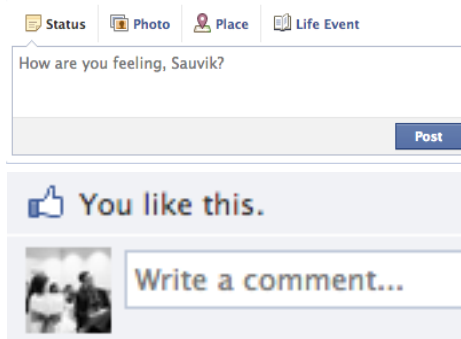


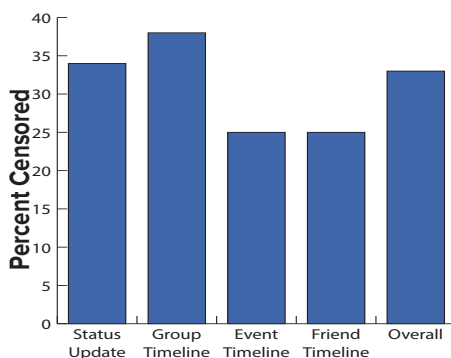
Sauvik Das

Research Statement

HCI researcher
working on
security and
privacy.



When people started writing content into these UI elements, we kept track of whether or not they ultimately decided to share.



Percent of censored posts across different sharing contexts.

Security and privacy help realize the full potential of computing. Without authentication and encryption, for example, few would use digital wallets, social media or even e-mail. The struggle of security and privacy is to realize this potential without imposing too steep a cost. Yet, for the average non-expert, security and privacy are just that: costly, in terms of time, effort and social capital. It is unsurprising, therefore, that for many laypeople, security and privacy tools are begrudgingly tolerated if not altogether subverted. This cannot continue. As computing encompasses more of our lives, we are tasked with making increasingly more security and privacy decisions. Simultaneously, the cost of every breach is swelling. Today, a security breach might compromise sensitive data about our finances and schedules as well as deeply personal data about our health, communications, and interests. Tomorrow, as we enter the era of pervasive smart things, that breach might compromise access to our homes, vehicles and bodies.

In my research, I aim to empower end-users with novel security and privacy systems that mitigate costs of time, effort and social capital in order to promote their use and social spread. To do so, I pursue two high-level research thrusts. First, I construct empirical models of people's security and privacy behaviors, particularly as they relate to social technologies and social use-cases of technology. Second, I construct novel security and privacy systems that better match people's behaviors and capabilities. While I identify primarily as a computer scientist, my work across these two research thrusts draws from a variety of other intellectual traditions including computational social science, cognitive psychology, usable security, ubiquitous computing and applied machine learning.

Modeling People's Security and Privacy Behaviors

To build useful systems for end-users, we must first understand those users. In my user modeling work, I employ a mixed-methods approach to do just that: empirically model people's security and privacy behaviors, preferences and capabilities. Here, I discuss two illustrative examples: modeling people's self-censorship decisions on social media as well as how social influence affects security behaviors.

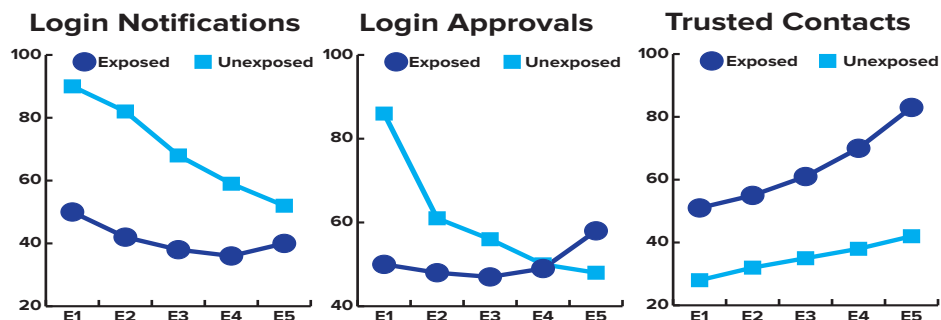
Last-Minute Self-Censorship on Facebook

Last-minute self-censorship refers to situations where people begin the process of sharing content on social media but ultimately decide against sharing. Understanding last-minute self-censorship is important: people use social media to stay connected to loved ones, but their desire to share can conflict with their desire to keep their lives hidden from the wrong people. Thus, understanding self-censorship can surface tensions between current privacy affordances and people's preferences.

