The 2nd Symposium on

Applications of Contextual Integrity

SYMPOSIUM REPORT

AUGUST 19-20 2019 | BERKELEY, CA

The 2nd Symposium on Applications of Contextual Integrity

Conference Co-Chairs

Serge Egelman (ICSI & UC Berkeley)

Helen Nissenbaum (Cornell Tech)

Yan Shvartzshnaider (New York University & CITP, Princeton)

Report

Jessie G Taft (Cornell Tech)

Compiled from notes taken by Amanda Stanhaus and Jessie Taft

More Information

http://privaci.info/ci_symposium.html

Generously Sponsored By

With additional support from

Program Committee

Sebastian Benthall (New York University)

Louise Barkhuus (The IT University of Copenhagen)

Marshini Chetty (University of Chicago)

Anupam Datta (CMU)

Pieter De Leenheer (Collibra)

Vicky Froyen (Collibra)

Jake Goldenfein (Cornell Tech)

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

Darakhshan Mir (Bucknell University)

Kirsten Martin (George Washington University School of Business)

Mainack Mondal (IIT Kharagpur)

Xinru Page (Bentley University)

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

Madelyn Sanfilippo (CITP, Princeton)

Luke Stark (Microsoft Research)

Andrew Selbst (Data & Society)

Eran Toch (Tel Aviv University)

Jessica Vitak (University of Maryland)

Pamela Wisniewski (University of Central Florida)

Primal Wijesekera (University of California, Berkeley)

Ding Xiaodong (Renmin University of China Law School)

Contents

Executive Summary

Berkeley's International Computer Science Institute and Cornell Tech's Digital Life Initiative hosted the Second Symposium on Applications of Contextual Integrity on August 19-20, 2019 in Berkeley, CA. The event brought together researchers from computer and information science, communication, political science, and public health to discuss research using the theory of privacy as contextual integrity (CI).

CI defines privacy not as control over information but as information flow that is appropriate given contextual information norms. Norms are defined by the information's subject, sender, recipient, type, and transmission principles. The theory provides a rigorous framework for determining people's perception of appropriateness and ethical implications of modern technologies. As such, its application is an important part of the process of evaluating the privacy risks posed by new and emerging technologies.

Applying the Theory of Privacy as Contextual Integrity is an important part of the process of evaluating the privacy risks posed by new and emerging technologies.

The researchers and practitioners who gathered at the Symposium have each incorporated CI into their work, whether as a foundational principle or as an analytical tool. Over the course of the two-day event, with five panels of talks and multiple collaborative activities, attendees explored how CI can inform policy and system design, and how the theory can be refined, operationalized, and applied to emerging technologies and heretofore unexplored social contexts.

CI: Theoretically Speaking

These presentations combined the CI theory with other theories from across information science, sociology, and organizational behavior. The three presenters focused on Hirschman's exit and voice framework, social exchange theory, and Bourdieu's Field theory. Talks explored the ways in which these theories inform and overlap with CI, and ways in which the theories can work together to create real-world outcomes.

Through the Lens of CI

Each presentation in this panel investigated real-world privacy concerns by applying the CI framework. Presentations explored privacy regulation in internet-connected toys, childrens' mental models of privacy, privacy in disaster contexts, and applications of CI in enterprise systems. While seemingly disparate, each presentation used CI to augment

existing research, resulting in better understandings of information flow in communication, computer science, and organizational behavior.

Norms Discovery

A complete CI analysis needs a full understanding of the norms at play in each context. This session presented research on discovering norms, expectations, and preferences in varying contexts: among older adults, on Facebook, for domestic employees working in smart homes, and in community-based research. The results of each of these studies points the way toward more precise CI analyses. Other outcomes include advances in research methods and data management, and suggestions for the design of privacy-preserving technologies.

CI and System Design

Presentations in this session demonstrate how CI can inform the design of technical systems. Researchers presented their work on using CI to explore privacy in the design of integrated electronic health records, surgical robots, and recommendation systems. Another presentation addressed how context can be best represented in technical systems generally. These projects suggest ways that the abstract concepts involved in CI can be made concrete for use in technical systems, and how technical systems that require considerations of privacy can benefit from the results of CI analyses.

CI and Reasonable Expectations

CI assumes that users of a technology have certain expectations, values, and norms about how information flows in each context. The presentations in this session explored privacy expectations in marginalized communities and in national social credit systems, proposed computational modeling of context with values as predicates, and developed a privacy paradigm that encompasses reasonable expectations and can be incorporated into existing engineering lifecycles. Each paper concluded with calls for further engagement with the specific contexts of the research, from global society to small communities.

Open Mics and Breakouts

At several points during the symposium, attendees were given the opportunity to raise their own questions relating to applications of CI, or to apply CI to a new context in a focused breakout session. The “CI Open Mic” on Monday provided a forum for feedback on new research ideas and open questions in the CI research community. The CI and Regulation breakout session on Tuesday gave attendees the task of understanding how to regulate new technologies via understanding of information flows.

Common Themes

Defining Context. Throughout the symposium, attendees raised the question of defining context. Privacy scholars use the word in different ways across disciplines and fields, and even within discussions of CI, no specific definition of context has been completely agreed upon. Many talks also surfaced the idea of agency in defining context, noting that context may be different from the perspectives of different actors within an information flow.

Technical Operationalization. A number of presentations explored how CI can be operationalized in research methods, technical systems, and interfaces. While better definitions of context, norms, values, and transmission principles may assist in this effort, no clear "best practices" for incorporating CI into new technologies has emerged. Future work will need to bring these research findings together to create guidelines for those wishing to apply CI to real-world projects.

