

Privacy by Design-Privacy Enabling Design

Workshop 2 Report

Executive Summary

For years, lawmakers, advocates and engineers have touted the potential benefits of Privacy by Design, of integrating privacy throughout the entire technical design process rather than after-the-fact. Nonetheless, we still struggle with how to practice Privacy by Design, whether it is how to conceptualize privacy, how to build privacy in the engineering process, how to present those privacy designs to users or how to incentivize practice of and compliance with Privacy by Design.

In order to identify a shared research vision to support these different facets of the practice of Privacy by Design, the Computing Community Consortium (CCC) is sponsoring a series of four workshops throughout 2015. The second workshop took place in May, 2015 in Atlanta, Georgia, focusing on privacy principles in the design process and how designers can take a more active role in the conversation? A group of over 50 collaborators were in attendance, representing various parts of industry, academia, government and civil society, including legal, philosophy, and computer science academics; researchers and members of industry; and designers from various major design firms.

Here we describe the outcomes of the second workshop, highlighting the key insights, questions, themes, disagreements, and further barriers to actionable progress.

Key insights

Designers lack adequate heuristics to follow when designing applications that may affect users' privacy. While there are principles and truisms, like 'make it usable,' disclose the most salient information necessary without explaining, and use words people understand, these principles are particularly difficult to translate to design in nontraditional interfaces. There is currently no set of heuristics that can be applied across the universe of potential products, and translating design heuristics into privacy spaces is difficult, as compelling design is not necessarily privacy preserving, and privacy design often does not generate as much revenue. However, user-centered design, rather than designing to mitigate risk, can help build trust and better align designs with users' evolving privacy interests. There was disagreement on whether usable heuristics existed or could be articulated, and whether user-centered design is a better approach than risk mitigation.

Users want modular privacy for different personal relationships. In looking at designs built to support elderly parents living alone, it was found that 1) parents generally appreciated being able to share data with their children, but wanted to be able to either hide certain information (like when their significant other was visiting), or exclude some children while sharing with others, and 2) granular data was less vital to the children than knowing whether or not their parent had a normal day. Additionally, patients seem happy to use applications to share data

with health care educators or doctors but this type of sharing leaves a “Loneliness Gap” felt by the patient. Patients were hesitant to share feelings of loneliness, isolation, and depression with their physicians or caretakers because they felt their healthcare professionals either wouldn’t care or that it wasn’t relevant to their treatment. Ultimately, these cases showed that user’s mental models of how applications work affect the decisions they make about their privacy in the context of these applications.

Designing for trust is a good framework, but requires a data-driven and research-driven process to make sure users’ expectations and needs are met. While studies have shown that users will continue to interact with applications performing data mining, even going so far as to rationalize the activity (e.g., a flashlight app using location data may be trying to determine whether or not it’s currently daytime for the user, or a dictionary app may be using location data to tailor what dictionary files it uses), continued use is not the same as continued trust. Even though users may keep interacting with something like retargeting, they do not like feeling spied upon and will have less trust in the system, or just seek to game it for better prices. As a working model, 1) there should be transparency, so users are aware of how information is collected and used; 2) users should be able to make choices regarding with whom and for what purposes their information is shared; and 3) users have to be actively engaged to find out what they care about. Unfortunately, many users don’t understand the difference between PII and how much aggregated data can say about them. Even though it is unlikely that there will be a one-size-fits-all solution, companies can be more transparent in what they collect and how they use it in order to build lasting trust with consumers, and better understand the contextual nuances of users’ privacy concerns. Companies that do so consistently will inspire trust, building more meaningful, lasting relationships with their customers

The workshop also heard reports from companies and designers on how they have been implementing privacy principles: the group received insights from designers in a variety of settings including healthcare applications, browser design and research on how to communicate privacy information to users. Key insights included:

Healthcare

- **Companies trade off privacy protections for user satisfaction.** Many end-clients want continued and uninterrupted access to user data so that they may provide users with experiences that are easy to move through and receive consent from them quickly while still meeting regulatory constraints. This makes it difficult to generate meaningful conversations between users and companies, who are resistant to change they feel might affect the quality and effectiveness of the experiences they are creating. Creating meaningful conversations with users about how their data will be collected and used is resisted, and the push is to meet regulations while getting users to finish assessments and grant consent as quickly as possible.
- **Lack of trust leads to user drop-off.** A perceived lack of privacy protection will cause users, and especially at-risk users, to avoid care and services, even when healthcare applications are incentivized by employers and provide clear benefits to the user.

