Exploring Users' Contextual Privacy Perspectives to Support GenAI Governance

Alisa Frik
International Computer Science Institute
Berkeley, CA, USA
University of California, Berkeley
Berkeley, CA, USA
afrik@icsi.berkeley.edu

Mada Alhaidary
King Abdulaziz City for Science and Technology
Riyadh, Saudi Arabia
malhaidary@kacst.gov.sa

Julia Bernd
International Computer Science Institute
University of California, Berkeley
jbernd@icsi.berkeley.edu

Basel Alomair
King Abdulaziz City for Science and Technology
University of Washington
Seattle, WA, USA
alomair@uw.edu

Serge Egelman
International Computer Science Institute
University of California, Berkeley
egelman@cs.berkeley.edu

Abstract

The vast amounts of data used to train Generative AI raise critical questions about data governance and privacy. Effective governance requires an approach that accounts not only for technical capabilities but also broader social factors shaping use of GenAI, including user needs and preferences and social norms related to privacy. Thus it is urgent to understand users' privacy concerns specific to GenAI and how they vary based on context, and to identify how privacy norms are developing around new GenAI capabilities. We plan a study to investigate these questions, and what privacy protections GenAI users desire. We will conduct a survey with GenAI users, including factorial vignettes in which participants will evaluate scenarios that vary in purpose, sphere of use, data sources and formats, recipients, and privacy/security protections. Our findings aim to inform governance efforts by identifying key contextual factors that shape privacy expectations and influence trust in GenAI technologies.

CCS Concepts

• **Social and professional topics** → Government technology policy; • **Human-centered computing** → **Empirical studies in HCI**; • **Security and privacy** → **Privacy protections**; *Social aspects of security and privacy*; Information flow control.

Keywords

Privacy, Artificial Intelligence, Generative AI, Contextual Norms

STAIG '25, April 27, 2025, Yokohama, Japan

ACM Reference Format:
Alisa Frik, Mada Alhaidary, Julia Bernd, Basel Alomair, and Serge Egelman. 2025. Exploring Users' Contextual Privacy Perspectives to Support GenAI Governance. In *Proceedings of the Sociotechnical AI Governance Workshop, at the ACM CHI Conference on Human Factors in Computing Systems (STAIG '25).* ACM, New York, NY, USA, 8 pages.

1 Introduction

As GenAI models are becoming increasingly sophisticated and rely on vast amounts of data, concerns about users' privacy have grown significantly [3, 41, 42]. Researchers and policymakers thus recognize the need for human-centered design principles and transparency in AI development and governance [10, 21, 27, 37, 52].

Generative AI gathers, uses, and emits information in ways that are quite new compared with previous technologies—even previous AI technologies [8]. For example, when a user inputs data to accomplish a task, the output derived can then become inputs or training data for the tool that are now disconnected from, but still containing some of the information from, the original input [9]. This iterative process assumes that later outputs generated by the tool will be sufficiently new that using and sharing them can't be viewed as violating the privacy of those who provided the original inputs—but this does not necessarily jibe with the assumptions of the users who provided that input data. (And, in the worst case, generated content may include information whose source is individually identifiable [41, 42].)

The evolving landscape of generative AI has spurred significant interest in governance and policy implications. Recent work, such as the Draft Report of the Joint California Policy Working Group on Frontier AI Models [11], highlights the urgent need to carefully consider the risks associated with advanced AI models and proposes various recommendations for responsible development and deployment. These policy discussions, while often focused on broader social impacts, underscore the importance of establishing frameworks that can influence user trust and perceptions of the safety and ethical considerations surrounding generative AI technologies.

The new capabilities of generative AI mean that users' normative expectations based on understandings of similar technologies (even existing AI technologies) may not longer be viable [40]. At the same time, new norms may be developing around new understandings and concerns. We plan to capture what people's current views are on this developing technology, with a particular focus on aspects and factors where we believe views may diverge significantly from views on other technologies—or where views may still be based on older understandings, whether those considerations apply or not.

In general, privacy attitudes have been repeatedly shown to be context dependent [e.g. 15, 19, 28, 30, 44, 53]. The Theory of Privacy as Contextual Integrity (CI) [29, 39, 40] provides a detailed framework for analyzing information flows and evaluating how technology is understood according to social norms and expectations within specific contexts—and thus how they may disrupt or align with societal values. Unfortunately, AI development may disregard nuances of context [36] by focusing on abstract technical considerations and general purposes of AI systems rather than specific application domains [7, 35]. Therefore, understanding people's contextual views of privacy is crucial for responsible AI development and AI governance frameworks that account for nuanced user needs.

Prior work has explored public perceptions of AI [5, 12, 20, 23, 31–33, 55], including perceptions about privacy and data governance [6, 17, 22, 24, 25, 43, 48, 49]. A significant contribution from the UK Department for Science, Innovation, and Technology [51] provides valuable insights into public attitudes towards data and artificial intelligence, tracking awareness, understanding, trust, and concerns related to these technologies, thus offering a broad perspective on user perceptions relevant to AI.

