EMILY TSENG

RESEARCH STATEMENT

My goal is to make digital technology safe and affirming for everyone. To do this, I work at the intersection of human-computer interaction (HCI) and computer security and privacy (S&P) to measure and mitigate digital technologies' role in harm: how platforms, devices, and algorithms are increasingly weaponized to control and harass the most vulnerable among us. I focus on high-stakes and real-world contexts where technology is central to people's experiences of vulnerability, and where power dynamics make it difficult for individuals to overcome the harms they face.

My research approach combines methods from multiple disciplines, spanning mixed-methods studies of large-scale observational data like social media and clinical records [5,17,31], community-partnered design research [1,20,30,33], and system-building and evaluation [32,34,35]. First, to understand the role of technology in harm, I **partner with affected communities** to map the social and technical ecosystems around the vulnerabilities they experience. Conscious of the power dynamics in these settings, I develop **creative and rigorous approaches to measurement**, including systems for ethical data analysis at scale, and design provocations that enable safe exploration of how technologies might upend social structures. Then, towards mitigating these harms, I create **principled frameworks for how to improve technology's impact** in these contexts, grounded in tenets of social responsibility (e.g., feminist ethics, trauma-informed care, participation). I use these frameworks to **build systems that positively impact safety and well-being** for affected communities, and **evaluate** their utility in real-world deployments. Through this practice, I develop **transferable design principles** for how to center equity and responsibility in digital technology, via real-world impact. In my PhD, I have applied this approach to investigate:

1. **Interpersonal abuse in intimate partner violence (IPV).** Digital technologies have created new avenues for targeted and persistent abuse from an intimate partner. I pioneered the delivery of S&P assistance to survivors of IPV, resulting in the real-world deployment of a clinic that has helped 500+ survivors since 2020 [3,9,29,30,33].

2. **Structural exploitation in workplace technologies.** In the age of distributed work, platforms, devices, and algorithms structure labor in ways that can improve conditions for some at the expense of others. With frontline health workers, I designed systems to ensure workplace technologies center workers' well-being [1,18,19,26,31].

3. **Research-related harm.** Successful interventions require understanding highly personal experiences of harm; but gathering such knowledge can itself cause harm. I have created scientific infrastructure for safer research with affected communities, by balancing individual and collective decision-making in data stewardship [4,6,28,32].

I publish at top-tier venues in HCI and S&P, including CHI and CSCW, USENIX Security, and IEEE Security and Privacy (Oakland). My work has also appeared at the highly selective Privacy Law Scholars Conference (PLSC). Across 18 publications, 7 first-authored, I have earned **three Best Paper Awards (top 1%)**, one Honorable Mention (top 5%), and third place in the 2020 Internet Defense Prize, as well as competitive fellowships from Microsoft Research, Cornell, and Rising Stars in EECS. In addition to advancing scientific progress, my research has created real-world impact in industry and policy. I am working with Google to translate my research into internal trainings for engineers, and I worked with community health organizations to study their security needs in a 6-month internship at Microsoft Research. I also consulted on U.S. Senate legislation creating protections for IPV survivors, resulting in the Safe Connections Act, signed into law in 2022.[1]

Protecting IPV survivors by reimagining security & privacy as care infrastructure

In my PhD, I deeply engaged with how to measure and mitigate one particularly high-stakes form of harm: technology's role in intimate partner violence (IPV). Affecting 1 in 3 women, 1 in 4 men, and 1 in 2 transgender or non-binary people in the U.S. [26], IPV encompasses many forms of mistreatment—including physical and emotional abuse, stalking, and harassment—levied against a victim by a current or former intimate partner. Digital technologies have worsened IPV's prevalence and severity, by creating more and more avenues for abuse. Smartphones and social networks leak location information to stalkers; spyware is easily found online; and attackers with an intimate partner's level of access and knowledge can easily bypass authentication. Intervening in IPV thus requires computer security and privacy (S&P) to revamp its canonical approaches to mitigating digital threats. To help survivors, who are already managing the traumatic effects of abuse, we need a sociotechnical approach that does not add to their burdens. Towards this vision, I have contributed (1) a measurement pipeline for understanding abusers [3,29]; and (2) a direct intervention assisting survivors [9,30,33].

