The Saucer

The Newsletter of the
CyLab Usable Privacy and Security (CUPS) Laboratory

Issue 4, Summer 2013

Contents

CUPS research sponsors

- Army Research Office (ARO)
- Fundacao para a Ciencia e Tecnologia (FCT)
- Department of Homeland Security (DHS)
- Google
- Microsoft
- National Institute of Standards and Technology (NIST)
- National Science Foundation (NSF)
- National Security Agency (NSA) Science of Security Lablet

Carnegie Mellon University
CyLab

Privacy Engineering Masters Program to Begin in August 2013

Carnegie Mellon University has announced a new masters degree program in privacy engineering, co-directed by CUPS faculty members Norman Sadeh and Lorrie Cranor. The co-directors have admitted the program's inaugural class of students, who will arrive on campus in August 2013.

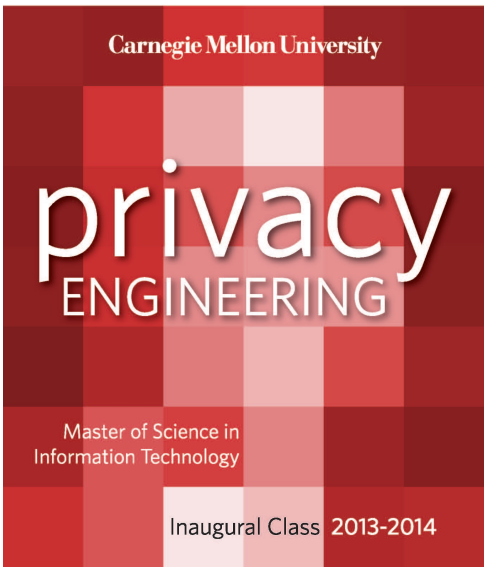

The Master of Science in Information Technology—Privacy Engineering (MSIT-PE) degree is a one-year program designed for computer scientists and engineers who wish to pursue careers as privacy engineers or technical privacy managers. Designed in close collaboration with industry and government, this program is intended for students who aspire to play a critical role in building privacy into future products, services, and processes. This first-of-its kind program responds to the rapidly growing need for technical privacy expertise. As organizations develop new products, services, infrastructures and business processes that facilitate the collection and management of an ever-wider range of customer data, they are discovering that privacy issues need to be addressed from the very beginning of the design process. Over the past several years, both industry and government organizations have created positions for people responsible for ensuring that privacy is an integral part of the design process. These people are brought in as in-house consultants who work as part of multi-disciplinary teams. They have to understand technology and be able to integrate perspectives that span product design, software development, cyber security, human computer interaction, as well as business and legal considerations. Today organizations are already reporting a significant shortage of people who are adequately trained to play this increasingly crucial role, while demand is continuing to increase.

Offered jointly by the School of Computer Science and College of Engineering, the MSIT-PE program includes two semesters of courses taught by leading academic privacy and security experts. Required semester-long courses include Privacy Policy, Law and Technology; Information Security and Privacy; Foundations of Privacy; Usable Privacy and Security; and Engineering Privacy in Software. The program concludes with a summer-long learning-by-doing, capstone project, where students will be brought in as privacy consultants to work on client projects. Students who complete this program will be well prepared for jobs as privacy engineers and technical privacy managers.

For more information, see **http://privacy.cs.cmu.edu**.

CyLab
Usable
Privacy *and*
Security
Laboratory

http://cups.cs.cmu.edu

Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213

Director
Lorrie Faith Cranor
lorrie@cmu.edu

Core Faculty
Alessandro Acquisti
Lujo Bauer
Nicolas Christin
Julie Downs
Jason Hong
Norman Sadeh
Marios Savvides

Supporting Faculty
Travis Breaux
David Brumley
Kathleen Carley
Laura Dabbish
Anupam Datta
Baruch Fischhoff
Greg Ganger
Virgil Gligor
Jim Herbsleb
Robert Kraut
Ramayya Krishnan
George Lowenstein
Brad Myers
Michael Shamos
Rahul Telang

CUPS IGERT students gather on campus for a photo in Spring 2013.

CUPS Doctoral Training Program Offers Unique PhD in Usable Privacy & Security

We are getting ready to welcome the fifth class of the Carnegie Mellon Usable Privacy and Security Doctoral Training Program in August. Students in PhD programs from throughout CMU have the opportunity to participate in the CUPS doctoral training program and earn a CyLab Usable Privacy and Security Meritorious Achievement Certificate. Thanks to an NSF Integrative Graduate Education and Research Traineeship (IGERT) grant, students who are US citizens can also apply for an IGERT traineeship, which provides two years of tuition and stipend support for students in the doctoral training program.

The CUPS doctoral training program offers students a cross-disciplinary training experience that prepares them to produce the key research advances necessary to reconcile ostensible tensions between security, privacy, and usability, moving away from an "either-or" view of these goals to a deeper understanding of underlying tradeoffs and eventually towards solutions where security, privacy, and usability are configured to reinforce each other. The goal of this program is to serve as a catalyst to shape the field of usable privacy and security by developing and training a new generation of researchers in methodologies, principles, and approaches that can be applied across systems and applications.

