1 Research Program

I work on understanding and finding solutions to problems that arise in sociotechnical systems as they become an increasingly invisible and indispensable part of everyday life. Sociotechnical systems involve people, technology, and information; these parts all interact and influence each other, and create complex behaviors and outcomes. They have great potential to help people and improve their lives; however, they also have the potential for harm. In particular, I focus on sociotechnical systems in which the technology and information components constitute a *black box* for the user: a "system whose workings are mysterious; we can observe its inputs and outputs, but we cannot tell how one becomes the other" [15].

For example, Google's Nest thermostat aggregates motion sensor data across time and from multiple households to infer more accurately whether users are not at home, enabling the system to conserve energy and save money by entering "away" mode sooner [13]. The conditions triggering away mode are based on data about household activity and usage patterns that the user cannot directly access or evaluate for accuracy. The Facebook News Feed ranking algorithm helps users cope with information overload by using data about users' past interactions with content to predict which new posts they are more likely to want to see [3]. The details about how the algorithm works and why it chooses to rank posts the way it does are not available to the user. And, many computer security tools and mechanisms, like automatic software updates and spam email filters, are designed with the assumption that systems are more secure when users are prevented from making security-related decisions. These tools are intended to hide the complexity from the user by taking security-related actions without any human intervention; the user is not always even aware that an update has taken place or an email has been filtered.

As computing becomes more embedded in the environment, systems become more like black boxes. This is by design; the easiest computer to use is one you don't even know you are using. In 1991, [redacted] called this *disappearance* [28], predicting that technologies would "weave themselves into the fabric of everyday life". Disappearance means that not only is the *computation* that transforms input to output invisible, but the *inputs* to the system may be invisible also. In one sense, this is a good thing, because it means that computing is increasingly capable of collecting data and making decisions on the user's behalf, approaching [redacted] vision of man-computer symbiosis [12]: offloading tasks that computers are good at and freeing humans to focus on problems computers are unable to solve. However, now that computing is becoming integrated into all sorts of everyday activities, users are starting to experience different kinds of harm that come along with these benefits – harms that system designers do not yet know how to prevent.

For example, [redacted] uses data collected from each household to send monthly reports to users informing them about how their energy consumption matches up against other households in their area [1, 2]. This social comparison is intended to motivate users to be more energy conscious. But, the same information could also be used to shame people who are perceived as wasteful. (See the Twitter hashtag #droughtshaming for examples of people using public shaming to enforce water conservation norms.) The underlying data may seem to users to be non-sensitive and harmless on its own; however, aggregation produces "derived data" consisting of new insights that are not obvious and can be surprising, unsettling or harmful when used for purposes users do not expect [17].

On Facebook, the News Feed algorithm makes choices to promote certain posts which necessarily cause other posts to be less visible. This helps users see more content that matches their preferences, but it also creates a "threat of invisibility" [7]: users interact with visible posts which then become more popular, but posts that are not displayed to users cannot receive the attention necessary for the algorithm to rank them highly. It can also result in a "filter bubble" [14], or a progressive decrease in the diversity of content a user is exposed to over time. Artificially biased access to information online can prevent minority opinions from being expressed and deliberated [6], and may mean that people miss content and opportunities for social interaction they would have wanted to see [9, 19].

Automatic software updates install security patches that fix vulnerabilities in the background, without the user necessarily being aware that this is happening. Keeping systems updated this way is good for security and also easier for users, who tend to want to delegate security tasks to others anyway [8]. But, updates can also make unexpected changes to functionality or user interfaces that people have grown accustomed to [23]. This can alter users' beliefs about whether they should allow future updates to take place, especially if they share their experiences with each other [22]. Social sharing of security-relevant information can help people avoid threats, but it can also cause increased risk that computing devices will be compromised.

