

3.2 Privacy as Expectations, Decision Making and Practice

Scholars in Human Computer Interaction (HCI) and Access Control (we restrict ourselves to research on user-centric access control at the intersection of HCI and User Modeling—there is greater body of work on OSN access control models that focuses on the formal properties of these rather than on user needs) have taken up the challenge of tackling social privacy in OSNs. In this research, the privacy problems users' faces are investigated through qualitative and quantitative studies. The users are consumers of OSN services whose concerns may show variety depending on demographics like gender, age, education, urbanity and technical skills. The results of these studies help to explore design mechanisms and principles that enable users to establish appropriate privacy practices.

In HCI research it is assumed that technical solutions that equate privacy with concealment are too rigid to accommodate the users' practices. Information concealment does not necessarily imply privacy, and disclosure is not inevitably associated with (undesirable) accessibility. Daily practices, such as making explicit that you do not want to be disturbed, illustrate that a disclosure can be used to negotiate privacy boundaries. Further, studies show that users develop their own strategies to maintain their privacy and manage their identity while benefiting from participating in OSNs. For example, some users create multiple accounts at a given service. These may be pseudonymous, obscured or transparent accounts. While these 'obscured' profiles may not conceal the users' profile effectively, users find that the protections they offer are sufficient for their daily needs.

Researchers perform user studies that are contextualized and are conducted iteratively. These studies observe how, given an OSN design, users negotiate and reconfigure their social boundaries. Hence, this research avoids focusing on one-off disclosure and concealment decisions without contextualization. Further, the researchers explore whether and how practices change when privacy design principles are applied by iterating user studies with enhanced prototypes. In addition to studying privacy practices, researchers have focused on the role of decision making in social privacy problems. A number of studies in behavioral economics point to failures in individual or social decision-making as the source of many social privacy problems in OSNs. These show that users systematically fail to correctly estimate privacy risks [1] and to match their privacy preferences to their actual behaviors [5]. These failures motivate the exploration of design mechanisms that aid users in making better privacy decisions—especially when they lack complete information, are subject to cognitive and behavioral biases, and are uncertain with respect to the outcomes of their decisions.

Specifically, contextual feedback mechanisms may aid users in making better disclosure decisions. These feedback mechanisms, also called privacy nudges, can help users to become aware of and overcome their cognitive biases. For example, if the users are experiencing harms or regrets with respect to emotional outbursts, they can be sent alerts before posting messages that use emotional language [21]. Such

feedback can be used to trigger reflection and self-censorship, instead of the desire for immediate gratification through disclosure. Users may also negotiate their boundaries by "skillfully" using their OSN privacy settings. However, there are major problems associated with privacy settings. A variety of decision-making problems re-appear when users utilize their OSN privacy settings. Users may be subject to social influence or may fail to predict future preferences. They may have a tendency to compromise in the present in order to get immediate gratification. In other cases, users may give greater prevalence to not-so-close friends (weak ties) and may experience difficulty in estimating trust towards these. All in all, given the multitude of decisions, users may simply experience cognitive overload.

To counter some of these problems, researchers have proposed corrective feedback mechanisms as well as a number of interface improvements to current privacy settings. In addition to decreasing the cognitive load of the user, these solutions aspire to make the potential effects of an action more visible in context. In one solution, users are able to view their effective permissions as they change their privacy settings [13].

Another major problem is that users encounter great difficulties to effectively configure their privacy settings. In order to successfully use their settings, users need to first locate them and understand their semantics. Further, the settings need to be at a meaningful granularity to express the users' disclosure preferences.

The response from the access control community, informed by research in user modeling, has been to develop privacy settings that are more expressive and closer to the users' mental models of OSNs. A number of the proposed access control models leverage users' 'attributes'. These attributes, e.g. relationships, roles, or other contextual information, can be used to aid users in configuring their settings to express their actual preferences. Other models propose using artificial intelligence to assist users in keeping their privacy settings up to date [18].

User studies have been successfully leveraged to re think social privacy and its evolution with OSN design. These studies have made the importance of the user factor visible to other privacy researchers, to policy makers and to regulators. Even further, some of their results have already found an audience in commercial OSNs. This illustrates that, in contrast to solutions developed to address surveillance concerns, the emphasis on OSN 'consumers' aligns well with the incentives of companies to design systems that are comfortable for their customers.

