The Privacy Policy Analysis Canvas

HAUKE SANDHAUS, Cornell Tech, Cornell University, USA
HELEN NISSENBAUM, Cornell Tech, Cornell University, USA

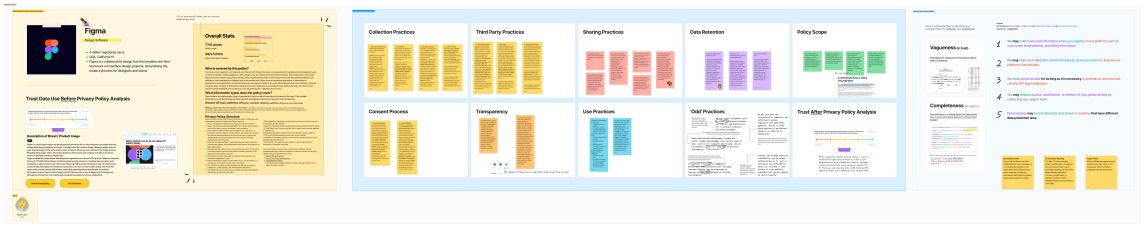


Fig. 1. A student-completed analysis canvas of the privacy policy for the design collaboration tool Figma [Link removed].

Privacy policies are the primary method used by digital service providers to inform users about their privacy practices. The privacy policy is also the mechanism employed by those providers to comply with the regulatory demand of public notice. Research has shown that users typically do not read privacy policies and often struggle to understand them due to their length, vagueness, complicated language, and incompleteness. To encourage deeper engagement with complex privacy practices, we developed a structured privacy policy analysis canvas for educational purposes. This allows students to critically examine and reflect on corporate privacy practices while revealing the fundamental insufficiency of notice-based approaches for privacy protection. We tested this canvas in a privacy class with 59 students who reviewed 31 privacy policies from diverse digital application domains. We report insights from students' use of the canvas and make it available for wide adoption as a customizable design template.

CCS Concepts: • **Security and privacy → Usability in security and privacy**; • **Human-centered computing → HCI design and evaluation methods**; • **Social and professional topics → Information systems education**.

Additional Key Words and Phrases: Privacy, Design, Education, Industry Practices

1 MOTIVATION AND RELATED WORK

Privacy policies serve as the primary mechanism used by digital service providers to inform users about their data practices and to comply with regulatory notice requirements. Yet research consistently shows that these policies fail at their ostensible purpose. Users rarely read privacy policies [10, 13] and, when they do engage with them, struggle to extract accurate information due to length, complex language, ambiguity, and incompleteness [14, 16]. Javed and Sajid [6] recently conducted a systematic review of privacy policy literature, analyzing 202 papers and finding persistent challenges in accessibility, readability, and completeness across diverse sectors.

Perhaps most concerning is what Martin describes as the "tabula rasa" phenomenon—users often project their own privacy expectations onto policies, reading into them what they expect to find rather than what the text actually contains [9]. This disconnect between expectation and reality undermines the entire premise of notice-and-consent frameworks, as users cannot meaningfully consent to practices they misunderstand or fail to identify.

Privacy Policies: Challenges and Limitations. The inherent challenges of privacy policies extend beyond length. Reidenberg et al. [14] identified widespread ambiguity in policy language that further complicates comprehension, while Shvartzshnaider et al. [16] demonstrated how missing contextual details in privacy statements hinder readers'

Authors' addresses: Hauke Sandhaus, hgs52@cornell.edu, Cornell Tech, Cornell University, New York, New York, USA; Helen Nissenbaum, hn288@cornell.edu, Cornell Tech, Cornell University, New York, New York, USA.

ability to evaluate data collection practices. These challenges persist despite years of research identifying the problems, suggesting the need for alternate approaches.

Visual and Canvas-based Approaches. Researchers explored various approaches to improve privacy policy presentation and comprehension. Visual representations [15] and interactive features [7] have shown promise, while standardized patterns [17] and automation tools [20] attempt to address structural issues. Kotut and McCrickard [8] proposed speculative approaches to demystify privacy policies, and other develop games [1] and toolkits [11] for privacy literacy.

