The 6th Annual PrivaCI Symposium Report

September 27-28, 2024 in New Jersey, US, **Rutgers University**.

The report was compiled by Yan Shvartzshnaider based on the notes taken by
John Oluwaseye Adebayo, Alexis Shore Ingbe, Olivia Figueira, Grace McGrath, Michael Beauvais, Kyra Milan Abrams, Rishi Jha, Shiva Mayahi, Smirity Kaushik, Min Cheong Kim, Michael Khavkin, Jake Chanenson, Rohan Grover, Madison Pickering, Kenan Kamel Alghythee, Laurent Haoyu Wang, Synthia Wang, Kyrie Zhixuan Zhou, Jason Brady (Brad) Stenger.

Special thanks to our sponsors: National Science Foundation, Rutgers's School of Communication and Information, Google and Digital Life Initiative

Table of Contents

Executive Summary

In September 2024, the CI community gathered at The Heldrich Hotel & Conference Center in New Jersey for the 6th annual PrivaCI Symposium, hosted by Rutgers University. With over 70 participants, the symposium offered a full agenda of discussions on the application of the Contextual Integrity (CI) theory and framework.

The symposium kicked off with a session exploring the use of CI to investigate remote patient monitoring systems and the post-pandemic adoption of COVID-19 mitigation mobile applications. This was followed by CI Community Feedback Session #1, which featured proposals on applying CI to governance challenges and designing frameworks for collective civic reporting of privacy harms.

Session 2, focused on CI as a framework, included discussions on developing a roadmap for applying CI to qualitative privacy research and efforts to standardize contextual integrity. Session 3 addressed using CI for privacy assessments, with presentations on exploring perceptions and acceptability of data sharing in virtual reality, as well as data-handling practices related to athletes' health and performance.

In Session 4, presenters used CI to audit privacy practices of online services aimed at children and adolescents, and to examine the concerts around synthetic content on video-based social media platforms. Session 5 continued the exploration of CI by discussing how large language models can be used to annotate privacy policies and introduce evidence-based privacy policies.

Day two began with CI Community Feedback Session #2, which featured talks on privacy concerns regarding Smart TVs from a user perspective, as well as the "Hidden Cam Inspector" tool designed to detect and locate hidden Wi-Fi cameras.

Session 6 explored the intersection of CI and differential privacy (DP). The first talk investigated the impact of DP on users' contextual data disclosure decisions. The second talk focused on understanding people's contextual choices regarding differential privacy.

In Session 7, the "CI and Society" session, speakers presented a case for using contextual integrity assessment of the ethics of incorporating derogatory data in clinical suicide prediction algorithms, along with a discussion on the contextual dimensions of data.

The final session, focused on CI and AI, featured a paper on AI detection for screenshot prevention and a position paper on CI for AI assistants.

Session 1: CI as a Lens

Notetakers: Michael Beauvais & Grace McGrath

Contextual Integrity in Remote Patient Monitoring for Managing Chronic Conditions at Home

D. Ruben Tjhie (University of Toronto)

The main points from both presentations revolve around privacy concerns in digital health applications. The first presentation from D. Ruben Tjhie discussed remote patient monitoring (RPM) systems, particularly for heart failure patients. RPM systems increase the accessibility of health data and the reliance on algorithms to assist in or take on caregiving roles, potentially leading to privacy issues.

Post-Pandemic Contextual Acceptance of COVID-19 Mitigation Mobile Applications in the US (published work)

Yuanyuan Feng, Brad Stenger (University of Vermont), Shikun Zhang (Carnegie Mellon University)

The second presentation from Yuanyuan Feng, Brad Stenger, and Shikun Zhang examined the acceptance of COVID-19 mitigation mobile applications in the US, focusing on data sharing and retention practices. Both studies utilize the theory of privacy as contextual integrity to analyze how privacy norms evolve in these new digital health contexts and to identify potential privacy violations.

Common themes between the two presentations included the exploration of how digital health interventions in both the clinical and public health contexts impact privacy and the application of contextual integrity theory to understand these impacts. Both studies considered the changing dynamics of health-related information flows, whether it's between patients and algorithms in RPM or between individuals and various recipients in COVID-19 mitigation apps. The research in both cases aims to understand how contextual factors influence privacy norms and user acceptance of health-related data sharing and retention.

Future work in this area could focus on several directions. First, there's a need to explore who is best placed to be responsible for information practices in digital health applications – clinicians, app developers, or other stakeholders. Second, further research could investigate how disease specificity factors into information flows and user sensitivity to certain types of health data. Finally, studying the role of intervention efficacy in shaping user acceptance of privacy trade-offs in digital health applications could provide valuable insights for future app

development and deployment strategies. These directions would help to refine our understanding of contextual integrity in evolving digital health landscapes and inform more privacy-aware design of health technologies.

CI Community Feedback Session #1

Notetakers: Olivia Figueira and Kyrie Zhixuan Zhou

Governance seam integrity, contextual integrity, and the integrity of socially meaningful contexts

Brett Frischmann (Villanova University)

This work asserts that there is a relationship between governance seams and contextual integrity, and the goal is to analyze this relationship with respect to the integrity of socially meaningful contexts. This goal includes defining governance seams, and their degradation or erosion, in the context of privacy through contextual integrity, to show how certain actions in society may degrade the integrity of socially meaningful contexts and how/why this is harmful.