4. Discussion

We showed in the previous sections that the two approaches frame and address the OSN privacy problem very differently. Each community of researcher's abstracts away some of the complexity associated with the OSN privacy problem through their framing, in the same way as we abstracted away institutional privacy in this article. Given the complexity of addressing privacy in OSNs, this is a necessary step to breakdown the problem into more graspable parts. The issue

is, however, that the surveillance and social privacy approached say actually have come to systematically abstract each other away. As a result, even though they speak about the same phenomenon, i.e., privacy in OSNs, they end up treating the surveillance and social privacy problems as independent of each other.

We argue that given the entanglement between surveillance and social privacy in OSNs, privacy research needs a more holistic approach that benefits from the knowledge base of the two perspectives. A first step for developing such a holistic approach lies in juxtaposing their differences. In doing so, we can understand the ways in which they are complementary well as identify where the gaps lie. Specifically, we find that the approaches tend to answer the following questions differently:

- Who has the authority to articulate what constitutes a privacy problem in OSNs?
- How is the privacy problem in OSNs articulated?
- Which user activities and information in OSNs are within the scope of the privacy problem?
- What research methods should be used to approach privacy problems in OSNs?
- What types of tools or design principles can be used to mitigate the issues associated with OSN privacy problems and why?
- How should these tools and design principles be evaluated?

In the following, we tackle some of the questions mentioned above: namely, the who, the how and the scope. We believe that a more thorough analysis of the different answers will pave the way to a possible integration of the two perspectives and to a more comprehensive approach to addressing user's privacy problems in OSNs.

4.1 Who has the authority to articulate the privacy problem?

While in PETs research "security experts" articulate what constitutes a privacy problem, in HCI, it is the "average OSN user" who does so.

With PETs, the emphasis is on the privacy risks that may arise when adversaries exploit technical vulnerabilities: this puts the "security experts" in the driver's seat. This has positive and negative consequences. On the positive side, expertise in analyzing systems from an adversarial view point is key to understanding the subversive uses of information systems; be it their repurposing for surveillance or the circumvention thereof. On the negative side, by formulating the problem as a technical one, the researchers bracket out the need to consider social and political analyses of surveillance practices. This introduces the risk of over-relying on techno-centric assumptions about how surveillance functions and what maybe the most appropriate strategies to counter it. Moreover, the focus on improving security guarantees and on designing tools that behave predictably in every context inevitably plays down the importance of the social context and the users' talents in subverting technical boundaries in unexpected ways. It also deemphasizes the importance of considering the difficulties users may face in integrating these tools into their everyday life.

In social privacy research, individual users are the actors articulating privacy concerns. This research makes evident that technologies are open-ended: their use in practice may differ from the use cases devised by the designers. However, the focus on contextual practices inevitably results in small intensive studies. Surveys have a greater reach, but they have in common with small studies a focus on the perceptions and concerns of individual users. Hence, such studies do not provide much insight into collective privacy practices of established OSN communities, e.g., specific interest groups.

Moreover, while user studies explore the correlations between demographics and privacy concerns, they rarely consider surveillance practices and how they may shape the privacy problem for specific populations. For example, under privileged groups that are subject to greater surveillance may have other (social) privacy problems. This may require examining other demographic criteria in user studies, e.g., immigrants or lower income communities. Further, most of the studies are done with users in North America and Europe; few consider the needs of users elsewhere. For example, it is unclear if a study focused on activists or users in contexts with limited ICT access would surface the same privacy concerns. Conducting such studies remains however extremely challenging: researchers do not always have easy access to these groups of users, and the design of the studies would need to take into account their specific socio-political context.

Finally, as OSNs become integrated into everyday life, users tend to take them as a given, and are likely to report on how they make do with the given design. This further constrains what can be discovered through user studies. For example, a study that asks users to critically engage in the values and ideologies embedded into a particular OSN design, or to imagine radical design alternatives, may overwhelm participants and fail to provide results. In order to address this limitation, we may have to introduce other methods, e.g., workshops in which experts explore designs together with users.

4.2 How is the privacy problem articulated?

'Who' has the authority to articulate the privacy problems inevitably determines how these problems are defined. In the two approaches, it determines whether privacy problems are mapped to technology-induced risks or to the harms perceived by users.