- **Engaging with clients on the topic of privacy is difficult.** Clients and partners often don't want users to know what data is being shared as it may encourage users to revoke (or discourage them from renewing) consent to use some or all of their data. Because clients associate wider data collection with more successful products, there is generally no room for raising privacy concerns in discussions between clients and designers, as they believe it might limit collection, and resist continued opportunities for users to renew or revoke consent during use. Often there are no voices for privacy concerns at the table for discussion between clients and vendors.

Browsers

- **Users have a strong mental model of advertisers and 3drd parties tracking online activity to target ads.** Pushing the use of add-ons like Ghostery in the first run experience for users, as well as testing of a Tracking Protection feature for browsers, left users feeling that their online behavior was more private and secure when trackers were actively blocked, and their experience was improved by the removal of heavy-handed and obvious ads, as well as lots of website clutter. However, users often mistakenly thought the browser was broken when ads would disappear, and some video playback and social interaction simply does not work when trackers are blocked.
- **Tracker blocking has hard to quantify benefits, but easily quantified pitfalls.** The value to users is difficult to articulate, but advertisers and small businesses want to be able to serve targeted ads and mine data, and resist these kinds of protections. Partnerships with other companies can also influence the adoption of tracker blocking.
- **Privacy onboarding made users more aware of tracking.** Users who tried the beta kept using the nightly builds to maintain the feature because they felt more protected, and those who did not understand tracking previously had stronger mental models after using the system.

Designers

- **Clients lack financial incentives to concern themselves with privacy.** The business of design has incentives to meet client requests and to do so quickly, but not to focus on privacy matters. If data aggregation leads to more personalized products and services, and a general rise in adoption of the product or service, companies have extensive financial incentives to dismiss privacy concerns. For this reason, among others, designers are compelled to concern themselves with designs that will get users to adopt the service while also dismissing privacy concerns.
- **Users understand process better than data.** Studies showed users had trouble synthesizing different presentations of what kinds of data points were being collected, and showed poor recall of those representations over time.
- **Privacy is not just a business issue, but also a design issue.** Increasingly, designers are being consulted early on by business stakeholders to make decisions that can affect privacy. There is an opportunity for designers to convince those in the c-suite, who block privacy protections, of their potential value.

- **Mobile clients have fast and aggressive timelines.** As a result, privacy related interactions are often deprioritized quickly by clients and stakeholders, and they choose to instead rely on built-in operating system mechanisms for determining privacy settings. Often, clients simply stuff high-incomprehensible terms of services, EULAs, and more into footers and about sections.
- **Automotive manufacturers come from a tradition of designing not to harm the user.** These clients have been slow to adopt design around data flowing in and out of cars and have a broad history of trying to avoid any potential harm to customers, including harm related to privacy and security concerns. As such, they are more inclined to attempt to limit data flow to and from vehicles. However, these manufacturers are in fact starting to look into the advantages of this data flow. Unfortunately, their longstanding relationship with designers is one in which designers are removed from initial strategy conversations making it difficult to take a holistic approach to design that incorporates security and privacy principles.
- **Users seem disproportionately concerned about smart home goods.** Users express greater concerns about data shared by connected home goods than they do about more intrusive browser tracking. Companies make efforts to address these concerns, but the conversation is often focused on how to overcome or alleviate privacy concerns and fears. This includes implementing strategic speed bumps that prompt users to think about the data they are sharing and attempt to make them comfortable sharing it. In this space, at least, a conversation addressing privacy concerns has been initiated.
- **Design agencies lack privacy expertise.** Stakeholders often place privacy at the bottom of their list of priorities in creation of a product and design agencies have thus not traditionally had to engage meaningfully with the issue. Often, designers have no recourse if they see an aspect of the design that raises red flags in regards to privacy.

“Encroaching Externalities” limit the freedom to design systems. Underlying infrastructure and building blocks are, understandably, not necessarily designed to reflect a user’s mental model of how a system works. . However, designing systems to better align with users’ mental models, or that can be more easily understood, should improve a system’s readiness and ability to communicate privacy issues to users. Unfortunately, user privacy concerns are often not in alignment with the interests of those creating systems. If a business model is based on creating the richest, densest possible social graph of user activity, any attempts seen as limiting the potential of the graph to display information will be met with resistance by the company, regardless of the user’s best interests.

Users trust themselves most to protect their own privacy, and believe that self-management is an effective way of doing this. Even though people are not good rational decision makers, especially when weighing immediate gratification against unspecified future risks, they believe they can best manage their own privacy. Although this is not surprising, considering the fact that people manage their privacy in their daily, physical lives rather effectively themselves, in their digital lives, their actions and decisions do not reflect their expressed privacy concerns. On top of that, those decisions can be influenced by an application’s design, including UI design,

default settings, and phrasing. It seems clear, in this context, that transparency controls alone will not be sufficient in and of themselves.