As an intern at Facebook, I investigated the last-minute self-censorship behavior of 3.9 million Facebook users [1]. I defined and implemented self-censorship as beginning to write content, but deciding not to actually share that content within 10 minutes. I modeled the frequency of that behavior with person-level demographics and behavior as well as interface-level context. Over the 17 days that we followed our sample, **71% of the 3.9 million users self-censored at least once** (a simple statistic that became the talking point of many popular press articles). I also found, among many other things, that original content ("posts") was self-censored more than comment replies to extant posts (33% vs. 13%), that men self-censored more than women, and that people part of more social circles self-censored more. To better understand what content people self-censored as well as why they did so, I complemented this large-scale "big-data" exploration of self-censorship with a "small-data" week-long diary study [2]. From this data, one of the most actionable insights I found was that people self-censored content that they wanted to share with dynamically defined groups: such as people who like pizza or people who were taking a particular exam. However, **existing privacy controls are static and preclude one's ability to share with dynamically defined audiences.** As a result of my work, Facebook has released a number of new features that help people target dynamic groups (e.g., through public conversations and Facebook groups).

Building a Theory for Social Cybersecurity

To date, little is known about the role of social influence in security and privacy. "Users" are treated as isolated actors whose security and privacy decisions are



Matched propensity adoption rates of Facebook security tools for people who were and were not exposed to friends. Different tools had different patterns.

Keep Your Account Safe
108 of your friends use extra security settings. You can protect your account and make sure it can be recovered ever lose access.

[Improve Account Security](#)

Keep Your Account Safe
You can use security settings to protect your account make sure it can be recovered if you ever lose access.

[Improve Account Security](#)

A social (top) and non-social (bottom) announcement tested in my experiment.

made in a vacuum, independent of external social influences. For example, we know little about how Alice’s decision to use two-factor authentication affects her friend Bob’s security choices. Likewise, little is known about the social consequences of security system design. For example, if Bob secures his home WiFi with a meticulous 30 character password, does that influence how he is perceived by his guests? My social cybersecurity work bridges these gaps in our understanding.

In formative work, I used semi-structured interviews to ask people of various ages and backgrounds about their security-related behavior changes and communications (SOUPS’14) [3]. Next, in a complementary large-scale quantitative exploration, I used matched propensity sampling analysis to understand how three security tools diffused (or not) through the social networks of 1.5 million people on Facebook (CSCW’15) [4]. Third, I empirically validated one of the key insights gained from the prior two exploratory studies, that observability and social proof are key to the widespread adoption of security behaviors, in a large-scale experiment with 50,000 Facebook users (CCS’14, *honorable mention for the NSA’s Best Scientific Cybersecurity Award*) [5]. Fourth, as an intern at Google I participated in a cross-cultural study about how social influences affected the adoption of “secure” instant messengers (SOUPS’16) [6]. Finally, in an upcoming submission to ICWSM’17, I quantitatively modeled how people learn about and re-share breaking-news events about security and privacy breaches.

The key takeaways from all of this work are three-fold. First, **social influence is a key driver of security and privacy behavior change**. In my interview study, half of all reported behavior changes were due to social influence. In my exploratory quantitative analysis, social influence significantly affected (both positively and negatively) the adoption of all tested security tools. Finally, in my experiment, an announcement with social proof significantly increased the awareness and adoption of promoted security tools compared to a control with no social proof (14.4% versus 10.5% clicked the announcement, and 4.8% vs. 3.7% adopted one of the promoted security tools). Second, **people feel accountable for the security of their loved ones but have few options to act on these feelings**. Indeed, in my interview study, people reported that they did not talk much about security for fear of being boring, but would break their silence to warn, teach or learn from others they cared about. Third, **the design of a security tool affects its potential for social diffusion and most existing security tools are designed in a way that precludes their social spread**. From my interview study, we found that the single most common catalyst for security-related behavior change was observing other people engage in security behaviors or use security tools. From my quantitative analysis of how security tools diffuse through social channels, we found empirical evidence to support this finding: security tools that were more observable and socially inclusive (i.e., included friends in the process of providing security) were much more likely to spread through social proof. In contrast, security tools that were non-observable and non-inclusive were negatively influenced by social proof: i.e., we observed a negative correlation between exposure to friends who use a tool and likelihood to adopt the tool. Existing security tools like two-factor authentication, however, are designed to be both invisible and non-inclusive. Taken together, my work on social cybersecurity broadly suggests **we should be creating security and privacy systems that are more observable, inclusive and stewarded** in order to promote their use and social spread.

