The 8th PrivaCI Symposium Report

June 24-25, 2026 at the School of Information Sciences,
University of Illinois Urbana-Champaign

Editors: Madelyn Rose Sanfilippo and Yan Shvartzshnaider

Notetakers: Kyra Milan Abrams, Jake Chanenson, David Eby, Isaac Haizel, Muhammad Hassan, Muhammad Hussain, Inyoung Jang, Felix Oke, Mubarak Raji, Pratyasha Saha, Sophia Wushanley, Yaman Yu

We are grateful for the support: the School of Information Science, University of Illinois Urbana-Champaign; the Grainger College of Engineering, University of Illinois at Urbana-Champaign; the AMD Center of Excellence and AMD HACC at UIUC; Digital Life Initiative, Cornell Tech; PrivaCI Institute for Technology Policy and Social Integrity

Table of Contents

Executive Summary

The 2026 Symposium brought together an interdisciplinary group of scholars to discuss current Contextual Integrity (CI) research and explore urgent privacy challenges with the support of the CI framework.

Beginning with the opening community discussion panel and continuing throughout the event, a core theme was that existing privacy norms are being strained by AI, automation, and cross-context data flows. The difficulty of maintaining contextual boundaries in modern digital systems is only growing more complex, though it is not necessarily producing entirely new challenges. Agentic AI, foundation models, workplace AI tools, remote patient monitoring systems, educational platforms, and IoT devices all create opportunities for data to move across contexts, institutions, and time horizons in ways that are difficult for information subjects and users to anticipate. Many presentations focused on downstream uses, secondary uses, context collapse, overlapping contexts, or future uses of data. AI was often described as amplifying uncertainty by obscuring who is acting, who benefits, who receives information, and how information will be reused. Across domains, symposium participants repeatedly questioned whether existing consent mechanisms, privacy regulations, and even CI itself are sufficient to address increasingly dynamic and interconnected information ecosystems.

Another core pattern was the emphasis on the need to enact CI in practice. Privacy governance must move beyond individual choice models, grounded in alternative conceptions of privacy, and toward system-level, collective, and adaptive approaches supported by the CI framework. Presentations on youth AI use, educational technology, athlete monitoring, workplace AI disclosure, geriatric care, biomedical AI, and remote patient monitoring all highlighted situations where users have limited ability to meaningfully control information flows. Economic necessity, institutional requirements, workplace pressures, health needs, and technical complexity often constrain individual choice. As a result, many projects explored alternatives: community norms, organizational governance, contextual rules, privacy-enhancing technologies, differential privacy, adaptive system design, professional obligations, and collective decision-making processes. The broader message across the symposium was that privacy should be understood not merely as an individual preference but with respect to social norms and as governance challenge requiring coordinated technical, legal, and institutional responses.

Community Discussion Panel

Notetakers: Inyoung Jang & Pratyasha Saha

The Symposium opened with a community discussion of around the impact of AI on contextual integrity. Moderator Madelyn Rose Sanfilippo posed a series of questions first to panelists Helen Nissenbaum, Yan Shvartzshnaider, and Noah Apthorpe, then to the full array of participants, before fielding additional questions from the audience.

The first question explored how agentic AI may fundamentally alter information flows compared with earlier forms of AI by creating new ambiguities around agency, responsibility, and privacy. A recurring theme in responses was that the concept of an "agent" carries social and philosophical assumptions that can obscure who is actually acting, whose interests are being served, and where data is flowing. Participants highlighted concerns that anthropomorphic framing may encourage users to overestimate autonomy, control, or empowerment while masking commercial incentives, conflicts of interest, and underlying corporate control. The conversation also examined how agentic systems complicate traditional understandings of senders, receivers, permissions, and accountability, potentially disrupting established privacy expectations and social norms around information sharing.

The discussion also emphasized the value of the Contextual Integrity (CI) framework for analyzing these challenges because it focuses on information flows and community norms rather than individual consent alone, by comparison to other privacy perspectives. Participants suggested that agentic AI may require new ways of mapping information flows, understanding transmission principles, and evaluating conflicts between user interests and organizational incentives. Several participants questioned whether existing regulatory, legal, and ethical frameworks are sufficient for highly capable agents operating with broad permissions, while others pointed to opportunities for future research on governance, liability, workplace disclosures, privacy boundaries, default settings, delegation of tasks, and the normative constraints needed to ensure agentic systems advance beneficial goals without creating new forms of manipulation, surveillance, or harm.

Session 1: CI and User Perceptions

Notetakers: Yaman Yu & Felix Oke

This session, chaired by Noah Apthorpe (Colgate University), included three papers exploring how users perceive information flows, as well as the divergence between data practices and social norms.

Alexis Shore Ingber (Syracuse University, S.I. Newhouse School of Public Communications) and Nazanin Andalibi (University of Michigan School of Information)
"The Role of Data Distance on Perceptions of Data Privacy: Distinguishing AI Inferences and Predictions"

AI has expanded the concept of personal data by generating inferences and predictions that are temporally distinct from the original information used to create them. Drawing on communication and psychological theories of distance, this work introduces the concept of "data distance," defined as the perceived psychological distance between an individual and an AI-generated output. The central argument is that existing privacy and legal frameworks do not adequately distinguish between user-generated data and AI outputs that may describe a person's present state, near future, or distant future. Because people tend to perceive future events as more abstract and psychologically distant, AI-generated predictions may warrant different normative and legal treatment than more temporally proximate inferences. Contextual Integrity (CI) provides the foundation for this work by examining how information flows are evaluated within specific social contexts, while data distance extends the framework by adding a temporal dimension to those evaluations.

To investigate these questions, the study employed a vignette-based online survey (n = 598) that varied information type, data distance, and recipient across contexts including advertisers, government agencies, technology firms, and medical teams. Findings suggest that participants view the flow of temporally proximate AI-generated data as more appropriate than the flow of temporally distant outputs, particularly when information concerns identity or major life events. Similarly, participants reported greater feelings of ownership over AI-generated data that was perceived as closer in time to their current circumstances. An unexpected finding was that higher levels of digital resignation were associated with viewing AI output flows as less appropriate overall. The study also highlighted methodological challenges, particularly the difficulty of assigning a clear purpose to AI-generated outputs, whose intended uses are often ambiguous. To address

this issue, participants evaluated AI outputs framed around public-interest purposes, helping isolate the effects of recipient and data distance.

Future work will explore how data distance may inform analyses of disparate impacts in algorithmic decision-making and how the concept applies to interactions with AI chatbots. More broadly, the project suggests that temporal characteristics of AI outputs shape perceptions of appropriateness, ownership, and privacy in ways that existing frameworks do not fully capture, raising important questions about identity, change over time, and the governance of AI-generated information.