CUPS students are required to take our Usable Privacy and Security course, plus three approved full-semester courses from a selection of courses in security, privacy, human computer interaction, social and decision sciences, and other areas. Students are also expected to participate in the weekly CUPS research seminar and be engaged in usable privacy and security research.

Prospective students should apply directly to a CMU PhD program of their choice, and also send a letter of interest to the CUPS program administrator indicating which CMU doctoral program they have applied to and describing their interest in CUPS-related research. See the CUPS website for more information **http://cups.cs.cmu.edu**.

From the Director

I've been on sabbatical this year, only coming by the CUPS Lab one day per week (well, that was the plan anyway). In my absence, the rest of the CUPS Lab folks continued their high level of productivity, and you can read about what everyone was up to in the rest of this issue. So, I will use my space to tell you a little bit about my sabbatical.

My sabbatical has been a "staybattical." Family needs prevented me from relocating, but that doesn't mean I couldn't do something different, exciting, mentally liberating, intellectually restorative, relaxing, and totally awesome. I have been spending my sabbatical as a fellow at the STUDIO for Creative Inquiry in the CMU School of Art, pursuing my interests in privacy and security, quilting, and technology.

While other faculty and students in the STUDIO spend the day creating new concepts from behind computer screens, I set up shop with an old sewing machine, an ironing blanket, a and a huge pile of colorful fabric. At the beginning of my fellowship, I smiled politely every time someone suggested ways of attaching the old sewing machine to a robotic arm, and spent days with needle and thread hand quilting colorful lines. Hand quilting is a process that offers one a lot of time to think, and I did spend a lot of time thinking about the art and craft of quilting, and how I might use technology in my work.

I appreciate the added value that technology can bring to my art, enabling me to create in ways that would be difficult or impossible for me unassisted. It is not my goal to use technology to eliminate the need for me to participate in the fabrication process. Part of my attraction to quilting and fiber arts is the tactile nature of the medium. Part of the fun is manipulating fabric and thread with my hands. I want to use technology to enhance my skills – let me sew straighter, faster, better – or, better yet, to let me create in ways I otherwise could not.

When I started a series of quilts, which I call "Interleave," I sketched the quilt designs in pencil and did some design experimentation with scissors and paper. As I started to design the third quilt in the series, I began using PowerPoint to sketch out some ideas involving sine waves. STUDIO director Golan Levin saw what I was doing and suggested I write a program using an arts engineering toolkit called Processing to draw my design. The program I wrote allowed me to generate the sorts of designs I had been struggling with, and it included sliders to allow me to experiment with sine waves of different frequencies and amplitudes. Using this program, I was able to rapidly iterate through large numbers of design possibilities before selecting one to actually fabricate. I did some engineering to figure out how to actually construct the quilt I designed, and then adapted my program to produce full-scale templates that I could print on paper and use to cut out my fabric.

In my sabbatical proposal, I said I would explore visualizing privacy and security concepts through art. It sounded like a plausible way to tie my research interests to my sabbatical plan, but I wasn't entirely sure how I was going to do that. But over the course of my sabbatical I actually created two privacy themed quilts, a privacy-enhanced self portrait and a quilt constructed from digitally printed photos called "De-identification." I have ideas for another privacy quilt as well as for a "security blanket" that I hope to create some time soon.

I was excited to have several of my quilts displayed in a solo exhibit at the Pittsburgh Children's Museum this summer. I was also commissioned to create two original quilts and I've been invited to give a lecture and teach a workshop at a local quilt guild. If you would like to read more about my sabbatical projects, check out my blog at **http://lorrie.cranor.org/blog/**

Lorrie

Students

Idris Adjerid
Hazim Saleh Almuhimedi
Rebecca Balebako
Sekhar Bhagvatula
Cristian Bravo-Lillo
Justin Cranshaw
Pedro Giovanni Leon
David Gordon
Hanan Hibshi
Eiji Hayashi
Peter Klemperer
Saranga Komanduri
Shing-hon Lau
Pedro Leon
Michelle Mazurek
Emmanuel Owusu
Rich Shay
Manya Sleeper
Blase Ur
Timothy Vidas
Tatiana Vlahovic
Jason Wiese

Post-Docs & Staff

Alain Forget
Mandy Holbrook
Eyal Peer
Tiffany M. Todd
Shomir Wilson

Quilts by Lorrie Cranor on display at the Pittsburgh Children's Museum. From left to right the quilts shown are Self Portrait, Lying on the Floor of the Pittsburgh Children's Museum Staring at the Ceiling, De-Identification, and Interleave #1.

CUPS alumni Patrick Gage Kelley speaks at the SOUPS 2012 conference.

CUPS PhD student Rebecca Balebako presents a poster at the 2012 CyLab Corporate Parters Conference

Pedro Leon, Lorrie Cranor, Rich Shay, and Blase Ur at the Future of Privacy Forum's Privacy Papers for Policy Makers Reception.