As a result of my work, tens of thousands of Facebook users are now aware of and use optional security tools like two-factor authentication. Facebook also now uses social cues to promote these security tools. In addition, this work has spawned a burgeoning field of research within usable security: social cybersecurity.

Inventing and Evaluating Novel Security and Privacy Systems

Ultimately, I model people’s behaviors and preferences to build new security and privacy systems that better match those behaviors and preferences. Here, I discuss a sampling of the security and privacy systems I have built: a context-aware authentication framework, a tool that helps people learn very strong passwords, and a socially-inclusive form of authentication.

Context-Aware Authentication

Different contexts naturally carry different risk profiles, yet despite the fact that most people spend most of their time in low-risk contexts (e.g., going back and forth between home and work), conventional security advice suggests that people should always use the type of complex, interruptive authentication strong enough for high-risk contexts. To address this mismatch, we developed CASA (context-aware scalable authentication, SOUPS’13) [7]. CASA assesses contextual risk through a Bayesian model of location and recency of device use and then varies how users are required to authenticate into their phones.

For example, in low risk situations (e.g., when one is at a familiar location and has recently used her phone), one may not be required to authenticate at all. In higher risk situations (e.g., when one has not recently used her phone and is at an unfamiliar location), one might have to enter a strong password. In a field-study evaluation, CASA reduced the number of explicit authentications participants required by 68% and motivated many participants to use mobile authentication more generally.

However, one issue with CASA was that while risk assessments were continuous, authentication strength was discrete—either no authentication, PIN entry or password entry. I next developed Autobiographical Authentication [8] (UbiComp’13, *best paper*) to better map on to this space of continuous risk assessments without requiring users to explicitly memorize any secrets. AutoAuth authenticates people based on their answers to questions about their day captured by their smartphone sensors and logs (e.g., “Who did you call around 4pm yesterday?”). One particularly unique attribute of AutoAuth is that authentication is not contingent only upon answering these questions correctly. AutoAuth also models expected memory lapses and incorrect answers. The output of AutoAuth, therefore, is not a binary “yes/no” decision but a confidence score based on the posterior probability that the attempting authenticator is the user given how she answered these questions. We tested AutoAuth in a field study over several weeks and found that users could reliably answer their own questions and that even many strong adversaries (e.g., those who had access to all of the correct answers) failed to authenticate.

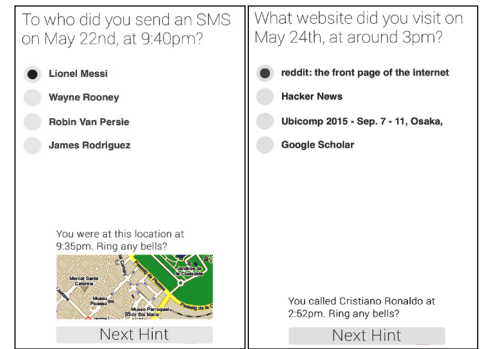
Facilitating the Memorization of Strong Secrets

Even laypeople sometimes require very strong authentication: for example, to protect online bank accounts, or to access encrypted drives with sensitive personal information. Accordingly, I developed two mnemonic training systems to incrementally help people learn 56.4 bit passwords composed of six words selected randomly from a set of 676 (e.g., ‘parcel cave turn head duck rhythm’) [9] (USEC’16). My story mnemonic trainer required participants to write two sentences, each of which contained three ordered words of their secret, while my peg-word mnemonic trainer required participants to write a sentence for each secret word, each also containing a non-secret word to assist participants in later recalling their secret words. I experimentally evaluated these different trainers in a randomized, controlled online experiment. I found that the story trainer performed best: 100% of over 50 participants remembered their secrets after 3 days and 84% remembered their secrets after 2 weeks. At the end of the study, a participant emphatically mentioned that my system helped him realize that he *was* capable of learning strong passwords.