Some recent studies have focused on the nuances of user privacy perceptions within the burgeoning field of *generative* AI, but in-depth empirical work on GenAI users' privacy perspectives, and especially variation according to specific contexts, is currently limited. Analysis of public sentiment reveals a generally favorable view of GenAI [34, 47]. At the same time, a recent report from the UK's AI Safety Institute [2] delves into the complexities of designing AI systems to behave like humans. While not directly focused on privacy, it highlights the crucial role of user trust and the potential for misinterpretations of AI capabilities, which can indirectly shape user's comfort and risk assessments when interacting with AI, including *generative* AI.

Such research serves to demonstrate that understanding these nuances will be crucial. In a more privacy-focused example, Zhang et al. [56] conducted a small interview study exploring user interactions and potential privacy risks of LLM-based conversational agents, and found that participants made complex, context-specific risk/benefit tradeoffs based on (often impoverished or incorrect) mental models of genAI. A small survey by Alkamli and Alabduljabbar [4] found that ChatGPT users focused more on concerns about unauthorized access than use of input data by the tool itself. The increasing adoption of AI assistants in various domains, such as software development, underscores the importance of understanding security and privacy practices and concerns among professional as well as casual users [26]. Our research aims to explore and quantify such contextual nuances.

By investigating concerns and normative judgments about data-sharing by GenAI tools, we aim to contribute to an empirical understanding of how people are evaluating this emerging technology, and what are their needs and preference for AI governance and privacy protections. By developing a deeper understanding of users' contextual privacy concerns regarding GenAI, we aim to inform the development of more responsible, privacy-protective GenAI systems, and provide guidance on what governance efforts are most important to users in various contexts.

Our research aims to answer the following questions:

- **RQ1:** What are users' privacy concerns about use of their data by GenAI algorithms?
- **RQ2:** What norms do users hold about sharing of their data by GenAI tools?
- **RQ3:** What contextual factors affect users' privacy concerns and norms about use of their data by GenAI?
- **RQ4:** What privacy protections, controls, and other governance mechanisms do users want for GenAI tools?

2 Planned Methods

To answer our research questions, we plan to run a factorial vignette-based study with GenAI users, in which participants will rate their levels of concern and normative judgments about presented vignettes, followed by a survey that will gather additional insights about their perspectives on GenAI.

2.1 Vignette Design

The first portion of our survey will use factorial vignettes to explore the impact of contextual factors on both *concerns* about GenAI tools' collection and use of users' data, and their *normative judgments* about information flows from those tools to other entities. Generally, vignettes present concise descriptions of situations with systematically varied parameters, used to compare respondents' views about aspects of the presented scenarios [16, 30, 46, 54]. To reduce participant fatigue, we will use a hybrid between/within subjects approach.

First, we will explain what we mean by Generative AI (see Appendix §A.1), and ask participants to imagine that they are users of a hypothetical GenAI tool, ProdigyHub. Then we will present each participant with one randomly assigned scenario with two parameters that will vary *between* subjects, including one of two *Spheres of Use*, and one of five *Purposes* for which the GenAI tool is used.

Then we will present each participant with several vignettes in random order, in which we will vary four parameters *within* subjects. Each vignettes will include additional description with one of four *Data Formats* the GenAI could use, then we will ask participants to respond to two types of question, using 5-point Likert scales. To explore how different privacy and security protections affect participants' concerns about, and thus potentially their willingness to use, GenAI tools, we present three types of *Protections* that could be provided for two *Data Sources*, and ask how concerned would they be about the privacy of their data in such a situation. To explore participants' views on sharing of the data beyond ProdigyHub, we ask about the appropriateness of ProdigyHub sending data from the two Data Sources to five data *Recipients*.

Table 1: Scenario and vignette parameters. Parameters marked with * are fixed per participant.

Parameters	Values
Purpose of the Tool*	*Imagine that you are using ProdigyHub, a Generative AI tool for:* - increasing productivity (e.g. analysis, programming, writing documents), - creating artistic content (e.g. creating music, visuals, fiction, etc.), - finding information (e.g. answers to questions, tips, learning materials), - engaging with entertainment content (e.g. games, jokes), - facilitating social connections (e.g. communication, emotional support),
Sphere of Use*	*for:* - work or school - personal use
Data Format	*To produce content, it uses:* - text data, including text from you - audio data, including audio from you - image data, including images from you - video data, including videos from you
Protection	*ProdigyHub provides the following protections:* - allows you to choose whether the [SOURCE] will be used to train and improve ProdigyHub - allows you to choose whether it will store or delete the [SOURCE] - provides built-in security safeguards (e.g., encryption, optional multi-factor authentication)
Data Source	- data provided by you as the user (including prompts, uploaded files) - outputs that were generated by the system for you
Recipient	*ProdigyHub share[s] [SOURCE] with the following recipients:* - the US government - your local/city government - third-party companies in the US - third-party companies outside the US - human content moderators for ProdigyHub

Scenarios and vignettes will be constructed from frame sentences filled in with parameter values from Table 1, as follows :

