RESEARCH STATEMENT Hyunsoo Lee

My research in Human-Computer Interaction (HCI), Usable Privacy, and Software Engineering focuses on designing interactive systems that support user agency, privacy, and responsible engagement in data-rich environments (e.g., smart homes, human-AI interaction). As computing systems increasingly collect and rely on personal data—through mobile sensing, wearable devices, smart homes, and AI-driven technologies—people are often left without clear ways to understand how their data is used, or how to control what's shared. My work addresses these gaps by designing and studying systems that help users reflect on their data practices, navigate trade-offs, and retain control over their personal information.

Grounded in a human-centered approach, I conduct empirical studies to understand people's privacy and security needs and expectations in everyday contexts. I use methods including interviews, vignette studies, prototyping, and real-world deployments to examine how people engage with complex systems for data collection and sharing, and where current designs fall short. Building on these insights, I build interactive technologies that surface data practices, enhance data literacy, and support both individual and collective privacy decisions in multi-user environments (e.g., smart homes).

As an interdisciplinary researcher, my research connects themes across ubiquitous computing, usable privacy, AI ethics, and sociotechnical design. My recent research includes two main threads. One explores **multi-user privacy and data sharing** in sensor-based environments. In this line of work, I investigate how people navigate shared data practices in everyday settings like the home, and how interactive visualizations and system design can support **contextual privacy reflection**, **data literacy**, and **collective decision-making**. Another thread focuses on **privacy risks in LLM-driven systems**. I study how large language models can be exploited to infer personal information, and how we might design more transparent and safer AI interactions. This includes empirical work to surface privacy vulnerabilities and conceptual frameworks for evaluating AI safety through a user-centered lens. Ultimately, I aim to develop systems that move beyond individualistic notions of privacy and toward more relational, user-driven models of data interaction.

My work has been published in top-tier venues such as CHI, UbiComp, CSCW and IEEE Pervasive Computing and has contributed to both academic knowledge and practical system designs. I also collaborate closely with industry partners to bridge research and practice, and engage with interdisciplinary teams to explore how interaction design can support care, ethics, and autonomy in the age of pervasive sensing and intelligent systems.

UNDERSTANDING USER ATTITUDES AND PRIVACY CONCERNS IN PERSONAL DATA COLLECTION

As I initiated my privacy research in 2020, I found that potential privacy issues and end-users' concerns in everyday sensing were largely unexplored despite the rising popularity of sensing devices, applications, and sensor-driven research. Particularly, the increasing use of diverse sensor types—such as physiological sensors,motion trackers, and environmental monitors—and their fusion to generate rich, context-aware data posed new, complex privacy challenges that remained underexamined.

Motivated by this gap, my UbiComp 2022 paper [1] focused on understanding individual users' attitudes and privacy concerns related to sharing personal sensor data collected through mobile and wearable devices. I conducted a large-scale, four-week in-the-wild study with 100 college students using mixed methods, including behavioral privacy surveys, weekly reflections, and follow-up interviews.

The findings revealed that users' privacy concerns were highly data-dependent: physiological signals and GPS elicited significantly higher sensitivity than other ambient data. Interestingly, it was not just the raw data that worried participants—what concerned them more was how such data could be combined to infer hidden states like mood, stress, or health status. Sensor fusion amplified perceived risks, particularly when the potential for unintended inferences or algorithmic profiling became visible. However, despite these concerns, most participants expressed a general willingness to share their data under vague or open-ended conditions—often due to a lack of awareness about what could be inferred or a sense of resignation about data collection practices.

This highlighted a crucial gap between users' abstract concerns and their actual data-sharing behaviors, shaped by limited understanding and agency. These findings pointed to the need for new tools that support informed and deliberate decision-making—enabling users not only to see what kinds of sensor data are being collected but also to understand their implications. This motivated the next phase of my work, which explores how interactive systems can help users reflect on and manage their sensor data through dynamic and contextual consent mechanisms.

DESIGNING DYNAMIC CONSENT TOOLS FOR USER-CENTERED DATA MANAGEMENT

After identifying key gaps in users' awareness and control over data sharing in the first phase, I focused the second phase on developing tools that help people navigate sensor data decisions in more informed and context-aware ways. I focused on *dynamic consent* as a model that allows for ongoing, context-aware control over data access—moving beyond one-time, static agreements.