Canvas-based tools have emerged as effective frameworks for structuring complex analysis across domains. Thoring et al. [18] mapped the design space of innovation canvases, identifying six key parameters that define their utility. Daou et al. [3] developed the EcoCanvas for circular economy business models, while Fritscher and Pigneur [4] showed how canvases can visualize evolution over time, supporting dynamic analysis rather than static snapshots. Chivukula et al. [2] surveyed ethics-focused design methods, showing that canvas-based tools are common in early design phases, can be applied throughout the entire process, but are uncommonly used post-technology design.

The Gap: Privacy Education and Analysis Tools. Privacy education frameworks provide strategies for enhancing critical privacy literacy. Wong and Mulligan [19] argued for broadening "design" in "Privacy by Design" through HCI perspectives, emphasizing the need for approaches that foreground social values. More recently, Gray et al. [5], found that UX practitioners often lack structured methods to engage with legal and regulatory knowledge in their everyday work practices, highlighting the need for frameworks that bridge design practice with policy understanding.

Despite extensive research on privacy policies and canvas-based methodologies, a significant gap exists in tools specifically designed for structured analysis of privacy policies in educational contexts. Our canvas differs from most design canvases in that it is specifically an analysis canvas applied outside the design process—to audit and scrutinize existing privacy policies. This approach fills a critical gap in both privacy education and analysis methodologies.

2 OUR APPROACH

We created the Privacy Policy Analysis Canvas as a tool that coalesces ethical, legal, and technical perspectives within a graduate-level course on digital privacy. The canvas was implemented after students had gained familiarity with privacy definitions—particularly contextual integrity [12]—but before covering technical tools or detailed legal frameworks. This timing ensured students had sufficient theoretical grounding without specialized knowledge that might narrow their focus. The canvas was shaped by over a decade of analyzing privacy policies in class activities. While core analytical questions remained consistent, new elements were introduced to address emerging issues such as AI inferences and regional regulations. Only the Information Flow Analysis section requires specific background from Reidenberg et al. [14] and Shvartzshnaider et al. [16]. For practical implementation, we used FigJam's digital environment to streamline logistics and ensure documentation of student work.

3 THE CANVAS

We created the Privacy Policy Analysis Canvas as a visual framework to systematically analyze privacy policies. The canvas integrates established privacy theories within three interconnected panels that guide students through a comprehensive evaluation process. Although the primary template was developed in English, we translated the canvas and had it proofread by native speakers in Chinese, Dutch, German, Hindi, Japanese, Korean, Portuguese, and Spanish to support global privacy education initiatives.

3.1 Company Context and Initial Trust Assessment (Left Panel)

The leftmost panel establishes context by capturing baseline information about the company and service. Key elements include company branding, service description, and quantitative metrics such as policy word count compared to industry benchmarks. This contextual grounding helps students situate the policy within its industry environment and recognize how external factors might influence privacy practices.

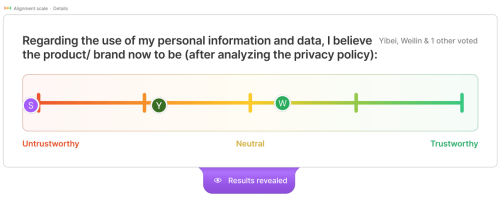


Fig. 2. Trust widget showing disagreement following policy analysis.

A distinctive feature of this panel is the pre-analysis trust assessment, which prompts students to document their initial perceptions before detailed analysis, creating a reference point for later comparison.

Interactive elements include a privacy policy length counter with comparative benchmarks, the trust ranking widget, company identification elements, and an itemized policy structure overview. These components familiarize students with the collaborative environment while establishing baseline knowledge needed for deeper analysis.

3.2 Privacy Practice Analysis Grid (Center Panel)

The central matrix presents 9 key privacy dimensions that operationalize distinct facets of information privacy practices:

- **Collection Practices**: Documents methods of data acquisition (direct input, cookies, inferences) and conditions under which information gathering occurs. Example questions include "How is information collected?" and "Does the policy mention data inferences?"
- **Third Party Practices**: Examines information flows between the service provider and external entities, including specificity of third-party identification. Questions explore whether third parties are explicitly named and what information types flow between entities.
- **Sharing Practices**: Analyzes how user data is distributed beyond the service provider, what types of information are shared, and how the sharing is characterized. Students evaluate whether certain data, such as location or sensitive data, receive special treatment.
- **Data Retention**: Investigates temporal aspects of information storage, deletion, and data portability. Questions address how long information is retained and whether users can move their data to other services.
- **Policy Scope**: Defines boundaries of policy application, including which users and interactions fall under its purview. Students identify whether unregistered users, dependents, or indirect users are covered.
- **Consent Process**: Evaluates mechanisms for obtaining, managing, and revoking user permission, including opt-out provisions. Questions explore how consent is obtained and whether partial opt-outs are possible.
- **Transparency**: Assesses tools and approaches for explaining data practices, including visual aids, explanatory resources, or specific privacy tools.
- **Use Practices**: Examines stated purposes for collected information, including potential AI training applications. Direct questions address whether user data is used for AI training and default settings.
- **Odd Practices**: Highlights unusual or concerning elements that warrant special attention from the students' point of view. This open-ended section often generates the most animated discussion.

The grid format enables comparative analysis across dimensions while visually organizing findings. Students use virtual sticky notes to document their discoveries, creating a collaborative knowledge map. The "Trust After" assessment creates a reflective conclusion that enables students to quantify shifts in perception resulting from their detailed analysis.

3.3 Information Flow Analysis Framework (Right Panel)

In the rightmost panel, students select and annotate five representative information flows for vagueness and completeness. This structured analysis applies two complementary theoretical frameworks:

Following Reidenberg et al. [14]'s approach to identifying ambiguous policy language, students highlight terms and phrases (particularly modal verbs like "may" and "might") that create uncertainty regarding actual practices. This vagueness analysis reveals how policies can technically disclose practices while maintaining strategic ambiguity about what actually occurs. Using Shvartzshnaider et al. [16]'s contextual integrity framework for privacy policy analysis, students color-code flows to indicate whether they adequately address all five parameters: subject, sender, recipient, information type, and transmission principle. This analysis helps identify incomplete flows and "parameter bloating"—where multiple semantically different values appear within a single parameter without clear relationships. Figure 3 shows an example of student annotations, demonstrating how the visual coding system reveals patterns across multiple information flows.

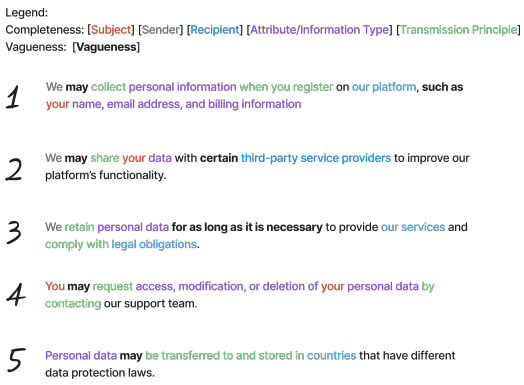


Fig. 3. Student annotation of information flows using contextual integrity framework.

4 CANVAS IMPLEMENTATION AND STUDENT ENGAGEMENT

4.1 Trust Analysis

Our analysis of student pre- and post-analysis trust evaluations (N=38 of individual students who filled in the trust widget before and after the analysis) showed that privacy policy analysis significantly affected students' perceptions of digital services. We converted student responses to a trust question on the canvas widget (see Figure 2) to a Likert scale score ranging from 1 (untrustworthy) to 5 (trustworthy), with 3 being neutral.

Though the aggregate trust change was modest (M_{diff} = -0.27, SD = 1.03), grouping policies by technology maturity revealed that emerging technologies (N=8, founded/widely available post-2015) experienced substantial trust increases (M_{diff} = +0.63, SD = 0.69) after policy analysis, while established technologies (N=30, pre-2015) showed trust decreases (M_{diff} = -0.51, SD = 0.98). This difference was highly significant ($t(36) = 3.96$, $p < 0.001$, $d = 1.25$), suggesting fundamentally different approaches to privacy communication between newer and older companies.

For example, Euki's privacy policy (a period tracking app) produced a trust increase of +1.15, while Slack's policy resulted in a substantial decrease of -2. Instagram (-1.2) and Washington Post (-0.65) similarly diminished trust, while newer services like Replika (+0.5) and Figma (+0.9) enhanced it.