Next steps. Previous symposia raised the need for a CI 2.0 - one that would incorporate updated definitions and clarify understanding of norms, values, information flows, and contexts. While the 2019 symposium made efforts to unify these diverse interpretations, as many new questions were raised as were answered. Future symposia will build on this year's enthusiasm for collaboration to further refine the CI framework and promote understanding of CI across domains.

Future CI Symposia will build on this year's enthusiasm for collaboration to further refine the CI framework and promote understanding of CI across domains.

Symposium Overview

Day 1, Session 1: CI: Theoretically Speaking

Chair: Jake Goldenfein

Legitimacy in Context

Ashley E. Gorham, Helen Nissenbaum, Madelyn R. Sanfilippo, Katherine Strandburg, **Mark Verstraete**

Contextual Integrity's focus on informational norms raises questions about how those norms are constructed. This talk describes three cases of norms that may not be sufficiently addressed by CI: bad norms, conflicting norms due to overlapping contexts, and yet-to-be-established norms due to new technologies. While CI generally offers robust privacy prescriptions, these three cases raise questions of legitimacy and governance.

To address these gaps, the authors use Hirschman's exit, voice, and loyalty framework, which describes the actions available to those who wish to express dissatisfaction with an organization. Their analysis focused on the prevalence of notice and choice, and how exit - leaving a situation or group as an expression of disagreement, or here, choosing not to use a system - is insufficient in some cases. They describe similar situations in homeowner associations and landlord/tenant agreements.

Q&A

Discussion of this work addressed the gaps the researchers noted in the CI framework. It was noted that the question becomes one of new practices vs. entrenched expectations – which are legitimate in light of new technology? Attendees also added that loyalty is voluntary and makes a distinction between personal loyalty and public accountability.

Another point of discussion addressed a way to tie the examples (homeowner's association and landlord/tenant) in ways that would strengthen their work. The authors responded that the commonality of the examples is that the exit options are limited, power imbalance present, and include private property.

Blending Contextual Integrity and Social Exchange Theory: Assessing Norm Building Under Conditions of "Informational Inequality"

Jennifer King, Andreas Katsanevas

This work addresses issues of institutional privacy - privacy as it applies to information exchange between an individual and an institution. These situations often involve power asymmetry, which must be accounted for in privacy analyses.

The work's research questions are:

- How does social structure affect privacy and personal disclosure?
- How do the imbalances between individuals and companies affect decisions to disclose information?
- Does examining personal disclosure from a perspective that accounts for social structure reframe the privacy paradox? That is, are privacy decisions rational given available options?
- How are "appropriate" norms determined in the context of what Nissenbaum describes as information inequality? What kind of social processes are constructed when entities wield power over individuals?

This work utilizes social exchange theory (SET), a theory from sociology that focuses on the benefits people obtain from, and contribute to, social interaction. It accounts for how power is distributed, and contributes to a network view of the interactions between people and institutions. For example, it may help describe how norms of information exchange between people and companies evolve over time.

While useful when considering disclosure, SET does not define privacy. Thus, CI may be helpful when disclosure relations, and their accompanying norms, are formed under extreme power imbalances. As an example, the researchers discuss responses to Facebook's evolving privacy practices, from treating users' information as public by default at its beginning to facing a 2011 FTC consent decree for sharing information without user consent. The researchers argue that these practices have forced Facebook's own evolving information norms onto the Facebook-using public, in contrast with Mark Zuckerberg's statements that social norms have evolved naturally over time.

Applying CI to Facebook's actions, the researchers argue, shows a violation of contextual integrity. Though Facebook eventually walked back some of its public defaults, as a powerful actor it manipulated norms of exchange for its own advantage, undermining the meaning of contextually appropriate norms.

What are the implications for contextual integrity of entities with power influencing norms in digital context? For example, companies offering consumer genetic testing services are defining norms of what it means to share genetic data online before consumers themselves have the chance to develop their own values around that information.

Q&A

The resulting discussion focused on norms vs. practices, with attendees noting that Mark Zuckerberg's quote about social norms used in the presentation really meant *practices*. Outstanding research questions include how norms are entrenched and the power dynamics in play as norms are formed and practices are implemented.

Contexts are Political: Field Theory and Privacy

Sebastian Benthall and Bruce Haynes

The present research aims to address the question of how non-dominant and dominant norms can co-exist using Bourdieu's Field theory. The researchers note that CI has two gaps: cross-context information flows and social adaptation to sources of normativity.

Bourdieu views social outcomes as predictable because power distributions are predictable. Individuals' relationships to social others matter when privacy is considered, and the power relations that happen within fields structure the relational world. These power relationships structure how participants in the field think of legitimacy, institutions, capital, and other concepts that have implications for privacy.

This complex interconnectivity between institutions, relationships, power, and privacy has implications for the design of real-world systems such as smart infrastructure.

Q&A

The following discussion focused on the European approach, with its utilization of purpose-driven norms. Under the GDPR, data providers are obligated to use data for the purpose it was collected, however context is not relevant. Who decides what purposes are legitimate?

Day 1, Session 2: Through the Lens of CI

Chair: Yafit Lev-Aretz

Evaluating the Contextual Integrity of Privacy Regulation: Parents' IoT Toy Privacy Norms Versus COPPA

Noah Apthorpe, Sarah Varghese and Nick Feamster

In 2017, the FTC updated its Children's Online Privacy Protection Act (COPPA) guidelines to address privacy issues caused by connected toys and other IoT devices. However, the extent to which COPPA-mandated data handling practices align with parents' privacy norms is unclear. More generally, research is needed to identify whether population norms align with regulation.