> Imagine that you are using ProdigyHub, a Generative AI tool for [PURPOSE*], for [SPHERE*].
>
> To produce content, it uses [DATA FORMAT]. The data is processed and stored on ProdigyHub's servers.
>
> ***[Concern]*** How concerned are you about the privacy of your data in such a situation, if ProdigyHub provides the following protections?
> [Presented as grid: PROTECTION (SOURCE)]
>
> ***[Normative judgment]*** How appropriate would it be for ProdigyHub to share [SOURCE] with the following recipients?
> [Presented as grid: RECIPIENT]

Selection of Vignette Parameters and Values The vignette parameters and values are selected based on factors identified in prior work as affecting privacy concerns and expectations, including some of the information flow parameters identified in CI theory [29, 38, 40]. For example, our within-subject parameters Data Format (text, audio, image, video) and Data Source (provided by the user as input vs. generated by the tool as output) reflect aspects of the *information type* (or *attribute*) parameter in CI theory. Recipient corresponds to the *recipient* in CI, while the *sender* is here fixed as ProdigyHub. The privacy and security protections constitute types of *transmission principles* that we vary in the vignettes, while storage of the data on a cloud server is a fixed transmission principle. In addition, we vary (between subjects) the Purpose of the software and what general Sphere it is being used in (work/school vs. personal) . Contextual purpose and sphere of use are aspects that may be views as evoking different *domains* with different sets of privacy expectations and norms.[1]

Example Vignette In all, participants will be shown four vignettes with three questions about each. For example:

> Imagine that you are using ProdigyHub, a Generative AI tool for **increasing productivity (e.g. analysis, programming, writing documents)**, for **work or school**.
>
> To produce content, it uses **text data, including text from you**. The data is processed and stored on ProdigyHub's servers.
>
> How concerned are you about the privacy of your data in such a situation, if ProdigyHub provides the following protections?
> – Allows you to choose whether **the data provided by you as the user (including prompts, uploaded files)** will be used to train and improve ProdigyHub

[1] See, e.g., Nissenbaum [40] for discussion of the relationship between the concepts of *context* and *domain*.

– [remaining Protections, as a list]
Answers presented in a grid: Very unconcerned, Somewhat unconcerned, Neutral, Somewhat concerned, Very concerned

How appropriate would it be for ProdigyHub to share **the data provided by you as the user (including prompts, uploaded files)** with the following recipients, in this situation?
– The US government
– [remaining Recipients, as a list]
Answers presented in a grid: Completely inappropriate, Somewhat inappropriate, Neutral, Somewhat appropriate, Completely appropriate

How appropriate would it be for ProdigyHub to share **outputs that were generated by the system for you** with the following recipients, in this situation?
– The US government
– [remaining Recipients, as a list]
[same answers as above]

2.2 Post-Vignette and Exit Surveys

In the *post-vignette survey* (Appendix §A.2), we will ask about additional factors that we do not plan to vary systematically as vignette parameters (to keep the complexity of the survey manageable), but which could be important to how participants reason about data flows in GenAI. This includes the appropriateness of scraping publicly available data for model training, use of watermarking to identify AI-generated content (AIGC), views on copyrighting of AIGC, and preferences about privacy protections and controls. In addition, we ask about views on open-source vs. proprietary models; while prior work has explored the risks and opportunities of open-source AI [e.g., 14, 18], public views (if any) on this topic are less understood.

In the *exit survey* (Appendix §A.3), we will ask about experience with GenAI tools and privacy/security violations when using GenAI, and measure participants' attitudes about AI [45] and personal IT innovativeness [1].

2.3 Deployment and Analysis

A Qualtrics survey will be deployed on Prolific, recruiting US-based participants who use generative AI weekly; Prolific will also supply demographic information.

Statistical analysis will examine the impact of contextual factors and participant characteristics on privacy concerns and normative judgments, potentially using Cumulative Link Mixed Models (CLMMs) [13, 50] to accommodate the non-independence of repeated observations within groups or subjects and inclusion of fized and random effects. Open-ended responses will be analyzed using thematic analysis with open coding, and code occurrences compared across conditions, e.g. using Chi-square or Fisher's exact tests.

2.4 Limitations

We focus on parameters that we expect to present interesting interactions or trade-offs, as suggested by prior research or common sense, in particular where we think views on generative AI may differ from views on other technologies. To maintain a reasonable number of parameters and values for quantitative analysis, we will have to exclude some possible dimensions that impact participants' views. Dimensions such as company size and type, open- vs. closed-source models, and details of storage and access control could be explored in future work, as could real-world GenAI user behavior in in-situ experiments or observational studies.

3 Discussion

By presenting this work-in-progress at the STAIG Workshop, we hope to bridge empirical research with governance design. We aim to address key challenges with usable GenAI governance, including understanding users' privacy expectations, concerns, and normative judgments, identifying their governance priorities, and developing recommendations for designing user-centric interventions. Our insights will support the development of governance frameworks that align with user expectations and promote more trustworthy and ethical GenAI systems.