Investigating abusive tactics by leveraging social media. S&P requires knowledge of an attacker's goals and capabilities to improve a system's robustness to attack. In IPV, such knowledge is difficult to obtain: People who abuse are unlikely to respond to interviews or surveys. However, when they seek resources or community, they turn to the anonymity offered by

[1] Safe Connections Act of 2022 and accompanying FCC Report and Order in October 2023 implementing the law.

the Internet. I thus built **a measurement pipeline crawling and analyzing social media** to understand abusers: the capabilities they seek, the vulnerabilities they exploit, and how they are enabled by online communities. I built a novel dataset of 200k posts from five public forums where people in relationships seek advice on how to surveil their partners. Through an in-depth content analysis, I showed these forums provide potential abusers with detailed knowledge of tools for digital surveillance, and escalate the severity of attacks [5,31]. This work earned a Distinguished Paper Award (top 1%) at USENIX Security (a top S&P venue) and 3rd place in the 2020 Internet Defense Prize ($40k), as well as a Best Paper Award (top 1%) at ACM CSCW (a top social computing venue). I presented this work to practitioners at Facebook and Google; and S&P scholars have used my approach to study hard-to-reach adversaries in deepfakes [29] and smart homes [15].

Directly assisting survivors via care infrastructure. My work on abusers' tactics underscored the need for sociotechnical approaches to handle the complexity of today's security and privacy problems. Where S&P traditionally centers on assuring the confidentiality, integrity and availability of *systems*, we now need to attend to the holistic needs of the *people* facing targeted attacks. To bridge these traditions, I proposed an approach called care infrastructure [33]. Inspired by the feminist ethic of care and the science studies notion of infrastructure, care infrastructure calls for S&P to create systems oriented towards ongoing relations of care between survivors and S&P experts: a new terrain for how to mitigate digital insecurity.

I tested the ideas behind care infrastructure in a multi-year collaboration with the Clinic to End Tech Abuse (CETA)[2], which provides IPV survivors with personalized 1:1 security and privacy assistance from volunteer S&P experts ("consultants") trained in the specific threat model of IPV. Since 2018, I have worked as a researcher-practitioner at CETA, where I personally consult survivors with high-risk needs. **At CETA, I have created new approaches to security that have helped 500+ survivors since 2020**. For example, at the height of the COVID-19 pandemic in 2020, I investigated the safety challenges of helping survivors who might be locked down with their abuser [30]. At the time, CETA paired survivors with consultants for one-time, one-hour consultations, held via synchronous audio or video calls. But this model meant consultants struggled to fully investigate survivors' complex social and technical vulnerabilities.

I hypothesized that if we redesigned CETA's service towards care infrastructure, we could better support survivors. I thus designed and implemented a new care model that emphasizes long-term connection between survivors and consultants: aligning S&P assistance more towards the role of a family doctor than tech support. I implemented a secure case record system, and a referral system matching survivors to consultants with specific expertise. In an 8-month deployment, I found the new system was able to reach survivors with higher-risk concerns; route survivors to specialized assistance; and increase survivors' confidence in managing their digital privacy. However, delivering S&P as care also imposed new burdens on consultants: the need to negotiate boundaries with survivors, who expected them to be "always on"; the need to assure safety for survivors; and the need for methods to evaluate the success of their work. These new burdens, I argue, open a fruitful new S&P subfield, on how technology enacting the ethic of care can help achieve holistic digital security.

The care model I developed is still in use by CETA today. Survivors come away from CETA more informed about technology, and in many cases, their digital security is improved via removal of spyware or account compromise. Awarded a Best Paper (top 1%) at ACM CHI 2022, this work has already been influential in the field: the notion of *care infrastructure* has been used to consider emerging S&P issues from educational tech [34] to everyday security practices in Lebanon [17]. I was also invited to share this work at Google Trust and Safety, and at the Grace Hopper Celebration.

Protecting frontline workers from structural harm in workplace technologies

Beyond enabling interpersonal abuse, technology can also create structural harms in organizational contexts. Data-driven technologies scaffold and distribute labor, and thus determine which workers benefit from gains in efficiency, and which are burdened or exploited. I explored tech-mediated harm with frontline health workers: skilled caregivers like doctors, nurses, and home health aides. Frontline health is fundamentally collaborative and empathetic—but also increasingly mediated through workplace platforms, and thus routinized, stratified, and isolated. With several groups of workers, I designed tools for decision support and productivity in distributed care [20,21,28,33] and personal sensing for workplace well-being [1].