A previous project used survey methods to study users' privacy norms for home IoT devices. The researchers modified this method to compare parents' privacy norms to regulation. Their process is as follows:

1. Choose parameter values from the CI model that are related to regulation. For example, data senders are a subject or attribute; the recipient is a manufacturer or service provider; and transmission principles are derived from the FTC's COPPA compliance plan.
2. Generate combinations of parameters and assemble them into information flows. Generate vignettes from each flow, and create surveys that ask about the acceptability of each flow.
3. Test surveys using cognitive interviews using crowdworkers who are expert survey-takers.
4. Deploy surveys to parents of children of COPPA-relevant ages.
5. Group responses by individual parameters, pairs of parameters, and participant demographics.
6. Compare likert acceptability scores across groups.

Results showed that information flows allowed by COPPA are acceptable on average, while those disallowed were unacceptable on average. The researchers are careful to note that this does not mean that COPPA is sufficient to protect privacy, but that COPPA's requirements align with the general direction of parental privacy concerns.

Other findings include:

1. Parents dislike targeted advertising.

2. Notification and consent transmission principles are more influential than those related to confidentiality and security, pointing to issues surrounding consent. Parents may assume that security is built in to a toy, or is not a problem.
3. Parents are more accepting of data collection by first than third parties.
4. Respondents familiar with COPPA are more comfortable with data collection. The researchers posit that this may indicate a false sense of security provided by the regulation.
5. Older parents see data collection as less acceptable than younger parents do. This may be due to differences in norms, or in mental models.

The researchers suggest that future work should apply this survey method to other privacy regulations. The results on the alignment between population norms and policy should then be used to draft more informed regulations. Finally, the survey should be re-run every few years to capture the evolution of norms surrounding IoT and connected toys.

The researchers posed several questions relating to research on contextual integrity:

- Are null parameter controls contrary to theory? The researchers used a 4-parameter CI model rather than the full 5 parameters. Adding parameters may be useful as experimental methodology, but could also confound the results as participants make assumptions with regard to the missing parameters.
- Context of context: what parameters do people assume? How does this affect the reported appropriateness of an information flow?
- Transmission principles: COPPA guidelines have post-flow criteria, or they may imply other or multiple information flows. How these principles should be constructed in CI is not obvious.
- Beyond norm discovery: In addition to norms, research must also measure functions, purposes, and values; and their distribution across communities.
- Norms and mental models: How do norms and mental models interact? How is the appropriateness of a flow affected by mental models?
- Research recommendation: Mixed effects models and pre-deployment testing with crowdworkers can improve results.

The researchers concluded by noting that they would propose that the FTC to improve specificity of information recipients, closing the spaces within “third party” definitions that allow for a variety of acceptability. They also call for more technical specificity in consent and notification processes, moving toward a model that acknowledges that notice and consent does not work well.

Understanding Children's Mental Models of Privacy based on the Theory of Contextual Integrity

Hoda Mehrpouyan, Jerry Alan Fails, Dhanush kumar Ratakonda

This research group's approach is to extend CI to provide mathematical models and algorithms that enable creation of privacy management norms. They aim to include environmental and temporal variables to answer the following question: What happens when you want information to be shared in a certain sequence?

To address this question, they use the User Information Sharing Model (UISM) and the Privacy Preservation Model (PPM) of behavior. Their goal is to map these models onto each other using privacy verification.

The UISM is a model based on entities (senders), agents (receivers), attributes, and knowledge (what recipient knows). It includes a defined set of abstractions for roles and information types. This allows for graphical modeling of attribute types, which can then be learned based on context and user actions. Role abstractions allow agents to be assigned to multiple roles, which are then ordered. UISM is represented as a 4-tuple transition system.

PPM is a model designed to manage and govern user information sharing activities at runtime, where it either allows a transition to happen or not. It defines access permissions and attributes, environmental and temporal conditions using positional logic with predicates and variables. For temporal conditions, regular expressions can be used to find sequences.

While contexts may change, norms must remain consistent. The researchers propose a verification engine that ensures a mapping between UISM and PPM. It verifies compliance by comparing regular expressions, and in case of violation, does not allow information sharing behavior; rather, it asks the user whether they are comfortable sharing.

To test these models, the researchers conducted semi-structured interviews with children about what they share, with whom they share, and in what contexts they share information. Children were asked about their knowledge of security and privacy, and the social roles they occupy. Differences between stated behavior and actual behavior were also tested, revealing a considerable gap between stated privacy practices and actual behavior.

Current privacy management systems do not address dynamic requirements of privacy, since they're not designed from users' perspectives. CI allows privacy management

systems to be customized to individual user needs. Further work is required to allow automatic norm creation based on user behavior, rather than building norms on stated behavior.

Results show some idiosyncrasies in children's understanding of privacy and security. They are sometimes over-trusting, and sometimes under-trusting when it might not make sense. More research is needed to understand behavior changes over time and after additional exposure to privacy concepts.

Q&A

Attendees commented that this research brings back the hidden variable missing from the previous talk, and emphasizes the importance of qualitative methods in understanding preferences and expectations. While some preferences can be uncovered by observing and learning behavior, there are also risks in becoming captured by the appeal of machine learning based on observational data.

Disaster Privacy/Privacy Disaster

Madelyn R. Sanfilippo, Yan Shvartzshnaider, Irwin Reyes, Helen Nissenbaum, Serge Egelman

In disaster contexts, saving lives depends on getting the right information to the right people at the right time. However, in times of disaster, norms change: where normally people do not want to share their location, this information is widely shared during disasters. Technologies entering a disaster context often change as well; for example, platforms encourage using social media to inform others of well-being during disasters.

Government agencies also share information during disasters. Federal disaster privacy governance focused on information types, rather than overall information flows. Guidelines restrict what can be collected and who has access. Transmission principles are limited to "need to know", some norms exist for what is expected for those who receive information, but not much else. The researchers illustrate this with a CI analysis of FEMA's leak of the personally identifying information of disaster survivors.

What do disaster apps do? The researchers annotated app privacy policies using CI, looking for actors, attributes, and transmission principles. They then conducted a dynamic analysis with AppCensus. Results show that apps send data to third parties for unknown reasons not described in their policies.