Acknowledgments

This research is supported by funding from King Abdulaziz City for Science and Technology (KACST). Many thanks are due to Katharina Baum for substantial input on the survey design. We are grateful also to Noah Apthorpe, Eugene Bagdasarian, Sebastian Benthall, and anonymous STAIG and PrivaCI reviewers for their suggestions.

References

[1] Ritu Agarwal and Jayesh Prasad. 1998. A conceptual and operational definition of personal innovativeness in the domain of information technology. *Information systems research* 9, 2 (1998), 204–215.
[2] AI Security Institute. 2024. Should AI systems behave like people? https://www.aisi.gov.uk/work/should-ai-systems-behave-like-people Accessed 2025-04-07.
[3] Mousa Al-kfairy, Dheya Mustafa, Nir Kshetri, Mazen Insiew, and Omar Alfandi. 2024. Ethical challenges and solutions of generative AI: An interdisciplinary perspective. In *Informatics*, Vol. 11. MDPI, 58.
[4] Shahad Alkamli and Reham Alabduljabbar. 2024. Understanding privacy concerns in ChatGPT: A data-driven approach with LDA topic modeling. *Heliyon* 10, 20 (2024).
[5] Theo Araujo, Natali Helberger, Sanne Kruikemeier, and Claes H De Vreese. 2020. In AI we trust? Perceptions about automated decision-making by artificial intelligence. *AI & society* 35, 3 (2020), 611–623.
[6] Sumit Asthana, Jane Im, Zhe Chen, and Nikola Banovic. 2024. "I know even if you don't tell me": Understanding Users' Privacy Preferences Regarding AI-based Inferences of Sensitive Information for Personalization. In *Proceedings of the 2024 ACM Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA. To appear.
[7] Abeba Birhane, Pratyusha Kalluri, Dallas Card, William Agnew, Ravit Dotan, and Michelle Bao. 2022. The Values Encoded in Machine Learning Research. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*. ACM. https://doi.org/10.1145/3531146.3533083
[8] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* (2021).
[9] Hannah Brown, Katherine Lee, Fatemehsadat Mireshghallah, Reza Shokri, and Florian Tramèr. 2022. What Does It Mean for a Language Model to Preserve Privacy?. In *2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) *(FAccT '22)*. Association for Computing Machinery, New York, NY, USA, 2280–2292. https://doi.org/10.1145/3531146.3534642
[10] Tara Capel and Margot Brereton. 2023. What is Human-Centered about Human-Centered AI? A Map of the Research Landscape. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 359, 23 pages. https://doi.org/10.1145/3544548.3580959

[11] Jennifer Tour Chayes, Mariano-Florentino Cuéllar, and Fei-Fei Li. 2025. *Draft Report of the Joint California Policy Working Group on AI Frontier Models*. Technical Report. Joint California Policy Working Group on AI Frontier Models. https://www.cafrontieraigov.org/ Accessed 8 April 2025.

[12] Chun-Wei Chiang, Zhuoran Lu, Zhuoyan Li, and Ming Yin. 2023. Are two heads better than one in ai-assisted decision making? comparing the behavior and performance of groups and individuals in human-ai collaborative recidivism risk assessment. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–18.

[13] Rune Haubo B Christensen. 2015. A Tutorial on fitting Cumulative Link Models with the ordinal Package. Retrieved from www.cran.r-project.org/package=ordinal.

[14] Francisco Eiras, Aleksandar Petrov, Bertie Vidgen, Christian Schroeder, Fabio Pizzati, Katherine Elkins, Supratik Mukhopadhyay, Adel Bibi, Aaron Purewal, Csaba Botos, Fabro Steibel, Fazel Keshtkar, Fazl Barez, Genevieve Smith, Gianluca Guadagni, Jon Chun, Jordi Cabot, Joseph Imperial, Juan Arturo Nolazco, Lori Landay, Matthew Jackson, Phillip H. S. Torr, Trevor Darrell, Yong Lee, and Jakob Foerster. 2024. Risks and opportunities of open-source generative ai. *arXiv preprint arXiv:2405.08597* (2024).

[15] Pardis Emami Naeini, Sruti Bhagavatula, Hana Habib, Martin Degeling, Lujo Bauer, Lorrie Cranor, and Norman Sadeh. 2017. Privacy Expectations and Preferences in an IoT World. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*. USENIX Association, Santa Clara, CA, 399–412. https://www.usenix.org/conference/soups2017/technical-sessions/presentation/naeini

[16] Janet Finch. 1987. The vignette technique in survey research. *Sociology* 21, 1 (1987), 105–114.

[17] Faraz Ghalebikesabi, Yang Zhang, Sahil Garg, Deepak Narayanan, Jiaqi Zhai, Iyad Rahwan, Soyeon Park, Diyi Yang, Cynthia Breazeal, and Amir Rahimi. 2024. Operationalizing Contextual Integrity in Privacy-Conscious Assistants. In *International Conference on Learning Representations*. https://arxiv.org/abs/2408.02373

[18] Riccardo Ghioni, Mariarosaria Taddeo, and Luciano Floridi. 2024. Open source intelligence and AI: a systematic review of the GELSI literature. *AI & society* 39, 4 (2024), 1827–1842.