Year in Review

July 2012

- Alessandro Acquisti testified at the US Senate Committee of the Judiciary Subcommittee on Privacy, Technology and the Law hearing on "What Facial Recognition Technology Means for Privacy and Civil Liberties"

August 2012

- CUPS password research was mentioned in the *Pittsburgh Post-Gazette* in the article "Password Length is more Beneficial than Complexity"
- Lorrie Cranor was appointed a Privacy by Design Ambassador by Ontario Information Privacy commissioner Ann Cavoukian
- Lorrie Cranor was interviewed by *The Wall Street Journal* for the article "New Rules on Kids' Web Ads"
- CUPS faculty member Travis Breaux and PhD student David Gordon received a distinguished paper award for the paper entitled "Reconciling Multi-Jurisdictional Requirements: A Case Study in Requirements Water Marking" at the 20th IEEE International Requirements Engineering Conference

October 2012

- CMU introduced a new privacy engineering master's program
- Two CUPS papers were selected to appear in the Future of Privacy Forum's annual Privacy Papers for Policy Makers. "Smart, Useful, Scary, Creepy: Perceptions of Online Behavioral Advertising" by Blase Ur, Pedro G. Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang; and "Will Johnny Facebook Get a Job? An Experiment in Hiring Discrimination via Online Social Networks" by Alessandro Acquisti and Christina Fong.
- Alessandro Acquisti and Christina Fong were awarded the Privacy Law Scholar Conference/IAPP award at the IAPP Privacy Academy for their research on the impact of social networks information on firms' hiring behavior
- Janice Tsai, Serge Egelman, Lorrie Cranor, and Alessandro Acquisti won the prestigious Information Systems Research (ISR)'s 2012 Best Published Paper Award for the best paper published in 2011. The award was for the paper "The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study"
- Jason Hong was interviewed for *The New York Times* in the article "Data Gathering via Apps Presents a Gray Legal Area"

December 2012

- "Q&A: Privacy engineers could hold the key" with Lorrie Cranor was published in the *Pittsburgh Tribune Review*
- Jason Hong was interviewed for a December 26 CBS Morning News segment on "Smartphone Snoops: How your Phone Data is Being shared"

January 2013

- Jialiu Lin, Norman Sadeh, and Jason Hong were interviewed by *Pittsburgh Tribune Review* in the article "Many Free Apps Come at a Price" January 16
- Jason Hong was interviewed by *Live Science* in the article "10 Surprising Apps That Are Sharing Your Data" January 16
- Jason Hong was interviewed by *NBC News* in the article "A shock in the dark: Flashlight app track your location" January 16

- Lorrie Cranor was interviewed by *Pittsburgh-Post Gazette* for the article "Electronic devices and services monitor us 24/7, but there are few ways to block them" January 29
- Norman Sadeh moderated the session "Will the Mobile Web and social Networking Mark the End of Privacy?" during Data Privacy Day at CMU

February 2013

- Lorrie Cranor was interviewed by *MIT Technology Review* for the article "This Dopey, Decade-Old Tech Privacy Video Actually Got a Lot of Things Right" February 8

March 2013

- "A Guide to Facebook's privacy options" By Lorrie Cranor appeared in *The Wall Street Journal* March 11
- Rebecca Balebako received the "Best Presentation Award" for her talk about student-parents and family leave at the regional conference for the National Association for Graduate-Professional Students March 13
- Alessandro Acquisti was interviewed by *The New York Times* in the article "Letting Down Our Guard With Web Privacy"
- Justin Cranshaw was awarded a 2013-14 Facebook Graduate Fellowship

April 2013

- Alessandro Acquisti served as keynote speaker at the USEC Workshop on Usable Security at financial Cryptography 2013 in Okinawa, Japan
- Alessandro Acquisti was interviewed by *Salon.com* for the article "Why Facial Recognition Failed"

May 2013

- Rebecca Balebako received the Google *Anita Borg Memorial Scholarship*
- Alessandro Acquisti and Ey'al Pe'er were interviewed by the *New York Times* for the article "Brain Interrupted" which focused on their research about the impact of gadgets on our brains May 3
- Alessandro Acquisti served as keynote speaker for the Proskauer Conference on Privacy in New York
- Alessandro Acquisti and Marios Savvides were interviewed for a segment on *60 Minutes* entitled "'Big Brother' is Big Business?"
- Blase Ur received the Best Presenter Award at PSOSM 2013 for "A Cross-Cultural Framework for Protecting User Privacy in Online Social Media"

June 2013

- Lorrie Cranor was interviewed by *NPR Radio* for "Poll: Majority of Americans Comfortable with Surveillance" June 12
- Alessandro Acquisti was a speaker at the IFIP Summer School on Privacy & Identity Management in The Netherlands

July 2013

- Lorrie Cranor was the keynote speaker at the 2013 Privacy Enhancing Technologies Symposium in Bloomington, Indiana
- SOUPS 2013 will be held July 22-24 at Northumbria University, UK

Alain Forget was awarded a Carelton University Senate Medal for Outstanding Achievement.

