

Foundations and Trends® in Privacy and Security
Vol. 2, No. 1 (2017) 1–69
© 2017 S. Benthall, S. Gürses and H. Nissenbaum
DOI: 10.1561/3300000016



Contextual Integrity through the Lens of Computer Science

Sebastian Benthall
NYU Steinhardt MCC and UC Berkeley School of Information

Seda Gürses
FWO/COSIC, KULeuven and CITP, Princeton

Helen Nissenbaum
Cornell Tech and NYU Steinhardt MCC

Contents

1	Introduction	2
2	Privacy and Context in Computing	7
2.1	Contextual Integrity	8
2.2	Context in Computing	12
2.3	Context in Privacy Policy	16
3	The Study	18
3.1	Research Questions	19
3.2	Study Methodology	20
3.3	Three exemplars of analysis	22
4	Results	33
4.1	(RQ1) Architecture	35
4.2	(RQ2) Capturing Context	39
4.3	(RQ3) Source of Normativity	42
4.4	(RQ4) Expanding Contextual Integrity	45
5	Findings and Discussion	51
5.1	Architecture: Towards a Modular Contextual Integrity . . .	52
5.2	Diverse concepts of context	55
5.3	Privacy as a Moral Imperative: Between bits and norms . .	57

	iii
5.4 Expanding and Sharpening Contextual Integrity	60
Acknowledgements	63
Appendices	64
A Research Template	65
References	66

Abstract

The theory of Privacy as Contextual Integrity (CI) defines privacy as appropriate information flow according to norms specific to social contexts or spheres. CI has had uptake in different subfields of computer science research. Computer scientists using CI have innovated as they have implemented the theory and blended it with other traditions, such as context-aware computing. This survey examines computer science literature using Contextual Integrity and discovers: (1) the way CI is used depends on the technical architecture of the system being designed, (2) ‘context’ is interpreted variously in this literature, only sometimes consistently with CI, (3) computer scientists do not engage in the normative aspects of CI, instead drawing from their own disciplines to motivate their work, and (4) this work reveals many areas where CI can sharpen or expand to be more actionable to computer scientists. We identify many theoretical gaps in CI exposed by this research and invite computer scientists to continue exploring the horizons of CI.

1

Introduction

Privacy is both an elusive moral concept and an essential requirement for the design of information systems. The theory of Contextual Integrity (CI) is a philosophical framework that unifies multiple concepts of privacy – as confidentiality, control, and social practice [Gürses and Diaz, 2013] – and has potential as a systematic approach to privacy by design [Nissenbaum, 2009]. Indeed, over the last decade, computer scientists in a variety of subfields such as security, HCI, and artificial intelligence have approached the challenge of technical privacy design by applying CI.

This is a structured survey and review of this body of work. This survey has threefold aims: 1) to characterize the different ways various efforts have interpreted and applied CI; 2) to identify gaps in both contextual integrity and its technical projection that this body of work reveals; 3) perhaps most significant, it aims to distill insights from these applications in order to facilitate future applications of contextual integrity in privacy research and design. We call this, “making CI more actionable for computer science and computer scientists.”

Over the last 20 years, privacy by design [Cavoukian, 2009, Federal Trade Commission, 2012, Regulation (EU), 2016] and privacy

engineering [Gürses and del Alamo, 2016] have become research topics that span multiple sub-disciplines in computer science [Rubinstein, 2010, Danezis and et al, 2015]. Prior work has shown that computer scientists often stick to a single definition of privacy, for example, confidentiality, secrecy, or control over personal information. Although reducing privacy to a narrow definition has generated interesting work, it has limitations in addressing the complexities of privacy as an ethical value. In the wild narrow definitions offer analytic clarity, yet they may stray too far from a meaningful conception of privacy, that is, a conception that people actually care about.