Some apps conformed both to government policies and to their own privacy policies. For others, actual information flows are not consistent with either app privacy policies or

government policies. Some apps pose as government services but are not, creating situations in which the user expects regulatory compliance but the app shares information with third parties.

In general, analyses showed that more third parties have access to disaster app data than governance models account for. Not all parameters of disaster information are governed, and we lack a clear understanding of these contexts.

Q&A

The discussion for this presentation explored the meaning of trust in disaster contexts. Disasters activate community and increase interpersonal trust, allowing said trust to be taken advantage of. How can these tradeoffs be managed without alienating people? In addition, how can risks be managed when humanitarian aid delivery depends on data collection?

Use Case: Using Contextual Integrity in an Enterprise Context

Vicky Froyen

Collibra provides data management and governance services for large enterprise clients. Dr. Froyen presents it as a potential use case for evaluating CI; as it balances creating data frameworks for enterprises with growth and adaptation to upcoming privacy regulations. Researchers and developers within Collibra want to be accountable innovators, with understanding of data governance, user expectations, and real world contexts.

Research questions that could be addressed with CI analyses in this context include:

- Can CI be a universal information flow language?
- How can CI be employed to monitor data flows and estimate risk?
- How can CI adjust policy and educate users/others?

At Collibra, governance of data requires shared understanding – people doing something with data need to understand what a flow is to effectively place an internal regulation on it. Collibra's position as a growing and privacy-responsive company make it an ideal testbed for applying CI principles.

Day 1: CI Open Mic

Several new collaborative projects were born from last year's CI symposium. The organizers hoped to use this time to discuss open questions in CI research and provide feedback on ongoing research projects, possibly even sparking new collaborations.

Attendees raised questions of the subjectivity of context. Timothy Kariotis noted that CI research often fails to ask about peoples' lived experiences. In the healthcare context, for example, the practice of "active pre-interpretation" should be more carefully examined. The subsequent conversation centered on the question of who gets to decide what context is, and ways to ensure that those involved in each context have agency.

Another line of discussion was suggested by Bernadette Boscoe, who sought attendees' advice on using CI to manage qualitative research data. She noted that researchers adhere to rules set by their IRBs, or to FERPA or HIPAA rules. However, they may fail to recognize privacy violations that might adhere from re-identification. attendees provided literature review suggestions along these lines.

Finally, attendees discussed the question of norms vs. expectations. It was noted that expectations are individual, while norms apply across groups. Some wondered how many people must have an expectation for it to become a norm, and whether norm differences across sub-populations might affect understandings of context. They noted as a larger point that a clearer understanding of the term "context" as it is used in CI is needed, to avoid confusion across disciplines.

How can we create a clearer understanding of *context* as it is used in CI to avoid confusion across disciplines?

Day 2, Session 1: CI and Norms Discovery

Chair: Kirsten Martin

The Contextual Preferences of Older Adults on Information Sharing

Alisa Frik, Julia Bernd and Noura Alomar

As the population of older adults increases and the number of available caretakers decreases, many are coming to rely on technology to supplement care for the aging. New technologies focus on both independence and health for older adults, and many rely on communication to families and caregivers. These devices are novel and complex, and collects a lot of information. Older adults themselves often have limited technology literacy and experience, declining abilities, less awareness of privacy risks, and greater susceptibility to security concerns.

The goal of the present research is to inform the design of effective systems for older adults, allowing them to make informed decisions, have better control over their data, and maintain better security practices.

The researchers conducted semi-structured interviews with older adults, with questions such as: What information do they expect to be collected, what is ok to collect and share? With whom is it ok to share? How can information be used and misused? They asked about specific devices: fitness trackers, smart home things, and virtual reality devices.

It is important to note that the researchers did not begin this project intending to use CI as a framework. However, the results of their interviews showed that the dimensions that matter to older adults when making privacy decisions - decision-maker, data, recipients, purpose/benefit, risk, system, environment - bear many parallels to CI, although the present research is more explanatory and does not rely as heavily on norms.

The study found that transmission principles affect privacy decision-making most heavily. Previous experiences with privacy violations also impact privacy attitudes and behavior, as does environment - the ecosystem of norms, media stories, friends' experiences, and alternatives in which information is shared. The researchers note that there are many complex interactions between different variables, making context difficult to define.

One lesson to be drawn from this research is the need for an expanded understanding of transmission principles. Their salience in participants' decision-making makes this understanding essential to creating design recommendations for privacy protection. The researchers also suggest incorporating other privacy theories, such as privacy calculus.

Future research areas include explicating stated vs. actual privacy behaviors. Participants mention paradigm examples when explaining privacy behaviors, but tend to mean more granular preferences. These findings could help develop better intervention strategies.

Q&A

Questions after this talk addressed the interplay between environment and context. If environments fit within context, how can context be deconstructed? And what is the relationship between an observer and the environment or context they are observing?

Using Long-Lived Facebook Accounts to Understand Implicit Norms of Consent in Contextual Integrity

Mainack Mondal, Zhou Jin, Tamara Babaian, Xinru Page, Blase Ur

Social media contexts present plentiful opportunities for misunderstanding. In many situations, the sender and data subject are the same entity, but this isn't always the case. Facebook, for instance, has possibilities for explicit consent – users can review posts before they appear on their timelines. But there's also implicit consent – people implicitly believe that there's a social norm that makes it ok to share information about others. CI could be used to untangle these norms of consent.

Research questions in this study were:

- What are the perceived CI norms when the sender does not equal data subject?
- When do sender-subject mismatches lead to violated expectations?
- What are potential mitigation strategies for these violations?