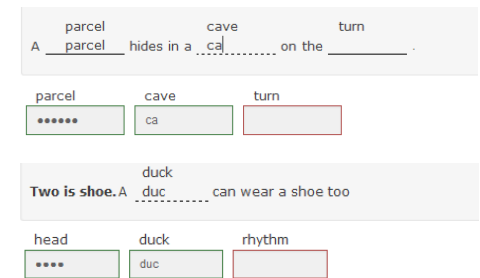
Socially-Inclusive Authentication

Authentication, today, is not inclusive: it is predicated on the notion of one secret per person. But, as my social cybersecurity work suggests, this notion is inappropriate for many small, local groups who share access to devices and spaces: Is it ideal, for example, for a father and son to each have their own secret password to in order for the shared family iPad to recognize them as distinct?


To make authentication more inclusive, I created Thumprint [10] (CHI’17, Qualcomm Innovation Fellowship project winner). Thumprint authenticates groups based on each member’s expression of a shared, 3-second knock on a surface instrumented with an accelerometer and microphone. As the secret knock is shared, group members need not maintain their own individual secrets. However, because individual expressions of the knock are variable, Thumprint can still identify individuals. Thumprint works by clustering expressions of the secret knock as expressed by actual group members during training. Authentication then becomes a question of comparing unlabeled authentication attempts against those learned clusters. I implemented Thumprint on Android and evaluated it over a multi-session lab study with over 30 participants. I found that (1) different people who enter the same thumprint can be reliably recognized, (2) people can consistently enter their thumprints over time-separated sessions, and (3) thumprints are reasonably secure against casual, motivated adversaries who know the exact knock and have ten attempts to replicate the secret. Thumprint is a promising first step towards the vision of more socially compatible cybersecurity tools and systems. In my future work, I hope to create and evaluate a suite of social cybersecurity systems like Thumprint.



Example AutoAuth questions. Users do not necessarily have to answer all questions correctly, as even their memory lapses are modeled.



My two mnemonic password trainers. Top: Users wrote a sentence linking three randomly assigned secret words; Bottom: Users wrote a sentence linking each secret word to a public “peg” that could later assist in recall.



Thumprint authenticates small, local groups and identifies individual group members through shared, 3-second secret knocks.

16 conference publications

3 paper awards

4 fellowships awarded

10 h-index

3 program committees
(WWW'17, ICWSM'16-17)

Recognition & Impact

Recognition: My work as a Ph.D. student has resulted in over sixteen publications in high-impact HCI and security venues (e.g., CHI, CSCW, UbiComp, SOUPS, CCS, ICWSM). I have earned three paper awards: a best paper award at UbiComp 2013 (top 1% of submissions), a best paper honorable mention at CHI 2016 (top 5% of submissions), and an honorable mention for the NSA's Best Scientific Cybersecurity Paper Award in 2014 (top 3 out of 50 anonymous nominations among all published security papers in 2014). I have also won four prestigious academic fellowships (the NDSEG, NSF EAPSI, Qualcomm Innovation, and Stu Card fellowships) and was a finalist for two others. I was one of five students nominated for the inaugural John Karat Usable Privacy and Security Student Research Award. I have been invited to serve on three technical program committees. And, my work on self-censorship and social cybersecurity has been extensively covered by the press, including features on, for example, The Atlantic, Slate, The Financial Times and Mashable.

Impact: According to Google Scholar, my work has been cited over five hundred times (with an h-index of 10). In addition, my work on self-censorship, social cybersecurity and context aware authentication has been incorporated into many classroom curricula and textbooks on usable privacy and security. In terms of impact outside of academia, my social cybersecurity work was selected for inclusion in the SERENE-RISC Smart Security Network's quarterly cybersecurity knowledge digest: a publication that summarizes influential academic work on cybersecurity for policy makers and industry practitioners. My work on self-censorship and social cybersecurity has helped change Facebook's approach to security and privacy: Facebook now uses the metrics and approaches I introduced to evaluate and promote new privacy and security tools. In fact, my experiment on using social proof to increase security sensitivity resulted in thousands of additional Facebook users using or becoming aware of optional-use security tools provided by Facebook.