Some might argue that people use black-box-like systems any time they use computing technology. For example, editing a document with Microsoft Word might actually resemble using a black box system, since most Word users

are not software developers and would not be able to understand the underlying code even if they had access to it. However, the basic functionality of Microsoft Word stays the same no matter how many documents a user edits. A key characteristic of black box sociotechnical systems is that they are constantly evolving. Facebook's news feed algorithm shapes conversations online based on what kinds of posts are shown [11], and thereby affects characteristics of new posts which are subsequently ranked for display by the algorithm. As smart home technologies become more diverse and interoperable (e.g., "Works with Nest" integration), additional inferences and derived data about users and their families can be generated from the new inputs. And it is hard for even security experts to keep up with the changing landscape of security threats.

In other words, a black box sociotechnical system can produce different outputs and outcomes for end users as the input data change. Black box sociotechnical systems are challenging for HCI researchers and designers to create and study because fundamental aspects of what they do change over time as the systems evolve. These conditions create *sociotechnical bugs* – incorrect, unexpected or unwanted behaviors that result from the interactions between people, information and technology in the system – that can constrain access to information, compromise security, and violate users' privacy.

My research program is focused on analyzing interactions between people, technology, and information in sociotechnical systems like these, in order to contribute generalizable scientific knowledge that characterizes how the harmful aspects of these systems arise. My work also identifies and evaluates solutions, such as ways to minimize the effects of bias, and help users manage their privacy and make better security-related decisions. I am currently the Principal Investigator on three projects funded by grants from the US National Science Foundation that are focused on these issues. Since 2011, I have been awarded over $1.2 million to carry out this research.

1.1 Project One: Managing Privacy of Derived Data

I recently began work (in Fall 2015) on a project in which I am studying the privacy implications of derived data in ubiquitous computing systems involving networked sensors. Sociotechnical systems like these involve interaction and interdependence between human, computational, and information components, which creates complex behaviors and outcomes. For example, my Fitbit activity tracker is able to more accurately infer whether I am asleep by analyzing data from all the other users who choose to wear a Fitbit while sleeping, walking, and riding in cars over potholes and bumpy railroad tracks. Therefore, by contributing my data, I am helping to improve how everyone's Fitbit works. But, using my Fitbit also enables potentially sensitive inferences about myself and others. Sensor data like accelerometer readings and GPS locations collected by an activity tracker device and app and aggregated across users could be used to infer information like how many times a day the user goes to the bathroom, or the likelihood that the user is a parent of young children. People adhere to social rules and norms for offline privacy-related behaviors [10, 16]. But because interdependence in ubiquitous computing is invisible, people can't coordinate with each other on privacy-related matters in those systems. They cannot currently develop rules or norms for deciding which uses of derived data are acceptable, such as those that make the system perform better, and those that are unacceptable, like making sensitive inferences that may be unrelated to system operation.

This project aims to identify norms for the use of derived data in a ubiquitous computing system, and design and evaluate a mechanism for coordination among users of the system such that they can jointly manage the derived data as a common-pool resource. Using the system creates the resource: the derived data. Like common-pool resources in social-ecological resource systems, derived data is *non-excludable* in that it is difficult for users to prevent others from joining the system and contributing data; one user cannot prevent another from buying and using a Fitbit. It is also *subtractable*, in that as more users contribute more sensor data, negative privacy-related inferences about users from a larger and more diverse dataset become more likely. This represents a shift in the way we think about information privacy problems from the self-management model to a collective governance model, which may lead to new avenues of research and development.

The work to identify norms is currently underway. In an initial semi-structured interview study, I found that activity tracker users were more aware of the derived data aspects of the system (e.g., step counts) than the sensors that are actually measuring their activity (e.g., accelerometers, GPS). They were also more positive about new kinds of derived data that would clearly help them gain additional awareness of activities they would like to monitor, or that would provide a positive benefit to themselves and others. This indicates that rules or norms about benevolence [5] might be applicable to derived data. However, without being aware of what is actually being measured, it was challenging for these users to imagine what other kinds of inferences might be possible. These findings show how the sensor inputs to the activity tracker black box are invisible to users, and they are only aware of the derived data outputs related to their reasons for adopting the activity tracker in the first place.

