Eric Zeng Research Statement

As computers have become a ubiquitous part of daily life, people increasingly encounter threats to their computer security and privacy. Everyday threats like online scams and misinformation, pervasive data collection for targeted advertising, and abuse of smart devices for interpersonal surveillance harm people and weakens trust in technology. My research investigates how to **design technology and inform policy** to address **computer security and privacy issues that impact end users' experience**. In my work, I use three complementary approaches to develop these solutions:

1. **User research:** I inform the design space for technical and policy solutions by conducting user research on peoples' security and privacy needs and the sociotechnical context that underlies those needs. For example, my studies of smart home users uncovered that privacy threats from other people in the home were often more salient to users than threats from hackers or device vendors, and that smart home platforms often lacked features like access controls that could mitigate these threats [8,10].
2. **Internet Measurement:** I provide transparency into widespread security and privacy issues that impact users by conducting large-scale empirical measurements. For example, I conducted measurement studies of online advertising to identify the websites, advertisers, and platforms that were disproportionately responsible for serving misleading and deceptive advertisements [3,5,6,7].
3. **Building Tools and Interventions:** I evaluate the effectiveness of user-facing tools and interventions by building high-fidelity prototypes and testing them in rigorous experiments with users. For example, I ran an experiment to evaluate the effectiveness of security notifications for HTTPS misconfigurations with over 20,000 websites [9] and I conducted an in-home study to evaluate whether novel access controls for smart homes could address users' concerns about control and surveillance [8].

To enable my research, I use a breadth of qualitative and quantitative methods ranging from interviews to field experiments, I implement infrastructure and prototypes, and I build on state-of-the-art techniques from across computer science including natural language processing [5], explainable AI [1], cryptography [11], and formal methods [2].

My work has been published at top conferences in security and privacy, human-computer interaction, and internet measurement, such as USENIX Security, CHI, and IMC (where my work was the runner-up for Best Paper Award). My research on smart homes was among the first to identify interpersonal privacy as a key concern of users and laid a foundation for research on safety for bystander users in smart homes. My work on online advertising has been of significant interest to stakeholders in industry and government, and I have presented my findings at venues including the US Federal Trade Commission's PrivacyCon and the Ad-Blocker Dev Summit. Additionally, my measurement tools, datasets, and taxonomy have been used by other researchers to audit online advertising platforms [12].

Past Work

Next, I will illustrate how I use these approaches to address complex sociotechnical security and privacy challenges facing users, through two of the research topics I worked on during my PhD.

Deceptive Online Advertising

Deceptive and misleading online advertising can negatively impact users' security, privacy, and overall experience. For example, ads can promote scams, deceive people into installing spyware, and broadly irritate, misinform, or waste users' time. Anecdotally, such ads remain common on the web, but online ad platforms are notoriously opaque, making it difficult to assess how large of an impact these ads have on users, and which actors in the online advertising ecosystem are responsible. To systematically study this problem, I developed a qualitative taxonomy to describe what users dislike about online advertising [6], and I conducted large-measurement studies of online advertising on the web to identify the websites, advertisers, and platforms that were disproportionately responsible for serving misleading and deceptive advertisements, in case studies of misinformation websites [7], political advertising [5], and the economics of ad auctions [3].

Creating a User-Centric Taxonomy of Problematic Advertising [6] To enable systematic, quantitative audits of harmful ad content, I conducted a user study to define a taxonomy of users' perceptions of ads. First, I conducted an initial qualitative survey with 60 people to describe what they liked or disliked about a sample of 30 ads and synthesized their responses into a taxonomy of 15 reasons for liking or disliking ads. Second, to identify which kinds of ad content elicited these reactions, I surveyed 1025 participants and asked them to label a dataset of 500 ads with their perceptions of the ads using our taxonomy, assigning 10 participants to each ad to capture subjective disagreements. Our taxonomy of ad perceptions included categories such as "deceptive", "distasteful", "clickbait", "ugly" and "not relevant". We found that people perceived ads that sell supplements, health products, and software downloads to be particularly deceptive and "clickbaity". Additionally, we found that participants reacted negatively to "native advertising", a deceptive format of ad that mimics the

look and feel of content on the page, common on news websites. I publicly released our taxonomy and dataset[1], which has been used in recent studies auditing the targeting of deceptive advertising [12].