Brad Stenger and Yuanyuan Feng (University of Vermont)

"Evidence of College Athletes' Trust and Self-belief Attitudes in the Gray Area of Data Privacy Expectations"

Personal tracking technologies (PTTs), such as wearables and athlete monitoring systems, shape information flows within collegiate athletics. As data is collected and shared on sleep, nutrition, injuries, recovery, training load, performance, and other aspects of athlete wellbeing, complex information flows arise among athletes, coaches, medical staff, teammates, and institutions. Using Contextual Integrity (CI) as an analytical framework, the project seeks to identify "gray areas" where norms governing data sharing are contested, unclear, or in tension. The authors considered how athletes navigate the competing demands of individual privacy and collective team accountability, and how expectations about appropriate information sharing vary across different types of data and use cases.

The study draws on seven semi-structured interviews with college athletes and a personal tracking technologies questionnaire. Interview data were used to map information flows across multiple scenarios, including training load, sleep, injuries, setbacks, nutrition, menstrual health, and team leaderboards. Analysis revealed that athletes generally described their experiences and data practices in strongly individual terms, despite operating within a fundamentally collective team environment. This suggests that athletes often understand themselves as distinct information subjects whose data remains personally meaningful, even when shared for team purposes. At the same time, attitudes toward information sharing varied substantially by data type. Injury and setback data elicited more negative responses, while sleep and nutrition data were viewed more

positively, indicating that contextual norms are shaped not only by recipients and purposes but also by the nature of the information itself.

The project also highlights several conceptual and methodological challenges. High levels of trust in coaches and athletic institutions may lead athletes to retrospectively justify data-sharing practices that could otherwise be perceived as privacy violations, complicating efforts to identify normative boundaries. In addition, the distinction between “I” and “we” raises questions about how CI should be applied when individuals operate within highly coordinated groups. Team sports blur the line between personal and collective contexts, making it difficult to determine whether information-flow norms belong primarily to individual athletes, the team, or both simultaneously. Ongoing survey research and vignette-based studies will further investigate these tensions by systematically examining how data type, sender, recipient, purpose, and transmission principles shape perceptions of appropriate information flows in athletic settings.

Jake Chanenson and Marshini Chetty (University of Chicago)

“Privacy in the Classroom: A Study of Student Perspectives of Educational Technology in School”

Students are often forced to use edtech platforms and systems that routinely collect academic, behavioral, and interaction data as part of everyday educational activities. This study examines how high school students perceive privacy on educational technologies that are ubiquitous in their education. Using Contextual Integrity (CI) as a framework, the project investigates students’ mental models of educational data flows, including what information they believe is being collected, who has access to it, and which uses they consider appropriate. A typical educational data flow involves a student submitting coursework through a learning management system to a teacher for assessment, but the increasing complexity of digital learning environments raises questions about how students understand and evaluate these exchanges.

Drawing on nineteen semi-structured interviews with students in grades 9 through 12, the study found that most participants were generally comfortable with the collection and sharing of academic data within what they perceived to be a conventional educational situation. Students frequently viewed educational data as fundamentally different from social media data and often assumed that platforms adopted by schools were trustworthy and safe, reflecting what the authors called a broader "school halo effect." Participants appeared to evaluate data-sharing practices primarily through the lens of purpose rather

than information type, suggesting that educational objectives often legitimized information flows that might otherwise raise concerns. Comfort levels were not uniform, however, and varied depending on the recipient and intended use of the data. Students were generally most comfortable with parents accessing educational information, often more so than teachers or administrators, while expressing considerably greater concern about peers gaining access to the same information.

The findings also reveal important challenges to educational privacy norms. While students' intuitions that academic data should serve academic purposes aligns well with existing norms, monitoring systems such as device-tracking and behavioral surveillance tools generated substantially greater discomfort as they diverged from norms. Approximately half of participants reported feeling uneasy about these technologies, distinguishing routine data collection from active forms of surveillance. These systems can blur the identity of recipients, particularly when AI-enabled monitoring tools route information through opaque technical infrastructures, making it difficult for students to determine who ultimately accesses their data. The study further suggests that apparent acceptance of educational data practices may sometimes reflect compliance rather than genuine normative approval, as participation in many edtech systems is mandatory. This raises broader questions about whether students' responses reflect established contextual norms, emerging expectations, or simply adaptation to environments in which meaningful alternatives are unavailable.

Summary:

The three presentations in this session collectively examined how contextual integrity (CI) extends — and strains — in data environments shaped by AI inference, institutional tracking, and mandatory technology adoption. Shore Ingber and Andalibi introduced "data distance" as a CI-compatible conceptual tool that adds a temporal dimension to the standard CI parameters: people hold temporally proximate AI outputs to different normative standards than distant ones, with perceived appropriateness and ownership both sensitive to how far removed an output is from the moment of data collection. Stenger and Feng applied CI to the gray areas of college athlete tracking, finding that athletes' privacy expectations are shaped by trust in their institutions and by a tension between individual identity and collective team accountability — a tension CI's standard parameters do not cleanly resolve. Chanenson and Chetty examined high school students' mental models of educational technology, finding that student intuitions broadly track CI's contextual norms but that two conditions — diffuse recipient identity in monitoring

tools and acceptance under coercion — create situations that superficially satisfy CI norms while undermining them.

The Q&A surfaced three recurring tensions across all three papers, regarding context specificity and what "ownership" means within CI. Further, a question about emerging norms pointed to a productive gap: when contextual norms are still being negotiated (as with educational technology), it becomes difficult to determine whether CI is describing what norms are or prescribing what they should be.

These papers point toward productive new directions for CI research. The data distance concept offers a temporal axis that the standard CI parameters (sender, recipient, information type, transmission principle) do not currently capture, and the forthcoming extensions to disparate impact claims and AI chatbot interactions suggest it has broad applicability. The gray area work in athletics and education both point toward the need for CI frameworks that can account for institutional power asymmetries — cases where the appearance of norm satisfaction may reflect coercion or resignation rather than genuine contextual alignment. More broadly, AI and algorithmic outputs pose challenges for traditional conceptualization of information flows, raising important new questions about recipient identity, temporal origin, and purpose.

Session 2: CI and Theory

Notetakers: Jake Chanenson & Isaac Haizel

This session, chaired by Helen Nissenbaum (Cornell Tech), included two papers exploring conceptual privacy challenges via the lens of contextual integrity.

Noah Apthorpe (Colgate University)
"Integrating Contextual Integrity and Data Minimization"

This presentation explored what types of actions would uphold both Contextual Integrity (CI) and Data Minimization (DM), defined by the GDPR principle that personal data should be "adequate, relevant and limited to what is necessary." While the two approaches share a common concern for privacy protection, they evaluate practices according to different criteria. CI focuses on whether information flows align with contextual norms, social expectations, and appropriate transmission principles, whereas DM centers on whether data collection and use are necessary for a specified purpose.

This comparison highlighted instances in which a practice may appear acceptable under one framework while raising concerns under the other, as well as how CI and DM could serve as complementary frameworks to evaluate the legitimacy of data practices.