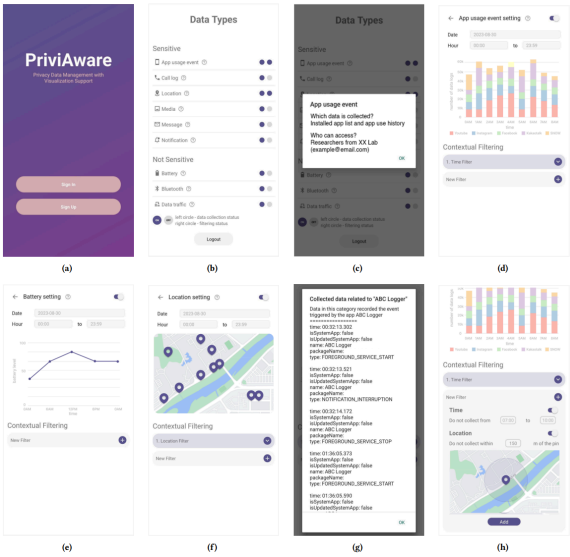

In the first study, I developed and deployed a dynamic consent prototype in a four-week in-the-wild study with 23 participants [2, 3]. The system provided frequent, context-sensitive consent prompts along with interactive visualizations of their collected data. This setup allowed participants to adjust their sharing decisions on the fly while also reflecting on what kinds of data they were comfortable sharing in different situations. While many defaulted to broad sharing settings, certain data types—like location or physiological signals—were selectively withheld depending on context. Participants reported that the visual feedback and contextual filters helped them feel more in control and thoughtful about their privacy, demonstrating the

value of dynamic consent in real-world scenarios. In a follow-up study, I explored how to support users' nuanced control over data sharing by designing and evaluating a contextual filtering interface with interactive data visualizations [4]. The system, **PriviAware**, enabled users to set granular sharing preferences based on context variables such as location, activity, and time, and to preview how their sensor data would be shared under different conditions. Participants (N = 58) interacted with the interface and explored various filtering scenarios during the study. The data visualizations played a critical role in helping users understand what their data revealed and how different settings would impact data exposure. Many participants expressed that seeing their own data in a visual and tangible way made them more reflective and intentional about their sharing decisions. They found contextual filters particularly empowering for managing sensitive information like location or audio in private situations. These findings highlight the importance of combining dynamic consent with visual and contextual tools to address the interpersonal and situational complexities of data sharing in sensor-driven environments.

INTERPERSONAL DATA SHARING AND PRIVACY IN MULTI-USER SENSING ENVIRONMENT

Building on my earlier work designing tools for individual control over sensor data, I next turned to examine how data sharing plays out in multi-user, co-present contexts. While prior studies highlighted the importance of context-aware and granular controls, they left open how people interact with one another's data in shared environments—where sensor information is not only personal but also relational.

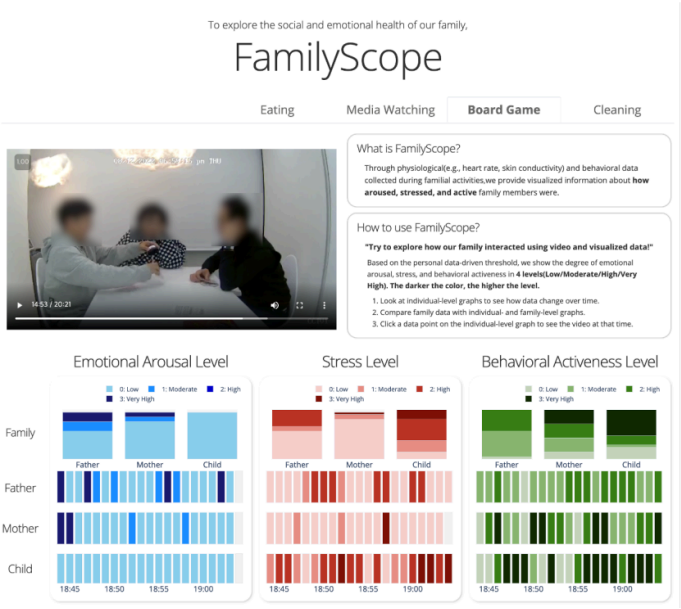

To investigate this, I conducted a three-hour, lab-based field study with 10 families (30 participants) in a smart home–like setting. Participants engaged with an interactive system called **FamilyScope** [5] that visualized multimodal sensor data—such as audio activity, physical activeness, emotion-labeled behaviors—in real time. This setting allowed me to observe how families collectively made sense of the data, responded to perceived patterns, and navigated emergent privacy boundaries.

The study revealed that sensor data often prompted scrutiny: participants interpreted others' behaviors, inferred emotional states, and even confronted each other based on the visualized traces. These interactions surfaced both moments of empathy and discomfort, highlighting that privacy in shared sensing is deeply interpersonal. People's concerns weren't just about what *they* revealed, but also how *others* might misread or over-interpret the

data. Some participants requested hiding or altering parts of the display, while others sought more transparency for coordination or care.

These findings suggest that supporting interpersonal data sharing requires systems that can mediate not just access, but interpretation—offering mechanisms for negotiated visibility, relational context, and co-constructed meaning in sensor-driven environments.

RETHINKING DATA IN SHARED SENSING ENVIRONMENTS: FROM INDIVIDUAL OWNERSHIP TO RELATIONAL AWARENESS

As sensing technologies increasingly capture data from shared routines, conversations, and social interactions, traditional notions of privacy—as personal, discrete, and individually controlled—prove insufficient. In sensor-rich environments like smart homes, a voice command may record multiple people's speech; a motion log may reflect entire household rhythms. Such data is not just about individuals—it is *OurData*: relationally situated and collectively implicating.