Contextual integrity will be used to analyze the norms and appropriateness of information flows and context those to the definition of and analysis of governance seams and seam erosion, such as analyzing cases of dataset aggregation, redesign of interfaces, and surveillance technologies.

The presentation discussed the problem, motivation, and goals for this work, as it is a work in progress presented for feedback. The speaker discussed a case study, namely school-issued laptops and how this may diminish the integrity of the school and/or the home contexts, which are both socially meaningful and have their own norms, if there are or aren't privacy violations. CI is applied to analyze appropriateness of data flows, considering the norms of these contexts, which may be different, and to show how governance seam degradation. Considering this use case, a seam may be needed to protect the family's determination.

The speaker discussed how this work can be used to show privacy harms, and it may even be useful to show how deterioration of governance seams around socially meaningful contexts can be harmful even if there aren't any privacy violations. Developing and analyzing use cases in this work was noted as a challenge.

The main next step noted is to dig deeply into the antecedent social theories layed out in the original CI theory to analyze the relationship between governance seams and CI.

Q&A:

Q: When two contexts experience a degradation of a governance seam between them, it may have harm to both contexts. What are the normative assumptions that go into that? Do we need to know anything else about a governance seam degrading to make a claim that it's wrong?

A: It could be a harm, but not necessarily. The governance seam maintains the norm as a socially meaningful context, and sometimes the degradation is bad,
but not always.

Q: (comment) We should try to show how the degradation of a governance seam is a harm, and we should persuade others of this, because it may show why things such as aggregation of datasets for ML purposes causes such degradation. This way, we can show why it's bad or harmful.

A: Speaker agrees; the seam is not necessarily the harm, but the impact of degrading the seam on socially meaningful contexts, where they lose their integrity, is the harm, and the means by which you get there is the degradation of the seam.

Q: In the context of school-issued devices going into family's homes, how about when people in those spaces don't understand what goes on/in that machine? The people at the school setting them up likely haven't read all the privacy policies, but they have done the required security analysis. The people in the home don't have those skills either. What kind of understanding is possible when the people setting it up and the people receiving it don't have the capacity to fully understand what they are getting?

A: This is what can lead to seam erosion. The recipients, such as the teachers, are under-informed about what's on the device, and this makes it more frictionless for the seam to erode. Causal dynamics are why the seam erodes, and we don't have resistance or governance of these properties.

Design(ing) Fictions for Collective Civic Reporting of Privacy Harms

Yuxi Wu (Georgia Institute of T echnology), William Agnew (Carnegie Mellon University), W. Keith Edwards (Georgia Institute of T echnology), Sauvik Das (Carnegie Mellon University)

This work is trying to design a way for society at large to be able to report privacy harms at scale, which can enable us to define privacy harms more concretely based on users' experiences. Existing platforms do not work well, and the impacts of evidence-gathering attempts are poorly-understood and not well-recognized. Related work has largely focused on privacy harms through targeted advertising.

Contextual integrity is applied in the formulation of the research questions/study, as the core of this work is to address what makes an information flow harmful in the context of user privacy.

The presentation described the researchers' user study conducted to analyze design requirements for such a system and to analyze the cultural ideals at play, which relate to their two main research questions. Their findings thus far include what the form should look like, what it should take in, and participants' conceptions about how privacy risks should be protected with respect to reporting these issues. Cultural ideals for the participants centered around duty to help others and pride in volunteering such information for the good of society.

Challenges encountered include applying these findings in the task of developing such a system in the long term. The lessons learned in this work will help to inform future studies with the goal of developing the system and using its results to define privacy harm more specifically.

Future directions include continuing this work to design such a system based on their user study findings. This work will be published in CSCW 2025, and thus future continuations may follow.

In this session, two researchers presented their work for feedback from the contextual integrity community. The first speaker presented their preliminary work about defining the relationship between governance seams and contextual integrity, where the goal is to analyze this relationship with respect to the integrity of socially meaningful contexts. This work includes defining governance seams, and their degradation or erosion, in the context of privacy through contextual integrity, to show how certain actions in society may degrade the integrity of socially meaningful contexts and how/why this is harmful. Empirically, privacy harms are not well-understood and hard to measure. Existing platforms for reporting privacy incidents primarily operate as intake forms – people can only voice, but have no other participatory stakes.

The second speaker discussed their work about designing a system for society to be able to report privacy harms at scale, which can enable us to define privacy harms more concretely based on users' experiences. Their latest work includes conducting a user study to better understand users' expectations and ideals regarding the reporting of privacy harms to inform the design of their system, such as protections for volunteers, cultural ideals, and duty to help others. The Q&A was focused mainly on the first talk, due to time, and the topics of the Q&A included both how the presenter approached the definition of privacy harm with respect to seam degradation and erosion and comments surrounding the work and case studies presented.

Common themes between these presentations include incorporating contextual integrity with the goal of attempting to better define privacy harms, either from a theoretical point of view in context with governance seams and case study analyses, or via large-scale analyses of privacy violations in society through the development of a user-centered privacy harm reporting system. These works demonstrate the growing need for a definition of privacy harm that extends beyond the usual legal definition and is adaptable to both emerging technologies, such as large language models and artificial intelligence, and the increasingly diverse environments in which people use such technologies, such as work, school, and their homes.

The first project presented is in its preliminary stages and thus the future work is to continue developing their theories and case studies surrounding governance seams, contextual integrity, and their relevance to socially meaningful contexts' integrity. The second project discussed work that is to be published, and the future work includes ideating and iterating on the design of their proposed privacy harms reporting system based on their user study findings.