[19] Weijia He, Nathan Reitinger, Atheer Almogbil, Yi-Shyuan Chiang, Timothy J Pierson, and David Kotz. 2024. Contextualizing Interpersonal Data Sharing in Smart Homes. *Proceedings on Privacy Enhancing Technologies* (2024).

[20] Pegah Karimi, Jeba Rezwana, Safat Siddiqui, Mary Lou Maher, and Nasrin Dehbozorgi. 2020. Creative sketching partner: an analysis of human-AI co-creativity. In *Proceedings of the 25th international conference on intelligent user interfaces*. 221–230.

[21] Davinder Kaur, Suleyman Uslu, Kaley J. Rittichier, and Arjan Durresi. 2022. Trustworthy Artificial Intelligence: A Review. *Comput. Surveys* 55, 2 (Jan. 2022), 39:1–39:38. https://doi.org/10.1145/3491209

[22] Patrick Gage Kelley, Celestina Cornejo, Lisa Hayes, Ellie Shuo Jin, Aaron Sedley, Kurt Thomas, Yongwei Yang, and Allison Woodruff. 2023. "There will be less privacy, of course": How and why people in 10 countries expect AI will affect privacy in the future. 579–603. https://www.usenix.org/conference/soups2023/presentation/kelley

[23] Patrick Gage Kelley, Yongwei Yang, Courtney Heldreth, Christopher Moessner, Aaron Sedley, Andreas Kramm, David T. Newman, and Allison Woodruff. 2021. Exciting, Useful, Worrying, Futuristic: Public Perception of Artificial Intelligence in 8 Countries. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (AIES)*. 627–637. Poster presentation.

[24] Sage Kelly, Sherrie-Anne Kaye, Katherine M. White, and Oscar Oviedo-Trespalacios. 2023. Clearing the way for participatory data stewardship in artificial intelligence development: a mixed methods approach. *Ergonomics* (Nov. 2023). https://www.tandfonline.com/doi/abs/10.1080/00140139.2023.2289864 Publisher: Taylor & Francis.

[25] Kimon Kieslich, Birte Keller, and Christopher Starke. 2022. Artificial intelligence ethics by design. Evaluating public perception on the importance of ethical design principles of artificial intelligence. *Big Data & Society* 9, 1 (Jan. 2022), 20539517221092956. https://doi.org/10.1177/20539517221092956 Publisher: SAGE Publications Ltd.

[26] Jan H Klemmer, Stefan Albert Horstmann, Nikhil Patnaik, Cordelia Ludden, Cordell Burton Jr, Carson Powers, Fabio Massacci, Akond Rahman, Daniel Votipka, Heather Richter Lipford, et al. 2024. Using ai assistants in software development: A qualitative study on security practices and concerns. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*. 2726–2740.

[27] Ansgar Koene, Chris Clifton, Yohko Hatada, Helena Webb, and Rashida Richardson. 2019. A governance framework for algorithmic accountability and transparency. (2019).

[28] Hosub Lee and Alfred Kobsa. 2016. Understanding user privacy in Internet of Things environments. In *2016 IEEE 3rd World Forum on Internet of Things (WF-IoT)*. IEEE, 407–412.

[29] Nathan Malkin. 2023. Contextual Integrity, Explained: A More Usable Privacy Definition. *IEEE Security & Privacy* 21, 1 (Jan. 2023), 58–65. https://doi.org/10.1109/MSEC.2022.3201585 Conference Name: IEEE Security & Privacy.

[30] Kirsten Martin and Helen Nissenbaum. 2016. Measuring privacy: an empirical test using context to expose confounding variables. *Columbia Science & Technology Law Review* 18 (2016).

[31] Jon McCormack, Toby Gifford, Patrick Hutchings, Maria Teresa Llano Rodriguez, Matthew Yee-King, and Mark d'Inverno. 2019. In a silent way: Communication between ai and improvising musicians beyond sound. In *Proceedings of the 2019 chi conference on human factors in computing systems*. 1–11.

[32] Jon McCormack, Patrick Hutchings, Toby Gifford, Matthew Yee-King, Maria Teresa Llano, and Mark D'inverno. 2020. Design considerations for real-time collaboration with creative artificial intelligence. *Organised Sound* 25, 1 (2020), 41–52.

[33] Christian Meurisch, Cristina A Mihale-Wilson, Adrian Hawlitschek, Florian Giger, Florian Müller, Oliver Hinz, and Max Mühlhäuser. 2020. Exploring user expectations of proactive AI systems. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 4 (2020), 1–22.