Users intuitively recognize causality when their OSN activities lead to concrete harms in interpersonal relationships. However, they cannot be reasonably expected to articulate concerns with respect to the more "abstract" privacy risks, derived from surveillance that often motivate the need for PETs. These may be risks that affect parts of the OSN population. For example, users deemed as not fitting societal 'norms' may be discriminated or repressed as a result of inferences made from their data. Other abstract risks affect society as a whole rather than individual users. For example, the greater intrusion in the private life of citizens that is enabled by OSN surveillance may result in an erosion of basic rights and freedoms.

Often, even the experts struggle to articulate how the abstract risks associated with OSN surveillance may materialize into actual harms. In practice, it may even be impossible to establish the link between personal data disclosures and their ultimate consequences. This is because the decision making processes of the organizations holding the data are complex and opaque. These processes involve multiple entities and sources of data, as well as sophisticated data processing algorithms.

For example, studies have shown that friendship relations in OSNs can be analyzed to infer sensitive personal preferences, such as sexuality and political orientation, even if the users have not disclosed this information. The inferred preferences may or may not be correct, and we do not know if OSN providers employ such inference mechanisms. If they do employ them, we do not know which decisions are made based on them, or who else has access to the inferences.

Understanding how decisions are made on the basis of which data, however, would require access to algorithms and management decisions that are typically not available for scrutiny by either users or independent experts. The opacity of OSN providers poses an enormous challenge to both research in PETs and in social privacy.

PETs designers can only guess which data is collected and how it could be exploited to the disadvantage of the user. Without information on actual OSN surveillance practices, it is hard to establish the capabilities and objectives of the adversaries, or the accuracy of the risk analysis. In such cases, the researchers prefer to study ‘worst case scenarios’. While this is technically sensible, it may not reflect the most pressing practical concerns posed by surveillance. In social privacy, one challenge lies in determining the appropriate mechanisms through which OSN users can be exposed to complex and opaque privacy issues. This may empower users to find their positions on matters that do not seem to directly impact them. How to conduct studies that surface the user perspective on abstract risks and harms remains however an open question.

4.3 What is in the scope of the privacy problem?

On the other hand, in social privacy self-censorship is explored as a strategy. For example, some solutions aim to avoid regrettable disclosures by cautioning users when they are about to disclose sensitive content. Privacy practices are hence associated with silence as much as with expressing oneself. This raises the question of who has the authority to decide on the norms that underlie privacy nudges, e.g., who decides what constitutes sensitive content?

Finally, users may benefit from being able to question norms asserted through design. There are situations in which OSN providers make certain actions invisible in order to avoid conflict, e.g., in Facebook users are not informed when their friends delete their relationship. These norms set by OSN providers enable certain interpersonal negotiations but disable others. This begs a greater question that is missing in social privacy research and that is only partially addressed with PETs: what can we offer users to enhance their ability to

The first difference between the approaches lies in the way they treat explicit and implicit data disclosures. In the social privacy perspective, the privacy problems are associated with boundary negotiation and decision making. Both aspects are concerned with volitional actions, i.e., intended disclosures and interactions. Consequently, user studies are more likely to raise concerns with respect to explicitly shared data (e.g., posts, pictures) than with respect to implicitly generated data (e.g., behavioral data). In contrast, PETs research is mainly concerned with guaranteeing concealment of information to unauthorized parties. Here, any data, explicit or implicit, that can be exploited to learn something about the users is of concern.

Shedding light on users’ perception of implicit data may benefit both approaches. Studies showing how far users are aware of implicitly generated data may help better understand their privacy practices. The results of such studies may also provide indicators for how PETs can be more effectively deployed. If users are not aware of implicit data, it may be desirable to explore designs that make implicit data more visible to users.

The content of the data shared by the user with trust entities is out of the scope of PETs. Researchers only consider the disclosure of data with respect to the ‘adversary’, and PETs offer no protection to data disclosures made at the discretion of the user, e.g., to ‘trusted friends’. Further, the actual semantics of the data shared by the user are also out of the scope. Social privacy studies however reveal that the privacy concerns of users include the semantics of intentional data disclosures towards ‘trusted friends’. This points to possibly irreconcilable difference between the two approaches concerning what ‘privacy’ actually entails.

The two approaches have a fundamentally different take on censorship. In PETs research, privacy technologies are often instrumental for free speech and eluding censorship. They can enhance the user’s ability to express themselves shielded from pressure by peers and authorities. PETs can conceal who is speaking and what is being said in a content-agnostic manner.

say what they want – including expressions that contest design, as well as social norms?