Session 2: CI as a framework

Notetakers: Laurent Haoyu Wang and Jason Brady (Brad) Steng

A Roadmap for Applying the Contextual Integrity Framework in Qualitative Privacy Research

Priya C. Kumar (Pennsylvania State University), Michael Zimmer(Marquette University) and Jessica Vitak (University of Maryland)

The paper seeks to improve the accessibility of the Contextual Integrity framework, and in the process see it move from something that is seen as conceptual to something that is more of applied tool.

The roadmap in the paper is straightforward guidance for researchers on how to use the CI framework to set up and then execute a privacy analysis. "We said, here's a goal or an aim for you, the research team to guide you, the research team in each step. And then here are a set of guiding questions to help you figure out what specifically to look for in your data." At each step in the roadmap there are also examples from the real user case of Fitbits and privacy.

The guidance provided by the roadmap is rigorous, but not prescriptive. The fitbit examples offer an analysis that is not meant to show "how to do the analysis" but instead explains

"here's what you might be able to get out of this analysis" and that doing this analysis can clearly advance privacy research questions.

Standardizing Contextual Integrity

Sebastian Benthall (New York University), Darra Hoffman(San Jose State University) ,Ido Sivan-Sevilla (University of Maryland)

The paper is trying to introduce clarity and consistency will help Contextual Integrity to achieve wider adoption. Standards should help, and therefore so should standardization working groups.

The standards working groups include: "formal standards working group" and it thinks about the core rules, like existing rules for privacy in information security situations. The "working with industry on compliance" is taking all of the questions about standards and determining "what the parameters should be"

Edge cases seem to be the challenge and the modus operandi for creating standards around CI. An edge case prompts a new consideration for CI and work by a CI standards group proposes an ontology that fits the new case with existing practice. Examples are regional governments and Native American where cultures and cultural norms have to be reconciled. GDPR and Differential Privacy are also edge case examples.

Any future work that is planned, or other future directions of interest for the CI community. The standards groups are looking for knowledgeable people to help with the work.

Summary:

Each of the two presentations in "Session 2: CI as a framework" could have gone one step further and called the relationship to CI an "applied framework." The first presentation, "A Roadmap for Applying the Contextual Integrity Framework in Qualitative Privacy Research," is a how-to that used rules and examples to illustrate and not to prescribe any specific type of analytical problem solving solution. The second presentation, "Standardizing Contextual Integrity," is an update and call to action to help make CI clearer and more consistent, and so useful to more and more people who are interested in privacy.

The two presentations both featured significant consideration of edge cases. In the Roadmap presentation, edge cases (non-human actors, interaction media, GKC) are research questions that will scale up the complexity of future roadmaps, and at the same time, present research opportunities. In the Standardizing presentation, edge cases (indigenous cultures, GDPR, Differential Privacy) are source material for evolving current practice into robust standards.

After a couple of questions that were forward-looking and related to edge cases (marginalized groups, inductive analysis, IoT devices) there was a hard-to-pin-down discussion of law, politics methods and power relationships. The last few minutes were a fitting reminder that public policy, like human-computer interaction and computational social science, is the domain of applied CI. If a roadmap is a how-to and standards are a what-for then law, politics and power are good reasons why CI and why research using CI are critically important.

Session 3: CI and Assessment

Notetakers: Kenan Kamel A Alghythee and Smirity Kaushik

PrivaCI in VR: Exploring Perceptions and Acceptability of Data Sharing in Virtual Reality Through Contextual Integrity

Emiram Kablo, Melina Kleber (Paderborn University) and Patricia Arias-Cabarcos (Paderborn University and KASTEL/KIT)

There exist many sensors in a typical virtual reality set, and as such, it may propose a privacy violation. The current trend of user studies focuses on XR/VR users' and developers' privacy concerns based on what type of data is being shared. The current research does not consider recipients under the condition data is being shared or privacy needs. The researchers included VR non-users and created data flows using textual integrity. Furthermore, the researcher aims to answer three research questions, Under which conditions do VR users and non-users consider sharing VR data flows acceptable? How aware are VR users and non-users of VR data collection? What are VR users' privacy-protecting behaviors when using VR?

The researchers used contextual integrity theory as an instrument to construct the information flows and allowed the participants in the study to rate the flows that they found acceptable, and the team generated 1,387 different information flows.

The current state of the research is still ongoing, and the results from the pilot study of which it includes 10 participants, 5 surveys with each 2 participants, and an average time of 10 minutes. The team learned to include examples for user controls and privacy features. Another point is the consideration of bystander as a subject of which it will present a paper extension; Also, the lack of a complete list of flows is another challenge to be addressed in their future work/progress.

Information Flows for Athletes' Health and Performance Data

Brad Stenger and Yuanyuan Feng (University of Vermont)

2% of college athletes become professional. However, there is no policy for data collected from athletes. Additionally, there is huge interest in collecting data and helping coaches mark progress. This leads to privacy risks for athletes. In future work, authors will develop scenarios for the Contextual Integrity vignette study.

Athletes can't set what their security means to them. So athletes who are complying have a different context compared to those who are not. Also, there is a case of purple save athletes, data, and teams. Use of different privacy without impacting the competitive advantage.

The first parts of an NSF study are underway include Literature review and Technology survey. Next steps include mapping concepts and terminology to technologies and developing scenarios for Contextual Integrity vignette study. Additional initiative look into collaborating with colleagues working on Differential privacy and applying Privacy Enhancing Technology to solve problems in sports and athlete development.