Future Research Agenda

As I look forward, I intend to use the expertise I have developed in my Ph.D. to implement usable and social solutions to some of the most pressing security and privacy problems of the present and near future. Three domains of interest are:

Scalable security and privacy: How can we make security and privacy scale? I want to develop computational and social approaches that help people keep up with the ever-increasing security and privacy decisions they must make. Computationally, I plan to use advancements in deep learning to create intelligent, personalized security assistants. These assistants will scour vulnerability databases and popular press for security and privacy attacks that affect their owners and make or suggest changes automatically (e.g., forcing a password reset). Socially, I plan on implementing an expert-sourced social network where laypeople can follow trusted experts and loved ones who can send personalized, context-specific notifications about security to their followers in the wake of relevant security breaches.

Usable security for the Internet of Things: How can we make IoT secure? Connected, pervasive smart things (e.g., the Nest thermostat and electronic smart locks) are here. These devices provide tremendous utility, but are insecure. The Mirai IoT botnet that takes over and enrolls insecure IoT devices to perform DDoS attacks, for example, is allegedly responsible for many service failures on the Internet today (including that of Twitter in October of 2016). In my future work, I intend to create usable security systems and controls suitable for IoT devices. I am particularly keen on applying the social cybersecurity principles of observability and inclusivity to facilitate the viral spread of the approaches I develop.

Security for and with programmable matter: What should security look like for the the computing ecosystems of tomorrow? For example, physical computing interfaces that can re-configure physical matter programmatically (e.g., through shape-memory alloys or swarms of robots) are on the horizon. For the most part, however, these interfaces are vaunted for their interaction affordances: their security, and their implications for security, remain largely unexplored. In future work, I am interested in both inventing approaches to secure these futuristic interfaces as well as exploring how they can be used to create novel, intuitive security controls.

References

- [1] **Sauvik Das** and Adam Kramer. *Self-Censorship on Facebook*. ICWSM'2013.
- [2] Manya Sleeper, Rebecca Balebako, **Sauvik Das**, Amber McConohy, Jason Wiese, and Lorrie Cranor. *The Post That Wasn't: Examining Self-Censorship on Facebook*. CSCW'2013.
- [3] **Sauvik Das**, Tiffany Hyun-Jin Kim, Laura Dabbish and Jason I. Hong. *The Effect of Social Influence on Security Sensitivity*. SOUPS'2014
- [4] **Sauvik Das**, Adam Kramer, Laura Dabbish and Jason I. Hong. *Increasing Security Sensitivity with Social Proof: A Large Scale Experimental Confirmation*. CCS'2014.
- [5] **Sauvik Das**, Adam Kramer, Laura Dabbish and Jason I. Hong. *The Role of Social Influence in Security Feature Adoption*. CSCW'2015.
- [6] Alexander de Luca, **Sauvik Das**, Martin Ortlieb, Ben Laurie and Iulia Ion. *Expert and Non-Expert Attitudes Towards (Secure) Instant Messaging*. SOUPS'2016.
- [7] Eiji Hayashi, **Sauvik Das**, Shahriyar Amini, Jason Hong and Ian Oakley. *CASA: Context-Aware Scalable Authentication*. SOUPS'2013.
- [8] **Sauvik Das**, Eiji Hayashi and Jason Hong. *Exploring Capturable Everyday Memory for Autobiographical Authentication*. UbiComp'2013.
- [9] **Sauvik Das**, Jason Hong and Stuart Schechter. *Testing Computer-Aided Mnemonics and Feedback for Fast Memorization of High-Value Secrets*. USEC'2016.
- [10] **Sauvik Das**, Gierad Laput, Chris Harrison and Jason Hong. *Thumbprint: Socially-Inclusive Local Group Authentication Through Shared Secret Knocks*. CHI'2017 (to appear).