CUPS faculty member Marios Savvides gets ready to launch at the CyLab holiday party live Angry Birds tournament at the National Aviary.

Lorrie Cranor with one of her quilts at the STUDIO for Creative Inquiry.

New Additions

Former CUPS postdoc **Yang Wang** and wife Yuan Huang welcomed their second son, David H. Wang in January. Yang is now an assistant professor in the School of Information Studies (aka iSchool) at Syracuse University. He continues to conduct research in usable privacy and security and collaborate with his CUPS colleagues.

Henry cuddles his little brother David

CUPS PhD student **Hanan Hibishi** and husband welcomed a baby girl named Layla in October.

CUPS PhD student **Rebecca Balebako** and husband Eric welcomed their second daughter, Dissi Dula Balebako in June. Dissi (pronounced D.C.) is a Losso name meaning "Blessing."

CUPS Ph.D. student **Hazim Almuhimedi** and his wife welcomed a newborn baby girl Lana Hazim Almuhimedi in June.

CUPS PhD student **Tim Vidas** and his wife Sheila welcomed a newborn baby girl named Catherine Claire in July.

Weddings

Congrats to CUPS PhD student **Saranga Komanduri** who got married this past May to fiancée Pradipta at the Hindu Temple of Minnesota. Best wishes to the newlyweds!

CUPS PhD student **Jason Wiese** married Eliane Stampfer in Portland, OR on June 16. Eliane is a PhD student in HCII. Congratulations!

Recent Publications

Most CUPS publications are available on the CUPS website. The following are a selection of publications from the past year.

Privacy Decision Making

Privacy as Part of the App Decision-Making Process

Patrick G. Kelley, Lorrie Faith Cranor, and Norman Sadeh

CHI 2013

Smartphones have unprecedented access to sensitive personal information. While users report having privacy concerns, they may not actively consider privacy while downloading apps from smartphone application marketplaces. Currently, Android users have only the Android permissions display, which appears after they have selected an app to download, to help them understand how applications access their information. We investigate how permissions and privacy could play a more active role in app-selection decisions. We designed a short "Privacy Facts" display, which we tested in a 20-participant lab study and a 366-participant online experiment. We found that by bringing privacy information to the user when they were making the decision and by presenting it in a clearer fashion, we could assist users in choosing applications that request fewer permissions

Necessary But not Sufficient: Standardized Mechanisms for Privacy Notice and Choice

Lorrie Faith Cranor

Journal of Telecommunications and High Technology Law, Vol. 10

For several decades, "notice and choice" have been key principles of information privacy protection. Conceptions of privacy that involve the notion of individual control require a mechanism for individuals to understand where and under what conditions their personal information may flow and to exercise control over that flow. Thus, the various sets of fair information practice principles and the privacy laws based on these principles include requirements for providing notice about data practices and allowing individuals to exercise control over those practices. Privacy policies and opt-out mechanisms have become the predominant tools of notice and choice. However, a consensus has emerged that privacy policies are poor mechanisms for communicating with individuals about privacy. With growing recognition that website privacy policies are failing consumers, numerous suggestions are emerging for technical mechanisms that would provide privacy notices in machine-readable form, allowing web browsers, mobile devices, and other tools to act on them automatically and distill them into simple icons for end users. Other proposals are focused on allowing users to signal to websites, through their web browsers, that they do not wish to be tracked. These proposals may at first seem like fresh ideas that allow us to move beyond impenetrable privacy policies as the primary mechanisms of notice and choice. However, in many ways, the conversations around these new proposals are reminiscent of those that took place in the 1990s that led to the development of the Platform for Privacy Preferences ("P3P") standard and several privacy seal programs.

In this paper I first review the idea behind notice and choice and user empowerment as privacy protection mechanisms. Next I review lessons from the development and deployment of P3P as well as other efforts to empower users to protect their privacy. I begin with a brief introduction to P3P, and then discuss the privacy taxonomy associated with P3P. Next I discuss the notion of privacy nutrition labels and privacy icons and describe our demonstration of how P3P policies can be used to generate privacy nutrition labels automatically. I also discuss studies that examined the impact of salient privacy information on user behavior. Next I look at the problem of P3P policy adoption and enforcement. Then I discuss problems with recent self-regulatory programs and privacy tools in the online behavioral advertising space. Finally, I argue that while standardized notice mechanisms may be necessary to move beyond impenetrable privacy policies, to date they have failed users and they will continue to fail users unless they are accompanied by usable mechanisms for exercising meaningful choice and appropriate means of enforcement