5. Conclusion and Future Work

By juxtaposing their differences, we were able to identify how the surveillance and social privacy researchers ask complementary questions. We also made some first attempts at identifying questions we may want to ask in a world where the entanglement of the two privacy problems is the point of departure. We leave as a topic of future research a more thorough comparative analysis of all three approaches. We believe that such reflection may help us better address the privacy problems we experience as OSN users, regardless of whether we do so as activists or consumers.

References

- [1] Alessandro Acquisti and Jens Grossklags. Privacy and rationality in individual decision making. *IEEE Security and Privacy*, 3(1):26 – 33, January/February 2005.
- [2] J. Anderson, C. Diaz, J. Bonneau, and F. Stajano. Privacy-Enabling Social Networking over Untrusted Networks. In *ACM Workshop on Online Social Networks (WOSN)*, pages 1–6. ACM, 2009.
- [3] Miriam Aouragh and Anne Alexander. The Egyptian Experience: Sense and Nonsense of the Internet Revolutions. *International Journal of Communications*, 5:1344 – 1358, 2011.
- [4] F. Beato, M. Kohlweiss, and K. Wouters. Scramble! Your social network data. In *Privacy Enhancing Technologies Symposium, PETS2011*, volume 6794 of LNCS, pages 211–225. Springer, 2011.
- [5] B. Berendt, O. Günther, and S. Spiekermann. Privacy in E-Commerce: Stated Preferences vs. Actual Behavior. *Communications of the ACM*, 48(4):101–106, 2005.
- [6] E. De Cristofaro, C. Soriente, G. Tsudik, and A. Williams. Hummingbird: Privacy at the time of twitter. In *IEEE Symposium on Security and Privacy*, pages 285–299. IEEE Computer Society, 2012.
- [7] A. Cuttillo, R. Molva, and T. Strufe. Safebook: A privacy-preserving online social network leveraging on real-life trust. *Communications Magazine*, 47(12):94–101, 2009.
- [8] R. Dingledine, N. Mathewson, and P. Syverson. Tor: The second generation onion router. In *USENIX Security Symposium*, pages 303–320, 2004.
- [9] FTC. Ftc charges deceptive privacy practices in Google’s rollout of its buzz social network. Online, 03 2011.
- [10] Glenn Greenwald. Hillary Clinton and internet freedom. *Salon (Online)*, 9. December 2011.
- [11] James Grimmelmann. Saving facebook. *Iowa Law Review*, 94:1137–1206, 2009.
- [12] Kevin D. Haggerty and Richard V. Ericson. The Surveillant Assemblage. *British Journal of Sociology*, 51(4):605 – 622, 2000.
- [13] Heather Richter Lipford, Jason Watson, Michael Whitney, Katherine Froiland, and Robert W. Reeder. Visual vs. Compact: A Comparison of Privacy Policy Interfaces. In *Proceedings of the 28th international conference on Human factors in computing systems, CHI ’10*, pages 1111–1114, New York, NY, USA, 2010. ACM.
- [14] Evgeny Morozov. Facebook and Twitter are just places revolutionaries go. *The Guardian*, 11. March 2011.
- [15] Deirdre K. Mulligan and Jennifer King. Bridging the gap between privacy and design. *Journal of Constitutional Law*, 14(4):989 – 1034, 2012.
- [16] Leysia Palen and Paul Dourish. Unpacking”privacy” for a networked world. In *CHI ’03*, pages 129 – 136, 2003.
- [17] Kate Raynes-Goldie. Privacy in the Age of Facebook: Discourse, Architecture, Consequences. PhD thesis, Curtin University, 2012.
- [18] Rula Sayaf and Dave Clarke. Access control models for online social networks. In *Social Network Engineering for Secure Web Data and Services. IGI - Global*, (in print) 2012.
- [19] Fred Stutzman and Woodrow Hartzog. Boundary regulation in social media. In *CSCW*, 2012.
- [20] Irma Van Der Ploeg. Keys To Privacy. Translations of “the privacy problem” in Information Technologies, pages 15–36. Maastricht: Shaker, 2005.
- [21] Yang Wang, Saranga Komanduri Pedro Giovanni Leon, Gregory Norcie, Alessandro Acquisti, and Lorrie Faith Cranor. “I regretted the minute I pressed share”: A Qualitative Study of Regrets on Facebook. In *Symposium on Usable Privacy and Security*, 2011