To explore how people perceive and navigate this complexity, I conducted a qualitative study with 10 families (N = 40). Using speculative vignettes grounded in real-world sensing contexts, I examined how participants defined shared data, expressed concerns, and imagined fairer ways to manage access and consent. Scenarios varied by sensing modality (e.g., voice, emotion, location) and social context (e.g., parent-child dynamics, guests, external platforms).

Participants did not view group data as an aggregation of individuals' inputs. Instead, they described it as emotionally situated and socially constructed—shaped by roles, routines, and relationships. Tensions emerged in households where one member held more control (e.g., parents, tech-savvy users), exposing gaps in systems that assume static roles or individual decision-making. Many participants emphasized the need for ongoing discussion, transparency, and mutual understanding in managing shared data.

In response, I developed **OurData** (submitted to IMWUT'25) as a sociotechnical framework to conceptualize data not as property to be owned, but as a shared artifact to be negotiated. This reframing sees privacy not as a one-time individual choice, but as an ongoing, relational practice embedded in everyday life.

This perspective calls for a shift in ubiquitous computing—from designing for individual context-awareness to relational awareness. Sensing systems should reflect multi-user dynamics, support co-consent, and recognize how data practices are shaped by evolving relationships and social norms. By embedding these insights into system design, OurData offers a path toward more equitable and accountable data governance in shared sensing environments.

FUTURE RESEARCH DIRECTIONS

My future research centers on designing systems that embed social accountability, privacy, and interpersonal awareness into data-driven technologies. Across three interconnected directions—sensing informatics, multi-user data governance, and socially reflective LLMs—I aim to develop technologies that help people understand, negotiate, and manage their data in ways that reflect real-world relationships and social dynamics.

1. Privacy-Aware Sensing Informatics

Building on systems like LifePensieve (under review for IMWUT'25), which enables users to review and manage temporally organized sensing data for mental health management, and SELaD, which visualizes multi-user sensor data in domestic environments, I develop sensing informatics systems that are scrutable, context-sensitive, and privacy-aware. My work focuses on designing system logic and interaction models that reflect the relational and shared nature of sensor data, particularly in spaces where data boundaries, access, and control must be negotiated among multiple users.

2. Multi-User Data Governance and OurData Systems

I will continue developing data governance frameworks and systems for shared or jointly produced data, particularly in family or collaborative environments. Through my empirical work, I proposed the concept of *OurData*: data that is relationally entangled and collectively implicating. Traditional systems often assume data to be individually owned and managed, but sensor-rich environments challenge that logic. My research has shown how interpersonal dynamics—like power asymmetries, mutual dependencies, and shared routines—demand new approaches to privacy, access, and consent. Building on this foundation, I aim to design multi-user data systems that reflect collective expectations, support co-consent and joint control, and enable socially negotiated access. These systems will respond to real-world tensions in environments like homes, workplaces, or care settings, moving beyond the limitations of single-user consent models.

3. Designing SAFE (Socially Aware, Fair, and Explainable) LLM Systems

My research extends beyond sensing informatics into human-AI interaction, focusing on privacy risks and safeguards in complex AI environments. I am investigating how privacy breaches can occur through interactions between multiple AI agents that share or process sensitive data, analyzing vulnerabilities in these multi-agent systems. At the same time, I am conducting empirical studies on privacy risks arising from users' natural prompts when interacting with large language models (LLMs). This research explores how everyday user inputs can unintentionally expose sensitive information or create privacy vulnerabilities. The goal is to understand these risks in real-world settings and inform future development of privacy-aware LLM systems that better protect users' data and respect social contexts. Together, these efforts aim to develop socially aware and trustworthy AI systems that respect user privacy in dynamic, multi-user contexts.

REFERENCES

[1] Lee, Hyunsoo, Soowon Kang, and Uichin Lee. "Understanding privacy risks and perceived benefits in open dataset collection for mobile affective computing." Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies 6.2 (2022): 1-26.

[2] Lee, Hyunsoo, and Uichin Lee. "Toward dynamic consent for privacy-aware pervasive health and well-being: A scoping review and research directions." IEEE pervasive computing 21.4 (2022): 25-32.

[3] Lee, Hyunsoo, and Uichin Lee. "Dynamic consent for sensor-driven research." 2021 thirteenth international conference on mobile computing and ubiquitous network (ICMU). IEEE, 2021.

[4] Lee, Hyunsoo, Yugyeong Jung, Hei Yiu Law, Seolyeong Bae, and Uichin Lee. "PriviAware: Exploring Data Visualization and Dynamic Privacy Control Support for Data Collection in Mobile Sensing Research." In Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems, pp. 1-17. 2024.

[5] Lee, Hyunsoo, Yugyeong Jung, Youwon Shin, Hyesoo Park, Woohyeok Choi, and Uichin Lee. "FamilyScope: Visualizing Affective Aspects of Family Social Interactions using Passive Sensor Data." Proceedings of the ACM on Human-Computer Interaction 8, no. CSCW1 (2024): 1-27.