I am a ***human-centered computing researcher*** with strong interdisciplinary expertise in human-computer interaction, usable privacy & security, privacy-enhancing technologies, applied machine learning/artificial intelligence (ML/AI), ubiquitous computing, and health informatics. My research aims to ensure ***the use of people's personal information is appropriate, fair, and meaningful*** by emerging technologies.

My work spans ***the full research and development cycle of human-centered computing systems around people's personal information***: I examine how ***emerging computing technologies*** impact the ways people interact with their personal information (e.g. the Internet of Things, wearable devices); I seek to identify underlying personal and social challenges around these technologies (e.g. data privacy transparency, trustworthy ML/AI); I design and develop real-world systems that address these challenges (e.g. ML/AI-enabled personalized privacy assistants). The ultimate goal of my research is to improve people's personal wellbeing and achieve greater social good. In this research statement, I outline my current and previous work in the following aspects, upon which I will build my future research program.

• Increase Data Privacy Transparency of the Internet of Things (IoT)

In the digital world, users often do not know the full extent of which their personal data is being collected, used, or shared, because the dominant privacy notice mechanism (e.g., privacy policies) is ineffective to convey key privacy-related information to users. Also, the choices (e.g., privacy settings) available for users to control their data privacy are often absent or difficult to locate. These data privacy issues are particularly concerning with the plethora of IoT technologies (e.g., Bluetooth beacons, smart speakers, motion sensors, facial recognition cameras). To provide exciting context-aware services like smart cities and smart homes, IoT rely on ubiquitous collection and use of potentially privacy-sensitive data. IoT data practices are often opaque due to the fact that most IoT devices and sensors lack traditional interfaces to communicate with users. Therefore, it is extremely challenging to ensure data privacy transparency and support individual privacy choices in IoT context.

My current work focuses on resolving the technical challenges to increase data privacy transparency of IoT. I manage the development efforts for the ***IoT Privacy Infrastructure (IoTPI),*** a real-world solution to facilitate data privacy notice and choice in the particularly challenging IoT context (see Figure 1 for system architecture). IoTPI enables owners of various IoT to disclose the data practices of their IoT

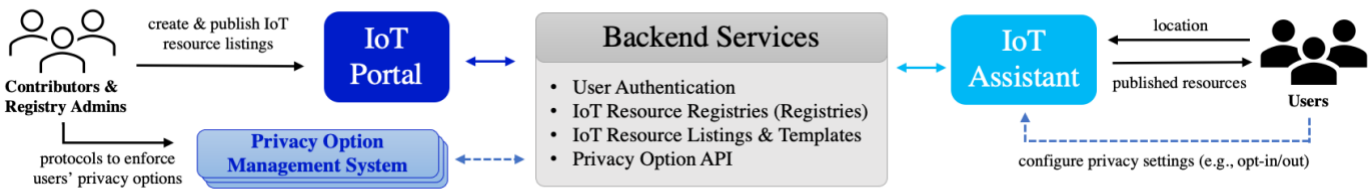

Figure 1: System Architecture of the Internet of Things Privacy Infrastructure (IoTPI). IoTPI has two major user-facing components: (1) the IoT Portal (www.iotprivacy.io) website for stakeholders of IoT to publicize the data practices of their IoT devices and services; (2) the IoT Assistant mobile app for the general public, available on both iOS and Android.

devices and services, which helps them achieve compliance with the new data privacy regulations such as the General Data Protection Regulation (GDPR). IoTPI also empowers individual users to discover nearby IoT data collection and to manage their IoT data privacy via available privacy options. In summary, by involving multiple stakeholders (e.g., owners, manufacturers, vendors, and users of IoT), our infrastructure provides a novel mechanism to increase data privacy transparency and facilitate privacy choices in IoT context. Also, I lead all user research that collects in-the-wild data from real users of IoTPI to shed light on various IoT data privacy issues. My recent work outlining the design space for privacy choices inspired by our design process of IoTPI is accepted by the 2021 ACM CHI Conference [1].

- ## Design Usable and Trustworthy AI-enabled Privacy-Enhancing Technologies

"Notice and choice" is the dominant mechanism to data privacy under most regulations worldwide. Under this mechanism, the burden to manage personal data privacy unfairly falls on individual users, where they have to spend time and efforts to configure privacy settings. With countless mobile and IoT devices in everyday environments, if privacy choices are available (as increasingly required by new data privacy regulations), people would have to configure an unrealistic number of privacy settings to effectively manage their data privacy. Research repeatedly shows that high user burden often leads privacy fatigue and privacy resignation, turning people away from effective personal data privacy management.

My current research on ***personalized privacy assistants (PPAs)*** aims to address this usability problem by reducing people's burden to manage their personal data privacy in mobile and IoT contexts. Specifically, I conduct human subject research studies that employ ML/AI techniques to model people's privacy preferences in both mobile and IoT contexts [2][3]. In a large-scale online survey exploring how different data usage purposes affect people's decisions for Android app permissions, we built predictive models of people's privacy preferences that can be integrated into mobile PPAs to recommend privacy settings based on individual preferences [2]. In a 10-day experience sampling study, we collected participants' in-situ privacy attitudes towards possible deployments of different facial recognition applications in the real-world places they actually visited [3]. We identified patterns in people's privacy preferences towards different purposes for facial recognition in public places. We demonstrated that AI-based PPAs could reduce people's burden to configure various facial recognition consent settings if such privacy choices become available. This research also offers public policy insights towards a potential consent mechanism for the controversial facial recognition applications [4].