Session 4: CI and Policies (Part 1)

Notetakers: Shiva Mayahi and John Oluwaseye Adebayo

DiffAudit: Auditing Privacy Practices of Online Services for Children and Adolescents

Olivia Figueira, Rahmadi Trimananda, Athina Markopoulou, and Scott Jordan

The Children's Online Privacy Protection Act (COPPA) in the United States (US) governs online service providers' practices regarding acquiring and disseminating personal data about minors under 13. Although previous research has established frameworks and methodologies for privacy auditing and measurement across various digital platforms, including mobile applications, web browsers, and virtual reality environments, there has been no comprehensive audit of general audience services concerning both COPPA and CCPA about children and adolescents, particularly examining their behaviors before and after the provision of consent and the disclosure of user age.

Employing the GPT-4 classification methodology, the Contextual Integrity parameters were utilized to scrutinize the data flows categorized by age across website and mobile platforms, focusing on age-specific traces, privacy policy examination, and data linkability across

diverse platforms, with particular attention to the classification of data types and their respective attributes.

The research revealed that all services analyzed through DiffAudit engaged in collecting and disseminating personal information pertaining to child and adolescent users in a manner that raises significant concerns regarding compliance with COPPA and CCPA. When conducting a comparative analysis of the data flows about children, adolescents, and adults utilizing a differential analysis approach, the findings indicated that these services do not adjust their data collection and sharing practices to the extent anticipated for child and adolescent users. Furthermore, every service also collected and shared information regarding users before ascertaining their age and obtaining consent (i.e., during the logged-out state), including interactions with third-party advertising tracking services, a practice deemed inappropriate under COPPA and CCPA regulations.

There exist limitations associated with applying machine learning techniques to automatically categorize data types derived from network traffic data, as well as challenges in analyzing network traffic, which often needs to be presented as fully comprehensible text. The researchers opted to manually engage with the services to enhance the quality and comprehensiveness of the dataset.

The researchers intend to advance their data type classification methodologies in tandem with the ongoing enhancements of large language models by scrutinizing the privacy policies of additional platforms for a comparative analysis of network traffic through differential analysis, employing an automated and customizable data type classification approach to detect anomalous behaviors in network traffic across any platform, with a particular emphasis on safeguarding the privacy of children and adolescents under COPPA and CCPA.

Q&A:

The questions asked: There was a question on whether the children's behavior would still be the same if examined within the context of other privacy acts that differ from CCPA. A question was also asked about the rationale for the use of GPT-4 model for the study.

Main points of the answer: The authors replied that although such direction was not captured in the study, there might be slight differences considering the scopes of different privacy acts.

Governing Manipulative and Synthetic Content on Video-based Social Media Platforms

Smirity Kaushik and Madelyn Sanfillipo

Problem Statement: Deepfake ads are on video-based social media platforms targeted at specific user groups to influence their behavior through online advertisement feeds. As technology advances, young adults, being the most users and active participants on social media platforms are at risk of fraudulent ads, which can influence specific behavior and may erode trust in content on digital platforms. As technology develops, users of digital platforms find it challenging to identify between authentic and artificially generated content.

Application of Contextual Integrity Approach to Privacy: The researchers used the Contextual Integrity parameters because of their connection with the two research questions that guided the study. Specifically, the study employed the GKC-CI Framework, a combination of Contextual Integrity (CI) parameters and the Governing Knowledge Commons frameworks to explore how deepfake ads influence online behaviors among young adults.

Current Progress and Preliminary Results described in the presentation: The authors have completed a qualitative analysis of 25 policy documents from 3 significant stakeholders and a review analysis from 19 audiences based on the CI framework. Preliminary results show the institutional structure of the policies, as content norms and strategies are the major priorities in privacy policies that shape people's behavior towards ads on digital platforms. This finding implies that it further established the need for more assessment of the effectiveness of privacy policies on various platforms regarding deepfake advertisements targeting young adults.

Future Plans: The preliminary findings will help the authors expand their work's scope by exploring young adults' experiences with deepfake on various platforms and how such can influence contextual policy reviews.

The questions asked: The presenter was asked if consequences and norms are the strategies to evaluate the content of deepfakes and GenAI in advertising when privacy policies become more mature.

Main points of the answer: The presenter responded that the findings of their study will still play significant roles in analyzing contents used in deepfake and GenAI advertising to shape internet governance.

Session 5: CI and Policies (Part II)

Notetakers: Kenan Kamel A Alghythee and Rohan Grover

Qualitative Analysis of Governing Knowledge Commons and Contextual Integrity (GKC-CI) Privacy Policy Annotations with Large Language Models

Jake Chanenson (University of Chicago), Madison Pickering (University of Chicago), Noah Apthorpe (Colgate University)

The authors presented a study in which they sought to automate detecting governing knowledge commons (GKC) and contextual integrity (CI) parameters in the texts of privacy policies. It was motivated by prior research that has demonstrated the utility of the unified GKC-CI framework to interpret privacy policies and normatively evaluate them. However, relying on manual annotation limits the scale at which the GKC-CI framework can be applied. Thus, the authors were interested in exploring whether a large language model (LLM) could be useful in automating annotating privacy policies using the GKC-CI framework. In the study, they fine-tuned 50 off-the-shelf LLMs on 21,588 GKC-CI annotations from 16 ground truth privacy policies and examined their performance and errors. The best performing model produced GKC-CI annotations with above 90% accuracy.