Silent Listeners: The Evolution of Privacy and Disclosure on Facebook

Fred Stutzman, Ralph Gross, Alessandro Acquisti

Journal of Privacy and Confidentiality: Vol. 4: Iss. 2, Article 2

Over the past decade, social network sites have experienced dramatic growth in popularity, reaching most demographics and providing new opportunities for interaction and socialization. Through this growth, users have been challenged to manage novel privacy concerns and balance nuanced trade-offs between disclosing and withholding personal information. To date, however, no study has documented how privacy and disclosure evolved on social network sites over an extended period of time. In this manuscript we use profile data from a longitudinal panel of 5,076 Facebook users to understand how their privacy and disclosure behavior changed between 2005---the early days of the network---and 2011. Our analysis highlights three contrasting trends. First, over time Facebook users in our dataset exhibited increasingly privacy-seeking behavior, progressively decreasing the amount of personal data shared publicly with unconnected profiles in the same network. However, and second, changes implemented by Facebook near the end of the period of time under our observation arrested or in some cases inverted that trend. Third, the amount and scope of personal information that Facebook users revealed privately to other connected profiles actually increased over time---and because of that, so did disclosures to ``silent listeners" on the network: Facebook itself, third-party apps, and (indirectly) advertisers. These findings highlight the tension between privacy choices as expressions of individual subjective preferences, and the role of the environment in shaping those choices.

Privacy

"I Read my Twitter the Next Morning and was astonished": A conversational perspective on Twitter regrets

Manya Sleeper, Justin Cranshaw, Patrick G. Kelley, Blase Ur, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh
CHI 2013

We present the results of an online survey of 1,221 Twitter users, comparing messages individuals regretted either saying during in-person conversations or posting on Twitter. Participants generally reported similar types of regrets in person and on Twitter. In particular, they often regretted messages that were critical of others. However, regretted messages that were cathartic/expressive or revealed too much information were reported at a higher rate for Twitter. Regretted messages on Twitter also reached broader audiences. In addition, we found that participants who posted on Twitter became aware of, and tried to repair, regret more slowly than those reporting in-person regrets. From this comparison of Twitter and in-person regrets, we provide preliminary ideas for tools to help Twitter users avoid and cope with regret.

The Post that Wasn't: Exploring Self-Censorship on Facebook

Manya Sleeper, Rebecca Balebako, Sauvik Das, Amber Lynn McConahy, Jason Wiese, Lorrie Faith Cranor
CSCW 2013

Social networking site users must decide what content to share and with whom. Many social networks, including Facebook, provide tools that allow users to selectively share content or block people from viewing content. However, sometimes instead of targeting a particular audience, users will self-censor, or choose not to share. We report the results from an 18-participant user study designed to explore self-censorship behavior as well as the subset of unshared content participants would have potentially shared if they could have specifically targeted desired audiences. We asked participants to report all content they thought about sharing but decided not to share on Facebook and interviewed participants about why they made sharing decisions and with whom they would have liked to have shared or not shared. Participants reported that they would have shared approximately half the unshared content if they had been able to exactly target their desired audiences.

What Matters to Users? Factors that Affect Users' Willingness to Share Information with Online Advertisers

Pedro G. Leon, Blase Ur, Yang Wang, Manya Sleeper, Rebecca Balebako, Richard Shay, Lujo Bauer, Mihai Christodorecu, Lorrie Faith Cranor
SOUPS Proceedings 2013

Much of the debate surrounding online behavioral advertising (OBA) has centered on how to provide users with notice and choice. An important element left unexplored is how advertising companies' privacy practices affects users' attitudes toward sharing data. We present the results of a 2,912-participant online study investigating how facets of privacy practices—data retention, access to collected data, and scope of use—affect users' willingness to allow the collection of behavioral data. We asked participants to visit a health website, explained OBA to them, and outlined policies governing data collection for OBA purposes. These policies varied by condition. We then asked participants about their willingness to permit the collection of 30 types of information. We identified classes of information that most participants would not share, as well as classes that nearly half of participants would share. More restrictive data-retention and use policies increased participants' willingness to allow data collection, while policies allowing users to review and modify their data had a smaller impact. We discuss user-interface improvements and public-policy approaches to support users' privacy preferences.

Are They Actually any Different? Comparing Thousands of Financial Institutions' Privacy Practices

Lorrie Faith Cranor, Kelly Idouchi, Pedro G. Leon, Manya Sleeper, Blase Ur
WEIS 2013

Although large-scale comparisons of privacy practices across an industry have the potential to illuminate the state of consumer privacy and to uncover egregious practices, the freeform legalese of most privacy policies makes such comparisons time-consuming and expensive. Financial institutions in the United States are required by the Gramm-Leach-Bliley Act to provide annual privacy disclosures. In 2009, eight federal agencies jointly released a model privacy form for these disclosures. While use of the model privacy form is not required, it has been widely adopted. With so many financial institutions' policies available in a standard format, large-scale comparisons are now more readily achievable.