The researchers conducted interviews with people with long-lived Facebook accounts, with the goal of collecting information from those users who would have had a variety of experiences. Interviews included both audio recording and screen sharing, allowing the researchers to see users' Facebook feeds. The study sees Facebook as a data archive: in the activity log, the participant is a data subject. In the timeline, the participant is the sender. Part of the interview focused on each aspect. The interviewer walked the participants through CI parameters while browsing through their timelines to get separate thoughts on each area.

Participant responses included anecdotes describing information flows and information management. Some noted that posts by others that unearthed past posts constituted inappropriate flow. Others noted that management of self-disclosure is difficult or is overtaken by others. Finally, participants mentioned bowing to social pressure when they responded to posts that solicit information that they may not have wished to provide.

As a next step, the researchers seek feedback on the methods and framing of the study.

Q&A

Attendees suggested using findings from these studies to discover whether people have surprise or disappointment reactions in response to information flow violations.

Further discussion involved centering the technology in platform-specific research. Here, interactions are mediated by Facebook, but is Facebook the context? The context might be better described as a relationship between two people. Input from strangers might change the context entirely.

Smart Home Bystanders: Further Complexifying a Complex Context

Julia Bernd, Alisa Frik, Maritza Johnson, Nathan Malkin

Bystanders – here, accidental participants in information flows - are a non-canconical data subject. IoT devices, which collect data about environments, have privacy controls that are designed for end users. Neither data collection targets nor bystanders are end users: they have different knowledge, expectations, preferences, and attitudes about data collection. This study focuses in particular on smart homes. The presence of multiple contexts within a home make them complex in terms of determining privacy norms.

Bystanders in smart homes can be domestic employees, residents, visitors, or groups. This study investigates domestic employees, highlighting the differences between targets and bystanders. The domestic employment context also brings complicated employer-employee power dynamics, questions of accountability, and SES differences. Here, the location of interest is not just a home context with home-based privacy norms, but also a workplace context and a caregiving context. While this particular caregiving arrangement is recognized as a common situation, norms differ across the world.

The present study investigates differing attitudes between nannies and parents: what are their experiences, expectations, power dynamics, and points of intervention? The researchers reviewed posts about smart devices on the Nanny forum on Reddit in a preliminary investigation of these attitudes.

Results showed that nannies acknowledge conflicting contextual norms. They also report accepting being monitored in their workplace if they are also able to access those same monitoring techniques understand their purpose. As a next step, they hope to conduct additional surveys and interviews to better understand how CI can address the complexities of the smart home context.

Applying Contextual Integrity Framework for Community-Engaged Research Data Management

Jina Huh-Yoo

The present work is a case study describing how community obesity prevention researchers engage with privacy issues in designing health technologies and managing health data. The researchers take a holistic approach to child obesity prevention, improving home routines with informational awareness. In a move to shift away from mobile health and toward ambient awareness, they sought to provide families in a predominantly immigrant community with ambient monitoring devices that pick up acoustic information about the home environment. The researchers hoped to use a custom-designed device, while families preferred established consumer devices such as those made by Amazon.

This preference seemed paradoxical, even given the fact that families saw consumer goods such as Amazon's as more desirable than a device custom-built for research. The ongoing negotiation between researchers and families showed that community members trusted Amazon's treatment of their data, but lacked trust in researchers and university. In fact, the university in question did not have a structure in place to protect community data: participants gave sensitive information in interviews, but that data was only minimally protected.

This case study reveals the complicated context of community-engaged research data management. Questions for consideration include:

- To what extent is CI effective as a framework for community-engaged research data management?
- Is a community representative's information flow preferences the same as any other individual's in that community?
- How should intra-community dynamics be accounted for in research and data management?
- Should community advisors be considered human subjects with respect to data management and confidentiality requirements?
- How can communities be educated about over-disclosing personal information?

Q&A

The discussion for this talk centered around CI's impact on research institutions: Could CI help create a checklist for human subjects research projects that allows them to better account for participants' privacy concerns? And how can relationships of trust between institutions and communities of participants be built, or improved upon?

Day 2, Session 2: CI and System Design

Chair: Susanne Wetzel

Contextual Privacy by Design for Integrated Electronic Health Records: The Information Continuum Project

Timothy Kariotis, Megan Prictor, Kathleen Gray, Shanton Chang, Darakhshan Mir

The present work has its roots in efforts to design the integrated electronic health record (EHR) of the future, one that is integrated across services for a collaborative model of care. Such records are complicated and must have many types of information integrated within them. The researchers find that current EHR models breach information norms - but that the benefits of effective information sharing may still outweigh the risks. They hope to use CI to understand how norms adapt to the introduction of new technologies, and how the evaluative power of CI can help design future systems.

Other privacy-related theories and frameworks can assist in this analysis. Privacy by Design is useful, if applied in a way that takes context into account. Theories of appropriation of technology can help explain what happens when technology is put in context and when context is shaped by technology. Adaptive structuration theory, which states that technology has structural features (things you can do) and spirit (values and interpretation), is one such. When placed in an organizational structure, technology is appropriated in a way that is either faithful to the intentions of its designers, or is not.

When CI is applied to this understanding of technological appropriation, we begin to see emergent norms, which are apart from designed-in norms (the result of structural features and spirit) and context of design (designers' interpretations of a context). The tension between designed-in norms and actual norms in context lead to varying appropriations of the technology.

The authors propose using participatory and co-design methods to discover existing and emergent information norms for EHRs. Themes from stakeholder focus groups will inform design guidelines that help EHRs respect information norms.

Q&A

The discussion after this talk addressed variations in context. Do applications of CI vary in international contexts? How can CI help differentiate between the people involved in an information flow and the information transmitted? Attendees noted that norms

themselves are a recognition of the fact that contexts can vary widely. Norms can be both beneficial and toxic.