[34] Kunihiro Miyazaki, Taichi Murayama, Takayuki Uchiba, Jisun An, and Haewoon Kwak. 2024. Public perception of generative AI on Twitter: an empirical study based on occupation and usage. *EPJ Data Science* 13, 1 (2024), 1–20. https://epjdatascience.springeropen.com/articles/10.1140/epjds/s13688-023-00445-y

[35] Alexander Martin Mussgnug. 2022. The predictive reframing of machine learning applications: good predictions and bad measurements. *European Journal for Philosophy of Science* 12, 3 (2022), 55.

[36] Alexander Martin Mussgnug. 2025. Technology as uncharted territory: Contextual integrity and the notion of AI as new ethical ground. arXiv:2412.05130 [cs.AI] https://arxiv.org/abs/2412.05130

[37] Jessica Newman. 2021. Explainability won't save AI. https://policycommons.net/artifacts/4143628/explainability-wont-save-ai/4952717/ Policy Commons. Accessed 1 May 2024.

[38] Helen Nissenbaum. 2004. Privacy as contextual integrity. *Wash. L. Rev.* 79 (2004), 119.

[39] Helen Nissenbaum. 2009. Contextual integrity as a general conceptual tool for evaluating technological change. *Information, communication & society* 12, 5 (2009), 511–530.

[40] Helen Nissenbaum. 2019. Contextual Integrity Up and Down the Data Food Chain. *Theoretical Inquiries in Law* 20, 1 (Jan. 2019), 221–256. https://doi.org/10.1515/til-2019-0008 Publisher: De Gruyter.

[41] University of Illinois at Urbana-Champaign. 2024. Privacy Considerations for Generative AI. https://cybersecurity.illinois.edu/privacy-considerations-for-generative-ai/

[42] Norwegian Board of Technology. 2023. Generative AI brings new data privacy challenges. https://teknologiradet.no/en/generative-ai-brings-new-data-privacy-challenges/

[43] Hyanghee Park, Daehwan Ahn, Kartik Hosanagar, and Joonhwan Lee. 2021. Human-AI Interaction in Human Resource Management: Understanding Why Employees Resist Algorithmic Evaluation at Workplaces and How to Mitigate Burdens. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, 1–15. https://doi.org/10.1145/3411764.3445304

[44] Javad Pool, Saeed Akhlaghpour, Farhad Fatehi, and Leonard C. Gray. 2022. Data privacy concerns and use of telehealth in the aged care context: An integrative review and research agenda. *International Journal of Medical Informatics* 160 (April 2022), 104707. https://doi.org/10.1016/j.ijmedinf.2022.104707

[45] Astrid Schepman and Paul Rodway. 2023. The General Attitudes towards Artificial Intelligence Scale (GAAIS): Confirmatory validation and associations with personality, corporate distrust, and general trust. *International Journal of Human–Computer Interaction* 39, 13 (2023), 2724–2741.

[46] Yan Shvartzshnaider, Schrasing Tong, Thomas Wies, Paula Kift, Helen Nissenbaum, Lakshminarayanan Subramanian, and Prateek Mittal. 2016. Learning privacy expectations by crowdsourcing contextual informational norms. In *Fourth AAAI conference on human computation and crowdsourcing*.

[47] Hyo-Jeong So, Hyeji Jang, Minseon Kim, and Jieun Choi. 2023. Exploring public perceptions of generative AI and education: topic modelling of YouTube comments in Korea. *Asia Pacific Journal of Education* (2023), 1–20. https://www.tandfonline.com/doi/full/10.1080/02188791.2023.2294699

[48] Logan Stapleton, Min Hun Lee, Diana Qing, Marya Wright, Alexandra Chouldechova, Ken Holstein, Zhiwei Steven Wu, and Haiyi Zhu. 2022. Imagining new futures beyond predictive systems in child welfare: A qualitative study with impacted stakeholders. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*. Association for Computing Machinery, New York, NY, USA, 1162–1177. https://doi.org/10.1145/3531146.3533177

[49] Yuan Sun, Magdalayna Drivas, Mengqi Liao, and S. Shyam Sundar. 2023. When Recommender Systems Snoop into Social Media, Users Trust them Less for Health Advice. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–14. https://doi.org/10.1145/3544548.3581123

[50] Jack E Taylor, Guillaume A Rousselet, Christoph Scheepers, and Sara C Sereno. 2023. Rating norms should be calculated from cumulative link mixed effects