The theory of contextual integrity (CI) was offered as a rigorous philosophical account of privacy that reflected its natural meaning while also explaining its moral force. Generally CI characterizes privacy as appropriate information flow, and appropriate flow characterized in terms of three parameters: actors (subject, sender, recipient), information type, and transmission principles. This definition immediately sets it apart from definitions in terms of subject control or stoppage of flow. Besides allowing for a more expressive framing of privacy threats and solutions than other approaches, the additional factors allow for greater specificity – hence less ambiguity – in prescribing and prohibiting certain flows. Because CI allows formal representation of flow constraints, it may serve to bridge privacy needs experienced by humans, in situ, with privacy mechanisms in digital systems. Although CI’s account of privacy’s ethical importance plays a lesser role in the work we have surveyed it remains important as a normative justification for Privacy by Design (PbD) initiatives grounded in privacy as contextual integrity.

With this survey we aim for more than a description of leading scientific applications of CI; in addition, we seek an exchange of ideas. In assessing how these applications have engaged with CI we ascertain, in one direction of exchange, how true to the letter they have been and how the framework might have been better or more fully reflected in the work. Equally, in the other direction, we assess how these frontrunners may materially inform future developments of CI itself. Such insights are crucial to enhancing the capacity of CI both to challenge and inspire scientific work and technical design, thus making CI more actionable for

computer scientists. We conclude the survey by providing prescriptive guidance going forward.

On the one hand, our findings reveal that, for the most part, computer scientists engaged in technical design do not take up contextual integrity in its full theoretical scope. They often give specificity and depth to some of its concepts while bracketing others, rarely addressing normative dimensions of CI explicitly. Another common departure from CI is how researchers interpret context, which often maps onto their respective disciplinary assumptions and strategies, and literatures. In reviewing the literature, it was not our aim to declare any of these approaches “wrong” or “misguided”. Instead, our aim was to record our findings, identify different perspectives, and note discrepancies as opportunities to learn from, revise, expand, and improve CI, to guide future research practice, and, in our terms, to make CI more “actionable.”

On the other hand, the forays into implementing systems using contextual integrity lead to significant innovations and improvisation that we believe can inform contextual integrity theory. For example, the papers we have surveyed have elaborated on different types of contexts, most prominently those that arise as a consequence of interactions (with people and machines) or those that come to bear as changes occur in the environmental conditions surrounding users. In doing so, computer scientists tease out different socio-technical situations that may impact how informational norms play out in a social sphere.

Relatedly, the desire of computer scientists to design systems that observe and adapt to changes throughout time is common. This is reflected in the use of technical mechanisms that capture changes to contexts, social norms and environments and that respond to the evolving conditions. Finally, papers we surveyed often position users as central actors, highlighting their role and agency in engaging and transforming informational norms in a context and throughout time. These concrete instances of socio-technical contexts, adaptivity and user agency shed light on issues that, with some elaboration, could enhance the analytical power of CI for privacy by design.

The authors have also taken up known challenges to CI, as in the case of papers that propose solutions to applying the framework when

information flows from one social situation to the other or when multiple contexts are co-located. The question of multiple contexts is acute especially in technologies that act as infrastructures, e.g., platforms that host multiple applications or that host actors and actions from multiple social contexts. The solution space proposed by the authors include mechanisms to negotiate information flows across contexts or agents introduced to reason about multiple contexts in a single application.

Overall, when authors delegate responsibility of governing CI to technical elements that act semi-autonomously in an adaptive environment, they also raise novel questions for CI. It is not uncommon that researchers design agents that reason about contexts, auditing mechanisms that ensure informational norms are not violated, or apps that take active part in negotiating permissions. Can technical mechanisms be seen as actors in CI? If so, are they acting in the same context as they are “serving” or are they in a different context? These are some of the hard questions revealed by the works we studied, and that require input when considering future applications of CI in technical systems.

