Position Paper: Purpose-Specific Consent To Minimize Privacy Harms

RAKIBUL HASAN, Arizona State University, USA

We propose a reformulation of the consenting mechanism. Consent should be purpose-specific rather than data-specific; data collectors must explicitly list potential benefits and harms of consenting to disclose data (akin to the human-subject research consent mechanism); and consent must not be treated as a blanket license to use data as one pleases. Additionally, data must be deleted once the purpose is fulfilled, and consent may expire after a certain duration or when the harm landscape changes.

Moving the focus from "what data can be collected" to "how data can be used" enhances consumer protection from privacy harms. Explicit listing of benefits and harms increases transparency, aids informed decision-making, and allows built-in mechanisms to ensure purpose-specific data use or to prevent specific harms. Finally, the current binary notion of consent, where consenting legitimizes any and all use of the collected data, and encourages the data collectors to employ deceptive practices to obtain consent, should be replaced by what Solove termed *Murky consent*. Under this notion, consent is a continuous negotiation process, and Data collectors can be held accountable for the use of data in ways that harm the data subject beyond what was stated or beyond "reasonable harms," even when consent was acquired.

This consent framework requires more engagement from all stakeholders, and computational agents can be leveraged for automation. Agents can (continuously) negotiate consent after a detailed examination of data use purposes and harm-benefit tradeoffs. When a data subject's agent "decides" to "consent," it should be treated with caution as a limited permit to use data in "reasonable" ways. Data subjects need to have the legal right to hold the collectors accountable when the "reasonable harm" threshold is crossed, even if the agent "consented" (perhaps mistakenly). This setup resembles how corporations have been enjoying protections through disclaimers when their (AI) tools fail.

Additional Key Words and Phrases: Privacy, Privacy Harms, Informed Consent

ACM Reference Format:
Rakibul Hasan. 2018. Position Paper: Purpose-Specific Consent To Minimize Privacy Harms. *ACM/IMS J. Data Sci.* 37, 4, Article 111 (August 2018), 3 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Problem statement

Current consent methods—both "opt in" or "opt out"—have proven to be largely ineffective in protecting people's privacy and agency [17]. It is neither voluntary nor informed: people often have no choice but to consent since otherwise they cannot use a technology they need to participate in society [16], and they are not provided with details on the extent of data collection, use, and possible consequences. Further, meaningfully engaging with consent documents and terms of service is infeasible for most, if not all. We propose a new consent framework with the following properties:

Purpose-specific consent. Current consent mechanisms mainly focus on what data might be collected, with purposes stated vaguely (e.g., marketing or improving services). We propose that consent should primarily focus on the purposes, regardless of what data is being collected. This is because collected data can be used to derive new data not included in the consent form or privacy policy. For example, behavioral data can be used to

Author's Contact Information: Rakibul Hasan, Arizona State University, Tempe, USA, rakibul.hasan@asu.edu.

Manuscript submitted to ACM

profile or even uniquely identify people, which is consistent with a statement like "we do not collect identifying information" but harms data subjects as they are profiled or identified. To prevent such loopholes, either data collectors must specify that the behavioral data will be used for profiling (explicit) or abandon that practice. Note also that (inappropriate) use of data, rather than collection, most often causes harms; the same data can yield different harms depending on how it is used and in what context [3, 18]. For example, physical or cognitive disability status can be predicted from technology use patterns; this information in an education context can help identify students with learning disabilities to offer support, while in an employment context can lead to discrimination [5]. Purpose-specific consent will thus facilitate listing concrete benefits and harms in a given context. It will further facilitate preventive measures to mitigate harms. For example, before disclosing to the collecting party, data can be altered to allow certain types of use but not others. Concrete examples include blurring images and videos that allow activity recognition (e.g., for fall detection from cameras at an elderly care center) but hide identity and emotion [9, 11], and transforming behavioral data to allow legitimate use but prevent private data predictions [10].

Explicit listings of potential benefits and harms. The benefits of data-sharing are often implicit (e.g., to enable certain functionality), while the harms are entirely omitted. Following the consenting mechanism for human-subject research, we propose that both should be listed explicitly to help users contextualize and make an informed decision. Research has repeatedly shown that people lack knowledge about the extent of data collection and use; they make decisions in a vacuum where privacy is an abstract concept. Explicit mention of harms has been shown to improve decision-making; in a recent study (will appear at CHI 2026 [3]), we demonstrate how to reliably measure people's understanding of privacy with concrete and contextualized harm statements. Based on this study, we propose that privacy be conceptualized as a shield against harms that might arise when user data is used in a given context. This approach is superior to measuring other concepts like privacy concerns in helping people grasp and contextualize the consequences of data disclosing [3, 4].

In summary, data subjects must be informed how data will be used and how that use might benefit or harm them. To realize this consent framework, we would need a comprehensive taxonomy of privacy harms; we could start with the many taxonomies that exist in the literature (e.g., [2, 3]) and privacy incident databases [14], and extend them.

Automated consent management. Careful manual evaluation of privacy risk is already infeasible; a detailed list of data use and associated harms will certainly add to the burden. This is a great opportunity to use personalized privacy agents. If data use is specified in a predefined template, the user agent can evaluate the potential benefits and harms (either provided by the collector's agent or self-computed). The agent may communicate with the data subject when a manual evaluation is needed (e.g., in a high-risk situation, or when the use/harms are ambiguous). However determined, the final consent decision must not legitimize any use of the data. Rather, consent is "murky" [17]: data must be used with caution, in particular if the potential negative consequence could be unexpected or higher than what one might reasonably expect. This is also important due to the inherent uncertainty in estimating harms. The user agent makes a "best effort" decision. The user reserves the right to take actions against the data collector in such a case, even if the agent "consented," especially if the resulting harm was foreseeable. As Solove writes, legal guardrails must reflect reasonable expectations of individuals and embody principles of good faith and fair dealing [17]. This ambiguous consent is functionally equivalent to how corporations get away with disclaimers, such as the tools they develop may make mistakes, and they should not be held liable for it. In his recent work, Solove laid the groundwork of how such a consenting mechanism could be supported with a legal framework [17].