Measuring Problematic Advertising on the Web [5,7] Though problematic advertising is known to pervade the web, it is difficult to pinpoint which advertisers, websites, and ad platforms are primarily responsible. To better inform technical and policy approaches to reducing users' exposure to problematic ads, we conducted two measurement studies to quantify the prevalence of problematic advertising.

To conduct our measurements, I developed an infrastructure for scraping ads from the web. I built a web crawler based on Chrome's browser automation library (Puppeteer) to automate ad detection, scraping, and clicking on ads. Our infrastructure was used to scrape over 1.4 million ad screenshots and landing pages across several projects. To scale our qualitative analyses of the ad content we scraped, we built an automated analysis pipeline that used components like language model-based text classifiers to perform initial filtering and indexing of ads by their topic.

In our first case study [7], we investigated the phenomenon of clickbait advertising on news websites: were they more prevalent on unreliable news sources? And which ad networks were responsible? Using our ad crawling infrastructure, we collected ads from 100 most popular news websites and misinformation websites. We manually labeled 5414 ads, identifying categories of misleading content, and types of ad networks. We found that while mainstream news and misinformation websites hosted similar quantities of misleading ads, specific ad networks like Taboola, Outbrain, and RevContent were disproportionately responsible for serving misleading advertising. Our results suggest that stronger regulation of deceptive techniques like native advertising and more consistent content policies across ad networks is necessary to improve the overall quality of advertising on the web.

In our second case study [5], we investigated misleading political advertising during the 2020 U.S. Elections. Using our crawling infrastructure, we crawled ads on news and misinformation sites from six locations across the U.S. in toss-up, Republican and Democratic leaning states (through VPNs). We identified several categories of deceptive ads containing political content. Our findings reveal an ecosystem of advertisers that leveraged political controversies for monetary gain and polluted the information ecosystem with misleading political narratives. Of most concern were misleading political polls or petition ads, whose end goal were to get users to subscribe to political mailing lists, which solicited campaign donations, advertised questionable products, and spread political misinformation. Based on these findings, we made policy recommendation to ad platforms, such as expanding their definition of political ads beyond those from political action committees and banning certain deceptive political advertising techniques such as fake polls. This work was recognized as a runner-up to the Best Paper award at IMC 2021.

I publicly released the measurement tools that powered these studies on GitHub[2] and continue to maintain them, to enable other researchers to conduct audits and measurements of ad content on the web.

Interpersonal Privacy and Security in Smart Homes

As smart homes became increasingly popular, the research community initially focused on software security vulnerabilities in individual devices and app platform vulnerabilities. However, little research had characterized security and privacy concerns of smart home users themselves. To bridge this gap, I conducted some of the first studies with modern smart home users to learn about what security and privacy issues were most salient to users, and how smart home platforms could be designed to address their concerns.

Understanding Smart Home Users' Security and Privacy Concerns [10] To develop an initial understanding of smart home users' concerns, threat models, and protective behaviors, I conducted an interview study with fifteen smart home users. I found that while users were aware of common security and privacy threats like vulnerable devices or sensor data collection, the concerns more relevant to their lives were privacy concerns about how other people in the household used the smart home. For example, participants were concerned that other users could use the home to surveil the other household members using cameras and device logs or restrict access to devices or apps necessary to control the home. These risks were exacerbated in households where one tech-savvy household member (primary user) was responsible for setting up the smart home and other household members used the smart home more passively (incidental users), because the primary user could use their access and knowledge to exercise power and control over other household members.