Just as important to our analysis is what the authors have not attended to in applying the CI framework. We reference here CI’s account of privacy’s ethical legitimacy, which identifies stakeholder interests, societal values, and contextual ends, purposes, and values as the basis for such legitimacy. Although omission of this aspect of CI might not be problematic for the technical accounts of privacy given by computer scientists, it nevertheless warrants explanation and justification; we discuss these issues at the end of the paper. Finally, we found no reference to information theoretic advances in privacy technologies, e.g., differential privacy or privacy preserving machine learning. Given the growing role that machine learning and artificial intelligence is playing in information systems, we believe there is a great potential in exploring how CI may be applied in systems that have the ability to infer and reason using data.

Finding that these papers both narrow and expand the CI framework, our review concludes with issues that will be important to address if CI is going to be useful to a wider spectrum of computer

scientists. These include attending to questions such as: how to be more technically precise about the nature of contexts in order to relate the definition given in CI with more concrete notions used by computer scientists; how to advance normative concepts in CI – i.e. ends, purposes, and values – by taking advantage of well developed methods in the scientific study of privacy, including user studies, models, threat models and threat agents; how to use CI in systems where data not only flows, but also persists in a single place; and, how to apply contextual integrity to systems that function as infrastructure across multiple contexts.

2

Privacy and Context in Computing

In this section we provide theoretical background in privacy and context in computer science that situates our findings. Section 2.1 details Contextual Integrity as a philosophical theory that aspires to be robust to changes wrought by technology, rooted in legal, ethical, and social theory. Our study is primarily of how this framework has been used by computer scientists. As an inductive result we discovered that this work often drew from different conceptions of context as relevant to privacy that came from different subdisciplines in CS. In particular, we found computer scientists working at the intersection of CI, introduced in computer science literature with Barth et al. [2006], and in the tradition of ubiquitous computing that has given a central place to context and its implications for privacy design [Dey et al., 2001], [Dourish, 2004]. We detail the latter in section 2.2. We note in section 2.3 how the connection between context and privacy has become recognized by policy-makers [White House, 2012, World Economic Forum, 2012]. We speculate that this policy recognition was responsible for an uptick in the interest of computer scientists in context and privacy. The resulting creative synthesis of multiple traditions offers an opportunity for realizing new theoretical insights and opening new research

problems.

2.1 Contextual Integrity

The practice of privacy may be as old as social life itself but the contemporary need for a concept of privacy, rich enough to drive policy and precise enough to shape architecture follows in the wake of advances in technologies that have disrupted how we create, collect, communicate, disseminate, interpret, process, and utilize data. (It is worth noting, however, that rarely, if ever, is it raw technology that stirs agitation; instead agitation is a response to technology embedded within particular social practices and particular political ecosystems.) In the US the contemporary need to sharpen the concept and strengthening protection is often dated back to 1895 with Warren and Brandeis’s historic call to define a legal right to privacy in the wake of new photographic and printing technologies. In 1973, the landmark Report to the Secretary of Housing, Education, and Welfare issued Principles of Fair Information Practices (FIPs) following the rise of massive computerized databases.

In the 1990s and 2000s, such systems extended to video, audio, and online surveillance, RFID, and biometrics systems. Subsequently, public attention has turned to hyperbole over “big data” – database technologies, computational power, and scientific advances in information and data processing. Dramatically amplifying the privacy impacts of these technologies are transformations in the software engineering industry – with the shift from shrink-wrap software to services– spawning an agile and ever more powerful information industry. The resulting technologies like social media, user generated content sites, and the rise of data brokers who bridged this new-fangled sector with traditional industries, all contribute to a data landscape filled with privacy perils.

Approaches to privacy that depended on neat divisions of spaces or data into public and private have been severely challenged by these developments. Long entrenched definitions of privacy rights as rights to control information or rights to secrecy, that is, to block access, were

overly simplistic, either too easily challenged by those with competing interests or over-claiming on the part of data subjects. An account that captured the complex contingencies of legitimate privacy claims was needed – one that benefitted from conceptual building blocks of existing theories but offered a greater expressive agility, to resist incursions while allowing the positive potential of novel socio-technical systems to be realized. Contextual integrity intends to provide such an account. For one, it addresses gaps in prior entrenched conceptions allowing it to identify privacy threats to which other accounts were blind (e.g. “privacy in public”). It also offers a view on the nature and sources of disruptive information flows in order to distinguish that that constitute threats from those that do not.

