Data Privacy and the Elusive Goal of Empowering the User

Vijay Kothari
Dept. of Computer Science
Dartmouth College
vijayk@cs.dartmouth.edu

Sean Smith
Dept. of Computer Science
Dartmouth College
sws@cs.dartmouth.edu

Jim Blythe
Information Sciences Institute
University of Southern California
blythe@isi.edu

Ross Koppel
Dept. of Sociology
University of Pennsylvania
rkoppel@sas.upenn.edu

CHI 2018 Workshop on Networked Privacy, "Moving Beyond a 'One -Size Fits All' Approach: Exploring Individual Differences in Privacy," April 21, 2018, Montréal, Québec, Canada.

Abstract
A major aim of privacy research is to empower end users by giving them more control over their data privacy. However, precisely defining empowerment, let alone achieving it, can be challenging. In this paper, we examine key subgoals and challenges to achieving the grand objective of user empowerment, study the interconnectivity of these subgoals and challenges, and list open questions for future privacy research.

Author Keywords
privacy; ethics; risk perceptions; metrics; user interface

ACM Classification Keywords
H.5.2 [Information interfaces and presentation]: User Interfaces; H.5.m [Information interfaces and presentation (e.g., HCI)]: Miscellaneous; K.4.1 [Computers and Society]: Public Policy Issues - Privacy

Introduction
An overarching pursuit behind much of the privacy research is to empower end users by giving them better control over how their private information is used and shared by the services with which they interact, including, but not limited to, online social networks, search engines, and online marketplaces. However, achieving this grand objective involves so much more than simply providing privacy knobs. For some

examples: accurate but esoteric or confusing language accompanying privacy options may drive users toward settings they don't want; frightening warnings that appear over stringent privacy settings may nudge the malleable user toward undesirable lax settings; the service may provide useful privacy settings, but users may ignore them altogether because they feel they have already lost control over their privacy.

Empowering the user requires understanding the challenges and subgoals underlying this grand challenge, as well as how they are all intertwined. Doing so will enable us to develop better privacy solutions and more informative metrics to evaluate those solutions. To this end, in this paper, we provide a preliminary disaggregation of the grand challenge of empowering users into component subgoals and challenges, and we explore the interconnectivity of these components. We also provide further discussion of related challenges.

Key Challenges

We enumerate key challenges for achieving the goal of empowering users. This list is not, nor is it intended to be, exhaustive.

Users Do Not Correctly Conceptualize Ramifications of Privacy Decisions: Numerous researchers have shown that users hold misperceptions and employ flawed mental models when it comes to privacy, e.g., [9]. For one example, many users subscribe to the nothing-to-hide argument despite its logical flaws [15]. For another, Turow et al. [16] revealed, from a survey they conducted, the following staggering statistic, amongst others:

> 65% [of respondents] do not know that the statement "When a website has a privacy policy, it means the site will not share my information with other websites and companies without my permission" is false.

Indeed, recent work by Draper [5], as well as by others [7, 10, 14, 16, 17], call into question the view that most users are privacy pragmatists who make well-informed privacy decisions. Draper argues that many users feel their privacy is gone and so they resign to giving up control over their data privacy.

There are also secondary effects associated with configuring privacy settings that users may struggle to conceptualize. For example, configuring a privacy setting on an online social network fundamentally changes the interactions the user has with the service and may, in turn, influence how the user behaves, what information the user consumes, how the user thinks, whether the user becomes addicted to the service, the happiness derived from using the service, and so forth.

For another more nuanced point, the user's determination of a privacy option may, itself, leak information. For example, Lewis et al. note that both options not to share or share information correlate with other demographic information [11]. (This point provides further justification for an opt-in approach to privacy over an opt-out one, especially in situations where most users stick with defaults.)

Which Privacy Settings Truly Empower the User?: There is no consensus on what information to use as a source for the user's true privacy preferences. Users have beliefs, desires, and intentions— and when it comes to privacy, they needn't align. Indeed, much research has been devoted to understanding privacy paradoxes stemming from these disconnects, e.g., [10, 3, 4]. As mentioned earlier, the characterization of most users as privacy pragmatists has been debunked. Given this observation, what information should be elicited from the user to empower them? And how do we

use it?

There's a Cost to Configuring Privacy Settings: Actual and perceived time and effort to configure privacy settings may influence whether the user begins configuring them and whether they finish. For example, the user may be dissuaded from using an interface that appears illogical, complex, or hard to navigate.