A central contribution of the work was identifying vulnerabilities that can undermine each framework when applied in isolation. For CI, a key concern is "context gaming," in which organizations strategically shape user expectations or exploit contextual cues, such as through dark patterns, to make questionable information flows appear acceptable. For DM, vulnerabilities arise when organizations define purposes or claims of necessity in ways that justify broader data collection than users might reasonably expect. These weaknesses are largely orthogonal rather than overlapping: CI is well positioned to identify violations of social and contextual norms, while DM is better suited to scrutinize claims about functional purpose and technical necessity. Because functional requirements are not always reflected in contextual transmission principles, and contextual expectations do not always constrain assertions of necessity, each framework can detect privacy concerns that the other may overlook.

The presentation illustrated these dynamics through case studies involving Discord's age assurance mechanisms and Meta's use of AI-driven advertising personalization. These examples demonstrated how evaluations can diverge depending on whether the primary focus is contextual appropriateness, user expectations, technical necessity, or organizational purpose. User opinions and prevailing social norms may align with the conclusions generated by CI, while DM may reach different judgments based on whether data collection is genuinely required to achieve a stated function. The discussion concluded that CI and DM should be viewed not as competing approaches but as complementary tools. Combining them enables a more comprehensive assessment of privacy practices, increasing the likelihood of identifying problematic data uses and creating opportunities for closer collaboration between the two research communities in developing more robust and streamlined methods of privacy evaluation.

Sophia Wushanley (University of Michigan)
Contextual Integrity's Privacy Problem

Wushanley explored interpretations of Contextual Integrity (CI), ranging from viewing CI as the equivalent of privacy to a framework describing how appropriate flows should replace privacy. This presentation also examined privacy through what was termed a "stoppage problem," emphasizing that privacy disputes often center on decisions about

whether information should be permitted to flow or whether access should be restricted. Using the example of a physician deciding whether to disclose a patient's disease status to public health authorities or individuals who may have been exposed, the presentation highlighted how privacy questions frequently arise when important social goods depend on the sharing of information. The example illustrates that privacy is rarely an isolated value; instead, it exists alongside competing concerns such as public safety, justice, accountability, and individual liberty.

A central argument was that framing privacy exclusively as a barrier to information flow can contribute to privacy skepticism and the perception that privacy is inherently anti-innovation or obstructive. Rather than rejecting this characterization entirely, the presentation suggested accepting that privacy often functions as a constraint on information sharing while recognizing that societies simultaneously pursue multiple normative and informational goods. From this perspective, privacy losses may sometimes be justified when outweighed by other legitimate values. Decisions about information disclosure therefore require balancing privacy against considerations such as transparency, fairness, public health, or collective welfare, making privacy an "all-things-considered" concept rather than an absolute rule against information sharing.

The discussion also raised important conceptual questions about the distinction between privacy and secrecy. Audience members noted that the absence of information flow does not automatically constitute privacy, since privacy is fundamentally concerned with rights, protections, and legitimate barriers governing access to information rather than simple concealment. Unlike secrecy, which carries no inherent moral valence and may be viewed positively or negatively depending on context, privacy was characterized as a more neutral normative framework for determining when restrictions on information access are justified. This distinction reinforces the idea that privacy should not be understood as preventing all information flows, but rather as establishing appropriate conditions under which information may be shared, withheld, or protected.

Summary:

This session raised a number of themes that can help to guide conceptual contextual integrity research moving forward, as well as offered an opportunity during Q&A to explore the implications of their work for understanding norms around new tools that facilitate computation without transferring user data.

One important theme from this session was that both papers challenge the idea that privacy can be evaluated in isolation. The Data Minimization (DM) and Contextual Integrity (CI) comparison argued that privacy judgments require multiple evaluative lenses because contextual appropriateness and technical necessity do not always align. Similarly, the discussion of privacy as a "stoppage problem" emphasized that privacy must be balanced against other social values such as justice, liberty, transparency, and public welfare. In both cases, privacy is but one normative consideration among many.

Second, both emphasize the importance of justification rather than simple information disclosure or nondisclosure; the transmission principle parameter of the CI framework offers significant structuration to explore the conditions under which flows are or are not appropriate. Beyond considering whether a data practice is justified according to contextual norms, user expectations, purpose, or necessity, there are important implications competing values which may result in privacy loss or inappropriate flows. Both presentations therefore move beyond a binary "share versus don't share" view of privacy and instead ask what makes a particular information flow legitimate.

Third, both highlighted the strengths of contextual integrity over the limitations and vulnerabilities of competing privacy frameworks. Both suggest that many other privacy models can produce blind spots that require broader theoretical approaches and emphasized the importance of normative governance of information flows for legitimacy. Implicitly, both presentations ask the same underlying question: *What makes an information flow appropriate, and who gets to decide?*

Session 3: CI and Privacy Risks

Notetakers: Mubarak Raji & Inyoung Jang

This session, chaired by Alexis Shore Ingber, explored decision-making regarding and perception of privacy risks, both practically and empirically, via two papers.

Ece Gumusel and Madelyn Sanfilippo (University of Illinois Urbana-Champaign)
Multiverse Privacy Theory: Modeling Context-Dependent and Time-Sensitive Privacy Decisions

Modern data-driven systems, particularly large language models (LLMs), require individuals to make privacy judgments that are dynamic, situational, and often time-sensitive. This work aims to understand how the countless permutations of situations, contextual changes, and uncertainties shape people's privacy decisions and perceptions of risk by integrating the CI framework with multiverse theory. While CI provides a valuable normative framework for evaluating information flows, it treats privacy decisions as either relatively static and binary, or considers only one time slice. The project introduces Multiverse Privacy Theory (MPT) as a probabilistic framework designed to capture the uncertainty, individual variation, and contextual complexity that characterize contemporary privacy decision-making.

MPT conceptualizes privacy decisions as a set of alternative "privacy universes," each representing a different combination of situational factors that may influence perceptions of appropriateness and risk. Rather than assuming a single correct privacy judgment, the framework models how evaluations can vary across individuals and contexts and how uncertainty affects those evaluations. This approach is particularly relevant in environments where contextual boundaries are blurred, overlap, or conflict, raising questions about context collapse and the application of norms from one context to another. The framework also offers a way to examine how privacy norms may evolve over time as technologies, social practices, and user expectations change. Consistent with broader findings in communication and privacy research, the work recognizes that higher levels of uncertainty, such as those associated with future outcomes or unfamiliar technologies, make risk assessment more difficult and can alter perceptions of acceptable information flows.

The conceptual approach was evaluated through a vignette-based survey involving 500 participants across 96 LLM-related scenarios. These responses, coupled with more than 1.2 million posterior predictive simulations, revealed that perceived appropriateness is the strongest driver of privacy evaluations. The results also demonstrated substantial variation across individuals and contexts, supporting the need for approaches that move beyond binary judgments of acceptable and unacceptable information practices. MPT successfully distinguished among low-, medium-, and high-risk privacy outcomes within a unified probabilistic model, suggesting a promising avenue for understanding privacy under conditions of uncertainty. More broadly, the work points toward privacy governance approaches that can account for evolving norms, contextual complexity, and the diverse ways individuals interpret and respond to information flows in rapidly changing technological environments.