1.2 Project Two: Algorithmic Curation in Social Media

Sociotechnical systems provide access to ever-increasing quantities of information online. These systems often implement algorithmic curation: automated selection of what content should be displayed to users, what should be hidden, and how it should be presented. Virtually every Internet user who reads online news, visits social media sites, or uses a search engine has encountered algorithmic curation at some point, in many cases without even realizing it [19]. Personalization algorithms are a necessary and beneficial part of the infrastructure; but, users' own awareness and understanding of what the algorithms are doing could cause them to adjust their behavior in ways the algorithm isn't designed to handle properly, with consequences for individuals and the system as a whole.

For example, the Facebook News Feed records data about users' interaction with content, and uses algorithms that aggregate information collected from many users when making inferences about users' preferences and deciding how posts should be ranked for display on a per-user basis. In this project, I am investigating the relationship between social behavior and algorithmic curation, to find ways to help users and system operators identify, measure, and reason about system-level effects that arise from individual-level design interventions. In proposing this project, I was one of the first to use the phrase "algorithmic curation" to describe the effects of filtering and ranking algorithms. I have identified a disconnect between the motivations of content producers for creating posts and the interests of consumers in what posts they prefer to see [21]. I also have discovered that despite the black box characteristics of the Facebook News Feed ranking algorithm, users form mental models for how it works, and these mental models shape their behavior in the system [19]. And, my work describes how the algorithm might affect closeness of users' real-world relationships as a result of the way that it ranks posts for display [18]. In a paper under review, I report results of an analysis of a dataset released by researchers at Facebook in conjunction with a recent paper [4], and a simulation study based on the data. I found that the ranking algorithm interacts with user scrolling behavior to determine which posts users see, which means the influence of the algorithm is strongest at the "top" of the user's News Feed. Future work on this project includes applying these findings to the implementation of agent-based models that will allow us to examine the effects of different kinds of content prioritization schemes on the information users are exposed to.

1.3 Project Three: Mental Models of Computer Security

Many security vulnerabilities are the result of a human choice to act (e.g., clicking on a link in a phishing mail) or not act (e.g., not installing updates or security patches). My security work focuses on how people make choices about protecting themselves and their computing devices from security-related problems and threats that are very hard for them to be aware of and to understand. I have been working with a collaborator since 2011 on a third project in which we are studying the beliefs and behaviors of computer users who are not experts in computer security. The contexts in which security choices and behaviors must take place are expanding, and both the need to protect one's devices and information, and the complexity of that task, are increasing for non-expert users. The goal of this project is to find out how what non-experts learn, know, and believe is related to their security behavior, and to develop ways to take advantage of what people already know and do, and how they already learn about security, to help them behave more securely [25].

We have identified social aspects of learning about security [22] and ways that mental models for software updates based on past experience can actually cause computing devices to be *less* secure [23]. We also found that strategies for removing the user from the security loop, like automated software updates, can prevent learning from taking place and thereby cause users to be less able to make educated decisions [24]. To better understand differences in the security-related information available to users for learning about security, we compared three informal sources of computer security information, and found that security information from peers usually focuses on *who* conducts attacks, but expert advice focuses instead on *how* attacks are conducted. These differences may prevent users from associating protective measures with the generalized threats they are concerned about, and talk about with each other [20].

We are currently analyzing data we collected in a six-week field study using a survey instrument we developed to measure beliefs and self-reported behaviors related to computer security [26], and a software tool installed on participants' computers designed to measure a number of security-related behaviors. This tool recorded log data as participants used their computers, and detected behaviors such as when they updated their operating system, or clicked on a link in an email, or entered a password into a web page. Our analysis is focused on correlating the behavioral data with the survey results to discover how different mental models, and self-reports of behavioral intentions, are associated with actual behaviors. Our first paper from this unique dataset, published in 2016 [27], reported evidence that password composition policies interact with users' limited memory capacity: participants

tended to re-use their most frequently-entered password, which was also their most complex password, relatively speaking. This means that organizations that force frequent authentication are effectively training users to memorize those passwords, which are then more likely to be re-used on other accounts because they are easier for users to recall. This points to an unexpected and previously unknown interdependence between accounts; and, it is one way that sociotechnical security shows characteristics of a black box system.