models. *Behavior research methods* 55, 5 (2023), 2175–2196.
[51] UK Department for Science, Innovation and Technology. 2024. *Public Attitudes to Data and AI Tracker Survey: Wave 4 Report.* Government Statistical Report. GOV.UK. https://www.gov.uk/government/publications/public-attitudes-to-data-and-ai-tracker-survey-wave-4/public-attitudes-to-data-and-ai-tracker-survey-wave-4-report Accessed 2025-04-07.
[52] Inga Ulnicane, William Knight, Tonii Leach, Bernd Carsten Stahl, and Winter-Gladys Wanjiku. 2022. Governance of Artificial Intelligence: Emerging International Trends and Policy Frames. In *The Global Politics of Artificial Intelligence*, Maurizio Tinnirello (Ed.). Taylor & Francis. https://doi.org/10.1201/9780429446726-2
[53] Jessica Vitak, Yuting Liao, Anouk Mols, Daniel Trottier, Michael Zimmer, Priya C Kumar, and Jason Pridmore. 2023. When Do Data Collection and Use Become a Matter of Concern? A Cross- Cultural Comparison of U.S. and Dutch Privacy Attitudes. (2023).
[54] Lisa Wallander. 2009. 25 years of factorial surveys in sociology: A review. *Social Science Research* 38, 3 (2009), 505–520.
[55] Jack West, Bengisu Cagiltay, Shirley Zhang, Jingjie Li, Kassem Fawaz, and Suman Banerjee. 2025. "" Impressively Scary:"" Exploring User Perceptions and Reactions to Unraveling Machine Learning Models in Social Media Applications. *arXiv preprint arXiv:2503.03927* (2025).
[56] Zhiping Zhang, Michelle Jia, Hao-Ping Lee, Bingsheng Yao, Sauvik Das, Ada Lerner, Dakuo Wang, and Tianshi Li. 2024. "It's a Fair Game", or Is It? Examining How Users Navigate Disclosure Risks and Benefits When Using LLM-Based Conversational Agents. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–26.

A Survey Instruments

A.1 Introduction to Main Survey

Welcome to the survey about Generative AI.

To make sure we're on the same page, let us clarify what we mean by that. Generative AI refers to artificial intelligence tools that create text, images, video, audio, etc. by modeling it on very large datasets of human-created content. This modeling is often referred to as "training." Training means teaching generative AI to create human-like text, images, or other content by feeding it lots of examples so it can learn patterns and use them to generate new, similar content.

Generative AI tools can be used, for example, for entertainment, information search, artistic efforts, social connection, professional productivity, etc., as well as deceptive uses.

In this survey, we will ask you to imagine different situations involving Generative AI and ask your opinions about them. Most of those questions will be about your personal opinions. There are no right or wrong answers; we really just want to learn what you think. Although some of the information or scenarios in the survey will be hypothetical, please try to answer as close as you can to what your honest response in real life would have been.

[Scenarios/vignettes and evaluation questions go here.]

A.2 Post-Vignette Survey

Now let's talk about generative AI apps and tools in general, instead of focusing on ProdigyHub specifically.

As a reminder, "generative AI" refers to artificial intelligence tools that create text, images, video, audio, etc. by modeling it on very large datasets of human-created content. Modeling or "training" means teaching the generative AI to create human-like text, images, or other content by feeding it lots of examples so it can learn patterns and use them to generate new, similar content.

[appropriateness of scraping publicly available data] In general, how appropriate do you think it is for a company to train generative AI models using content that was scraped from **publicly available sources**?
1 - Completely inappropriate, 2 - Somewhat inappropriate, 3 - Neutral, 4 - Somewhat appropriate, 5 - Completely appropriate

[factors affecting the use of publicly available data] What factors (if any) affect your opinion about whether it is appropriate to use publicly available content for training generative AI?
[Free-answer]

For the next couple of questions, we'd like you to imagine you are advising a Copyright Office about what you think regulations **should** say about content created by generative AI (regardless of whether you think they already say that or not). It is ok if you are not sure about your answers, we just want to know your opinions.

[copyright] What **output** produced using generative AI should the **users** of the AI tools be able to copyright?
1 - All output, 2 - Some output, 3 - None of the output

[reason for not copyrighting anything] *[If "None of the content" is selected]* Why do you think none of the output produced using generative AI should be copyrightable?
[Free-answer]

[reason for copyrighting everything] *[If "All content" is selected]* Why do you think all output produced using generative AI should be copyrightable?
[Free-answer]

[what copyright should depend on] *[If "Some content" is selected]* What factors do you think should decide whether output produced using generative AI is copyrightable?
[Free-answer]

[competition vs privacy] How should United States laws balance different priorities in regulating generative AI?
Response options are presented as a slider with the following anchors: Ensure data privacy protection – Encourage innovation in generative AI software

[optional comments about competition vs privacy] Do you have any comments about laws balancing different priorities in regulating generative AI?
[Free-answer]

[user controls] Are there any particular privacy settings and user controls that you would like generative AI tools to offer to their users? If so, please specify.
[Free-answer]

[privacy protections - data types] In your opinion, what privacy protections should apply to the following types of data about users that a generative AI tool might use? Choose ALL that apply.
Data types: A - Location, B - Data about children under 13, C - Health information, D - Financial information (e.g. income, investment history), E - Political beliefs, F - Religious beliefs, G - Sexual orientation, H - Gender identity, I - Unique personal identifiers (e.g., social security number, ID or driver's license number, etc.), J - Personally identifiable information (e.g. first and last name, email address)
Response options, presented as a grid against data types: 1 - Allow users to choose whether it will be used to train and improve the generative AI tool, 2 - Allow users to choose whether the generative AI tool will store or delete it, 3 - Provide built-in security safeguards (e.g. encryption, optional multi-factor authentication)

for it, 4 - Other (please specify), 5 - No particular privacy protections should apply to this data type *[exclusive option]*