Zihao (Raio) Huang, Quinn Waeiss (Center for Biomedical Ethics, McCoy Family Center for Ethics in Society, Stanford University)

How Researchers Identify and Mitigate Downstream Privacy Risks: A Contextual Integrity Analysis

These researchers examined how privacy risks are recognized, articulated, and governed within the research process through analysis of proposals reviewed under Stanford's Ethics and Society Review (ESR) framework. The study focused on how researchers anticipate potential harms arising from data collection, sharing, and reuse, as well as how the rights and interests of downstream users are considered during project design. Information flows described in funding proposals were coded using contextual parameters, allowing the researchers to systematically map how data moves across actors, contexts, and stages of use. The relationship between original research purposes and the broader ecosystem of downstream applications that may emerge after data, models, or findings leave the immediate research setting was considered to identify emergent risks.

The analysis revealed that privacy risks were not consistently foregrounded in proposal narratives. Of the 52 proposals examined, only 19 explicitly discussed privacy concerns. Moreover, 17 of the 23 distinct information flows identified involved cross-contextual movement of information, often through secondary uses, downstream deployment, or combinations of both. Despite these patterns, privacy risks were frequently framed in narrow terms and often stopped at concerns related to identification. Commonly cited risks included identity theft, surveillance, harmful inference, misuse, data leakage, unauthorized access to sensitive information, mass data collection, invasion of privacy, and data breaches. The findings suggest that researchers may underestimate the broader implications of information flows that extend beyond their immediate projects, particularly when data or technologies are later adapted for new contexts and purposes. At the same time, the discussion emphasized that downstream actors should not automatically be viewed as malicious, as information can move through a spectrum of legitimate and unintended uses between the original research context and broader public deployment.

The presentation also highlighted the importance of understanding how institutional, temporal, and disciplinary factors shape privacy governance. Questions emerged about how privacy protections can be preserved as data and technologies travel across organizational boundaries, how downstream recipients may differ in their obligations and

capabilities, and how reproducibility requirements may interact with privacy considerations over time. Participants discussed the need to account for varying types of recipients, including potential roles for techniques such as differential privacy, while also considering how governance mechanisms affect day-to-day research practice. Overall, the work demonstrated that CI reveals a much richer picture of privacy risk than traditional assessments, particularly by drawing attention to cross-contextual information flows, secondary uses, and the evolving relationships between researchers, downstream users, and the broader public.

Summary:

This session was characterized by the theme that privacy decisions are increasingly shaped by uncertainty, complexity, and changing contexts, which necessitates change in governance and oversight. In both cases, privacy cannot be understood as a fixed property of a single information flow because it evolves as data moves across contexts, recipients, and future or downstream uses.

These downstream, and cross-contextual, information flows are important privacy challenges that are not well accounted for by most privacy governance in effect. Many policies and procedures don't well account for what happens after information leaves its original context. Both projects suggest that privacy risks often emerge not from the initial collection of data but from its movement into new environments where expectations, norms, and governance structures may differ from the original context.

Further, this session emphasized the importance of grounding privacy governance in the CI framework, as practical frameworks need to better account for variation in risk perception and governance. These works call for privacy governance frameworks that move beyond binary notions of acceptable versus unacceptable data use and instead account for uncertainty, diverse stakeholders, evolving norms, and the multiple possible futures that may emerge from a single information flow. Such changes can better support the complexity and dynamics of the modern information environment.

CI Community Feedback Session #1

Notetakers: Isaac Haizel & Felix Oke

In the first CI community feedback session, moderated by Christopher Muhawe (University of Illinois Chicago), three works in progress were presented.

Kyra Abrams (University of Illinois at Urbana-Champaign)
Contextual Implementation of Data Autonomy

This work explores the concept of data autonomy through the lens of CI, proposing autonomy as a key condition for determining whether information flows are appropriate. Drawing on philosophical work by Miller and Braumwich concerning the requirements for being "fully informed," the project examines whether individuals can meaningfully exercise autonomy over their data in contexts where choices may be constrained, poorly understood, or shaped by complex digital environments. The discussion highlighted how perceptions of data autonomy may sometimes reflect an illusion of control rather than genuine agency, raising important questions about whether individuals truly understand and can influence the collection and use of their information. Q&A emphasized the strong fit between this line of inquiry and Contextual Integrity's focus on appropriate information flows, while also suggesting future methodological enhancements, including supplementing qualitative recruitment and interview approaches with a concise survey to broaden empirical insights.

Yuxing Yang (University of Illinois Urbana-Champaign)
Contextual Integrity and Age: Rethinking Transmission Principles Across Age

This work explores whether attributes of the information subject, such as age, should be into the CI framework. The author suggests that age shapes expectations about appropriate information flows, particularly distinguishing perceptions of transmission principles. Treating age as an attribute of the information subject, the project investigates whether privacy norms vary systematically across age groups and whether technologies should account for such differences when evaluating the appropriateness of data sharing. Analysis thus far focuses on prompts by systems to re-post a user’s material across time or context. The discussion highlighted challenges in isolating age as the primary explanatory variable, particularly in examples such as social media features that encourage users to revisit or share past content. In these cases, expectations about disclosure may reflect assumptions about nostalgia, memory, platform design, or social context rather than age alone. Workshop participants also raised broader questions about the interaction between age and other contextual factors, including culture and social norms. Examples such as dating applications and differences between Korean and U.S. expectations illustrated how privacy norms may vary across cultural settings, creating tension with technological systems that seek to apply a single standardized set of rules

across diverse populations. More generally, the discussion underscored the difficulty of determining whether contextual parameters should be treated independently or understood as deeply intertwined with particular actors, cultures, and social environments.

Sam Hafferty, Daniel Flyer, Nitya Nadgir, Prateek Mittal, Mihir Kshirsagar, Mona Wang (Princeton University) and Shaanan Cohney (Melbourne University)

Between Safety and Surveillance: Online Activity Monitoring Products and Student

Hafferty presented preliminary results from a study of online student activity monitoring products structured by the CI framework, focusing on the tension between student safety and digital surveillance in educational settings. Using simulations of student, teacher, and administrator experiences on configured Google Chromebooks, the research analyzed monitoring features, network activity, data flows, system logs, and metadata generated by commonly used school surveillance technologies. By tracing how information is collected, transmitted, and made visible to different stakeholders, the study evaluated whether these flows align with educational privacy norms and expectations. The CI framework provided a way to assess the appropriateness of data sharing among human and technological actors, while also raising questions about how non-human systems should be represented within CI's parameters. Discussion further explored the relationship between contemporary digital monitoring and more traditional forms of surveillance, highlighting how automated collection, aggregation, and analysis of student activity may reshape the balance between protecting students and preserving their privacy.

Session 4: CI and Health AI

Notetakers: Sophia Wushanley & David Eby

This session, moderated by Masooda Bashir (University of Illinois Urbana-Champaign), included two papers that empirically apply the CI framework to information flows associated with Health AI.

