UsersFirst: A User-Centric Privacy Threat Modeling Framework for Notice and Choice

Tian Wang Xinran Alexandra Li Miguel Rivera-Lanas Yash Maurya Hana Habib
Lorrie Faith Cranor Norman Sadeh
Carnegie Mellon University

1 Background

In today's data economy, a large number of products, services, and business processes are powered by data [8]. The rapid adoption of AI is further fueling our dependence on the collection and use of personal data across increasingly complex and diverse dataflows [3]. Concurrently, new data privacy regulations impose increasingly stringent requirements on the collection and use of data. This includes more specific obligations about the disclosure of data practices and the need to provide data subjects with more comprehensive sets of choices or controls [10]. Penalties for not complying with these requirements have also become significantly steeper.

In this context, organizations are looking for guidance to help them organize the way in which they identify and mitigate potential privacy risks. A particularly important objective is to apply systematic and consistent approaches to handle potential risks and methodically document the processes involved in the identification and mitigation of threats. Privacy threat modeling frameworks, including LINDDUN [11], the NIST Privacy Framework [4], or more recently MITRE's PANOPTIC framework [9], have been introduced to provide a structured methodology to help organizations to systematically analyze and address potential privacy risks throughout the lifecycle of a system. While these frameworks provide general guidance about privacy threats that may exist in a system and are gaining initial adoption, their guidance when it comes to addressing the lack of or ineffectiveness of privacy notices and choices is limited to high-level considerations (e.g., "unawareness as data subject," "lack of rectification/erasure"). In contrast, regulations such as GDPR and CCPA have been mandating increasingly comprehensive sets of notifications and choices, as well as increasingly emphasizing the need for these choices to be user-centric. With the passage of CPRA, there is additional regulatory scrutiny surrounding the use of dark patterns in privacy notice and choice interfaces [2].

This poster presents work on the development of UsersFirst, a user-centric framework intended to supplement coarser threat modeling frameworks to help an organization's privacy team identify and remedy areas where their privacy notices[1] and choices[2] fall short. UsersFirst aims to reflect emerging trends in privacy regulations where perfunctory approaches to notices and choices are no longer sufficient, instead, they are expected to be noticeable, usable, unambiguous, and devoid of deceptive designs. The framework provides organizations with a systematic approach to identify and mitigate potential threats, while affording them the autonomy to determine their own acceptable risk thresholds and objectives.

2 Our Proposed Framework: UsersFirst

Our proposed UsersFirst framework is intended to support organizations and help them systematically address risks associated with the lack or ineffectiveness of notices and choices. This framework is designed to be flexible, allowing organizations to identify and implement notice and choice requirements that are appropriate for their specific context. The UsersFirst framework is informed by a growing body of previous research in usable privacy, including models of the design spaces associated with privacy notice and choice [1, 6, 7], which look at dimensions of notices and choices that include timing, channel, modality, and control with the objective of

USENIX Symposium on Usable Privacy and Security (SOUPS) 2024.
August 11–13, 2024, Philadelphia, PA, United States.

[1] A privacy notice is a presentation of terms, sometimes but not exclusively in the form of text in a privacy policy or terms of use agreement, intended to inform users about the data practices of a system and what rights, if any, a user of the system might be able to exercise.

[2] A privacy choice is a mechanism by which a user is allowed to control one or more practices associated with the collection or processing of data about them.

supporting the design of user-friendly privacy notices and choice controls. The product of this work is a taxonomy of usability threats that are commonly found in privacy notices and choices. The selected threats were based on a systematic analysis of the academic literature, regulatory documents, existing frameworks, and industry white papers, as well as knowledge from privacy professionals and researchers.

3 How UsersFirst Works

The UsersFirst framework revolves around an iterative approach with two phases, where organizations identify and design notice and choice interfaces (the Design Phase), then evaluate the resulting designs against the UsersFirst taxonomy of user-oriented threats (the Analysis Phase). Results of this analysis are used to inform design revisions intended and mitigate these threats.

The design phase itself consists of several steps. The first step is for the organization to determine those privacy notices and choices it will be present to data subjects. Identifications of notices and choices are based on three sets of considerations: 1) applicable laws and regulations that the organization needs to comply with, 2) corporate policies, which may mandate going beyond minimum regulatory requirements, and 3) out of a broader contextual integrity perspective [5]. After determining those specific notices and choices that will be presented to users, the next step consists of the identification of "touchpoints" through which notices and choices will be presented to data subjects (generally referred to as "users" in the framework). Each touchpoint involves an interaction with the user and may involve presenting the user with one or more notices and/or choices. As part of this step, each notice and choice requirement identified may be mapped onto one or more touchpoints. This is in general a many-to-many mapping, where a touchpoint might be used to present more than one notice or more than one choice to a user, and where each notice and each choice may also be accessible through multiple touchpoints.