We built an automated web crawler and document parser for the model privacy form and automatically evaluated thousands of financial institutions' disclosures. We found large variance in data-sharing practices, even among banks of the same class. While thousands of financial institutions share personal information without providing the opportunity for consumers to opt out, some institutions' practices are more consumer-friendly. Institutions' practices vary by region and by the size of the institution. Furthermore, we uncovered violations of financial regulation, such as failing to allow consumers to limit data sharing even when required to do so. We identify issues with the design and use of the model privacy form, ranging from poorly designed categories to institutions making self-contradictory statements. We discuss implications for privacy in the financial industry, as well as future directions for standardized privacy notices.

A Cross-Cultural Framework for Protecting User Privacy in Online Social Media

Blase Ur and Yang Wang
PSOSM 2013

Social media has become truly global in recent years. We argue that support for users' privacy, however, has not been extended equally to all users from around the world. In this paper, we survey existing literature on cross-cultural privacy issues, giving particular weight to work specific to online social networking sites. We then propose a framework for evaluating the extent to which social networking sites' privacy options are offered and communicated in a manner that supports diverse users from around the world. One aspect of our framework focuses on cultural issues, such as norms regarding the use of pseudonyms or posting of photographs. A second aspect of our framework discusses legal issues in cross-cultural privacy, including data-protection requirements and questions of jurisdiction. The final part of our framework delves into user expectations regarding the data-sharing practices and the communication of privacy information. The framework can enable service providers to identify potential gaps in support for user privacy. It can also help researchers, regulators, or consumer advocates reason systematically about cultural differences related to privacy in social media.

Tweets are forever: A Large-Scale Quantitative Analysis of Deleted Tweets

Hazim Almuhimedi, Shomir Wilson, Bin Liu, Norman Sadeh, Alessandro Acquisti

Proceedings of the 2013 Conference on Computer Supported Cooperative Work (pp. 897-908)

This paper describes an empirical study of 1.6M deleted tweets collected over a continuous one-week period from a set of 292K Twitter users. We examine several aggregate properties of deleted tweets, including their connections to other tweets (e.g., whether they are replies or retweets), the clients used to produce them, temporal aspects of deletion, and the presence of geotagging information. Some significant differences were discovered between deleted and undeleted tweets, namely in the clients used to post them, their conversational aspects, the sentiment vocabulary present in them, and the days of the week they were posted. However, in other dimensions for which analysis was possible, no substantial differences were found. Finally, we discuss some ramifications of this work for understanding Twitter usage and management of one's privacy.

Passwords

Optimizing Password Composition Policies

Jeremiah Blocki, Saranga Komanuri, Ariel Procaccia, and Or Sheffet

Proceedings of the fourteenth ACM conference on electronic commerce (EC'13)

A password composition policy restricts the space of allowable passwords to eliminate weak passwords that are vulnerable to statistical guessing attacks. Usability studies have demonstrated that existing password composition policies can sometimes result in weaker password distributions; hence a more *principled* approach is needed. We introduce the *first theoretical model* for optimizing password composition policies. We study the computational and sample complexity of this problem under different assumptions on the structure of policies and on users' preferences over passwords. Our main positive result is an algorithm that -- with high probability --- constructs almost optimal policies (which are specified as a union of subsets of allowed passwords), and requires only a small number of samples of users' preferred passwords. We complement our theoretical results with simulations using a real-world dataset of 32 million passwords.

Helping Users Create Better Passwords

Blase Ur, Patrick G. Kelley, Saranga Komanduri, Joel Lee, Michael Maass, Michelle L. Mazurek, Timothy Passaro, Richard Shay, Timothy Vidas, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor, Serge Egelman, and Julio Lopez

USEC 2013

Over the past several years, we have researched how passwords are created, how they resist cracking, and how usable they are. In this article, we focus on recent work in which we tested various techniques that may encourage better password choices. What we found may surprise you.

Despite a litany of proposed password replacements, text-based passwords are not going to disappear anytime soon. Passwords have a number of advantages over other authentication mechanisms. They are simple to implement, relatively straightforward to revoke or change, easy for users to understand, and allow for quick authentication; however, passwords also have a number of drawbacks. Foremost among these drawbacks is that it is difficult for users to create and remember passwords that are hard for an attacker to guess. Our research group at Carnegie Mellon University has been investigating strategies to guide users to create passwords that are both secure and memorable.

The Impact of Length and Mathematical Operators on the Usability and Security of System-assigned One-time PINS

Patrick Gage Kelley, Saranga Komanduri, Michelle L. Mazurek, Richard Shay, Tim Vidas, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor

USEC 2013

Over the last decade, several proposals have been made to replace the common personal identification number, or PIN, with often-complicated but theoretically more secure systems. We present a case study of one such system, a specific implementation of system-assigned one-time PINs called PassGrids. We apply various modifications to the basic scheme, allowing us to review usability vs. security trade-offs as a function of the complexity of the authentication scheme. Our results show that most variations of this one-time PIN system are more enjoyable and no more difficult than PINs, although accuracy suffers for the more complicated variants. Some variants increase resilience against observation attacks, but the number of users who write down or otherwise store their password increases with the complexity of the scheme. Our results shed light on the extent to which users are able and

willing to tolerate complications to authentication schemes, and provides useful insights for designers of new password schemes.