Jiunn-Tyng Yeh (Duke University) and Margot Hanley (Duke University), Yan Shvartzshnaider (York University)

Applying Contextual Integrity to Biomedical Foundation Models: An Analysis of a Multimodal Sleep Foundation Model

The governance of biomedical foundation models (FMs), which are increasingly trained on diverse forms of health-related data, including electronic health records, biometric information, medical archives, and data collected through clinical and consumer technologies, raises important unanswered questions. Biomedical foundation models are designed to support tasks such as diagnosis, prognosis, simulation, and clinical decision-making, but their development often involves combining data originating from different contexts, purposes, and consent regimes. The work focused on the challenge of understanding where training data comes from, how it moves across institutional boundaries, and what privacy and ethical obligations attach to data gathered under different regulatory and social expectations. The presentation explored how CI can support governance of downstream uses and appropriate flows in practice.

A central case study involved a sleep foundation model that integrates information from multiple sensors to support sleep diagnosis, analysis, and improvement. While such systems offer significant potential benefits for healthcare and research, the discussion emphasized that the same data and models may later be repurposed for objectives beyond their original intent. This creates uncertainty about consent, governance, and appropriate use, particularly because foundation models are inherently general-purpose technologies whose future applications may not be fully known at the time data is collected. Questions emerged about how information flows should be evaluated when data gathered for one health-related purpose becomes part of a model that can support many downstream uses, potentially extending far beyond the context in which individuals originally shared their information.

The presentation framed these challenges as an opportunity to apply CI to biomedical AI governance. Key research questions focused on how biomedical data should be governed when used to train foundation models, how downstream and secondary uses should be evaluated, and whether CI can provide a systematic framework for assessing the appropriateness of these evolving information flows. Uses that cross contextual and institutional boundaries raise risks, such as the possibility that health-related models or outputs could inform insurance risk assessments or other forms of decision-making unrelated to the original purpose of data collection. Overall, the discussion suggested that CI offers a promising approach for understanding and governing biomedical foundation models by highlighting the relationships among data sources, consent expectations, intended purposes, and the broader ecosystem of downstream applications.

Thalia Viranda, Jia Bo, Shupeng Wang, Dan Adler (Cornell University), Andrea Cuadra (Olin College of Engineering), Deborah Estrin (Cornell University), Tanzeem K. Choudhury (Cornell University),Walter Richard Boot (Weill Cornell Medicine), Sara Czaja (Weill Cornell Medicine)

From Adaptive Personalization to Accountable Care: Bridging the (Un)usable Privacy Gap in Voice-Based Health AI

Voice assistants and conversational agents play a growing role within geriatric care ecosystems, which involve complex interactions among older adults, family members, paid care partners, and clinicians. These technologies are increasingly being introduced to support health management, daily living, and social connection, particularly as concerns about social isolation among older adults continue to grow. At the same time, companies and healthcare organizations are beginning to develop AI-enabled tools specifically for patient care, including voice-based systems designed to support health-related communication and decision-making in home environments.

The project explored how large language model (LLM)-based voice assistants for home-based geriatric care might be integrated into clinical workflows through an AI health assistant delivered on a tablet device, with intentionally limited personalization to reduce privacy risks while still providing meaningful support. The discussion highlighted the increasing use of passive sensing technologies and health interventions in the home, raising important questions about how sensitive health information is collected, processed, and shared across multiple stakeholders. Because many current IoT devices,

voice assistants, and AI services were originally designed for consumer rather than clinical settings, their deployment in healthcare contexts introduces new challenges related to privacy, accountability, and appropriate information flows.

The presentation also emphasized a gap between privacy-aware sensing research, which often focuses on technical methods for protecting data, and usable privacy research, which examines how individuals understand and manage privacy in practice. Bridging this gap was presented as a key goal for future work on voice-based health AI systems. More broadly, the research argued that privacy considerations cannot be treated as secondary concerns in geriatric care technologies, particularly because these systems operate within a highly interconnected ecosystem involving patients, caregivers, clinicians, and technology providers. As AI-powered health assistants become more common, understanding how information flows across these relationships will be critical for developing systems that effectively support care while respecting the privacy expectations and needs of older adults.

Summary:

While focused on very different contexts of medical AI use, both papers emphasize the challenges associated with applying AI to sensitive health data, particularly when spanning contexts and stakeholders. Health information flows associated with these technologies involve many actors with different roles, responsibilities, and expectations.

The CI framework, therefore, offers a useful lens through which to govern AI-enabled healthcare technologies. Each project raises questions about appropriate information flows, contextual boundaries, and the alignment between technological capabilities and social expectations. Whether the issue is training foundation models on heterogeneous biomedical data or deploying voice-based assistants in home-care settings, both studies argue that privacy cannot be understood solely through technical safeguards. Instead, effective governance requires attention to context, stakeholder relationships, purposes of use, and the norms that determine when health-related information should appropriately flow across institutional and personal boundaries.

Further, echoing sessions from Day 1, these projects both also emphasize downstream uses and the problem of general-purpose AI systems. Biomedical foundation models may be trained for one purpose but later support very different applications, including secondary uses that participants did not initially anticipate. Likewise, voice assistants and

AI companions introduced into geriatric care may collect information that becomes useful for care coordination, monitoring, intervention, or other future purposes. Both projects highlight uncertainty about how data and AI outputs will be used over time and whether existing consent and governance mechanisms adequately account for those future possibilities.

Session 5: CI and Social Norms

Notetakers: Muhammad Hussain & Kyra Abrams

This session, moderated by Yan Shvartzshnaider (York University), presented two papers and one use case that all aimed to understand specific contextual norms.

Alex Chanhou Lou (Cornell Tech, Cornell University)
Contextual Integrity Theory in China: An Intellectual History from Interpretation to Reflexivity

This presentation examined the intellectual history of CI in China, tracing how the theory has been interpreted, translated, and adapted within Chinese scholarship, legal theory, and policy discussions. The theme of "theory in translation" highlighted that concepts developed in one cultural and legal setting may take on new meanings when introduced into another. The discussion explored how CI has gained influence in China through the translation and circulation of *Privacy in Context*, becoming an increasingly important framework for understanding privacy, information governance, and data protection. Rather than being adopted unchanged, CI has evolved through engagement with local legal traditions, social norms, and regulatory priorities.

The presentation also considered how cultural values shape the interpretation and application of contextual privacy norms. As Chinese society, technology ecosystems, and data governance regimes continue to evolve, scholars have used CI to examine the relationship between changing social expectations and emerging forms of data collection and use. Chanhou Lou explored whether privacy should be understood primarily as an individual right or as a value embedded within broader social and collective relationships. This raises important questions about how contextual norms differ across societies and whether CI requires cultural grounding or adaptation when applied beyond its original intellectual context.