Time of Configure, Time of Use: The user may configure their privacy settings once, when they begin using a service. However, over time, people may join or leave the service, leaving their privacy choices outdated. The available privacy options may also change over time. Gaw and Felten argue that the user may choose to reuse a password for an account (i.e., select a weak password before that account is associated with sensitive information), and, by the time that account has accrued information, "they're locked into their reused password" [6]. A similar phenomenon may be true with privacy settings. Namely, the user may choose privacy settings before sensitive information is tied to their account. By the time sensitive information is tied to their account, the user may no longer think about privacy. Moreover, in instances where the user does contemplate reconfiguring privacy settings, there's a possibility that the data in question may be perceived to already be lost and, therefore, not worth protecting.

On the other hand, some users may have already invested significant time or effort in selecting a piece of software, downloading it, and installing it before they configure their privacy settings, compelling them to continue using the service even if it does not meet their privacy needs. That is, they may fall victim to the sunk cost fallacy [2]. How then, do we ensure user privacy settings best serve the user at all times?

Users Have Limited information: Not only are users ill-equipped to make rational decisions pertaining to their privacy, but they lack the requisite information to make informed decisions, often because that information is simply not available. It is not always clear how services safeguard user data, nor the intricacies of how that data is used in practice. Privacy policies exist, but may be exorbitantly time-consuming to read and difficult to digest [13, 12]. Moreover, they are often vague and usually subject to change. A pessimist might argue that in practice many existing interfaces and privacy policies ensure users remain uninformed while presenting the veneer of informing the user, thereby persuading their users and other actors that user data is in good hands. Another concern is that primary services or third-party services may violate privacy policies, terms of service, or users' privacy expectations. This may even be compounded by a delay in reporting violations. Collectively, these and other factors support the argument that most users do not— and, at least in the current privacy landscape, cannot— have a concrete understanding of how their data is used.

Nudging: The user may be biased from the interface in a number of ways: Warnings may appear on certain privacy options to nudge the user toward select choices. Descriptions of privacy options may bias the user. Default selections may bias the user. The available options themselves may bias the user. For example, if there are k lax settings and l restrictive settings with $k >> l$, the user may be more inclined to choose a lax setting than they would under the condition $k \leq l$. Indeed, as Acquisti et al. [1] argue in their discussion on the ethicality of nudging, every design choice can be considered a nudge, it's hard to design interfaces that "optimize the benefits of nudges for all users," and evaluating the impact of nudges is hard. How then, do we ethically nudge?

Discussion

We present open questions and discussion topics that emerge from our enumeration of challenges.

User Interests vs. Corporate Interests: Corporations design the privacy policies and the interfaces for selecting privacy options. If we adopt the view that every design choice is a nudge as advocated by Acquisti et al. [1], then these corporations do quite a bit of nudging, and it is reasonable to question whether this nudging serves the user. Moreover, the corporations are the most knowledgeable about how user data is used or misused; and they often directly or indirectly influence the information channels upon which users rely to make privacy decisions. Yet the responsibility of reading and digesting privacy policies and selecting appropriate privacy options is often attributed to the user. Draper [5], in relaying an approach mentioned by Mayer-Schönberger and Cukier[19], states that "one possible approach to enhancing privacy protections would involve shifting responsibility from the data subject to the data user." An alternative approach may involve placing more responsibility on corporations to ensure end users are well-informed pertaining to their data privacy, while maintaining user control over their privacy selections. Indeed, there are other approaches to addressing this disconnect between user interests and corporat interests, e.g., other regulatory approaches, having independent watchdog organizations vet corporations, public shaming.

There's also the challenge of ensuring quality, cutting edge research from academia continues to be effective in the real world where they may be adopted by entities that are more concerned with meeting their own goals than doing what best serves the user.

Ethical Design: Vanderberghe and Slegers [18] advocate a core design philosophy espoused in the Ethical Design Manifesto by Ind.ie [8]: technology should not be designed for the other. Forward-thinking solutions that involve user segmentation may pose additional ethical concerns. In addition to the concerns regarding the ethicality of designing for others, it's difficult to validate segmentation-based solutions. On the other hand, segmentation-based privacy solutions do show promise in addressing some of the challenges we've mentioned, e.g., [20]. How do we weigh the ethical concerns of segmentation-based privacy solutions against the potential value provided by such solutions?