The authors also discussed findings from a qualitative analysis of the errors produced by the best-performing LLM. Notably, they found that many of the model's "errors" were in fact acceptable because the model's label and the ground truth label often differed only by a word that did not materially change the parameter annotation—something that is difficult to discern computationally. In addition, they found cases in which the model's parameter annotations were, in fact, more precise than the human annotation. At the same time, there were still cases in which the researchers confirmed that the model's annotation was erroneous, such as by adding irrelevant policy text to the annotation. This qualitative analysis was a key component of the study that can inform future model development and automation of GKC-CI analysis of privacy policies.

Evidence-based Privacy Policies

Allan Lyons (University of Calgary), Yan Shvartzshnaider (York University), Joel Reardon (University of Calgary)

The authors presented a study in which they audited companies' privacy policies. It was motivated by prior research that has identified deficiencies in privacy policies, such as being vague or difficult to understand. This study builds on such findings by comparing the text of privacy policies in the Google Play store with observed data transmissions in 20 Android apps. The authors annotated each app's privacy policy using components from the theory of contextual integrity (CI), analyzed their data safety statements, measured network traffic during use on an Android phone, and then compared the app's claims with its observed data transmissions.

One major finding is how widely privacy policies and data safety labels can differ. The authors noted that privacy policies and data safety statements are difficult to compare because they may be authored by different actors within the same company (e.g., lawyers vs. developers) and thus may use different terms or define the same terms differently. For example, they cited an example in which the app's data safety label said that no data was shared with third parties, but the corresponding privacy policy stated that data will be provided to third parties. This discrepancy could be explained by each policy's distinct scope: whether it applied only to the Android app or to the company overall.

In the future, they are interested in developing a tool or service to help developers understand the code in their applications. For example, if a developer uses a library for a particular functionality, they may not fully understand the implications for data transmission and thus potential violations of contextual integrity.

CI Community Feedback session #2

Notetakers: Min Cheong Kim and Madison Pickering

Privacy Concerns about Smart TVs from a User Perspective

Synthia Wang, Lan Gao, Marshini Chetty, Nick Feamster (University of Chicago)

While IoT privacy has been broadly studied, smart TV privacy remains underexplored due to the unique affordances of the devices. To understand the unique privacy needs and concerns inherent to these devices, the authors conduct semi-structured interviews using the CI framework. Some non-exhaustive but representative examples of questions asked include: who [does the user think] has access to their data? What info do they think is protected? And, what is their expectation for the data? The author's paper is in submission, so a more exhaustive discussion of the methods should be available in the future.

The authors find that users have only a vague perception of both what data is being collected as well as the information flows themselves. Further, users highlighted the existence of dark patterns such as multiple menus to adjust settings that limit data collection and transmission. In line with this, some participants reported feeling a lack of meaningful control as a result: "I think that you have to accept it and hope that it isn't used in a way that harms you."

Finally, the authors further notice tensions potentially unique to smart TVs as opposed to other smart devices. Namely, the device cannot be easily moved and requires one to log in, resulting in a potential exposure of watch history as a result. This was described as both inconvenient and annoying if a person moves out of their home environment and wants to use

a TV in another environment, e.g., a hotel.To this point, an audience member remarked that people have been doing semi structured interviews about user perceptions for IOT devices for a while, and if the authors observed differences in perceptions as smart TVs are no longer new as compared to 10 years ago? To this point, the author reiterated that smart TVs when compared to other IoT devices have unique properties (e.g., cannot be easily moved) and even when compared to themselves 10 years ago they are much closer to large tablets.

Hidden Cam Inspector: Usable Tool to Discover and Locate Hidden WiFi Cameras

Danny Y. Huang and Grace McGrath (New York University

The work is motivated by a concerning and rising trend of AirBNB hosts putting hidden cameras in air BnBs. As a result, potentially sensitive data leaves a home environment and is transmitted to an air bnb host as well as potentially the camera manufacturer without an individual's knowledge. To combat this inappropriate information flow, the authors want to make sure that people can find the hidden cameras themselves. As a result, the authors require a clear and intuitive human protocol and a UI, UX that supports an end-user walking through the requisite discovery steps themselves. To ensure that this is achievable by end-users, the authors utilize a participatory design process with the following stages: focus group, focus group with demos, in-person demos, and a real-world study where hidden cameras are put into a user’s space (i.e., their dorm).

The authors observe that when an individual moves into frame for a camera, the recording activity will create a burst of traffic on the network. Thus, one may be able to detect the presence of a hidden camera by monitoring the local network and looking for bursts of traffic. The authors utilize ARP spoofing and IOT Inspector to accomplish this. However, this work inherently assumes that the existence of a hidden camera is inherently a violation of privacy. Helen Nissenbaum noted that there could be problems with this. In particular, it is dangerous to do research with the assumption that “we” do the engineering work, and that the normative work is for “other folks” to do. One should always be focused on the normative aspects, and further, it is critical to acknowledge any assumptions that motivate work to ensure that produced work is of the highest possible quality. Finally, another individual asked if the authors had considered reporting to AirBnB if hidden cameras were found. To this, the authors noted it to be a useful area of work, but one they had not deeply thought about at this time due to the current stage of the project.