A key argument was that CI has become more than a descriptive theory within Chinese legal discourse; it increasingly functions as a doctrinal instrument that helps scholars and policymakers reason about appropriate information flows and legal protections. The presentation therefore traced a movement from interpretation to reflexivity, in which Chinese scholars not only apply CI to local problems but also use their experiences to reflect on and potentially extend the theory itself. By examining the interaction between cultural norms, legal institutions, and privacy values, this project highlights how the cross-cultural adoption of CI can deepen understanding of both privacy governance and the theory's broader applicability across diverse social contexts.

Glencora Borradaile, Alexandria LeClerc, Kelsy Kretschmer (Oregon State University)
Assessing privacy successes and failures through the lens of contextual integrity: Lessons from social movements

Borradiale presented this work, exploring privacy successes and failures within social movements from a CI perspective. Drawing on interviews with 40 activists representing 33 different organizations, the study explored organizational structures, communication practices, decision-making processes, and the technologies groups use to coordinate their activities. Results highlighted how the adoption of privacy-enhancing technologies (PETs) corresponded with specific and what factors influence intentions to protect privacy. Rather than treating privacy as a simple technological outcome, the project examined how privacy decisions emerge from the social, organizational, and strategic realities faced by activist groups.

A key insight was the distinction between privacy intentions and actual privacy practices. The research proposed a framework that considers both the level of intention to protect privacy and the extent to which groups use PETs, creating a spectrum of potential privacy successes and failures. This approach raises important questions about whether privacy should be measured primarily through technology adoption or through the motivations and goals that underlie organizational choices. The discussion suggested that some groups may intentionally avoid PETs in order to signal legitimacy, openness, or “mainstreamness,” while others may adopt such technologies to communicate concerns about safety, surveillance, or political risk. As a result, low adoption of PETs may not always indicate a privacy failure, just as high adoption does not necessarily reflect a well-developed privacy strategy.

The presentation also highlighted broader questions about collective privacy, norm formation, and group decision making. Symposium participants discussed how privacy expectations are negotiated within organizations, how consensus develops around communication practices, and how group-level privacy concerns may differ from those of individual members. The work suggested that both social movement research and Contextual Integrity are fundamentally concerned with information flows and contextual decision making, yet these perspectives have rarely been integrated.

Smirity Kaushik (George Washington University), Shiza Ali (George Washington University) and Adam Aviv (George Washington University)

Understanding AI disclosure norms at workplace using CI framework: A Use Case

Motivated by the rapid adoption of AI tools, with a substantial share of workers now using AI in their professional activities, the project focused not on whether employees should use AI, but on when, how, and to whom that use should be disclosed. The motivating scenario involved a newly hired software engineer who receives no guidance during onboarding about workplace expectations surrounding AI use. This ambiguity highlights a growing gap between the widespread availability of AI tools and the absence of clear organizational norms governing transparency about their use.

To investigate these issues, the study employed a vignette-based survey of 600 participants structured by the CI framework in order to assess what parameters of information flow shaped perceptions. The research centered on two primary questions: what information norms govern AI-use disclosure in workplace settings, and which Contextual Integrity parameters influence judgments about whether disclosure is appropriate. The project explored how factors such as the recipient of the disclosure, the purpose of the task, and the degree or granularity of AI involvement shape expectations. Rather than treating AI use as a binary condition, the research considered varying levels of assistance, raising questions about where people draw the line between acceptable undisclosed support and AI involvement that should be communicated to others.

A key motivation for the research is evidence that employees may face social consequences for acknowledging AI use. The presentation discussed findings suggesting a "competency penalty," in which workers who disclose AI assistance may be perceived less favorably, with reported penalties differing across gender. These dynamics create

tension between transparency and professional reputation, potentially discouraging disclosure even when organizations value honesty and accountability. By applying CI, the project seeks to identify the contextual factors that shape disclosure expectations and to better understand how workplace norms around AI use are developing as these technologies become routine components of professional work.

Summary:

All three projects emphasize that privacy cannot be understood through one universal definition, but rather privacy norms are socially constructed and context-dependent. The China presentation explored how Contextual Integrity (CI) is translated and adapted within different cultural, legal, and social traditions, demonstrating that privacy norms vary across societies. The social movements study showed that privacy decisions are shaped by organizational goals, collective identities, and activism contexts rather than by technological considerations alone. Similarly, the workplace AI disclosure project examined how expectations about disclosure depend on factors such as professional roles, task contexts, and organizational environments. These studies reinforce a core CI insight: judgments about appropriate information flows emerge from specific social contexts rather than from abstract principles.

Further, groups, communities, and institutions shape privacy expectations. Across all three projects, privacy is presented not merely as an individual preference but as something negotiated and maintained through collective norms and social relationships. These studies highlight the need to understand not only formal privacy rules but also how norms are interpreted, negotiated, and enacted in everyday practice. The Q&A session importantly highlighted opportunities for future directions and the need to be precise and systematic in applying CI parameters in survey work, both to support interpretation and synthesis.

CI and Community Feedback Session #2

Notetakers: Muhammad Hassan & David Eby

The second community feedback session, moderated by Smirity Kaushik (George Washington University), featured three works in progress.

Lili Dudas, Aljawharah Alzahrani, Satish Kumar Keshri, Yu-Jye Tung, Tanusree Sharma (Penn State)
Contextual Integrity in Youth AI Use: Privacy Norms Across Shared and Personal ChatGPT Account Contexts

This presentation explored privacy norms surrounding youth use of ChatGPT across both personal and shared account contexts, with particular attention to the constraints that shape privacy decision-making in different social and economic settings. Drawing on examples involving children, adolescents, and university students in the Global South, the work argued that many privacy choices are better understood as "impossible choices" rather than simple conflicts or paradoxes, since users may have limited ability to avoid sharing information when access to AI tools depends on shared devices, shared accounts, or economic necessity. The discussion highlighted cases from Bangladesh and Saudi Arabia where practical and financial considerations can compel account sharing, creating heightened privacy risks for sensitive queries and personal information. These dynamics raise important questions about responsibility and accountability, including whether the burden of protecting privacy should fall primarily on young users, families, institutions, or the providers of AI systems themselves, particularly when users have little meaningful choice over the conditions under which they access the technology.

Erin Gilmore (San Jose State University), Iori Khuhro (University of British Columbia), Darra Hofman (San Jose State University), Jim Suderman (Independent Expert), Frédéric Simard, Michel Barbeau, Mario Beauchamp, Nicholas Rivard (Carleton University)
Contextual Integrity and Age: Rethinking Transmission Principles Across Age

This presentation examined how CI can be applied to archival records and long-term information stewardship, with a particular focus on transmission principles, provenance, and privacy protections across the lifespan of records. Drawing on work associated with InterPARES Trust AI, the project explored the use of knowledge graphs and archival

metadata to model relationships among records, actors, and information flows. A central concern was how sensitive archival materials should be governed when records persist across generations and contexts, raising questions about appropriate transmission principles and privacy protections for both archivists and information subjects. By mapping provenance records into the CI framework, the work sought to better understand how the history, movement, and management of archival information can inform judgments about appropriate access, disclosure, and preservation practices over time.