Designing equitable tools for distributed care work. When workers are geographically and temporally scattered, managers need transparency into workers' productivity, and workers need support and recognition. Systems for time-tracking and real-time decision support could help make this work fairer and more equitable, but must be designed to preserve workers' autonomy. With 1199SEIU, the largest healthcare workers' union in the U.S., I explored how to design such systems with home health aides: formal caregivers who provide vital assistance to patients in their homes. Via several studies, I found aides' current tools serve to monitor their labor for insurance reimbursement, rather than streamlining their workflows, addressing their isolation, or accounting for their invisible work [20,21,33]. These inequities have created a labor crisis with grave social consequences: 1 in 2 workers leave the profession each year, even as more and more older adults seek aides to help them age in place. The field is projected to add 1.2 million openings over the next decade [2].

[2] Clinic to End Tech Abuse. CETA is a nonprofit that works with the NYC Mayor's Office to End Domestic and Gender-Based Violence.

I thus worked with the union to develop tools for secure communication, decision support, and productivity tools designed for aides' needs [33]. We found that intervening in distributed workplace dynamics requires **attending to the conflict** inherent to redistributing effort among a team. Both aides and their supervisors agreed that aides should be able to more easily reach their peers and supervisors, and receive credit for all their work. However, both parties also felt such tools could create new burdens, e.g., adding documentation to aides' day-to-day, or requiring supervisors to be "always-on". As a way forward, we propose technologists use design methods to concretize fairer workplace futures in near-term goals, as a way to elicit feedback from all stakeholders on how workplace technologies might fairly redistribute burden. Published at ACM CHI 2020 [32], this work laid the foundation for other HCI scholars to build interventions with aides [3,21–23], hospitality workers [26], and workforce development professionals [6]. My work was also cited in federal policy by the U.S. Centers for Medicare and Medicaid Services (CMS).[3]

Designing workplace well-being technologies. Re-distributing burden is not the only way workplace technologies can intervene in harm: technological interventions can also explicitly safeguard workers' well-being. The HCI subfield called Quantified Workplace has explored how data from personal devices ("personal sensing") can measure and mitigate workers' stress. With Northwell Health, a hospital system in New York City, I examined the promise and peril of personal sensing for mitigating burnout in medical residents: doctors-in-training who provide vital labor while working unpredictable and poorly paid shifts that can exceed 16 hours at a time. My study explored how measurements of residents' behavior and stress could be surfaced to their supervisors. In interviews with residents and supervising physicians, participants reflected on dashboards depicting real-time data on residents' sleep, physical activity, and productivity, alongside self-reports of residents' burnout.

We found residents see potential for personal sensing in self-management of their workplace well-being, and are especially open to aggregating sensed data to inform workplace-wide change. However, such programs will require addressing fundamental issues around the privacy of employees and the validity of measuring stress from personal sensing, as well as how to ensure supervisors are held accountable to improving workplace conditions: a rich terrain for further research. Published at ACM CSCW 2022 [1], this work has already informed studies on workplace well-being for software engineers at technology firms [7,11,35], and how ubiquitous computing scholars consider the harms of productivity tools [8]. I was also invited to share this work with health policymakers and innovation leaders at the 2023 Behavioral Health Next Summit.[4]

Creating scientific infrastructure for safe and consentful research on tech-mediated harms

The tools, workflows, principles and practices of research—our scientific infrastructure—deeply influence how we intervene in tech-mediated harms. To build appropriate interventions, researchers must engage with graphic and highly sensitive human experiences (e.g., accounts of IPV or workplace discrimination). But doing so takes a toll on researchers and participants. I thus built a line of research on frameworks [5,7] and tools [29,33] for safer digital-age research on technology's role in harm.

Guiding frameworks for researchers. With collaborators, I developed two frameworks to guide research on tech-mediated harm. The first, "Safer Digital Safety Research Involving At-Risk Users" [6], provides guidance for the burgeoning research community investigating digital safety for people at higher risk of targeted attacks ("at-risk users"). The number of papers on at-risk users at venues in HCI and S&P is steadily growing; but to-date, there has been little standard guidance on how to do this research in ways that minimize harm to participants and researchers alike. To fill this methodological gap, we conducted a systematization of knowledge (SoK) reviewing 196 recent papers in HCI and S&P and gathering oral histories from 12 expert researchers to arrive at 6 strategies for safer research, from planning and execution to publication. This work will appear at a top S&P venue, IEEE Symposium on Security and Privacy (Oakland), in 2024.