Once specific interfaces have been designed for each touchpoint, in the Analysis Phase, these notices and choices are evaluated according to an extensive taxonomy of usability threats with four categories: discovery and use, comprehension, appropriate choices, and nudging (Figure 1). These privacy threats were identified through a comprehensive literature review of usable privacy studies and informed by insights from privacy professionals and researchers. This taxonomy of threats is designed to capture elements of notices and choices that are ineffective, confusing, misleading and more generally inadequate. Mitigation of possible threats typically involves revisiting decisions made in earlier steps such as revisiting the specific set of notices and choices selected for presentation to users, adding or modifying some touchpoints, modifying mappings between notices/choices and touchpoints or modifying the design of specific notices or choices. The UsersFirst framework is intended to be iterative in nature. It may involve iteratively revisiting one or more steps in the process until all threats are deemed to have been satisfactorily resolved.

4 HealthWay: A Case Study

As part of our work, we are conducting a series of evaluations (e.g., case studies, comparison with other frameworks) intended to validate and refine our framework. In this poster, we introduce a use case study on HealthWay, a hypothetical multi-channel retail pharmacy inspired by some real-life retail chains, to examine how the UsersFirst framework can be applied in practical scenarios.

Starting from the Design Phase of the UsersFirst framework, our first step was to determine a set of notice and choice requirements. We assume HealthWay needs to comply with CCPA and other laws (the regulatory perspective), address its corporate policies related to data collection, and consider contextual integrity or user expectations. Next we considered ways HealthWay can implement effective touchpoints to deliver these requirements to customers. We assume that HealthWay implements a chatbot as a personalized shopping assistant that answers customers' questions, provides customer service and support, and recommends personalized products and deals based on customers' preferences. The chatbot will also be able to present users with their privacy choices upon request. By communicating with the chatbot, users can choose to share their chat history and other personal data for personalized services (promotions, recommendations, deals), and they can also choose to opt out the collection of their chat history and stop receiving any personalized deals in the future.

In the Analysis Phase, the next step is to elicit privacy notice and choice threats through the proposed UsersFirst threat taxonomy described in the framework. Examples of possible threats include "Less Privacy Protective Defaults" and "Manipulative Statements." By considering these threats, an analyst might observe that the chatbot will automatically collect the customer's chat history for promotions by default, and warn customers about missing out on personalized deals if they choose to terminate the sharing of their chat history. To mitigate these threats, the chatbot needs to have more privacy protective defaults, provide clear, accurate feedback on current privacy settings, and objectively notify customers about the effects of the privacy choices they made.

The above example represents part of the comprehensive analysis presented in the full case study, which includes more detailed explanations and a more comprehensive set of touchpoints. Overall, the case study offers insights into the framework's effectiveness, adaptability, and potential areas for improvement. By focusing on specific instances, it allows for an in-depth understanding of the framework's strengths and limitations, facilitating a comprehensive evaluation that can inform future refinements and implementations.

Acknowledgments

The researchers gratefully acknowledge the support of the Digital Transformation and Innovation Center at Carnegie Mellon University sponsored by PwC. We thank Seo Young Ko, Geetika Gopi, Ziping Song, and Yanhao Qiu for their initial work on the UsersFirst framework.

References

[1] Yuanyuan Feng, Yaxing Yao, and Norman Sadeh. A design space for privacy choices: Towards meaningful privacy control in the internet of things. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–16, 2021.

[2] Johanna Gunawan, Dave Choffnes, Woodrow Hartzog, and Christo Wilson. Towards an understanding of dark pattern privacy harms. In *Position Paper at the CHI 2021 Workshop: What Can CHI Do About Dark Patterns*, 2021.

[3] Matthew Humerick. Taking ai personally: how the eu must learn to balance the interests of personal data privacy & artificial intelligence. *Santa Clara High Tech. LJ*, 34:393, 2017.

[4] Naomi Lefkovitz and Kaitlin Boeckl. Nist privacy framework: An overview. 2020.

[5] Helen Nissenbaum. Privacy as contextual integrity. *Wash. L. Rev.*, 79:119, 2004.

[6] Florian Schaub, Rebecca Balebako, and Lorrie Faith Cranor. Designing effective privacy notices and controls. *IEEE Internet Computing*, 21(3):70–77, 2017.