2 Instruction and Advising

I have developed two new interdisciplinary courses for the Department of Media and Information, focusing on human computer interaction (master's level, MI845) and digital privacy (undergraduate level, MI239). In both courses, I combine hands-on activities, discussions of current topics, and thinking and writing assignments to help create a respectful and intellectually engaging learning environment in the classroom. For example, I created a time-use diary assignment for my digital privacy course, in which students record all of their activities for two days at 15-minute intervals, along with metadata about who they are with, what technologies they are using, and how public or private they feel at the time. They use their diaries for analysis and reflection in a writing assignment and in class discussion about their own beliefs and observations regarding privacy in their daily lives. I help students in this course form their own opinions about privacy in sociotechnical systems, by preparing multiple ways of explaining and illustrating the concepts, and by listening to students' concerns. This creates a supportive environment where students feel comfortable asking questions, voicing confusions, and talking about their own experiences, which helps me tailor my approach to a unique and diverse group of students each time I teach the course. This course began as a special topics course, and is now a regular yearly offering and fills a requirement in the department's new undergraduate curriculum.

I have also created a master's-level overview course about Human Computer Interaction (HCI). The course identifies themes in the field, integrates historical trends and current topics, and forecasts them into the future. For example, we watch excerpts from [redacted] 1968 demo of the very first computer mouse, discuss Fitts' Law which predicts the time it takes a human to move the mouse and click on a target, and expand that to consider interacting with touchscreens and gestural interfaces. Then, we imagine and discuss what the user interfaces of 20 years from now might look like, and what new forms of interaction they might enable. I designed the major assignment in this course to help the students develop writing and argumentation skills: they complete a writing project in which they brainstorm, plan, write, conduct peer evaluation, and revise an essay about a topic important to them that is also related to HCI. I combine the final essays into a single volume, which I then print as a library-quality paperback using Michigan State's Espresso Book Machine. Each student receives a copy of the book at the end of the semester, which makes the piece they are writing feel more real than a typical term paper assignment.

In addition to my teaching in the classroom, I also train and mentor graduate and undergraduate researchers. I have previously worked with 10 different MIS PhD students as research assistants, and have been a dissertation committee member or chair for two of those students. This Fall (2016), I have been assigned an incoming Information and Media (IM) PhD student as an advisee for the first time. Most of the students who matriculate in the IM program are not interested in sociotechnical systems research or HCI, and this has presented challenges related to staffing, leadership, mentoring, and project management on my sponsored projects. It has also presented opportunities to experiment with different ways of organizing the work. For example, I work extensively with undergraduate research assistants during the academic year and in the summer in the research lab I founded and lead, the BITLab (Behavior, Information and Technology Lab). During 2012-2015 I co-organized a very successful NSF-funded REU (Research Experiences for Undergraduates) program in which we recruited 5-7 exceptional students each year for 10-week summer internships. These undergraduate research assistants join active research projects as junior colleagues, and participate in all aspects of the research. Several of my publications include undergraduates I have worked with as co-authors [21, 24, 27], and many of the lab's undergraduate alumni have gone on to graduate school or careers in the technology industry. I find this work to be productive and rewarding, and plan to continue the REU program in the future.