Privacy with Surgical Robotics: Challenges in Applying Contextual Privacy Theory

Shishir Nagaraja, Ryan Shah

Surgical robots have high accuracy and precision, and high efficiency in a variety of simple surgical procedures. They also have implications for patient privacy, as the data needed to perform such procedures is highly sensitive.

Traditionally, patient privacy is based on the consent of the patient, and also of regulators and observers (eg. surgeons). Privacy relies on faith in the surgeon, medical staff, hospital standards, and legal guardians involved in the procedure. When robot surgeons are involved, the threat model changes. Unlike human surgeons, software can be untrusted.

The present work attempts to better understand this threat model by applying CI. The authors argue that CI is data-centric, and needs to focus on the services, subject, and operations involved in the surgical context. How can CI apply to managing these new uses of control? How can they manage the conflict introduced by the untrusted surgeon?

The authors propose the concept of privacy in calibration. The data required for calibrating a robot to perform surgery on a particular individual flows from a high-integrity source (secure medical data) to a low-integrity destination (a robot with untrusted software). Combining calibration traffic with identifiable patient information could compromise patient privacy. The authors call for future work in this area to express privacy in a way that's both enforceable and understandable, along with expert review to ensure that robot actions are not too complex for patients to understand.

Q&A

Discussion in the Q&A surrounded the normative "flavor" of the word "leak" - here, used to describe the transfer of information to an untrusted party.

The Limitations of Technical Mechanisms as Proxies for Context

Catherine Dwyer

This presentation addressed the challenge of representing context within technical systems. In these systems, context must necessarily be operationalized using ones and

zeros. However, context shifts constantly, and information flows occurs regardless of whether context has been established.

Computer science research conceptualizes context as concrete and technical, but CI calls it normative. Here, the author describes Hildebrandt's framework for the origin of personal data, used by the World Economic Forum. It characterizes data as being volunteered (deliberately provided; notice and consent implies that everything is in this category), observed (measurable behaviors that are machine readable), or inferred (an automated process combines data streams and comes up with something new).

These categories, the author argues, are being used in technical systems as a substitute for context, as a way to justify information flows. She argues that the assumptions inherent in these categories must be assessed through the lens of CI, and explores the specific contexts of social media data and location sharing. Information flow can disclose sensitive information (location, politics, etc). Incremental steps in the flow can be explained with the three aforementioned categories. However, there's no clear place to "fix" the flow, meaning any technical solutions rely on the three categories rather than on ethical norms.

Future work in this area will explore different ways of calculating location that are built into common apps, and beta testing tools to extract data.

Contextual Recommendation Sharing

Eugene Bagdasaryan

Recommendation systems help personalize users' online experiences. Using data traces from various sources, they can efficiently discover user tastes. Thus, greater access to user data means better recommendations, but also and more privacy and intellectual property risks.

In creating recommendation systems, larger platforms with more data sources and greater computing resources have an advantage over smaller ones. However, user preferences are rarely shared between platforms because of intellectual property and trade secret-related barriers. The present work explores ways to use CI to share recommendations across platforms without sharing data that may be protected. Specifically, can CI help limit exposure of user preferences? And can we vary exposure based on context?

As an example, can Google News help users get better recommendations from small or new sources? How can small platforms be prevented from inappropriately using data and models provided by larger platforms? Latent representations could be shared rather than

raw data or models, but these are different for different models and can be reverse engineered. Instead, the researcher proposes sharing the model's output - the recommendations themselves.

This approach allows the larger platform to use an "exposure controller" to filter out sensitive items from their recommendations, or otherwise distort the output, to protect user privacy. This has the added benefit of increasing the diversity of recommendations.

In sum, this approach to recommendation decentralizes insights from user data, provides privacy controls, and is relevant for many use cases.

Q&A

Discussion explored the nature of the large platform/small platform dichotomy. Why is the large platform always trusted and the small platform the attacker? Why is this adversarial model being used, rather than another?

Surprise Breakout Groups

Instead of the follow-up open mic described in the program, organizers used this time for an interactive activity. Attendees were divided into groups, and each assigned a technology, either extant or speculative. They were given the following prompt: You have all the power to regulate a technical system. How might you use CI to structure a regulatory or normative response to its uses? What would you need to know to do that?

How can Contextual Integrity structure a regulatory or normative response to a new technology?

Attendees generated lists of questions that would need to be answered to effectively regulate such a technology.

Facial Recognition

A system that identifies human faces and matches them with identities.

What is the purpose of the recognition? Who built it, and how, and who is accountable for it? Who holds power in this system? What are the pros and cons of the system for minorities? Technical questions - What does the flow from people to databases and back look like? Does detection happen on a camera or in the cloud? Is it always on, or activated selectively? How accurate is it, and how is ambiguity managed? What data was used to build the model, how was it acquired, and is it data accessible? How does Is there a human in the loop? Legal questions - Do data subjects understand the system? Do they have the right to opt out or dispute judgments? What are the implications of an incorrect identification? Is there purpose binding in the system?

Smart Home Device

An interactive in-home system such as a smart speaker.

Who is use for and by? What are the assumptions and mental models surrounding the device? What are the device's technical capabilities? What is the purpose, and the potential future use cases? Who has access? What else do they know about you? What is the benefit to the consumer vs. to the company? What is the permission structure for the technology? What are the reporting obligations, or is there a human in the loop? Who does the regulation? Is there human observation of the data? Are there concerns around manipulation, intellectual property, or anti-trust issues? Are there any other successful models that can be used to abstract this in a useful way?

Brain-wave Reader

A speculative technology that makes human thought legible.

