer. The work on this primer also benefited from supportive conversations with Jacob Metcalf, Patrick Davison, and Kiara Childs. Producing this primer would have been impossible without the support of the Communications and Research Management teams — Sona Rai, Chris Redwood, Eryn Loeb, Gloria Mendoza, Alessandra Erawan, and Siera Dissmore.

This material is partially based upon work supported by the National Science Foundation under Grant No. 1926174.

Data & Society is an independent nonprofit research institute that advances new frames for understanding the implications of data-centric and automated technology. We conduct research and build the field of actors to ensure that knowledge guides debate, decision-making, and technical choices.

www.datasociety.net

@datasociety

Layout by Hector Sandoval

June 2024