To conclude, this position paper combines the harm-centric framing of privacy [3] with the concept of murky consent to provide a consent framework that allows built-in, proactive privacy-enhancement and ensures greater accountability from the data collectors.

2 Author biography

Rakibul Hasan is an assistant professor of computer science at Arizona State University. He earned his PhD in CS from Indiana University, Bloomington, USA, in 2020. Prior to joining ASU, he was a postdoctoral research scientist at CISPA Helmholz Center for Information Security in Germany.

The author conducts interdisciplinary research in the area of data privacy and security. He applies both human-centric and computational approaches to understand privacy problems through experimental and observational studies [1, 6–8, 13, 15]. His solution approaches range from using machine learning [10] and computer vision techniques [9, 11, 12] to designing psychological interventions [1, 8, 13].

References

[1] Amon, M. J., Hasan, R., Hugenberg, K., Bertenthal, B. I., and Kapadia, A. Influencing Photo Sharing Decisions on Social Media: A Case of Paradoxical Findings. In *2020 IEEE Symposium on Security and Privacy (SP)* (May 2020), pp. 1350–1366.

[2] Citron, D. K., and Solove, D. J. Privacy harms. *BUL Rev. 102* (2022), 793.

[3] Gajavalli, S. H., Koizumi, J., and Hasan, R. What's Privacy Good for? Measuring Privacy as a Shield from Harms due to Personal Data Use, June 2025. arXiv:2506.22787 [cs].

[4] Gerber, N., Reinheimer, B., and Volkamer, M. Investigating people's privacy risk perception. *Proc. Priv. Enhancing Technol. 2019*, 3 (2019), 267–288.

[5] Glazko, K., Mohammed, Y., Kosa, B., Potluri, V., and Mankoff, J. Identifying and Improving Disability Bias in GPT-Based Resume Screening. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency* (New York, NY, USA, June 2024), FAccT '24, Association for Computing Machinery, pp. 687–700.

[6] Goenka, S., Prabhu, A., Sakurai, P., Ramachandran, M., and Hasan, R. Who's Watching You Zoom? Investigating Privacy of Third-Party Zoom Apps, Apr. 2025.

[7] Hasan, R. Understanding EdTech's Privacy and Security Issues: Understanding the Perception and Awareness of Education Technologies' Privacy and Security Issues. *Proceedings on Privacy Enhancing Technologies 2023*, 4 (Oct. 2023), 269–286.

[8] Hasan, R., Bertenthal, B. I., Hugenberg, K., and Kapadia, A. Your Photo is so Funny that I don't Mind Violating Your Privacy by Sharing it: Effects of Individual Humor Styles on Online Photo-sharing Behaviors. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (New York, NY, USA, May 2021), CHI '21, Association for Computing Machinery, pp. 1–14.

[9] Hasan, R., Crandall, D., Fritz, M., and Kapadia, A. Automatically Detecting Bystanders in Photos to Reduce Privacy Risks. In *2020 IEEE Symposium on Security and Privacy (SP)* (May 2020), pp. 318–335.

[10] Hasan, R., and Fritz, M. Understanding Utility and Privacy of Demographic Data in Education Technology by Causal Analysis and Adversarial-Censoring. *Proceedings on Privacy Enhancing Technologies 2022*, 2 (Apr. 2022), 245–262.

[11] Hasan, R., Hassan, E., Li, Y., Caine, K., Crandall, D. J., Hoyle, R., and Kapadia, A. Viewer Experience of Obscuring Scene Elements in Photos to Enhance Privacy. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (New York, NY, USA, Apr. 2018), CHI '18, Association for Computing Machinery, pp. 1–13.

[12] Hasan, R., Li, Y., Hassan, E., Caine, K., Crandall, D. J., Hoyle, R., and Kapadia, A. Can Privacy Be Satisfying?: On Improving Viewer Satisfaction for Privacy-Enhanced Photos Using Aesthetic Transforms. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow Scotland Uk, May 2019), ACM, pp. 1–13.

[13] Hasan, R., Weil, R., Siegel, R., and Krombholz, K. A Psychometric Scale to Measure Individuals' Value of Other People's Privacy (VOPP). In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg Germany, Apr. 2023), ACM, pp. 1–14.

[14] Katcher, S., Shapiro, S., Ballard, B., Isaacson, K., McEwen, J., and Slotter, S. Privacy threat modeling for everyone: MITRE PANOPTIC.

[15] Kelso, E., Soneji, A., Rahaman, S., Shoshitaishvili, Y., and Hasan, R. Trust, Because You Can't Verify: Privacy and Security Hurdles in Education Technology Acquisition Practices. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security* (New York, NY, USA, Dec. 2024), CCS '24, Association for Computing Machinery, pp. 1656–1670.

[16] Knowles, B., and Conchie, S. Un-Paradoxing Privacy: Considering Hopeful Trust. *ACM Trans. Comput.-Hum. Interact. 30*, 6 (Sept. 2023), 87:1–87:24.

[17] Solove, D. J. Murky Consent: An Approach to the Fictions of Consent in Privacy Law. *SSRN Electronic Journal* (2023).

[18] Solove, D. J. Data Is What Data Does: Regulating Based on Harm and Risk Instead of Sensitive Data, Jan. 2024.