Session 6: CI and Differential Privacy

Notetakers: Kyra Milan Abrams Synthia Wang

Investigating the Impact of Differential Privacy on Users' Contextual Data Disclosure Decisions

Michael Khavkin and Eran Toch (Tel Aviv University)

Michael Khavkin and Eran Toch from Tel Aviv University presented their work on "Investigating the Impact of Differential Privacy on Users' Contextual Data Disclosure Decisions." Presented by Michael, this paper attempts to solve the issue of combining Differential Privacy (DP) with CI. They recruited 588 individuals from the US, UK, and India to measure multi-dimensional preferences in decision making. Their results plotted indifference curves to measure the tradeoff between payment to users for data and DP protection level and they found a preferable tradeoff level between compensated users and DP level. They conclude that data analysts need to pay an appropriate amount to users if they are using data for revenue generation, user interaction exists between DP and CI, and understanding human decision-making is important to appropriately configure DP systems while fairly compensating users. Their future work plans to extend contextual framing to include other contextual factors such as privacy literacy and risk attitudes and how different modes of DP influence decision making.
A discussion followed with questions and recommendations. A question was asked surrounding correlation between attributes. Michael responded that the attributes were assigned to each data analysis randomly. According to each combination, a user had to make a decision. Other questions were asked about the sensitivity of the data and if it is worthwhile to present epsilon as greater than 1. Michael responded clarifying that DP was not applied to the data used and most companies present epsilon as greater than 1. Lastly, a recommendation was made to clarify aggregation and contextual factors highlighted and to look at recipients to not lose variation.

Understanding People's Contextual Choices of Differential Privacy

Kyrie Zhixuan Zhou and Madelyn Rose Sanfilippo (University of Illinois Urbana-Champaign)

This work focuses on the significance of privacy-data utility tradeoff in differential privacy (DP). With a proof-of-concept survey, it investigates how types of app and information receiver affect users' perception of appropriate information disclosure. Referencing the concept of contextual integrity (CI), the survey was constructed with questions under different scenarios that consisted of various data types and data recipients. User attitudes were explored from three aspects, namely acceptability, preference, and appropriateness.

Resonating with the concept of CI, the work found that people's privacy preferences are contextual and personalized. In addition, sequential symmetry existed between the perceived appropriateness of information flows and the desired privacy to data utility.

During the course of this research, the research group had difficulties recruiting participants, particularly as many participants quitted their survey halfway through it. To solve this problem, they experimented with different wordings in the survey. For future works, they were interested in deploying large-scale surveys across cultures. Other potential research directions include understanding people's contextual privacy choice in real-world settings and investigating people's perceptions of specific representations of DP.

Session 7: CI and Society

Notetakers: Rishi Jha and Jake Chanenson

The session on "CI and Society" featured two presentations exploring the ethical and privacy concerns surrounding the use of data in various contexts, with a common focus on how the theory of privacy as contextual integrity (CI) applies to modern data practices.

Considering the Ethics of Integrating "Derogatory Data" in Clinical Suicide Prediction Algorithms: A Contextual Integrity Assessment (use case)

Michael Zimmer (Marquette University)

Michael Zimmer's presentation focused on using CI to evaluate the ethical implications of integrating derogatory data, such as law enforcement and financial data, with electronic health records (EHR) for the purpose of suicide prevention. While this integration could help predict suicide risk years in advance, Zimmer raised concerns about the ethical issues surrounding the use of publicly available data without clear consent. Through surveys, his research revealed that while people generally support the use of data for public good, they are uncomfortable with certain types of personal data being shared, highlighting the tension between privacy and public benefit. In addition, as discussed in the Question-and-Answer period (QA), these findings challenge data scientists to selectively search for data features rather than straining privacy expectations by extracting correlations across different contexts.

Contextual Dimensions of Data Autonomy

Kyra Abrams (University of Illinois at Urbana-Champaign)

Kyra Abrams' presentation on data autonomy under power imbalance examined how CI can clarify the boundaries of control disenfranchised individuals have over their personal data. She explored situations where people are assumed to consent to data sharing but are not empowered to dictate how their data is being used. Abrams used the example of incarcerated individuals, whose electronic communications are monitored under the assumption that they have waived all privacy rights. Her work demonstrated the need for more nuanced interpretations of consent and autonomy in data practices, suggesting that CI could offer a framework to better understand the appropriateness of data flows. Future work will focus on understanding these issues more deeply and refining how CI can guide ethical data sharing. During the QA session, the community discussed the role of 'control' in Abrams' work and CI at large, to which Abrams suggested that different contexts, especially those with power imbalances, necessitate different amounts of control.

A common theme across the whole session was the application of CI to examine the appropriateness of data flows in society, especially in cases where data sharing occurs by default without sufficient consideration of consent or privacy expectations. Both speakers emphasized the need for further exploration of these issues, particularly as data continues to be integrated across different contexts in ways that may challenge traditional notions of privacy.

Session 8: CI and AI

Notetakers: Michael Khavkin and Alexis Shore Ingbe

AI Detection for Screenshot Prevention? Messaging Platforms and Beyond

Alexis Shore Ingber, University of Michigan School of Information

This use case is trying to solve the problem of screenshot collection and sharing of digital messages. This was presented as a use case, so the author provided an overview of the work that has been done on this topic and how CI could be applied. The author noted that use of the screenshot feature in the context of digital messaging can be classified as a violation of contextual integrity. CI could be used to further unpack the norms surrounding this behavior that are dependent on the sender, subject, recipient, transmission principles and information type. She also proposed using LLMs to train a model to recognize when to protect messaging platforms from the screenshot feature similar to Apple's Sensitive Content Warning. Challenges with this is that this protection cannot be one-size-fits-all, aligning with CI. The author plans to conduct future work on use of the screenshot feature using CI that provides broader implications for privacy management of others' information.