The theory of privacy as contextual integrity (CI) introduces three key concepts into the privacy vocabulary:

- 1) **Contexts.** These refer to social contexts, not formally constructed but, discoverable as natural constituents of social life. As theorized in sociology, social theory, and social philosophy, they have been assigned various labels, including, social domains, social spheres, fields, or institutions. (Throughout this survey, we will use the term *sphere* to denote this sense of context.) Societal recognition of distinct contexts, such as healthcare, family and home life, politics, religion, commercial marketplace, and education is clearly evidenced in distinctive structures of law and regulation.

For the framework of contextual integrity, contexts are formally characterized in terms of key elements, which include, paradigmatic activities, roles (or capacities), practices, and norms. Distinguishing contexts from one another, are contextual goals, ends, purposes, and values, around which activities and norms are oriented, and to which respective contexts are committed.

- 2) **Contextual informational (privacy) norms.** Among contextual norms, these govern information flows and, according to contextual integrity, are likely to map onto people’s privacy expectations. Informational norms are well-formed only if they refer

to five parameters: sender, recipient, and information subject, information types (topics, attributes), and transmission principle. The **parameters** of actors and attributes range over contextual ontologies, distinctive to respective social contexts, if not unique. Thus, in healthcare context, senders, recipients, and subjects range over **agents** acting in the capacities, such as, doctor, nurse, patient, surgeon, psychotherapist, etc. and **topics** may range over symptoms, diagnoses, and drug prescriptions. **Transmission principles** condition the flow of information from party to party, including those commonly associated with privacy, such as, *with permission of data subject*, *with notice*, or *in confidence*, in addition to those less salient, such as, *required by law*, *with a warrant*, and *entitled by recipient*.

Privacy as contextual integrity is respected when entrenched informational norms are followed. When these norms are violated (e.g. by disruptive information flows due to newly functioning technical systems) there is a prima facie case for asserting that privacy has been violated. The framework of contextual integrity allows, however, for the legitimacy of disruptive flows to be defended, as described below.

- 3) **Contextual ends, purposes, and values.** These may be considered the “essence” of a context, without which respective contexts would not be comprehensible. How would one properly describe a school, say, without indicating its purpose? These – let us call them – teleological factors are also important in defending the legitimacy of informational norms, particularly useful when comparing novel information flows against past expectations, or when no competing alternative is obvious, they are useful in evaluating the ethical legitimacy of given flows taken alone.

According to the CI framework, privacy norms can be assessed in terms of how they affect the interests of relevant parties (“stakeholders”) and how they impinge on societal values, such as equality, justice, fairness and political liberties. In addition to these considerations the norms governing flow can be evaluated in terms of their impacts on

the attainment of contextual ends, purposes, and values – either promoting or confounding them. For example, informational norms enabling (and enforcing) a secret ballot protects autonomous voting in elections and, as such, promotes ends and values of democracy.

This structure ensures that though CI is conservative, in the sense that it presumes in favor of entrenched norms, it nevertheless has built into it a set way for systematically evaluating and updating norms. This is done by examining balance of interests, general ethical and political values, and contextual values and purposes. It follows that norms adapt to their environment, crucial for an account of privacy to remain relevant in the face of advancing technologies of information and computational technologies. Societal and environmental shifts can destabilize entrenched privacy norms in many ways, either revealing that they are no longer optimal in achieving contextual ends and values or have nothing to say about disturbing information practices. Although such circumstances constitute challenges for ethicists and social policy makers, they do not necessarily constitute challenges to CI itself, which copes with novel or disruptive flows by presenting new norms for consideration.