3 Citizenship

In the Fall of 2012, which was the beginning of my second year at Michigan State, the department chair at the time asked me to become the Director of the department's Master's program. I began sitting in on committee meetings that semester, and formally took over beginning in January 2013. I was tasked with three goals: to resolve a long-standing conflict among the faculty regarding exit requirements for Masters students who choose not to do a thesis or a project to complete their degree, to analyze the program and make recommendations for longer-term

improvements, and to increase enrollment in the program. I accomplished the first two goals; however, working on these issues negatively affected my research productivity. I stepped down at the end of December 2013, and have worked hard since that time to make up for the gap in my publication record and continue to secure funding for my research program.

In addition to that leadership role, I have also worked to improve the quality of the IM PhD program. I have served on the Media and Information Department PhD committee, on a task force committee assembled in 2015 to analyze the PhD program and make recommendations, and I am also co-organizer of the yearly IM PhD Research Symposium. This is an event that a colleague and I initiated in the Spring of 2012 as an opportunity for students to showcase their work (and meet their practicum presentation requirement), while at the same time providing a venue for both students and faculty to become more familiar with the research that is going on in the IM program. Another objective of this activity is to promote a culture of research excellence among our doctoral students by providing incentives for good work. The symposium provides opportunities for recognition and feedback, and in addition, I contribute $750 each year from the AT&T endowment funds provided to me as part of my appointment at Michigan State, for awards given to the best "senior" and "junior" student presentations. Feedback from both faculty members and students about the Symposium has been very positive.

I have also served the department community by providing support for students and faculty members affiliated with the BITLab in the form of research infrastructure. This includes three components: space, technology, and personnel. I have transformed the research space that I was allocated from a jumbled mess of leftover furniture and abandoned equipment into a professional workspace, primarily using my AT&T endowment funds and my spare time. I have also used AT&T funds to provide computing equipment, servers, office supplies, and subscriptions to file sharing and backup services. I feel strongly that an investment in research infrastructure is also an investment in the people who work there, and that providing one's employees with the resources and freedom they need to do their jobs well is an important way to show them how valuable they are. I have also used AT&T funds to pay research-related expenses to help three MIS PhD students complete their dissertations, only one of whom was my advisee.

Finally, I am active in the wider academic community as a reviewer and program committee member. I have served on three NSF panels (and I have also reviewed single proposals as an ad-hoc reviewer), 8 conference program committees for HCI and security/privacy related venues, and I have been a reviewer for CHI annually since 2006 and CSCW since 2008. I have also reviewed for 11 other HCI-related journals and conferences. I have been recognized for writing exceptional reviews four times, for UIST 2014, CHI 2015 and 2016, and CSCW 2016.

4 Conclusion

I meet and exceed the requirements for tenure and promotion to Associate Professor in the Department of Media and Information at Michigan State University. I have a strong record of publishing in the human-computer interaction and usable privacy and security communities. I have published 19 peer reviewed conference publications and journal articles (12 since Fall 2011 when I began my appointment at MSU), which together have over 470 citations. I take a leadership role in all of my research projects, as indicated by the fact that I am the first or second author on all but one of my publications, and solo author on four. I have also been very successful at securing external funding for my research. Since 2011, I have received three research grants from the National Science Foundation as primary investigator (PI), totaling $1.25 million. I currently have a 100% success rate at the NSF. I receive strong student ratings of my teaching performance, averaging 1.58 for Instructor Involvement (on a 1 to 5 scale where 1 is 'Superior' and 5 is 'Inferior') across all of the courses I have taught. And, I am an active citizen both for my university and the wider academic community.

My research program is focused on problems that arise at the intersection of people, information and technology in sociotechnical systems with black box characteristics. I have found that benevolence norms may apply to sensitive derived data, that algorithmic curation can impact users' ability to maintain relationships in a social network, and that interdependence exists even in security systems that try to remove the user from the loop. As data collection and computation become more embedded in everyday infrastructure, and as *disappearance* spreads to new contexts, new systems that involve sociotechnical black boxes will continue to arise. My work contributes new understanding of these systems, and insight into how designers might begin to solve them. My future work will incorporate my findings and experience from these projects as I continue to investigate new instances of black box sociotechnical systems.