Some Privacy Goals are at Odds: Many of the subgoals we mentioned are useful to study in isolation, but it's worth noting that they may also conflict with other subgoals. For one example, ensuring that a privacy interface is deep enough to express users' "true" privacy preferences may lead a user to ignore the interface altogether if they perceive the effort of configuration as being too high. For another example, while in many circumstances an opt-in approach may empower the user, in select circumstances, giving up data privacy may be considered a social good and provide justification for an opt-out approach.

Metrics: Better metrics and methodologies to evaluate and compare proposed privacy solutions must be designed, adopted, and continually refined to meet real-world privacy needs. We believe these metrics should take into consideration the complexities of empowering the user that are mentioned in this paper as well as other papers.

Conclusion

In this paper, we explored the general aim of empowering users by providing them better control over their data privacy. We outlined key subgoals and challenges in this pursuit, touched on how they are intertwined, and outlined important research pursuits moving forward.

REFERENCES

1. Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, Yang Wang, and Shomir Wilson. 2017. Nudges for privacy and security: Understanding and assisting usersâĂŹ choices online. *ACM Computing Surveys (CSUR)* 50, 3 (2017), 44.
2. Hal R Arkes and Catherine Blumer. 1985. The psychology of sunk cost. *Organizational behavior and human decision processes* 35, 1 (1985), 124–140.
3. Naveen Farag Awad and Mayuram S Krishnan. 2006. The personalization privacy paradox: an empirical evaluation of information transparency and the willingness to be profiled online for personalization. *MIS quarterly* (2006), 13–28.
4. Susan B Barnes. 2006. A privacy paradox: Social networking in the United States. *First Monday* 11, 9 (2006).
5. Nora A Draper. 2017. From privacy pragmatist to privacy resigned: challenging narratives of rational choice in digital privacy debates. *Policy & Internet* 9, 2 (2017), 232–251.
6. Shirley Gaw and Edward W Felten. 2006. Password management strategies for online accounts. In *Proceedings of the second symposium on Usable privacy and security*. ACM, 44–55.
7. Eszter Hargittai and others. 2010. Facebook privacy settings: Who cares? *First Monday* 15, 8 (2010).
8. In.die. 2015. Ethical Design Manifesto. `https://ind.ie/ethical-design/`, (2015).
9. Ruogu Kang, Laura Dabbish, Nathaniel Fruchter, and Sara Kiesler. 2015. my data just goes everywhere:âĂİ user mental models of the internet and implications for privacy and security. In *Symposium on Usable Privacy and Security (SOUPS)*. USENIX Association Berkeley, CA, 39–52.
10. Spyros Kokolakis. 2017. Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & Security* 64 (2017), 122–134.
11. Kevin Lewis, Jason Kaufman, and Nicholas Christakis. 2008. The taste for privacy: An analysis of college student privacy settings in an online social network. *Journal of Computer-Mediated Communication* 14, 1 (2008), 79–100.
12. Aleecia M McDonald and Lorrie Faith Cranor. 2008. The cost of reading privacy policies. *ISJLP* 4 (2008), 543.
13. Jonathan A Obar and Anne Oeldorf-Hirsch. 2016. The biggest lie on the internet: Ignoring the privacy policies and terms of service policies of social networking services. (2016).
14. Kenneth Olmstead and Aaron Smith. 2017. Americans and cybersecurity. *Pew Research Center* (2017), 1–5.
15. Daniel J Solove. 2007. I've got nothing to hide and other misunderstandings of privacy. *San Diego L. Rev.* 44 (2007), 745.
16. Joseph Turow, Michael Hennessy, and Nora A Draper. 2015. The tradeoff fallacy: How marketers are misrepresenting American consumers and opening them up to exploitation. (2015).

17. Jennifer Urban and Chris Hoofnagle. 2014. The privacy pragmatic as privacy vulnerable. In *Symposium on Usable Privacy and Security (SOUPS)*. USENIX Association Berkeley, CA.

18. Bert Vandenberghe and Karin Slegers. 2016. Designing for Others, and the Trap of HCI Methods & Practices. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 512–524.

19. Mayer-Schönberger Viktor and Cukier Kenneth. 2013. Big data: A revolution that will transform how we live, work, and think. *Houghton Mifflin Harcourt* (2013).

20. Pamela J Wisniewski, Bart P Knijnenburg, and Heather Richter Lipford. 2017. Making privacy personal: Profiling social network users to inform privacy education and nudging. *International Journal of Human-Computer Studies* 98 (2017), 95–108.