Who has access and can you control it? How does the device actually work - does it detect any waveform, or just final intentions? What type of data – what image/wave, what part of the brain, what goes out in the world? Are there existing norms that should be accounted for, and what can be derived from them? Or would we need to start fresh with understanding the context? Issues of translation- content moderation around slang? How does inner thought change appropriate flow? Is it interceptable, and would we know if intercept occurs? Could thoughts be used in different contexts? Can thoughts be filtered? Is there a have and have not situation where people without are disadvantaged?

Conclusions

In order to effectively think about the regulation of each system, designers and policymakers need to think about flow. Asking these questions is the foundation of an effective CI analysis, and of regulation: each are easier when a system is well-specified.

Day 2, Session 3: CI and Reasonable Expectations

Chair: Michael Byrne

Contextual Integrity and Reasonable Expectations: Privacy Paradigm

Amin Rabinia, Daniel Nathan, Sepideh Ghanavati

Privacy is an interdisciplinary challenge, and as such is often unclear for developers and engineers. Privacy laws, where they exist, do not have clear guidelines for interpretation. Thus, engineering solutions to privacy problems are ad hoc. Improved developer understanding of privacy needs and practices is a necessity.

The authors propose privacy paradigms as a solution to this problem. They define a paradigm as a combination of theory, method, and technique. They propose a Privacy Paradigm (PriPa) in which CI is the theory and the Privacy by Design approach is the method. Here, the hallmark of privacy is consistency of the informational behaviors of a system with the reasonable expectations of its users.

To help guide technological development, PriPa can be integrated into engineering lifecycles in ways that make it understandable and usable by developers. The engineering lifecycle used here has five steps: extraction, refinement, modeling, design, and verification. The first three steps constitute the 'privacy requirements' phase, in which CI parameters are identified, conflicts are resolved, and models are documented. In the design phase, PbD's principles are incorporated. In the final phase, successful integration of privacy-preserving mechanisms is verified and reported.

In future work, the authors hope to answer the following questions:

- How well does CI capture users' expectations - is it sufficiently user-centric?
- How specifically can CI be implemented in the engineering lifecycle system?
- How can the paradigm be tested - how do we tell if the theory works?

Q&A

In the discussion after this talk, attendees challenged the assumption that the concept of a privacy paradigm is completely new, and suggested that the researchers incorporate or reconcile with existing privacy engineering models, such as NIST privacy engineering principles.

Attendees suggested that the authors look into literature on privacy compliance in engineering, and the different philosophies of privacy that underlie each of these theories and approaches.

Using MDPs to Model Contextual Integrity

Michael Tschantz

What does CI do well? it predicts why simpler theories fail to predict privacy outcries, and explains why events turn into privacy outcries. What does it not do? It is not predictive (it has too many degrees of freedom), it doesn't determine which contexts are relative, what their norms can be, or what to do when contexts conflict. In addition, many of its parts are too vague to program. Specifically, context needs to be made more precise.

The researcher presents the following two hypotheses, with the acknowledgement that they are too vague to be falsified:

> Hypothesis 1: The defining feature of context is a purpose.
>
> Hypothesis 2: Decision and game theory can be used to model context (because these define how purpose-driven agents should behave, forces precision, highlights gaps, makes assumptions explicit, and makes computer-aided creation and analysis of models easier).

The researcher gives the example of healthcare: Health decisions can be mapped into a probabilistic model of outcomes. This model can be used to decide what actions to take - it forms a Markov decision process.

Thus, context can be understood as purpose optimization, in which purpose is a scoring function to maximize satisfaction in each state. However, CI states that context is not just purpose: it also includes values. In this model, values are predicates stating whether the value is respected in each state. These factors create a constrained Markov decision process. Therefore, norms are understood as legitimate if for all states, the actions available are roughly optimal for purpose and value.

This model can be used to decompose purpose and context: purposes can have sub-purposes, some of which are mutually exclusive. Each sub-purpose can be optimized for individually.

Questions for further investigation include: Can we view society as a global context? If so, what is its purpose? Can the values of a society super-context be used to model a "deontological ethics" state? Can such modeling help us understand cross-context concerns?

Q&A

Are there values inscribed in purposes even before values are introduced? What happens when there's disagreement about what to optimize for? The author acknowledges that there are constraints.

How does this map on to CI? The arrows in this Markov model are more actions than flows. The author states that partially observable Markov models could be used instead. In addition, CI could be viewed as being about actions rather than just information flows?

Uncovering Privacy Norms in Marginalized Communities

Dylan Rogers, Desmond Dinkins, Gia Hayes, Richard Stover, Shin-Won Cho, Jennifer Silva, Evan Peck, Darakhshan Mir

Marginalized and rural populations are often most harmed by privacy violations. In areas like the coal regions of central Pennsylvania, communities marginalized by factors such as the worsening rural-urban divide experience hyper-visibility in systems and a host of other information-related harms.

The researchers present ongoing work that seeks to understand informational privacy norms in these communities. How do people define privacy? How do they understand technology use? What are their perceptions of information flow, values, and control? What are specific within-community information sharing and communication concerns? Focused discussion questions and thematic analyses uncovered these and other privacy values, experiences, and knowledge in rural Pennsylvania communities.

Results show that privacy is highly valued across these communities - rarely were participants apathetic about privacy, except insofar as they lacked agency - many felt that they did not have the self-efficacy and information literacy required to manage information flows. The researchers also found a lack of trust in government, law enforcement, and local media, feeling that their consent to participate in these systems is forced in order to receive assistance, access services, and obtain information.

Intra-community information sharing was a concern for many, who noted that rumors and personal information leaks were a concern. Participants also noted demographic change tensions, risks relating to social media, and the importance of mutual trust in small communities. They were interested in increasing their own literacy and engagement in information management and governance.

In future work, the authors hope to further explicate the folk model frameworks at work in these communities, and co-design educational materials with the goal of increasing agency to self-manage privacy.