Every successful non-traditional interface eventually becomes a traditional interface. New platforms need to build security and privacy protections in from the beginning, as any non-traditional platform has the chance of becoming widely adopted and ubiquitous. Having to build in security and privacy protections after wide adoption can present sometimes insurmountable challenges.

Key Statements and Research Topics

From the workshop survey, the following statements were widely agreed with:

1. **In the HCI tradition, we should conduct more research on the mental models that users have in connection with privacy, in a variety of settings for collecting user data.** This statement was also considered useful to research.
 2. **People act with an understood audience. There is a lot of anthropomorphism as applied to systems: “I’m talking with my counselor about my diabetes diet.” Research in Privacy by Design should focus on context – the audience with whom the user is explicitly or implicitly communicating.** This statement was also considered useful to research.
 3. **Privacy by Design is pervasively multi-disciplinary. Research should be done on how best to determine (1) when in the process designers should get involved to achieve privacy and other goals? (2) what team structure works best, in what contexts?** This statement was also considered useful to research.
-
1. **There is a tension between the complexity of data collection and use, on the one hand, and design goals such as simplicity, elegance, and usability. Research on Privacy by Design should more systematically address this tension.** This statement was also considered useful to research.
-
1. **Privacy is rarely a first- or second-level concern in design. Research should address how to design for privacy, understanding that it is often a third- or fourth-level concern for users.**
-
1. **We need to create infrastructure with building blocks that are closer to the way people think about things (e.g., whether true or not, people have certain beliefs about how an A/C control unit works).** This statement was also considered useful to research.

Future Workshops

The next workshop, the third in the series of four, will be in late August/early September in Pittsburgh, PA at Carnegie Mellon University and it will bring together engineers to discuss privacy research and practice. Finally, to wrap up the series, the fourth workshop will be in Washington, DC at Georgetown University to discuss legal and organizational research necessary to catalyze Privacy by Design.

Report Writers

Annie Anton	Georgia Tech
Justin Hemmings	Georgia Tech
Marie Le Pichon	Georgia Tech
Aaron Massey	Georgia Tech
Peter Swire	Georgia Tech

Workshop Participants

Brian	Anderson	IBM
Annie	Anton	Georgia Tech
Gabriela	Aschenberger	Create with Context
Travis	Breaux	Carnegie Mellon University
Kelly	Caine	Clemson U.
Sunny	Consolvo	Google
Sandra	Corbett	CRA
Lisa	Davis	microsoft
Nicola	Dell	Washington
Ann	Drobnis	CCC
Julie	Earp	NC State University
Keith	Edwards	Georgia Tech
Jennifer	Ehlers	18F
Johnathan	Fox	Intel/McAfee
Batya	Friedman	University of Washington
Patrick	Gage Kelley	Univ. of New Mexico
Vivian	Genaro Motti	Clemson University
Nathan	Good	Good Research

Elizabeth	Goodman	18F/GSA
Susan	Graham	University of California, Berkeley
Aislinn	Grigas	Mozilla
Seda	Gurses	NYU
Justin	Hemmings	Georgia Tech
Alina	Hua	Mozilla
Shawn	Kenney	IBM
Jen	King	Berkeley
Alfred	Kobsa	UCI
Neha	Kumar	USC
Marie	Le Pichon	Georgia Institute of Technology
Liana	Leahy	MeYouHealth
Keith	Marzullo	NSF
Aaron	Massey	Georgia Tech
Matt	Muller	Inflection
Deirdre	Mulligan	Berkeley
Beth	Mynatt	Georgia Tech/CCC
Heather	Patterson	Intel
Torrey	Podmajersky	microsoft
Wanda	Purinton	Georgia Tech
Evinn	Quinn	Critical Mass
Ira	Rubinstein	New York University
Clint	Rule	Frog
Richard	Rutledge	GaTech
Fred	Schneider	Cornell University
Katie	Skinner	Apple
Karen	Sollins	MIT
Peter	Swire	Georgia Tech
Ilana	Westerman	Create With Context
Lauren	Wilcox	Georgia Tech
Richmond	Wong	Berkeley
Helen	Wright	CCC
John	Yuda	18F
Michael	Zimmer	U Wisconsin
Susan	Landau	Privacy Ink

Katie Shilton UM College Park

This material is based upon work supported by the National Science Foundation under Grant No. (1136993).

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.