The second framework, "Trauma-Informed Computing" (TIC), describes how engineers and computer scientists can attend to the impact of trauma in technology [6]. Defined as the experience and aftermath of a distressing event (e.g., violence, illness, disaster), trauma affects >70% of people worldwide [12]. In our paper, we first used trauma as an explanatory lens, via vignettes on the tech-mediated harms faced by IPV survivors, victims of identity theft, and transgender people. We then used trauma as a generative lens: a guidance for how to create technologies that minimize trauma's impact, and avoid re-traumatization. We discuss how to adapt principles of trauma-informed care into all levels of computing: design, development, deployment, maintenance and support. This work appeared at ACM CHI 2022, and has already been taken up to consider social media [23], digital mental health [16], algorithmic welfare systems [24], and more.

Tools and techniques for safer research. As research on tech-mediated harm grows across computing, more and more researchers will be engaged in high-stakes and emotionally demanding studies. To concretely enact the frameworks and guidance provided for this growing community of researchers, I have designed and built systems for collecting and analyzing data about tech-mediated harms, and evaluated their use with researchers and data subjects.

[3] 86 FR 61555, a federal rule to establish COVID–19 vaccination requirements for healthcare workers.
[4] Behavioral Health Next Summit, headlined by Patrick J. Kennedy, former U.S. Representative; NYC Health Commissioner Dr. Ashwin Vasan; and former World Bank president Dr. Jim Yong Kim.

I focused first on data stewardship: the principles and practices of collecting and using data. Due to widespread availability of storage and compute, researchers increasingly look to amass observational datasets on social phenomena, like clinical records or archives of social media, and make them available to other researchers for secondary use. I myself used these methods to study abusers' surveillance tactics [3,29], IPV survivors' experiences of abuse [9,30,33], and healthcare workers' well-being during COVID-19 [15]. The drive towards big data enables more research at greater scale, but requires reimagination of fundamentals of research ethics, like consent, anonymity, and confidentiality.

I thus led a study using theories of participation to redesign systems for data stewardship, to better involve IPV survivors in digital-safety research [28]. I developed **a new heuristic framework for community participation in data stewardship** that emphasizes equitable division of burden and benefit between researchers and participants. Using this framework, I explored how to handle clinical records of digital-safety consultations: a dataset that could form a rich evidence base on survivors' security needs, if handled in appropriate and participatory ways. I developed **a technique for eliciting participants' data sharing preferences**, and deployed it in a field study with IPV survivors and their caseworkers at CETA. My findings show survivors and case workers want this data to be used to combat IPV, but differ widely in who they want to see this information, and for what purpose. Survivors in particular want assistance from a trusted steward to manage their records. To appear at ACM CSCW 2024, this work has already informed CETA's revision of its consent procedure.

I then focused on reimagining how we analyze collected data. Researchers seeking to understand experiences of tech-enabled harm typically use qualitative analysis to gain empathetic insight into human experience. But close study of disturbing content can cause secondary traumatic stress and re-traumatization in the researcher. I therefore designed and built **TIQA** [32], a system for **T**rauma-**I**nformed **Q**ualitative **A**nalysis that enables an analyst to **(1) leverage large language models (LLMs) towards personalizable content warnings; and (2) track their exposure to personally traumatic concepts.** To explore TIQA's utility, I built a functional prototype and tested it with 12 researchers who conduct empirical studies of IPV and online hate and harassment. My findings illuminate a broad design space around tools for managing traumatic exposure—and an equally broad space around how the high-level principles of trauma-informed computing can be realized in the low-level decision-making required to build software. This work is currently under review.

Future Research Agenda

Tech-mediated harms will impact more people, in more devastating ways, as digital technology expands its societal influence. To reap the benefits of computing without exacerbating harm for society's most vulnerable, we need rigorous and power-conscious study of how to measure and mitigate these harms. In the next five years, I envision this agenda to include:

Measuring and mitigating traumatic exposure across modalities. My work with TIQA showed there is an appetite for systems that can help the growing community of researchers studying tech-mediated harm with managing our well-being. To do this in standard scientific infrastructure, we need to address technical and design challenges in robust measurement and mitigation of traumatic exposure across data types. How can we operationalize personalizable content warnings in not only text (as I have done in TIQA), but also in images, audio, and video? One idea might be to use transfer learning to produce multimodal embeddings of a user's personally traumatic concepts, but this approach would still require the user to be exposed to traumatic content in the course of validation. We therefore need to understand how to evaluate measures of traumatic exposure, in ways that still minimize the user's exposure to the content they are trying to filter out. From this basis, we can explore mitigation: how can personalized ML-driven measurements of traumatic exposure be used to minimize researchers' stress reactions? I plan to answer these questions through iterative system-building and evaluation. Already, I have partnered with a group of researchers to implement TIQA in their practices, and build community around TIQA as an open-source tool.