The audience was wondering if watermarking could help protect private messages, where screenshots would have a watermark signaling that they were screenshots. Nissenbaum also pushed that "sensitivity" of data cannot be classified, and it depends on the relationship between actors. The moderator suggested further research on personalizing protection against the screenshot feature on messaging platforms.

Position Paper on Contextual Integrity for AI Assistants

Eugene Bagdasaryan (Google Research), Sahra Ghalebikesabi (Google DeepMind), Ren Yi (Google Research), Borja Balle (Google DeepMind), Leo Cheng (Google DeepMind), Diane Wan (Google DeepMind), Stefan Mellem (Google DeepMind), Octavian Suciu (Google Research), Helen Nissenbaum (Cornell Tech), Po-Sen Huang (Google DeepMind), Sarah de Haas (Google DeepMind)

In this presentation, the first author summarized three published papers on integrating CI into Large Language Models (LLMs) to serve as AI assistants. Similar to the first session's presentation, CI was applied in real-world scenarios.AI assistants are built on three core technologies: LLMs, multimodal I/O, and tools. The motivation for using CI in developing AI assistants is to ensure user privacy by aligning information flows with the users' privacy expectations. Challenges in applying CI to personal AI agents include adherence to norms, reasoning and generalization, robustness to adversarial attacks, and ensuring AI agents follow identified information flows.

The authors made three key contributions. First, the authors created CI benchmarks for assessing how well AI assistants protect personal information and evaluated them similarly to the evaluation process of LLMs.Second, the authors contributed to the direction of privacy-conscious AI assistants by defining an abstract model of information-sharing assistants that are useful for grounding benchmarks and designing evaluation metrics. In addition, such AI assistants can infer contextual attributes under the CI framework, thereby allowing to compare them across different LLMs. The third paper discussed the AI assistants' robustness against adversarial attacks. Inspired by the CI theory, a new adversarial attack named "context hijacking" was introduced, and an assistant that would protect against such attacks was proposed.

The main conclusion is that CI is crucial in designing AI assistants. By defining CI parameters and grounding LLM behavior in CI, we can enhance AI agents' performance. The presenter suggested future research on improving CI capabilities in existing LLMs and developing additional realistic benchmarks. During the Q&A session, several points were raised, including defining the utility of AI assistants and measuring it (e.g., how well LLMs map data to user norms). In addition, Nissenbaum questioned whether LLMs should be referred to as human agents, given they are designed by humans but their behavior is not always deterministic.

SYMPOSIUM CHAIRS

Louise Barkhuus (Rutgers University and The IT University of Copenhagen)
Ruobin Gong (Rutgers University)
Rebecca Reynolds (Rutgers University)
Yan Shvartzshnaider (York University)

PROGRAM COMMITTEE

Noah Apthorpe (Colgate University)
Sebastian Benthall (New York University)
Rachel Cummings (Columbia University)
Cathy Dwyer (Pace University)
Yuanyuan Feng (University of Vermont)
Brett Frischmann (Villanova University)
Kyle Jones (Indiana University-Indianapolis)
Bart Knijnenburg (Clemson University)
Priya Kumar (Pennsylvania State University)
Yafit Lev-Aretz (Zicklin School of Business, Baruch College)
Kirsten Martin (University of Notre Dame)
Lee James McGuigan (University of North Carolina at Chapel Hill)
Mainack Mondal (IIT Kharagpur)
Joel Reardon University of Calgary
Madelyn Sanfilippo (University of Illinois at Urbana-Champaign)
Ido Sivan-Sevilla (University of Maryland)
Luke Stark (Western University)
Katherine J. Strandburg (New York University School of Law)
Daniel Susser (Penn State University)
Eran Toch (Tel Aviv University)
Jessica Vitak (University of Maryland)
Primal Wijesekera (ICSI)
Shikun Aerin Zhang (TikTok)
Michael Zimmer (Marquette University)

STEERING COMMITTEE

Marshini Chetty (University of Chicago)
Helen Nissenbaum (Cornell Tech)

Friday Program

Friday **(PDF Version)**	
All Times are in Eastern Time Zone	
9:00 AM	**Registration, Coffee and Refreshments**
10:30 AM	**Welcome**
Session 1: CI as a lens ***Chair:*** **Sebastian Benthall (NYU)**	
10:40 AM	**Contextual Integrity in Remote Patient Monitoring for Managing Chronic Conditions at Home** D. Ruben Tjhie (University of Toronto)
10:50 AM	**Post-Pandemic Contextual Acceptance of COVID-19 Mitigation Mobile Applications in the US (published work)** Yuanyuan Feng, Brad Stenger (University of Vermont), Shikun Zhang (Carnegie Mellon University)
11:05 AM	**Discussion (10 mins)**
CI Community Feedback session #1 ***Chair: Madelyn Sanfilippo (University of Illinois at Urbana-Champaign)***	