The inter-annotator agreement, within students annotating the same policy, also improved markedly following structured analysis, with the average range between multiple evaluations of the same policy decreasing by 42% (from

0.62 to 0.36). The most dramatic agreement improvements occurred for BetterHelp (range decreased by 2.7) and Canvas (decreased by 1.2), demonstrating the canvas's effectiveness in promoting consistent, evidence-based privacy policy evaluations.

While the privacy policy analysis activity did affect students' trust in companies' data practices, this does not imply whether those practices are truly trust**worthy**. The difference in trust change between established and emerging services could be related not to better practices, but to several possible biases. Many established companies analyzed were from sectors where people typically are unaware of data processing practices (such as McDonalds, New York Times, Starbucks, GoodRX). Additionally, newer policies may carry less historical baggage, often using more default, template-style language without surprising concrete implementations, or simply being shorter and less complex.

4.2 Student Engagement and Feedback

The canvas proved highly effective as a classroom tool, with all students completing the template within the allocated time. All participants received full credit for completing the required elements, with many earning bonus points for additional analysis beyond the minimum requirements.

The activity generated substantial classroom discussion, with students actively comparing findings across different privacy policies. Notably, engagement with the canvas continued beyond the designated class period, with students referencing their discoveries in subsequent discussions and assignments. As one student noted: "I never realized how much information is missing from these policies until I tried to find whether BetterHelp is selling my data."

The "Odd Practices" section proved particularly illuminating, as students identified and shared unexpected data collection and processing mechanisms that might otherwise remain obscured in lengthy text policies. Examples included pixel tracking in email communications, voice analysis for emotion detection, and retention of deleted content for unspecified periods. These discoveries prompted deeper discussions about the boundaries of acceptable data practices and the adequacy of current regulatory frameworks.

The digital canvas format was effective for students working in small groups in person and remotely, as one student specifically highlighted, "I was just curious what other students are doing. I scrolled around to see their pointers and checked how comprehensive their answers are."

5 CONCLUSION AND FUTURE WORK

The Privacy Policy Analysis Canvas template is available on the Figma Community platform [link to be added after acceptance], allowing other educators to adopt and customize it for their privacy courses.

The Analysis Canvas helps students to critically evaluate privacy policies and highlight limitations of notice-based protection. Our in-class trial with 59 students analyzing 31 policies showed it promotes consistent, evidence-based evaluations and reveals systematic patterns in corporate privacy communication. The significant difference in trust impacts between established and emerging companies suggests meaningful engagement with privacy policies.

As AI tools in design canvases evolve, enhancing the canvas with automated analysis features while preserving critical reflection would further support privacy literacy development. Privacy policies are part of our digital lives, and unlikely to be replaced by something better soon. While common practices remain stable, certain data practices, such as AI, change over time and require our privacy analysis canvas to be kept up to date.

ACKNOWLEDGMENTS

We thank all students of the class CS 5436 / INFO 5303 (FA24): Privacy in the Digital Age for their participation. We thank Madiha Zahrah Choksi and Benjamin Laufer for their earlier contributions to the privacy policy analysis activity.

REFERENCES

[1] Jenny Berkholz, Aniqa Rahman, and Gunnar Stevens. 2025. Playing with privacy: Exploring the social construction of privacy norms through a card game. *Proceedings of the ACM on human-computer interaction* 9, 1 (10 Jan. 2025), 1–23. doi:10.1145/3701202

[2] Shruthi Sai Chivukula, Colin Gray, Ziqing Li, Anne C Pivonka, and Jingning Chen. 2024. Surveying a landscape of ethics-focused design methods. *ACM Journal on Responsible Computing* (17 July 2024). doi:10.1145/3678988

[3] Alain Daou, Camille Mallat, Ghina Chammas, Nicola Cerantola, Sammy Kayed, and Najat Aoun Saliba. 2020. The Ecocanvas as a business model canvas for a circular economy. *Journal of cleaner production* 258, 120938 (June 2020), 120938. doi:10.1016/j.jclepro.2020.120938

[4] Boris Fritscher and Yves Pigneur. 2014. Visualizing business model evolution with the business model canvas: Concept and tool. In *2014 IEEE 16th Conference on Business Informatics*, Vol. 1. IEEE, 151–158. doi:10.1109/cbi.2014.9