References

[1] About the Nest Home Report, November 2014. [redacted]

[2] How does the Nest Leaf work?, November 2014. [redacted]

[3] [redacted] News Feed FYI: A Window Into News Feed, August 2013. [redacted]

[4] [redacted] Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348(6239):1130–1132, 2015.

[5] [redacted] Activation of social norms in social dilemmas: A review of the evidence and reflections on the implications for environmental behaviour. *Journal of Economic Psychology*, 28(1):93–112, January 2007.

[6] [redacted]. Breaking the filter bubble: democracy and design. *Ethics and Information Technology*, 17(4):249–265, December 2015.

[7] [redacted] Want to be on the top? Algorithmic power and the threat of invisibility on Facebook. *New Media & Society*, 14(7):1164–1180, 2012.

[8] [redacted]. Security in the wild: user strategies for managing security as an everyday, practical problem. *Personal and Ubiquitous Computing*, 8 (6):391–401, 2004.

[9] [redacted] "I always assumed that I wasn't really that close to [her]": Reasoning about invisible algorithms in the news feed. In *CHI '15: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 153–162, 2015.

[10] [redacted]. *Behavior in Public Places: Notes on the Social Organization of Gatherings*. The Free Press, New York, NY, 1966.

[11] [redacted] Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, 111(24):8788–8790, 2014.

[12] [redacted] Man-computer symbiosis. *IRE transactions on human factors in electronics*, 1:4–11, 1960.

[13] [redacted]. The World's Best Thermostat Just Got Better, October 2012. URL `http://www.slate.com/articles/technology/technology/2012/10/nest_thermostat_the_ingenious_heating_and_cooling_system_keeps_getting_smarter_.html`.

[14] [redacted] *The Filter Bubble: What the Internet is Hiding from You*. The Penguin Press, 2011.

[15] [redacted] *The Black Box Society: The Secret Algorithms That Control Money and Information*. Harvard University Press, 2015.

[16] [redacted] *Boundaries of Privacy: Dialectics of Disclosure*. State University of New York Press, Albany, NY, 2002.

[17] President's Council of Advisors on Science and Technology. Big data and privacy: a technological perspective. United States Executive Office of the President, May 2014.

[18] [redacted]. Examining user surprise as a symptom of algorithmic filtering. *International Journal of Human-Computer Studies*, 98:72–88, 2017.

[19] [redacted] Understanding User Beliefs About Algorithmic Curation in the Facebook News Feed. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '15, pages 173–182, 2015.

[20] [redacted] Identifying patterns in informal sources of security information. *Journal of Cybersecurity*, 1(1):1–24, 2015.

[21] [redacted] Gap Between Producer Intentions and Consumer Behavior in Social Media. In *ACM International Conference on Supporting Group Work*, GROUP '12, pages 249–252, 2012.

[22] [redacted]. Stories as informal lessons about security. In *Proceedings of the Symposium on Usable Privacy and Security*, SOUPS '12, pages 1–16, 2012.

[23] [redacted] Betrayed by Updates: How Negative Experiences Affect Future Security. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, pages 2671–2674, 2014.

[24] [redacted] Out of the Loop: How Automated Software Updates Cause Unintended Security Consequences. In *Proceedings of the Symposium on Usable Privacy and Security*, SOUPS '14, pages 89–104, 2014.

[25] [redacted] Influencing mental models of security: a research agenda. In *Proceedings of the New Security Paradigms Workshop*, NSPW '11, pages 57–66, 2011.

[26] [redacted] Too Much Knowledge? Security Beliefs and Protective Behaviors Among United States Internet Users. In *Proceedings of the Symposium on Usable Privacy and Security*, SOUPS '15, pages 309–325, 2015.

[27] [redacted] Understanding password choices: How frequently entered passwords are re-used across websites. In *Symposium on Usable Privacy and Security*, pages 175–188, 2016.

[28] [redacted]. The computer for the 21st century. *Scientific American*, 265(3):94–104, 1991.