Privacy, Location, and Mobile Devices

QRishing: The Susceptibility of Smartphone users to QR Code Phishing Attacks

Timothy Vidas, Emmanuel Owusu, Shuai Wang, Cheng Zeng, Lorrie Faith Cranor

USEC 2013

The matrix barcodes known as Quick Response (QR) codes are rapidly becoming pervasive in urban environments around the world. QR codes are used to represent data, such as a web address, in a compact form that can be readily scanned and parsed by consumer mobile devices. They are popular with marketers because of their ease in deployment and use. However, this technology encourages mobile users to scan unauthenticated data from posters, billboards, stickers, and more, providing a new attack vector for miscreants. By positioning QR codes under false pretenses, attackers can entice users to scan the codes and subsequently visit malicious websites, install programs, or any other action the mobile device supports. We investigated the viability of QR-code-initiated phishing attacks, or QRishing, by conducting two experiments. In one experiment we visually monitored user interactions with QR codes; primarily to observe the proportion of users who scan a QR code but elect not to visit the associated website. In a second experiment, we distributed posters containing QR codes across 139 different locations to observe the broader application of QR codes for phishing. Over our four-week study, our disingenuous flyers were scanned by 225 individuals who subsequently visited the associated websites. Our survey results suggest that curiosity is the largest motivating factor for scanning QR codes. In our small surveillance experiment, we observed that 85% of those who scanned a QR code subsequently visited the associated URL.

Phoneprioception: enabling Mobile Phones to Infer Where they are Kept

Jason Wiese, T. Scott Saponas, A.J. Brush

CHI 2013

Enabling phones to infer whether they are currently in a pocket, purse or on a table facilitates a range of new interactions from placement-dependent notifications setting to preventing "pocket dialing". We collected data from 693 participants to understand where people keep their phone in different contexts and why. Using this data, we identified three placement personas: Single Place Pat, Consistent Casey, and All-over Alex. Based on these results, we collected two weeks of labeled accelerometer data in-situ from 32 participants. We used this data to build models for inferring phone placement, achieving an accuracy of approximately 85% for inferring whether the phone is in an enclosed location and for inferring if the phone is on the user. Finally, we prototyped a capacitive grid and a multispectral sensor and collected data from 15 participants in a laboratory to understand the added value of these sensors.

Home Storage

The Current State of Access Control for Smart Devices in Homes

Blase Ur, Jaeyeon Jung, Stuart Schechter

HUPS 2013

Although connected devices and smart homes are now marketed to average consumers, little is known about how access-control systems for these devices fare in the real world. In this paper, we conduct three case studies that evaluate the extent to which commercial smart devices provide affordances related to access control. In particular, we examine an Internet-connected lighting system, bathroom scale, and door lock. We find that each device has its own siloed access-control system and that each approach fails to provide seemingly essential affordances. Furthermore, no system fully supports user understanding of access control for the home.

Usable Cyber Trust Indicators

Operating System Framed in Case of Mistaken identity: Measuring the success of web-based spoofing attacks on OS password-entry dialogs

Cristian Bravo-Lillo, Lorrie Faith Cranor, Julie Downs, Saranga Komanduri, Stuart Schechter and Manya Sleeper

19th ACM Conference on Computer and Communications Security

When asking users to enter credentials, today's desktop operating systems often use windows that provide scant evidence that a trusted path has been established; evidence that would allow a user to know that a request is genuine and that the password will not be read by untrusted principals. We measure the effifficacy of web-based attacks that spoof these operating system credential-entry windows to steal users' device-login passwords. We recruited 504 users of Amazon's Mechanical Turk to evaluate a series of games on third-party websites. The third such website indicated that it needed to install software from the publisher that provided the participants' operating system: Microsoft's Silverlight for Windows Vista/7 users and Apple's QuickTime for Mac OS users. The website then displayed a spoofed replica of a window the participant's client operating system would use to request a user's device credentials. In our most effective attacks, over 20% of participants entered passwords that they later admitted were the genuine credentials used to login to their devices. Even among those who declined to enter their credentials, many participants were oblivious to the spoofing attack. Participants were more likely to confirm that they were worried about the consequences of installing software from a legitimate source than to report that they thought the credential-entry window might have appeared as a result of an attempt to steal their password.

Warning Design Guidelines
Lujo Bauer, Cristian Bravo-Lillo, Lorrie Faith Cranor, and Elli Fragkaki
CyLab Technical Report, February 2013
This document contains a set of guidelines aimed at helping software designers and developers in designing more effective warning dialogs. These guidelines were compiled from available literature on usable security and warnings research and from Human Interface Guidelines for three broadly used operating systems: Windows, MacOS, and Linux. The goal of this work is to help people cope with the shifting privacy landscape. While our work looks at many aspects of how users make decisions regarding their privacy, this dissertation focuses on two specific areas: the current state of web privacy policies and mobile phone application permissions. We explored consumers' current understandings of privacy in these domains, and then used that knowledge to iteratively design and test more comprehensible information displays. These prototyped information displays should not necessarily be seen as final commercially-ready solutions, but as examples of privacy notices that help users think about, cope with, and make decisions regarding their data privacy. We conclude with a series of design suggestions motivated by our findings.