[5] Colin M Gray, Ritika Gairola, Nayah Boucaud, Maliha Hashmi, Shruthi Sai Chivukula, Ambika R Menon, and Ja-Nae Duane. 2024. Legal Trouble?: UX Practitioners' Engagement with Law and Regulation. *Designing Interactive Systems Conference (DIS Companion '24), July 01â•fi05, 2024, IT University of Copenhagen, Denmark* 1, 1 (2024). doi:10.1145/3656156.3663698

[6] Yousra Javed and Ayesha Sajid. 2025. A systematic review of privacy policy literature. *ACM computing surveys* 57, 2 (28 Feb. 2025), 1–43. doi:10.1145/3698393

[7] Rhianne Jones, Neelima Sailaja, and Lianne Kerlin. 2017. Probing the design space of usable privacy policies: A qualitative exploration of a reimagined privacy policy. In *HCI 2017*. BCS Learning & Development. doi:10.14236/ewic/hci2017.50

[8] Lindah Kotut and D Scott McCrickard. 2022. The TL;DR Charter: Speculatively demystifying privacy policy documents and terms agreements. *Proceedings of the ACM on human-computer interaction* 6, GROUP (14 Jan. 2022), 1–14. doi:10.1145/3492842

[9] Kirsten Martin. 2015. Privacy notices as tabula Rasa: An empirical investigation into how complying with a privacy notice is related to meeting privacy expectations online. *Journal of public policy & marketing* 34, 2 (1 Sept. 2015), 210–227. doi:10.1509/jppm.14.139

[10] Aleecia M McDonald and L Cranor. 2009. The cost of reading privacy policies. (2009). https://lorrie.cranor.org/pubs/readingPolicyCost-authorDraft.pdf

[11] Vikram Mehta, Daniel Gooch, Arosha Bandara, Blaine A Price, and Bashar Nuseibeh. 2023. A card-based ideation toolkit to generate designs for tangible privacy management tools. In *Proceedings of the Seventeenth International Conference on Tangible, Embedded, and Embodied Interaction*. ACM, New York, NY, USA. doi:10.1145/3569009.3572903

[12] Helen Nissenbaum. 2009. *Privacy in Context*. Stanford University Press. doi:10.1515/9780804772891

[13] Jonathan A Obar and Anne Oeldorf-Hirsch. 2020. The biggest lie on the Internet: ignoring the privacy policies and terms of service policies of social networking services. *Information, communication and society* 23, 1 (2 Jan. 2020), 128–147. doi:10.1080/1369118x.2018.1486870

[14] Joel R Reidenberg, Jaspreet Bhatia, Travis Breaux, and Thomas B Norton. 2016. Automated comparisons of ambiguity in privacy policies and the impact of regulation. *SSRN Electronic Journal* (2016). doi:10.2139/ssrn.2715164

[15] Daniel Reinhardt, Johannes Borchard, and Jörn Hurtienne. 2021. Visual Interactive Privacy Policy: The Better Choice?. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21, Article 66)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3411764.3445465

[16] Yan Shvartzshnaider, Noah Apthorpe, Nick Feamster, and Helen Nissenbaum. 2019. Going against the (appropriate) flow: A contextual integrity approach to privacy policy analysis. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing* 7 (28 Oct. 2019), 162–170. doi:10.1609/hcomp.v7i1.5266

[17] Johanneke Siljee. 2015. Privacy transparency patterns. In *Proceedings of the 20th European Conference on Pattern Languages of Programs*. ACM, New York, NY, USA. doi:10.1145/2855321.2855374

[18] Katja Thoring, Roland M Mueller, and Petra Badke-Schaub. 2019. Exploring the design space of innovation canvases. *Conference Proceedings of the Academy for Design Innovation Management* 2, 1 (30 Nov. 2019). doi:10.33114/adim.2019.06.243

[19] Richmond Y Wong and Deirdre K Mulligan. 2019. Bringing Design to the Privacy Table: Broadening "Design" in "Privacy by Design" Through the Lens of HCI. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–17. doi:10.1145/3290605.3300492

[20] Razieh Nokhbeh Zaeem, Rachel L German, and K Suzanne Barber. 2018. PrivacyCheck: Automatic summarization of privacy policies using data mining. *ACM transactions on Internet technology* 18, 4 (30 Nov. 2018), 1–18. doi:10.1145/3127519