---

[1] https://badads.cs.washington.edu/ad-perceptions-dataset/table.html

[2] https://github.com/UWCSESecurityLab/adscraper

Designing and Evaluating Access Controls for Multi-User Issues [8] Our interview study indicated that a key sociotechnical challenge to address in smart home design was mediating conflicts and tensions between users relating to control, authorization, and privacy. To address these issues, we designed a set of fine-grained access controls for smart home platforms, with the goal of allowing users to codify rules for respectful usage of the smart home. I built a prototype smart home platform that included role-based access controls; proximity-based access controls, which restrict users from controlling devices that they are not physically close to, using Bluetooth beacons to localize mobile devices; supervisory access controls, which allow children to control the smart home when parents are nearby to supervise; and reactive access controls, which allow users to "ask for permission" when policies are too rigid.

We evaluated our prototype in a month-long in-home user study with seven households. We found that access controls were useful in situations with clear trust boundaries, such as restricting access for domestic workers, or in private rooms. However, due to the complexity of setting up access controls, the usability barrier often outweighed the (minimal) benefits to security and privacy for in high-trust households. In such cooperative households, existing norms and trust usually prevented misuse of the smart home, obviating the need for strict access control policies. For example, in two households, children were trusted to follow rules regarding control of lights, locks, and doors, regardless of whether it was via the smart home or the physical controls. Our findings suggest that to minimize tensions and conflicts between household members, smart homes should be designed to cultivate and promote positive household norms, rather than supplanting them with technical access controls.

Ongoing and Future Work

Looking forward, I hope to both address the risks of emerging technologies for end users' security and privacy, and leverage emerging technologies to build better tools to protect users. I am interested in a broad set of topics, but two directions I hope to pursue in the medium-term are 1) auditing targeted advertising for harms against users, and 2) developing machine-learning based tools for assisting users in making security decisions.

Building Tools for Auditing Targeted Advertising Researchers, regulators, journalists, and policymakers are interested in investigating the harms of targeted advertising: for example, are targeted advertisements harming people from sensitive demographic groups, through discrimination or deception? However, audits of targeted advertising must overcome many technical challenges. A large amount of data is required to make statistical inferences about the presence of targeting. But due to the opacity of the advertising ecosystem and the individualized granularity of targeting, it is difficult to collect sufficiently large datasets of ads across large samples of people. Additionally, it is difficult to label the content of ads at scale, especially for characteristics that typically require expert qualitative analysis to identify, like the presence of deceptive patterns. I am interested in developing tools for scalable, replicable measurements of targeted advertising, to enable other researchers and investigators to conduct rigorous audits of the practices of online advertisers and platforms.

I am interested in exploring three technical approaches to generalizing ad audits: 1) building infrastructure to enabling repeatable experiments of targeting using profiles seeded with real user histories, 2) leveraging large language models and to scale up qualitative categorization of ad content, and 3) incorporating additional tracking mechanisms beyond web tracking into profile construction, such as location and mobile app usage. I am currently conducting preliminary work at CMU to develop these tools. My collaborators and I are investigating whether people with chronic health conditions are being targeted with health-related advertising, and whether such ads contain misleading or fraudulent content. I am developing a distributed crawling infrastructure that creates advertising profiles based on real web histories provided by participants, and runs repeatable targeting experiments using these profiles, at the scale of hundreds of profiles in parallel. I am also working with experts in health misinformation and policy to develop a taxonomy and dataset of deceptive health-related claims in online ads, towards the goal of training a language model to recognize deceptive health ads at scale.

Assisting Security and Privacy Decisions with AI Machine learning-based tools are increasingly used to assist people in making security and privacy decisions. For example, in many network and operational technology environments, machine-learning based anomaly detectors can help alert security analysts to potential threats. However, in security-related scenarios; detection of threats is usually only the first step in an investigation and remediation process, and existing tools are often not trusted, or do not provide sufficient context for understanding their outputs. I plan to explore the sociotechnical gap between machine learning tools for security applications and their users.