Safety and participation in data-centric machine learning (ML). I am also interested in bringing lessons from my work on digital safety to the growing community in responsible ML. Increasingly, ML is best understood as a sociotechnical apparatus of mass data collection and reuse, where the quality of the data and feedback used to train a system is as important as the modeling and engineering ("data-centric ML"). In high-risk areas like content moderation, fraud detection and social risk scoring, human feedback—from both data workers [12] and end users—is especially needed for reliable and interpretable decisions around subjective constructs like toxicity. I want to study (1) how to make human feedback on these models safer for data workers, and (2) how to make these systems more participatory for affected communities.

Question (1) emerges from my work on scientific infrastructure for sensitive research, through which I observed many of the problems experienced by researchers may be amplified in data work. What are the challenges to health and well-being that occur when people are tasked with enriching data for generative AI, e.g., via reinforcement learning from human feedback (RLHF)? And how might approaches to mitigating workplace stress translate to the occupational context of data work? With respect to question (2), I am interested in addressing core problems in participatory design in high-risk systems reliant on human feedback. How do we give affected communities a meaningful voice in how these systems are built, without placing undue burden on them? Soliciting feedback from *all* members of an affected community is intractable; but how can we achieve representative governance in a community of people whose primary commonality is having experienced the same type of harm? I will start by blending my techniques for participatory data stewardship [30] with

techniques for human-centered machine learning. To begin this work, I have built links with interested partners in ML and AI: I co-led a series of workshops at 2023's ACM CSCW and FAccT on community collaboration in computing [16].

Digital safety in women's health. My work in IPV shows digital insecurity contributes to structural forms of gender-based violence that can have short- and long-term health consequences (e.g., trauma). Digital technologies' erosion of privacy also prevents vulnerable people from seeking care and participating in research: limiting their inclusion in modern biotechnology. To me, these insights point to **digital safety as an emerging health equity issue,** particularly in women's health.

Women's health encompasses not only individual factors like gynecological, maternal, and mental health, but also community-level and structural factors, like interpersonal violence, reproductive rights, and intersectionality. I also consider womanhood to be inclusive of transgender people. With this theoretical orientation, I want to explore **how interlocking forms of oppression affect women's trust in and access to digital health systems, and how the threat of digital insecurity affects their health outcomes**. My research agenda will include (1) empirical investigation of at-risk populations' access to and trust in digital health; (2) systems for sociotechnical safety in electronic health records and personal health apps; and (3) design principles for safely and equitably building technology with community health organizations. For example, I want to investigate: Since the U.S. Supreme Court overturned federal protection for abortion access in 2022, has there been a chilling effect on women seeking care for conditions not related to reproductive wellness? What are women's existing security practices and mental models around their health data, and do they align with those assumed by the builders of EHRs and health apps? And given results that show differential privacy disproportionately harms accuracy for minority groups in a dataset [3], what forms of privacy-enhancing technologies would support women's inclusion in research?

To execute this line of work, I will build off the techniques I used in my work with IPV survivors. I will begin by establishing collaborations with health organizations serving women on the front lines, and with the researchers and medical practitioners who center the under-studied health issues that deeply impact women, like endometriosis. Already I have begun to build this research area: with Dr. Mary L. Gray, I am investigating security and privacy in community health.

Conclusion

I create the systems, interventions, and design principles we need to measure and mitigate technology-mediated harm, towards making digital technology safe and affirming for everyone. By working in deep collaborations with affected communities, I demonstrate how technologists can positively impact safety and well-being for society's most vulnerable. I have made field-leading contributions across top-tier venues in human-computer interaction and computer security and privacy; informed U.S. federal policy; and enabled CETA to support 500+ IPV survivors since 2020. Together, my research activities advance **equity and justice in computing and data**, paving the way for engineers, policymakers, and affected communities to work together towards brighter technological futures.

References

* indicates equal contribution.