[other privacy protections - data types (optional)] In your opinion, are there any other types of data about users that should have special privacy protections when used by a generative AI tool? If so, please describe what privacy protections should apply and to what data.
[Free-answer]

[importance of watermarking (digital)] How important is it to you that content generated by AI contain a hidden digital marker showing that it's AI-generated, which can be detected using technical tools? (This is known as a machine-readable watermark.)
1- Very unimportant, 2 - Somewhat unimportant, 3 - Neutral, 4 - Somewhat important, 5 - Very important

[importance of watermarking (human-readable)] How important is it to you that content generated by AI is clearly marked as being AI-generated, so it's immediately obvious to humans?
1 - Very unimportant, 2 - Somewhat unimportant, 3 - Neutral, 4 - Somewhat important, 5 - Very important

[optional comments on importance of watermarking] Do you have any comments about the importance of machine-readable or human-readable markers of generative AI content?
[Free-answer]

[other signals] Besides knowing it was generated by AI, is there any other information about AI-generated content that you would like to know when consuming or interacting with such content? If so, please specify.
[Free-answer]

[awareness of open-source] How aware are you of the differences between open-source and proprietary generative AI models?
1 - Never heard of it, 2 - I've heard something about it but don't know the details, 3 - I know the basics, 4 - I'm well-informed about it

[open-ended open-source differences] *[If did NOT select "Never heard of it"]* Please tell us in a sentence or two what you know about this topic.
[Free-answer]

[preference between open-source and proprietary] In a nutshell, open-source and proprietary generative AI models can differ in:
- whether the training data, details of how the model is built, and/or the weights that influence the output are available to the public or not; and
- when use and modification of the model are allowed, through licensing.

Based on what you know, do you have a preference between using open-source or proprietary generative AI models?
1 - I prefer open-source generative AI models *[randomized]*, 2 - I prefer proprietary generative AI models *[randomized]*, 3 - I don't have a strong preference, 4 - I don't feel informed enough to have a preference

[reason for preferring open-source] *[If selected open-source]* Why do you prefer open-source generative AI models?
[Free-answer]

[reason for preferring proprietary] *[If selected proprietary]* Why do you prefer proprietary generative AI models?
[Free-answer]

A.3 Exit Survey

[experience with GenAI tools] What generative AI tools do you use (or have you used in the past)? Choose ALL that apply.
Options are presented in randomized order. 1 - ChatGPT, 2 - GitHub Copilot, 3 - Google Bard, 4 - Google Gemini, 5 - Bing Chat, 6 - Perplexity, 7 - Cohere Generate, 8 - Claude, 9 - Synthesia, 10 - DALL-E, 11 - Midjourney, 12 - Jasper, 13 - AlphaCode, 14 - Chat Sonic, 15 - Copy.ai, 16 - MetaLlama, 17 - Other (please specify)

[data format as input] What data formats do/did you typically provide as an **input** when using generative AI tools? Choose all that apply.
Options are presented in randomized order. 1 - Text, 2 - Images, 3 - Audio, 4 - Video, 5 - Code, 6 - 3D models, 7 - Other (please specify)

[data format as output] What data formats do/did you typically generate as an **output** when using generative AI tools? Choose all that apply.
Options are presented in randomized order. 1 - Text, 2 - Images, 3 - Audio, 4 - Video, 5 - Code, 6 - 3D models, 7 - Other (please specify)

[General Attitudes towards Artificial Intelligence Scale] *See Schepman and Rodway [45].*
Now we are going to ask about your experience with and views on **technology in general**, not just generative AI.

[Personal Innovativeness in the domain of IT (PIIT)] *See Agarwal and Prasad [1].*

[privacy attitudes] In general, how concerned are you about the privacy of your personal data?
1 - Very unconcerned, 2 - Somewhat unconcerned, 3 - Neutral, 4 - Somewhat concerned, 5 - Very concerned

[experience with prior privacy/security violations] Have you ever experienced an information privacy or security violation/incident?
1 - No, 2 - Yes

[timing of violation] *[If Yes is chosen]* When did you experience an information privacy or security violation/incident?
1 - Less than a month ago, 2 - 1-6 months ago, 3 - 6-12 months ago, 4 - More than a year ago

[description of violation] *[If Yes is chosen]* Please briefly describe the privacy or security violation/incident you experienced.
[Free-answer]

[experience with prior privacy/security violations in GenAI] Have you ever experienced an information privacy or security violation/incident in the context of using generative AI specifically?
1 - No, 2 - Yes

[timing of GenAI violation] *[If Yes is chosen]* When did you experience an information privacy or security violation/incident in the context of using generative AI specifically?
1 - Less than a month ago, 2 - 1-6 months ago, 3 - 6-12 months ago, 4 - More than a year ago

[description of GenAI violation] *[If Yes is chosen]* Please briefly describe the privacy or security violation/incident you experienced in the context of using generative AI specifically.
[Free-answer]

[study comments (optional)] Do you have any comments about the study?
[Free-answer]