11:15 AM	**Governance seam integrity, contextual integrity, and the integrity of socially meaningful contexts** **Brett Frischmann (VillanovaUniversity)**
	Design(ing) Fictions for Collective Civic Reporting of Privacy Harms **Yuxi Wu (Georgia Institute of Technology), William Agnew (Carnegie Mellon University), W. Keith Edwards (Georgia Institute of Technology), Sauvik Das (Carnegie Mellon University)**
11:25 AM	**Feeback (10 mins)**
Session 2: CI as a framework ***Chair: Louise Barkhuus (Rutgers University)***	
11:35 AM	**A Roadmap for Applying the Contextual Integrity Framework in Qualitative Privacy Research** Priya C. Kumar (Pennsylvania State University), Michael Zimmer(Marquette University) and Jessica Vitak (University of Maryland)
11:50 AM	**Standardizing Contextual Integrity** Sebastian Benthall (New York University), Darra Hoffman(San Jose State University) ,Ido Sivan-Sevilla (University of Maryland)
12:05 PM	**Discussion (10 mins)**
12:15 PM	**Lunch**
Session 3: CI and Assessment ***Chair: Danny Y. Huang (New York University)***	
1:45 PM	**PrivaCI in VR: Exploring Perceptions and Acceptability of Data Sharing in Virtual Reality Through Contextual Integrity** Emiram Kablo, Melina Kleber (Paderborn University) Patricia Arias-Cabarcos (Paderborn University and KASTEL/KIT)

1:55 PM	**Information Flows for Athletes' Health and Performance Data** Brad Stenger and Yuanyuan Feng (University of Vermont)
2:10 PM	**Discussion (10 mins)**
2:20 PM	**Break (30 min)**
Session 4: CI and Policies (Part 1) ***Chair: Noah Apthorpe (Colgate University)***	
2:50 PM	**DiffAudit: Auditing Privacy Practices of Online Services for Children and Adolescents** Olivia Figueira, Rahmadi Trimananda, Athina Markopoulou, Scott Jordan <sjordan@uci.edu> (University of California, Irvine)
3:05 PM	**Governing Manipulative and Synthetic Content on Video-based Social media Platforms** Smirity Kaushik and Madelyn Sanfillipo (University of Illinois at Urbana-Champaign)
3:20 PM	**Discussion (10 mins)**
Session 5: CI and Policies (Part 2) ***Chair: Eugene Bagdasaryan (Google / UMass Amherst)***	
3:30 PM	**Qualitative Analysis of Governing Knowledge Commons and Contextual Integrity (GKC-CI) Privacy Policy Annotations with Large Language Models** Jake Chanenson, Madison Pickering (University of Chicago), Noah Apthorpe (Colgate University)
3:45 PM	**Evidence-Supported Privacy Policies** Allan Lyons (University of Calgary), Yan Shvartzshnaider (York University), Joel Reardon (University of Calgary)

4:00 PM	**Discussion (10 mins)**
4:30 PM	**Panel: Reflections of the day**
5:00 PM	**Predinner break**
6:00 PM	**Dinner**

Saturday Program

Saturday (PDF Version)	
8:30 AM	**Registration, Coffee and Refreshments**
CI Community Feedback session #2 ***Chair: Priya Kumar (Pennsylvania State University)***	
9:00 AM	**Privacy Concerns about Smart TVs from a User Perspective** **Synthia Wang, Lan Gao, Marshini Chetty, Nick Feamster (University of Chicago)**
	Hidden Cam Inspector: Usable Tool to Discover and Locate Hidden WiFi Cameras **Danny Y. Huang and Grace McGrath (New York University)**
9:10 AM	**Feedback (10 mins)**
Session 6: CI and Differential Privacy ***Chair: Ruobin Gong (Rutgers Univesity)***	
9:20 AM	**Investigating the Impact of Differential Privacy on Users' Contextual Data Disclosure Decisions** Michael Khavkin and Eran Toch (Tel Aviv University)
9:35 AM	**Understanding People's Contextual Choices of Differential Privacy** **Kyrie Zhixuan Zhou and Madelyn Rose Sanfilippo (University of Illinois Urbana-Champaign)**
9:50 AM	**Discussion (10 mins)**
10:00 AM	**Break (20 min)**

Session 7: CI and Society ***Chair: Kirsten Martin (Notre Dame University)***	
10:20 AM	**Considering the Ethics of Integrating "Derogatory Data" in Clinical Suicide Prediction Algorithms: A Contextual Integrity Assessment (use case)** Michael Zimmer (Marquette University)
10:30 AM	**Contextual Dimensions of Data Autonomy (published work)** Kyra Abrams (University of Illinois at Urbana-Champaign)
10:45 AM	**Discussion (10 mins)**
Session 8: CI and AI ***Chair: Michael Zimmer (Marquette University)***	
10:55 AM	**AI Detection for Screenshot Prevention? Messaging Platforms and Beyond (use case)** Alexis Shore Ingber (University of Michigan School of Information)
11:05 AM	**Position Paper on Contextual Integrity for AI Assistants (published work)** Eugene Bagdasaryan (Google Research), Sahra Ghalebikesabi (Google DeepMind), Ren Yi (Google Research), Borja Balle (Google DeepMind), Leo Cheng (Google DeepMind), Diane Wan (Google DeepMind), Stefan Mellem (Google DeepMind), Octavian Suciu (Google Research), Helen Nissenbaum (Cornell Tech), Po-Sen Huang (Google DeepMind), Sarah de Haas (Google DeepMind)
11:20 AM	**Discussion (10 mins)**
11:30 AM	**Symposium Wrap Up**
12:00 PM - 1:00 PM	**Mentor's Lunch for NSF's Student Travel Grant Fellows**