Thesis Proposals

Understanding and Capturing People's Mobile App Privacy Preferences
Jialiu Lin
August 2012

Identifying and Communicating the Risks of Data Sharing by Smartphones and Smartphone Applications
Rebecca Balebako
February 2013

Improving Security Dialogs: an Exploration of Attention and Habituation
Cristian Bravo-Lillo
April 2013

Privacy Notice and Choice in Practice
Pedro Giovanni Leon
May 2013

Enabling an Ecosystem of Personal Behavioral Data
Jason Wiese
May 2013

Dissertations

Proximity Displays for Access Control:
A Doctoral dissertation
Kami Vaniea
September 2012
This thesis proposes the use of *proximity information displays* - small interface components spatially located near the data elements (or near a representation of data, e.g., file name in a file manager or thumbnail photo in a photo album) that contain information about who currently has access or who could access the data. These displays are intended to help users become more aware of how their data has been used in the past and how the data could be used in the future. We present empirical studies that test the hypothesis: Users of a system that includes proximity information displays of access control-information will implement policies that result in grant/deny actions that better match their preferences than will users of a system where access-control information is available only on a secondary interface. Participants who saw proximity displays that were more comprehensive and could be glanced at easily were better able to identify access-control policy errors. Participants who saw displays that were overly coarse-grained, on the sidebar, or showed information about who had previously viewed the photos, showed no improvement over those who saw permission settings only on a secondary interface. Our studies suggest that proximity displays for access control can help significantly the majority of users who do not normally check their access-control policies.

Designing Privacy Notices: Supporting User Understanding and Control
Patrick Gage Kelley
May 2013
Users are increasingly expected to manage complex privacy settings in their normal online interactions. From shopping to social networks, users make decisions about sharing their personal information with corporations and contacts, frequently with little assistance. Current solutions require consumers to read long documents or go out of their way to manage complex settings buried deep in management interfaces, all of which lead to little or no actual control.
The goal of this work is to help people cope with the shifting privacy landscape. While our work looks at many aspects of how users make decisions regarding their privacy, this dissertation focuses on two specific areas: the current state of web privacy policies and mobile phone application permissions. We explored consumers' current understandings of privacy in these domains, and then used that knowledge to iteratively design and test more comprehensible information displays. These prototyped information displays should not necessarily be seen as final commercially-ready solutions, but as examples of privacy notices that help users think about, cope with, and make decisions regarding their data privacy. We conclude with a series of design suggestions motivated by our findings.

In Other News

Congratulations to **Dr. Patrick Gage Kelley** and **Dr. Kami Vaniea!** Both completed their Ph.D. programs and graduated this past May. **Patrick** is an Assistant Professor of computer Science at the University of New Mexico. **Kami** is a Post Doctoral researcher at the Behavior Information Technology lab at Michigan State.

CUPS alumnus **Aleecia McDonald** was appointed director of privacy at Stanford's Center for Internet and Society.

CUPS alumni **Rob Reeder** started a new position at Google New York in April after being at Microsoft for 4.5 years.

CUPS alumni **Janne Lindqvist** was appointed as a tenure-track assistant professor of Electrical and Computer Engineering at Rutgers University starting September 2013. Janne has been an assistant research professor at Rutgers since September 2011. During his first year at Rutgers, Janne was awarded three NSF grants totaling nearly $1.3 million and a MobiCom best paper award. Janne leads the Human-Computer Interaction group at Rutgers University.

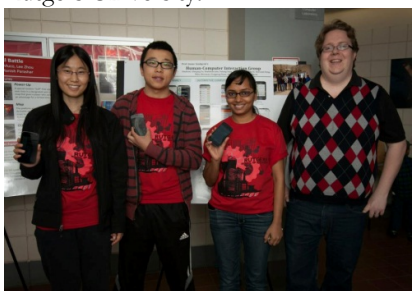

Alain Forget joined the CUPS Lab in October as a postdoc working on the Security Behavior Observatory project. Shortly after joining us, he graduated from Carleton University, where he did his thesis on "A World with Many Authentication Schemes," and was awarded a Senate Medal for Outstanding Achievement at the doctoral level.

CUPS PhD Student **Rich Shay** was one of several people from around the world invited to play "Magic" at the highest level, the Pro Tour. Magic is a trading card game for serious gamers only!

Below: Serious research: Members of the CUPS lab worked together to figure out the best items to use for a successful "Live Angry Birds" competition held at the CyLab Holiday Party, 2013

Above: Before graduation, Kami and Patrick show off their new lab coats and pose with their advisors, Lujo Bauer, Lorrie Cranor, and Norman Sadeh.

Right: CUPS Lab members gather in front of the Collaborative Innovation Center at CMU.