Q&A

What's the remedy to being asked for too much information? The presenter suggests that perceptions of requests for too much information may be based on a misunderstanding about why information was necessary. This could be solved by giving more clarity about why information is needed, or reforming systems so that they need less information.

Are there differences between information flows experienced by this population vs. other populations? Or are the norms just different? The presenter notes that the smaller community means that interpersonal relationships are more intimate, whether or not they are wanted, meaning that intra-community information flows are different.

Rethinking the Social Credit System: A Long Road to Establishing Trust in Chinese Society

Yuhao Zhong and Xiaodong Ding

This work uses CI to better understand the social credit system in China. This system of initiatives, at central and local levels and involving both corporate and government participation, is meant to encourage trust and sincerity in business and social interaction. It began with regional trials in 2009, with eight credit scoring firms proceeding to a national pilot in 2014. In 2018, these efforts were centralized under the Bank of China, with participation from other credit scoring firms.

The social credit system re-defines privacy, offering new opportunities for CI analyses. In breaking down barriers between "private" and "public" information, the social credit system changes what information flows might be considered inappropriate, making CI a useful privacy framework.

The authors offer the example of information flows and travel bans. Under the social credit system, people with low credit scores or those who have been convicted of undesirable behavior can be banned from flying and using high-speed trains. Instead of transmitting information about these infractions just between courts and transport companies, it is also transmitted to the media and to the public for purposes of public shaming.

CI may also assist in re-thinking the philosophies behind the social credit system. Currently, those "keeping trust" in the system maintain respect, those who break it are shamed. But these punishments are not based on context - people breaking the credit system's rules are shamed regardless of the severity of their acts.

The author concludes with a call for social credit system authorities to consider CI as a tool for discovering potential privacy violations. Remaining questions for future work include CI's utility for lawmakers considering social credit systems, and for finding flaws in existing legislation.

Q&A

Is there a reciprocity in this system, like with Uber? No, because this system involves an agency evaluating an individual, rather than individuals evaluating each other.

Other discussion surrounded the transparency of the score's input to the general population, and the logics of who within the system is entitled to have an impact on another person's score. Implications of merging financial and social data were also discussed.

Appendices

Program

Monday, August 19

2:15-3:15pm | CI: Theoretically Speaking
Chair: Jake Goldenfein

- Blending Contextual Integrity and Social Exchange Theory: Assessing Norm Building Under Conditions of "Informational Inequality"
  Jennifer King, Andreas Katsanevas
- Contexts are Political: Field Theory and Privacy
  Sebastian Benthall and Bruce Haynes
- Legitimacy in Context
  Ashley E. Gorham, Helen Nissenbaum, Madelyn R. Sanfilippo, Katherine Strandburg, Mark Verstraete

Short Break (15 min)

3:30-4:45pm | Through the Lens of CI
Chair: Yafit Lev-Aretz

- Evaluating the Contextual Integrity of Privacy Regulation: Parents' IoT Toy Privacy Norms Versus COPPA
  Noah Apthorpe, Sarah Varghese and Nick Feamster
- Understanding Children's Mental Models of Privacy based on the Theory of Contextual Integrity
  Hoda Mehrpouyan, Jerry Alan Fails, Dhanush kumar Ratakonda
- Disaster Privacy/Privacy Disaster
  Madelyn R. Sanfilippo, Yan Shvartzshnaider, Irwin Reyes, Helen Nissenbaum, Serge Egelman
- Use Case: Using Contextual Integrity in an Enterprise Context
  Vicky Froyen

Short Break (15 min)

5-6:00pm | Open Mic: CI Playbook for the Future of Privacy

7:30pm | Symposium Dinner

Tuesday, August 20

9:00am | Coffee and Refreshments

9:30-10:45am | CI and Norms Discovery
Chair: Kirsten Martin

- The Contextual Preferences of Older Adults on Information Sharing
  Alisa Frik, Julia Bernd and Noura Alomar
  Slides
- Using Long-Lived Facebook Accounts to Understand Implicit Norms of Consent in Contextual Integrity
  Mainack Mondal, Zhou Jin, Tamara Babaian, Xinru Page, Blase Ur
- Smart Home Bystanders: Further Complexifying a Complex Context
  Julia Bernd, Alisa Frik, Maritza Johnson, Nathan Malkin
- Applying Contextual Integrity Framework for Community-Engaged Research Data Management
  Jina Huh-Yoo

Short Break (15) min

11:00-12:15pm | CI and System Design
Chair: Susanne Wetzel

- Contextual Privacy by Design for Integrated Electronic Health Records: The Information Continuum Project
  Timothy Kariotis, Megan Prictor, Kathleen Gray, Shanton Chang, Darakhshan Mir
- Privacy with Surgical Robotics: Challenges in applying contextual privacy theory
  Shishir Nagaraja, Ryan Shah
- The Limitations of Technical Mechanisms as Proxies for Context
  Catherine Dwyer
- Contextual Recommendation Sharing
  Eugene Bagdasaryan

12:15-1:30pm | Lunch

1:30-2:15pm | Open Mic Follow Up/Lightning talks

Break

2:30-3:40pm | CI and Reasonable Expectations
Chair: Michael Byrne

- Using MDPs to Model Contextual Integrity
  Michael Tschantz

- Contextual Integrity and Reasonable Expectations: Privacy Paradigm
  Amin Rabinia, Daniel Nathan, Sepideh Ghanavati
- Uncovering Privacy Norms in Marginalized Communities
  Dylan Rogers, Desmond Dinkins, Gia Hayes, Richard Stover, Shin-Won Cho, Jennifer Silva, Evan Peck, Darakhshan Mir
- Rethinking the Social Credit System: A Long Road to Establishing Trust in Chinese Society
  Yuhao Zhong and Xiaodong Ding

Break

4:00-5:00pm | Discussion, Wrap up