Tian Wang (University of Illinois at Urbana-Champaign)

Restoring Contextual Integrity in Remote Patient Monitoring Systems via User-Level Differential Privacy

This presentation explored how CI can be operationalized through privacy-enhancing technologies in remote patient monitoring (RPM) systems. RPM technologies routinely collect sensitive health information, including heart rates, glucose levels, sleep patterns, and physical activity data, and share it among patients, clinicians, and healthcare organizations for clinical care purposes. The project examined whether these information flows remain appropriate when data is continuously collected and potentially repurposed beyond direct patient care, raising concerns about secondary uses that may not align with patients' expectations. Arguing that existing privacy norms often fail to adequately address the scale and persistence of RPM data collection, the work proposed user-level differential privacy (DP) as a mechanism for embedding CI principles directly into system design, echoing sessions from previous CI symposiums in which DP and CI were synthesized. By limiting the risks associated with continuous monitoring while preserving the clinical utility of health data, the approach aims to restore appropriate information flows and better align RPM systems with contextual healthcare privacy norms.

Summary:

These works in progress converge around three shared themes.

First, all three are concerned with privacy across time and changing contexts. Information often outlives its original context, creating challenges for maintaining appropriate information flows over time.

Second, all three emphasize the limits of individual control and responsibility. Privacy is framed as a collective governance challenge rather than solely an individual responsibility in all three projects.

Third, each explores how CI can be translated into practical governance mechanisms, connecting this feedback session to many of the prior sessions in the symposium. These presentations move beyond using CI as a descriptive framework and instead investigate how it can guide the design of institutions, policies, and technical systems that preserve appropriate information flows in complex real-world settings.

Session 6: CI and IoT

Notetakers: Kyra Abrams & Yaman Yu

This session, moderated by Quinn Waeiss (Stanford University), included two papers empirically studying information flows associated with IoT through the lens of the CI framework.

Isaac Haizel and Madelyn Sanfilippo (University of Illinois at Urbana Champaign)
Tradeoffs and Limitations of Wearable Location Tracking for Children

As parents increasingly leverage wearable and location-tracking technologies for children, including devices such as AirTags, smartwatches, Tile trackers, and platforms like Life360, information flows expand and increasingly diverge from common mental models. This project explored how parents use these technologies to monitor children's whereabouts and support feelings of safety. Rather than assuming that location data provides an objective or complete picture of children's activities, the project questioned how accurately these systems represent reality and how parents interpret the information they receive. Using CI as a foundation, the study investigated the information flows created by child-tracking technologies and the assumptions that underlie their use.

A central focus of the work was understanding parental sense-making and decision-making around location data. The research highlighted that location information is not neutral and often contains ambiguities that require interpretation. Real-time tracking can create a sense of certainty and control, yet data may be incomplete, imprecise, or disconnected from the broader social context needed to understand a child's situation.

Using a mixed-methods approach inspired by parental questions and everyday experiences, the authors explore norms surrounding child location tracking across contexts, situations, communities, and developmental stages, in comparison to the complex reality of information flows. The findings suggest that expectations about appropriate monitoring vary depending on factors such as a child's age, the environment, perceived risks, and family relationships. By foregrounding the relational dimensions of tracking technologies, the work aims to develop a framework for "relational precision tracking" that recognizes both the benefits and limitations of location monitoring.

Muhammad Hassan and Masooda Bashir (University of Illinois Urbana Champaign)
User Privacy Expectations in Connected Heart Devices

This presentation examined privacy expectations surrounding connected heart devices within the broader landscape of smart healthcare, the Internet of Medical Things (IoMT), and technologies that support aging in place. As connected cardiac monitoring systems become increasingly common, they generate continuous streams of health data that can be shared among a variety of stakeholders, including patients, clinicians, family members, device manufacturers, and insurers. The project sought to understand how users evaluate these information flows and what factors shape perceptions of their legitimacy and appropriateness. CI served as the primary analytical framework, providing a structured way to examine relationships among actors, data types, transmission principles, and healthcare contexts.

Using survey-based interviews centered on heart-monitoring scenarios, the study investigated how people evaluate different forms of data sharing associated with connected cardiac devices. The project distinguished between routine telemetry data and acute anomaly information, as these types of health data may carry different expectations regarding urgency, sensitivity, and disclosure. Findings indicated that the identity of the receiver plays a particularly important role in shaping perceptions of acceptable information flows. Participants differentiated among clinicians, family members, vendors, and insurers, assigning varying levels of legitimacy to each recipient depending on the context, purpose, and nature of the information being shared. These results reinforce the importance of context-sensitive information flows rather than one-size-fits-all approaches to health data governance.

The presentation highlighted several implications for the design of connected health technologies. Because privacy expectations vary according to circumstances and stakeholders, static privacy settings may be insufficient for managing health information appropriately. Instead, the findings suggest a need for adaptive, context-aware systems that can adjust data-sharing practices based on user preferences, clinical situations, and changing care needs. The discussion also raised broader questions about the relationship between perceived legitimacy and appropriateness, suggesting that users may evaluate information flows not only according to whether they fit existing norms but also according to whether particular actors are viewed as having a justified role in receiving health information. Overall, the work demonstrates how CI can help illuminate the nuanced privacy expectations that emerge as connected medical devices become increasingly embedded in everyday healthcare.

Summary:
This session raised important themes around IoT and privacy.

First, both projects emphasized that data does not speak for itself and must be interpreted within context, just as information flows themselves must be context-sensitive in practice. The value and implications of IoT data emerge through human interpretation rather than from the data alone. Further, IoT privacy cannot be governed effectively through universal rules because expectations change across situations and relationships.

Second, both reveal tensions between safety benefits and privacy concerns. Child-tracking technologies promise security and parental reassurance, while connected heart devices support health monitoring, aging in place, and potentially lifesaving interventions. Yet in both settings, increased visibility comes at the cost of increased surveillance and data sharing. The central question is not whether data should flow at all, but how to balance the benefits of monitoring against the risks associated with continuous observation and expanded access to personal information.

Third, both suggest adaptive design and policy solutions. IoT systems should be responsive to changing contexts, relationships, and user needs. Whether through relational approaches to child tracking or context-aware configurations for connected medical devices, the common goal is to align technological capabilities with evolving social norms and privacy expectations. This reflects a broader CI theme: privacy in IoT is fundamentally about managing appropriate information flows rather than simply limiting data collection.