To build trustworthy PPAs, I also conduct qualitative research to investigate people's attitudes towards AI-enabled PPAs [5]. We found that people weigh the desire for autonomy in privacy decision making against the burden of cognitive overload to arrive at different PPA automation preferences. We further discuss open issues like automated consent, privacy resignation, and algorithmic biases, contributing to the ongoing discussion on fairness, ethics, and trust in AI-enabled privacy-enhancing technologies.

- ## Leverage Ubiquitous Computing for Proactive Personal Health Management

Personal health management refers to an individual taking active responsibility for managing their own healthcare, which facilitates the transition from the predominantly reactive healthcare model to the overall lower-cost proactive healthcare model. My research in health informatics aims to leverage ubiquitous computing (e.g., health tracking technologies) to support people's long-term personal health management.

My doctoral work focused on ***activity tracking technology*** (e.g., fitness wristbands) that provides people new types of personal health information outside of traditional healthcare settings. My mixed methods research studies developed a deep understanding of activity tracker users' interaction with this technology in the long term (> one year) [6] [7]. I have identified two major groups of long-term activity tracker users, namely power users and consistent casual users, and their distinct rationales behind their usage patterns. Both groups are valid use cases that require different motivational design to facilitate their diverse health-related needs and support their long-term personal health management [6]. Further, I have also designed innovative solutions to promote physical exercises to support proactive personal health management, such as a mobile exergame app – ***StepQuest*** that uses Fitbit steps as game currency to motivate players to stay physically active in order to win team-based games [8]. This line of my research addresses human-centered health informatics, which is an increasingly crucial topic under the current global public health crisis.

- ## Expand Usable Privacy & Security Research in More Application Domains

I will continue my core research in usable privacy and privacy-enhancing technologies, particularly through the design of human-centered AI-enabled privacy tools to help people effectively and easily control their personal information in the digital world. I also seek to expand my research agenda into two domains: ***accessible privacy & security*** and ***the intersection of data privacy & health informatics***.

First, I strive to design accessible and inclusive privacy-enhancing technologies for all. Currently, I lead a qualitative research study examining the needs and concerns of people with vision impairment regarding privacy-enhancing technologies. This study aims to inform the design of an accessible natural language processing (NLP)-based privacy question answering tool. This tool could particularly benefit people with vision impairment by providing them with credible data privacy-related information, which helps protect this population against data privacy and security risks in the digital world. I will continue my research agenda in accessible privacy & security with a wider range of under-represented groups that are previously overlooked by privacy & security research community. Second, I am eager to integrate my health informatics expertise with my data privacy research. I will actively seek novel research opportunities and interdisciplinary collaboration to address data privacy issues in personal health tracking and public health domains (e.g., data privacy in pandemic mitigation efforts).

References

Note: Asterisk * denotes students mentored by me at the time of publication.

---

[1] **Feng, Y.,** Yao, Y., & Sadeh, N. (forthcoming). A design space for privacy choices: Towards meaning privacy control in the Internet of Things. Conditionally accepted to *the ACM Conference on Human Factors in Computing Systems* (CHI 2021).

[2] *Smullen, D., **Feng, Y.,** Zhang, S., & Sadeh, N. (2020). The best of both worlds: Mitigating trade-offs between accuracy and user burden in capturing mobile app privacy preferences. *Proceedings on Privacy Enhancing Technologies (PoPETS)*, 2020 (1), 195–215.

[3] *Zhang, S., **Feng, Y.**, Bauer, L., Cranor L.F., Das, A., and Sadeh, N. (forthcoming). Did you know this camera tracks your mood? Modeling people's privacy expectations and preferences in the age of video analytics. To appear in *Proceedings on Privacy Enhancing Technologies (PoPETS)*, 2021(2).

[4] *Zhang, S.., **Feng, Y.,** Das, A., Bauer, L., Cranor, L., & Sadeh, N. (2020). Understanding people's privacy attitudes towards video analytics technologies. *Proceedings of Federal Trade Commission's PrivacyCon* (FTC PrivacyCon 2020).

[5] *Colnago, J., **Feng, Y.,** Palanivel, T. et al. (2020). Informing the design of a personalized privacy assistant for the Internet of Things. *Proceedings of the ACM Conference on Human Factors in Computing Systems* (CHI 2020).

[6] **Feng, Y.,** & Agosto, D.E. (2019). From health to performance: Amateur runners' personal health information management with activity tracking technology. *Aslib Journal of Information Management*, 71(2), 217-240

[7] **Feng, Y.,** & Agosto, D.E. (2019). Revisiting personal information management through information practices with activity tracking technology. *Journal of the Association for Information Science and Technology*. 70(12), 1352-1367.

[8] Caro, K., **Feng, Y.,** Day, T., Freed, E. Fox, B., & Zhu, J. (2018). Understanding the effect of existing positive relationships on a social motion-based game for health. In *Proceedings of the 12th EAI International Conference on Pervasive Computing Technologies for Healthcare* (PervasiveHealth'18).