[7] Florian Schaub, Rebecca Balebako, Adam L Durity, and Lorrie Faith Cranor. A design space for effective privacy notices. In *Eleventh symposium on usable privacy and security (SOUPS 2015)*, pages 1–17, 2015.

[8] Andrea Sestino, Adham Kahlawi, and Andrea De Mauro. Decoding the data economy: a literature review of its impact on business, society and digital transformation. *European Journal of Innovation Management*, 2023.

[9] S. Shapiro. Mitre panoptic™ v1.0 tutorial. In *2nd Workshop on Privacy Threat Modeling (WPTM))*, 2023.

[10] Iris Van Ooijen and Helena U Vrabec. Does the gdpr enhance consumers' control over personal data? an analysis from a behavioural perspective. *Journal of consumer policy*, 42:91–107, 2019.

[11] Kim Wuyts, Laurens Sion, and Wouter Joosen. Linddun go: A lightweight approach to privacy threat modeling. In *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, pages 302–309. IEEE, 2020.

Notice and Choice Threats Taxonomy

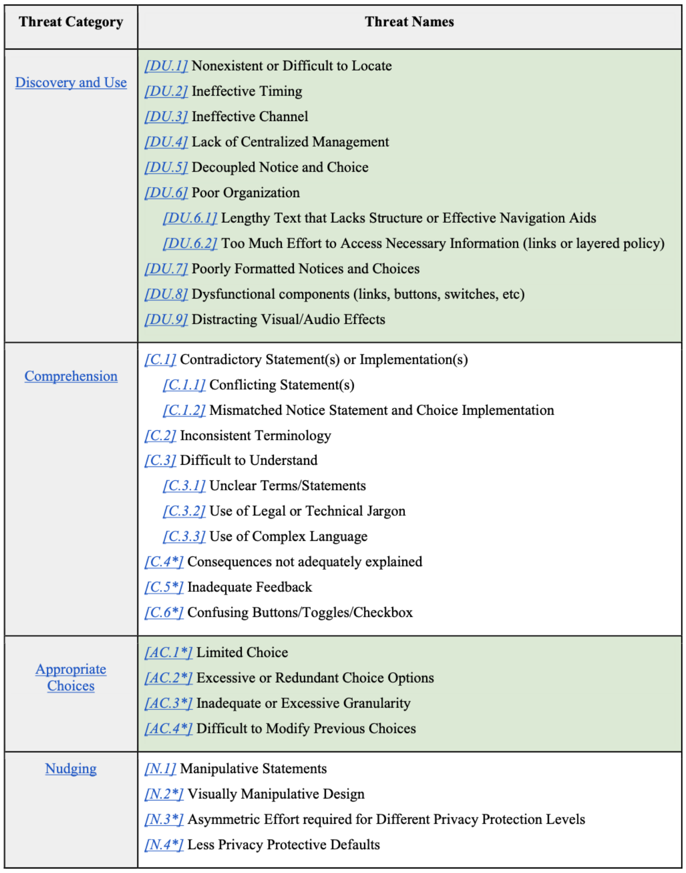

Threat Category	Threat Names
Discovery and Use	[DU.1] Nonexistent or Difficult to Locate [DU.2] Ineffective Timing [DU.3] Ineffective Channel [DU.4] Lack of Centralized Management [DU.5] Decoupled Notice and Choice [DU.6] Poor Organization [DU.6.1] Lengthy Text that Lacks Structure or Effective Navigation Aids [DU.6.2] Too Much Effort to Access Necessary Information (links or layered policy) [DU.7] Poorly Formatted Notices and Choices [DU.8] Dysfunctional components (links, buttons, switches, etc) [DU.9] Distracting Visual/Audio Effects
Comprehension	[C.1] Contradictory Statement(s) or Implementation(s) [C.1.1] Conflicting Statement(s) [C.1.2] Mismatched Notice Statement and Choice Implementation [C.2] Inconsistent Terminology [C.3] Difficult to Understand [C.3.1] Unclear Terms/Statements [C.3.2] Use of Legal or Technical Jargon [C.3.3] Use of Complex Language [C.4*] Consequences not adequately explained [C.5*] Inadequate Feedback [C.6*] Confusing Buttons/Toggles/Checkbox
Appropriate Choices	[AC.1*] Limited Choice [AC.2*] Excessive or Redundant Choice Options [AC.3*] Inadequate or Excessive Granularity [AC.4*] Difficult to Modify Previous Choices
Nudging	[N.1] Manipulative Statements [N.2*] Visually Manipulative Design [N.3*] Asymmetric Effort required for Different Privacy Protection Levels [N.4*] Less Privacy Protective Defaults