[1] Daniel A. Adler*, **Emily Tseng***, Khatiya C. Moon, John Q. Young, John M. Kane, Emanuel Moss, David C. Mohr, and Tanzeem Choudhury. 2022. Burnout and the Quantified Workplace: Tensions around Personal Sensing Interventions for Stress in Resident Physicians. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2 (November 2022), 1–48. https://doi.org/10.1145/3555531

[2] American Public Health Association. 2020. Supporting and Sustaining the Home Care Workforce to Meet the Growing Need for Long-Term Care. *APHA Policy Statements*. Retrieved November 14, 2023 from https://www.apha.org/policies-and-advocacy/public-health-policy-statements/policy-database/2021/01/13/supporting-and-sustaining-the-home-care-workforce

[3] Eugene Bagdasaryan, Omid Poursaeed, and Vitaly Shmatikov. 2019. Differential Privacy Has Disparate Impact on Model Accuracy. *In Proc. 33rd Conf. on Neural Information Processing Systems (NeurIPS)*. Curran Associates Inc., Red Hook, NY, USA, Article 1387, 15479–15488. https://doi.org/10.5555/3454287.3455674

[4] Vince Bartle, Liam Albright, and Nicola Dell. 2023. "This machine is for the aides": Tailoring Voice Assistant Design to Home Health Care Work. In *Proc. ACM CHI Conf. on Human Factors in Computing Systems*, 2023, Hamburg, Germany. https://doi.org/10.1145/3544548.3581118

[5] Rosanna Bellini, **Emily Tseng**, Nora McDonald, Rachel Greenstadt, Damon McCoy, Thomas Ristenpart, and Nicola Dell. 2021. "So-called privacy breeds evil": Narrative Justifications for Intimate Partner Surveillance in Online Forums. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW3 (January 2021), 210:1-210:27. https://doi.org/10.1145/3432909. **Best Paper Award (top 1%).**

[6] Rosanna Bellini, **Emily Tseng**, Noel Warford, Alaa Daffalla, Tara Matthews, Sunny Consolvo, Jill Palzkill Woelfer, Patrick Gage Kelley, Michelle L. Mazurek, Dana Cuomo, Nicola Dell, and Thomas Ristenpart. 2024. SoK: Safer Digital-Safety Research Involving At-Risk Users. *To appear at IEEE Symposium on Security and Privacy (Oakland)*, 2024. Retrieved October 28, 2023 from https://arxiv.org/abs/2309.00735v1

[7] Connie W. Chau, Kenneth Holstein, and Michael Madaio. 2022. Part of the Conversation: Workforce Professionals' Perspectives on the Roles and Impacts of Workforce Technologies. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2 (November 2022), 1–22. https://doi.org/10.1145/3555135

[8] Niu Chen*, Allison McDonald*, Yixin Zou*, **Emily Tseng**, Kevin A Roundy, Acar Tamersoy, Florian Schaub, Thomas Ristenpart, and Nicola Dell. 2022. Trauma-Informed Computing: Towards Safer Technology Experiences for All. In *Proc. ACM CHI Conf. on Human Factors in Computing Systems*, 2022, New Orleans LA USA. https://doi.org/10.1145/3491102.3517475

[9] Shreya Chowdhary, Anna Kawakami, Mary L Gray, Jina Suh, Alexandra Olteanu, and Koustuv Saha. 2023. Can Workers Meaningfully Consent to Workplace Wellbeing Technologies? In *2023 ACM Conf. on Fairness, Accountability, and Transparency (FAccT)*, 2023, Chicago IL USA. https://doi.org/10.1145/3593013.3594023

[10] Vedant Das Swain, Lan Gao, William A Wood, Srikruthi C Matli, Gregory D. Abowd, and Munmun De Choudhury. 2023. Algorithmic Power or Punishment: Information Worker Perspectives on Passive Sensing Enabled AI Phenotyping of Performance and Wellbeing. In *Proc. ACM CHI Conf. on Human Factors in Computing Systems*, 2023, Hamburg Germany. https://doi.org/10.1145/3544548.3581376

[11] Diana Freed*, Sam Havron*, **Emily Tseng**, Andrea Gallardo, Rahul Chatterjee, Thomas Ristenpart, and Nicola Dell. 2019. "Is my phone hacked?" Analyzing Clinical Computer Security Interventions with Survivors of Intimate Partner Violence. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW (November 2019), 1–24. https://doi.org/10.1145/3359304. **Honorable Mention (top 5%).**

[12] Mary L. Gray and Siddharth Suri. 2019. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Houghton Mifflin Harcourt.