Program Committee

Noah Apthorpe (Colgate University)

Karla Badillo-Urquiola (University of Notre Dame)

Ananth Balashankar (Google)

Masooda Bashir (University of Illinois at Urbana-Champaign)

Sebastian Benthall (New York University)

Sheena Bishop (University of Illinois at Urbana-Champaign)

Cathy Dwyer (Pace University)

Brett Frischmann (Villanova University)

Yuanyuan Feng (University of Vermont)

Kyle Jones (Indiana University-Indianapolis)

Priya Kumar (Pennsylvania State University)

Yafit Lev-Aretz (Zicklin School of Business, Baruch College)

Kirsten Martin (Carnegie Mellon University)

Lee James McGuigan (University of North Carolina at Chapel Hill)

Mainack Mondal (IIT Kharagpur)

Ido Sivan-Sevilla (University of Maryland)

Katherine J. Strandburg (New York University School of Law)

Daniel Susser (Penn State University)

Eran Toch (Tel Aviv University)

Vincent Toubiana (CNIL)

Salomé Viljoen (University of Michigan Law School)

Ine Van Zeeland (imec-SMIT-VUB, Universiteit Hasselt)

Jessica Vitak (University of Maryland)

Ben Zevenbergen (Google)

Michael Zimmer (Marquette University)

Program

Wednesday	
9:00 AM	**Registration, Coffee, Snacks and Refreshments**
10:00 AM	**Welcome**
10:15 AM	**Community discussion panel** ***Agentic PrivaCI: Challenges and Opportunities***
11:15 AM	**Break (15 min)**
Session: CI and User perceptions ***Chair: Noah Apthorpe, Colgate University***	
11:30 AM	**The role of data distance on perceptions of data privacy: Distinguishing AI inferences and predictions** **Alexis Shore Ingber (Syracuse University, S.I. Newhouse School of Public Communications) and Nazanin Andalibi (University of Michigan School of Information)**
11:45 AM	**Evidence of College Athletes' Trust and Self-belief Attitudes in the Gray Area of Data Privacy Expectations** Brad Stenger (University of Vermont) and Yuanyuan Feng (University of Vermont)
12:00 PM	**Privacy in the Classroom: A Study of Student Perspectives Of Educational Technology in School** Jake Chanenson (UChicago) and Marshini Chetty (University of Chicago)
12:15 PM	**Feedback (15 mins)**
12:30 PM	**Lunch**

Session: CI and Theory ***Chair: Helen Nissenbaum, Cornell Tech***	
2:00 PM	**Integrating Contextual Integrity and Data Minimization** Noah Apthorpe (Colgate University)
2:15 PM	**Contextual Integrity's Privacy Problem** Sophia Wushanley (University of Michigan)
2:30 PM	**Discussion (10 mins)**
	Break (15 min)
Session: CI and Privacy Risks ***Chair: Alexis Shore Ingber, Syracuse University***	
3:00 PM	**Multiverse Privacy Theory: Modeling Context-Dependent and Time-Sensitive Privacy Decisions** **Ece Gumusel and Madelyn Sanfilippo (University of Illinois Urbana-Champaign)**
3:15 PM	**How Researchers Identify and Mitigate Downstream Privacy Risks: A Contextual Integrity Analysis** **Zihao (Raio) Huang, Quinn Waeiss (Center for Biomedical Ethics, McCoy Family Center for Ethics in Society, Stanford University)**
3:30 PM	**Discussion (10 mins)**
	Setup buffer
CI Community Feedback session #1 ***Chair: Christopher Muhawe, University of Illinois Chicago***	
3:45 PM	**Contextual Implementation of Data Autonomy** **Kyra Abrams (University of Illinois at Urbana-Champaign)**

3:50 PM	**Contextual Integrity and Age: Rethinking Transmission Principles Across Age** **Yuxing Yang (University of Illinois Urbana-Champaign)**
3:55 PM	**Between Safety and Surveillance: Online Activity Monitoring Products and Student** **Sam Hafferty, Daniel Flyer, Nitya Nadgir, Prateek Mittal, Mihir Kshirsagar, Mona Wang (Princeton University) and Shaanan Cohney (Melbourne University)**
4:00 PM	**Feedback (15 mins)**
4:15 PM	**Panel: Reflections of the day**
5:00 PM	**Pre-dinner break**
6:00 PM	**Dinner**

Thursday	
9:00 AM	**Registration, Coffee and Refreshments**
Session: CI and Health AI ***Chair: Masooda Bashir, University of Illinois Urbana-Champaign***	
10:00 AM	**Applying Contextual Integrity to Biomedical Foundation Models: An Analysis of a Multimodal Sleep Foundation Model** **Jiunn-Tyng Yeh (Duke University) and Margot Hanley (Duke University), Yan Shvartzshnaider (York University)**

10:15 AM	**From Adaptive Personalization to Accountable Care: Bridging the (Un)usable Privacy Gap in Voice-Based Health AI** **Thalia Viranda, Jia Bo, Shupeng Wang, Dan Adler (Cornell University), Andrea Cuadra (Olin College of Engineering), Deborah Estrin (Cornell University), Tanzeem K. Choudhury (Cornell University),Walter Richard Boot (Weill Cornell Medicine), Sara Czaja (Weill Cornell Medicine)**
10:30 AM	**Discussion (10 mins)**
	Setup buffer
Session: CI and Social Norms ***Chair: Yan Shvartzshnaider, York University***	
10:45 AM	**Contextual Integrity Theory in China: An Intellectual History from Interpretation to Reflexivity** **Alex Chanhou Lou (Cornell Tech, Cornell University)**
11:00 AM	**Assessing privacy successes and failures through the lens of contextual integrity: Lessons from social movements** **Glencora Borradaile, Alexandria LeClerc, Kelsy Kretschmer (Oregon State University)**
11:15 AM	**Understanding AI disclosure norms at workplace using CI framework (Use Case)** Smirity Kaushik (George Washington University), Shiza Ali (George Washington University) and Adam Aviv (George Washington University)
11:25 AM	**Discussion (15 mins)**
11:40 AM	**Break (15 min)**
CI Community Feedback session #2 ***Chair: Smirity Kaushik, George Washington University***	
11:55 AM	**Contextual Integrity in Youth AI Use: Privacy Norms Across Shared and Personal ChatGPT Account Contexts** **Lili Dudas, Aljawharah Alzahrani, Satish Kumar Keshri, Yu-Jye Tung, Tanusree**

	Sharma (Penn State)
12:00 PM	**Contextual Integrity and Age: Rethinking Transmission Principles Across Age** **Erin Gilmore (San Jose State University), Iori Khuhro (University of British Columbia), Darra Hofman (San Jose State University), Jim Suderman (Independent Expert), Frédéric Simard, Michel Barbeau, Mario Beauchamp, Nicholas Rivard (Carleton University)**
12:05 PM	**Restoring Contextual Integrity in Remote Patient Monitoring Systems via User-Level Differential Privacy** **Tian Wang (University of Illinois at Urbana-Champaign)**
12:10 PM	**Feedback (15 mins)**
12:25 PM	**Lunch**
Session: CI and IoT ***Chair: Quinn Waeiss, Stanford University***	
1:55 PM	**Tradeoffs and Limitations of Wearable Location Tracking for Children** **Isaac Haizel and Madelyn Sanfilippo (University of Illinois at Urbana Champaign)**
2:10 PM	**User Privacy Expectations in Connected Heart Devices** **Muhammad Hassan and Masooda Bashir (University of Illinois Urbana Champaign)**
2:25 PM	**Discussion (10 mins)**
2:35 PM	**Symposium Wrap Up**
2:45 PM	**Ice Cream**