I plan to investigate two areas in this research direction: first, building on our preliminary work on anomaly detection for industrial control systems (ICS) at CMU [1], I will study how to make explainable AI methods into practical tools for cybersecurity analysts. In our recent work, I conducted a survey with ICS operators to understand whether explanations for machine learning-based anomaly detectors could help them investigate potential security breaches. We found that our

explanation outputs were helpful but insufficient for their larger root-cause analyses, and that more context from other sources was needed to aid in their decision-making. I plan to conduct research on understanding how cybersecurity analysts use different information sources during incident responses and how they evaluate and adapt to (imperfect) machine-learning based tools. Second, complementing my work on deceptive advertising, I plan to explore using language model-based tools to help end users interpret deceptive patterns online. Deceptive patterns often use subtle wording and language to manipulate users' decision making, exposing them to harms like higher costs. I will explore whether foundation models can be adapted to detect deceptive patterns, generate explanations to help users understand how they work, and evaluate the effectiveness of interventions in browsers and mobile apps.

Research Vision Stepping back, I plan to continue studying security and privacy threats from a human-centered perspective, with the goal of developing tools to make it easier for people to protect themselves, and informing policies to protect people when technological solutions and individual action are insufficient.

References

[1] Clement Fung, Eric Zeng, Lujo Bauer. *Attributions for ML-based ICS Anomaly Detection: From Theory to Practice*. Network and Distributed System Security Symposium (NDSS), 2024 (to appear).

[2] McKenna McCall, Eric Zeng, Faysal Hossain Shezan, Mitchell Yang, Lujo Bauer, Abhishek Bichhawat, Camille Cobb, Limin Jia, Yuan Tian. *Towards Usable Security Analysis Tools for Trigger-Action Programming*. Symposium on Usable Privacy and Security (SOUPS), 2023.

[3] Eric Zeng, Rachel McAmis, Tadayoshi Kohno, Franziska Roesner. *What Factors Affect Targeting and Bids in Online Advertising? A Field Measurement Study*. ACM Internet Measurement Conference (IMC), 2022.

[4] Miranda Wei, Eric Zeng, Tadayoshi Kohno, Franziska Roesner. *Anti-Privacy and Anti-Security Advice on TikTok: Case Studies of Technology-Enabled Surveillance and Control in Intimate Partner and Parent-Child Relationships*. Symposium On Usable Privacy and Security (SOUPS), 2022.

[5] Eric Zeng, Miranda Wei, Theo Gregersen, Tadayoshi Kohno, Franziska Roesner. *Polls, Clickbait, and Commemorative $2 Bills: Problematic Political Advertising on News and Media Websites Around the 2020 U.S. Elections*. ACM Internet Measurement Conference (IMC), 2021.

[6] Eric Zeng, Tadayoshi Kohno, Franziska Roesner. *What Makes a "Bad" Ad? User Perceptions of Problematic Online Advertising*. ACM CHI Conference on Human Factors in Computing Systems (CHI), 2021.

[7] Eric Zeng, Tadayoshi Kohno, Franziska Roesner. *Bad News: Clickbait and Deceptive Ads on News and Misinformation Websites*. Workshop on Technology and Consumer Protection (ConPro), 2020.

[8] Eric Zeng, Franziska Roesner. *Understanding and Improving Security and Privacy in Multi-User Smart Homes: A Design Exploration and In-Home User Study*. USENIX Security Symposium, 2019.

[9] Eric Zeng, Frank Li, Emily Stark, Adriana Porter Felt, Parisa Tabriz. *Fixing HTTPS Misconfigurations: An Experiment with Security Notifications*. Workshop on the Economics of Information Security (WEIS), 2019.

[10] Eric Zeng, Shrirang Mare, Franziska Roesner. *End User Security and Privacy Concerns with Smart Homes*. Symposium on Usable Privacy and Security (SOUPS), 2017.

[11] Ada Lerner, Eric Zeng, Franziska Roesner. *Confidante: Usable Encrypted Email - A Case Study with Lawyers and Journalists*. IEEE European Symposium on Security & Privacy (EuroS&P), 2017.

[12] Muhammad Ali, Angelica Goetzen, Alan Mislove, Elissa M. Redmiles, Piotr Sapiezynski. *Problematic Advertising and its Disparate Exposure on Facebook*. USENIX Security Symposium, 2023.