[13] Anna Kawakami, Shreya Chowdhary, Shamsi T. Iqbal, Q. Vera Liao, Alexandra Olteanu, Jina Suh, and Koustuv Saha. 2023. Sensing Wellbeing in the Workplace, Why and For Whom? Envisioning Impacts with Organizational Stakeholders. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW2 (September 2023), 1–33. https://doi.org/10.1145/3610207

[14] Rolf J. Kleber. 2019. Trauma and Public Mental Health: A Focused Review. *Front. Psychiatry* 10, (June 2019), 451. https://doi.org/10.3389/fpsyt.2019.00451

[15] Jingjie Li, Kaiwen Sun, Brittany Skye Huff, Anna Marie Bierley, Younghyun Kim, Florian Schaub, and Kassem Fawaz. "It's up to the Consumer to be Smart":

[16] Calvin Alan Liang, **Emily Tseng**, Akeiylah Dewitt, Yasmine Kotturi, Sucheta Ghoshal, Angela D. R. Smith, Marisol Wong-Villacres, Lauren Wilcox, and Sheena Erete. 2023. Surfacing Structural Barriers to Community-Collaborative Approaches in Human-Computer Interaction. In *Workshop Proc. ACM Conf. on Computer Supported Cooperative Work and Social Computing*, 2023, Minneapolis MN USA. https://doi.org/10.1145/3584931.3611294

[17] Matteo Malgaroli, **Emily Tseng**, Thomas D. Hull, Emma Jennings, Tanzeem K. Choudhury, and Naomi M. Simon. 2023. Association of Health Care Work With Anxiety and Depression During the COVID-19 Pandemic: Structural Topic Modeling Study. *JMIR AI* 2, 1 (October 2023), e47223. https://doi.org/10.2196/47223

[18] Gabriela Marcu and Jina Huh-Yoo. 2023. Attachment-Informed Design: Digital Interventions That Build Self-Worth, Relationships, and Community in Support of Mental Health. In *Proc. 2023 ACM Conf. on Designing Interactive Systems*, 2023, Pittsburgh PA USA. https://doi.org/10.1145/3563657.3596009

[19] Jessica McClearn, Rikke Bjerg Jensen, and Reem Talhouk. 2023. Security Patchworking in Lebanon: Infrastructuring Across Failing Infrastructures. https://doi.org/10.48550/arXiv.2310.16969

[20] Joy Ming, Elizabeth Kuo, Katie Go, **Emily Tseng**, John Kallas, Aditya Vashistha, Madeline Sterling, and Nicola Dell. 2023. "I Go Beyond and Beyond" Examining the Invisible Work of Home Health Aides. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW1 (April 2023), 1–21. https://doi.org/10.1145/3579492

[21] Fabian Okeke, **Emily Tseng**, Benedetta Piantella, Mikaela Brown, Harveen Kaur, Madeline R. Sterling, and Nicola Dell. 2019. Technology, home health care, and heart failure: a qualitative analysis with multiple stakeholders. In *Proc. ACM Conf. on Computing & Sustainable Societies*, 2019, Accra, Ghana. https://doi.org/10.1145/3314344.3332487

[22] Anthony Poon, Lourdes Guerrero, Julia Loughman, Matthew Luebke, Ann Lee, Madeline Sterling, and Nicola Dell. 2023. Designing for Peer-Led Critical Pedagogies in Computer-Mediated Support Groups for Home Care Workers. In *Proc. ACM CHI Conf. on Human Factors in Computing Systems*, 2023, Hamburg Germany. https://doi.org/10.1145/3544548.3580881

[23] Anthony Poon, Vaidehi Hussain, Julia Loughman, Ariel C. Avgar, Madeline Sterling, and Nicola Dell. 2021. Computer-Mediated Peer Support Needs of Home Care Workers: Emotional Labor & the Politics of Professionalism. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2 (October 2021), 1–32. https://doi.org/10.1145/3476077

[24] Anthony Poon, Matthew Luebke, Julia Loughman, Ann Lee, Lourdes Guerrero, Madeline Sterling, and Nicola Dell. 2023. Computer-Mediated Sharing Circles for Intersectional Peer Support with Home Care Workers. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW1 (April 2023), 1–35. https://doi.org/10.1145/3579472

[25] Carol F Scott, Gabriela Marcu, Riana Elyse Anderson, Mark W Newman, and Sarita Schoenebeck. 2023. Trauma-Informed Social Media: Towards Solutions for Reducing and Healing Online Harm. In *Proc. ACM CHI Conf. on Human Factors in Computing Systems*, 2023, Hamburg Germany. https://doi.org/10.1145/3544548.3581512

[26] Dilruba Showkat, Angela D. R. Smith, Wang Lingqing, and Alexandra To. 2023. "Who is the right homeless client?": Values in Algorithmic Homelessness Service Provision and Machine Learning Research. In *Proc. ACM CHI Conf. on Human Factors in Computing Systems*, 2023, Hamburg Germany. https://doi.org/10.1145/3544548.3581010

[27] Franchesca Spektor, Sarah E Fox, Ezra Awumey, Christine A. Riordan, Hye Jin Rho, Chinmay Kulkarni, Marlen Martinez-Lopez,

Betsy Stringam, Ben Begleiter, and Jodi Forlizzi. 2023. Designing for Wellbeing: Worker-Generated Ideas on Adapting Algorithmic Management in the Hospitality Industry. In *Proc. ACM Conf. on Designing Interactive Systems*, 2023, Pittsburgh PA USA. https://doi.org/10.1145/3563657.3596018

[28] Madeline R Sterling, **Emily Tseng**, Anthony Poon, Jacklyn Cho, Ariel C Avgar, Lisa M Kern, Claire K Ankuda, and Nicola Dell. 2020. Experiences of Home Health Care Workers in New York City During the Coronavirus Disease 2019 Pandemic: A Qualitative Analysis. *JAMA Intern. Med.* (2020).

[29] Brian Timmerman, Pulak Mehta, Progga Deb, Kevin Gallagher, Brendan Dolan-Gavitt, Siddharth Garg, and Rachel Greenstadt. Studying the Online Deepfake Community. *Journal of Online Trust and Safety*, 2(1). https://doi.org/10.54501/jots.v2i1.126

[30] **Emily Tseng**, Rosanna Bellini, Yeuk-Yu Lee, Alana Ramjit, Thomas Ristenpart, and Nicola Dell. 2024. Data Stewardship in Clinical Computer Security: Balancing Burden and Benefit in Participatory Systems. *To appear at ACM CSCW 2024*.

[31] **Emily Tseng**, Rosanna Bellini, Nora McDonald, Matan Danos, Rachel Greenstadt, Damon McCoy, Nicola Dell, and Thomas Ristenpart. 2020. The Tools and Tactics Used in Intimate Partner Surveillance: An Analysis of Online Infidelity Forums. In *29th USENIX Security Symposium (USENIX Security 20)*, August 2020. USENIX Association, 1893–1909. Retrieved from https://www.usenix.org/conference/usenixsecurity20/presentation/tseng. **Distinguished Paper Award (top 1%) and Internet Defense Prize, third place.**

[32] **Emily Tseng**, Diana Freed, Kristen Engel, Thomas Ristenpart, and Nicola Dell. 2021. A Digital Safety Dilemma: Analysis of Computer-Mediated Computer Security Interventions for Intimate Partner Violence During COVID-19. In *Proc. ACM CHI Conf. on Human Factors in Computing Systems,* 2021, New York, NY, USA. https://doi.org/10.1145/3411764.3445589

[33] **Emily Tseng**, Fabian Okeke, Madeline Sterling, and Nicola Dell. 2020. "We can learn. Why not?": Designing Technologies to Engender Equity for Home Health Aides. In *Proc. ACM CHI Conf. on Human Factors in Computing Systems*, 2020, Honolulu HI USA. https://doi.org/10.1145/3313831.3376633

[34] **Emily Tseng**, Thomas Ristenpart, and Nicola Dell. TIQA: A System for Trauma-Informed Qualitative Analysis. *In submission for an ACM conference*.

[35] **Emily Tseng**, Mehrnaz Sabet, Rosanna Bellini, Harkiran Kaur Sodhi, Thomas Ristenpart, and Nicola Dell. 2022. Care Infrastructures for Digital Security in Intimate Partner Violence. In *Proc. CHI Conf. on Human Factors in Computing Systems,* 2022, New Orleans LA USA. https://doi.org/10.1145/3491102.3502038. **Best Paper Award (top 1%).**

[36] Kelly B. Wagman, Elana B. Blinder, Kevin Song, Antoine Vignon, Solomon Dworkin, Tamara Clegg, Jessica Vitak, and Marshini Chetty. 2023. "We picked community over privacy": Privacy and Security Concerns Emerging from Remote Learning Sociotechnical Infrastructure During COVID-19. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW2 (September 2023), 1–29. https://doi.org/10.1145/3610036

[37] Novia Wong, Victoria Jackson, André Van Der Hoek, Iftekhar Ahmed, Stephen M. Schueller, and Madhu Reddy. 2023. Mental Wellbeing at Work: Perspectives of Software Engineers. In *Proc. CHI Conf. on Human Factors in Computing Systems*, 2023, Hamburg Germany. https://doi.org/10.1145/3544548.3581528