Adapting norms to novel or disruptive flows may involve adjusting any of the parameters. For example, the increasing digital mediation of transactions, communications, and interactions (including social media) creates new data recipients, which forces the reconsideration of norms. The same goes for increasing specialization and fracture of skills and functions within traditional contexts, healthcare being a prime example. The one-on-one physician-patient relationship paradigmatic of the distant past has been replaced by an immensely complex care and treatment ecosystem, involving specialists, insurance companies, pathologists, public health officials, wireless pacemaker service providers, and, with that, the emergence of new informational norms – in the ideal, to serve contextual ends and values.

A note on terminology: We refer to aspects of CI that deal with evaluating the *legitimacy* of norms as its normative, prescriptive, or ethical aspects. These aspects are contrasted with what we might call its *descriptive* or conceptual aspects, referring to the the structure of

informational norms.

2.2 Context in Computing

CI bridges two worlds. In one, it is an account of privacy as a term that has accrued meaning to describe alarm over wide-ranging, technology induced practices of surveillance, data accrual, distribution, and analysis. The alarm is due to disruptive practices that violate privacy expectations and create or amplify imbalances in power. CI posits contextual informational norms to model privacy expectations and explains when such expectations are morally legitimate and warrant societal protection. In the other, CI offers a formal structure for expressing rules of information flow (informational norms) and for building computational models of people’s privacy expectations in well-defined settings (contexts.) The first is a world inhabited by humanists, social scientists, lawyers, and regulators; the second is inhabited by mathematicians, computer scientists, and engineers. Perhaps because it seeks to map a meaningful conception of privacy onto a conception that strives for formal rigor contextual integrity has been taken up by computer scientists interested in privacy design and engineering.

Although philosophical versions of contextual integrity appeared in articles, dating back to 1998 [Nissenbaum, 1998] and, later, in the book, *Privacy in Context*, it was not represented in computer science literature until [Barth et al., 2006]. This paper, which we introduce as one of our survey exemplars in Section 3.3.1, formalized the fragment of CI known as context-specific (or, contextual) information norms. The authors, which include Nissenbaum, developed a logical framework for expressing and reasoning about norms and showed that this framework is adequate for expressing regulations drawn from U.S. sectoral privacy laws, such as, HIPAA, GLBA, and COPPA.

In fact, contextual integrity is but one source of influence that has drawn computer scientists to engage with the idea of context as it relates to privacy. Two others are worth discussing because they have roused the interest of computer scientists in context and shaped how they conceive of context with respect to privacy. We therefore find

in computer science loose interpretations of contextual integrity that consider these other forms of context. They are: the field ubiquitous computing and the Obama White House Bill of Consumer Privacy Bill of Rights [White House, 2012] (also the World Economic Forum and FTC Reports around a similar time [World Economic Forum, 2012]). Grasping these influences has been important in advancing our own ability to analyze the articles we have chosen for this survey.

2.2.1 Context in ubiquitous computing

Contextual integrity is not the only research tradition linking context and privacy in computer science. Many contemporary issues in human-computer interaction around mobile devices and IoT were anticipated in earlier waves of research into “ubiquitous computing”. This research program envisioned a world in which computation was not restricted to specialized workstations but, instead, was embedded in everyday objects and practices, enabling user interaction through sensors and actuators. Within ubiquitous computing research interest emerged in developing technologies that were responsive to social and environmental context, that is to say, ‘context-aware’ computing.

In their “anchor article” on context-aware computing, Dey et al. [2001] extensively analyze definitions of ‘context’ in the literature of their field and settle on the following for their own work:

Context: any information that can be used to characterize the situation of entities (i.e., whether a person, place, or object) that are considered relevant to the interaction between a user and an application, including the user and the application themselves. Context is typically the location, identity, and state of people, groups, and computational and physical objects. [Dey et al., 2001]

This definition of context, specifically referring to the concrete *situation* of persons and objects, starkly contrasts with the notion historically evolved, abstract, and normative social *spheres* of CI. In our survey of the computer science literature invoking contextual integrity, we found that several papers conceive of ‘context’ in ways that have more in

common with context-aware computing than with context as defined in CI. This has led to interesting synthetic work, in addition to incipient disjunctures.

That computer scientists have taken up the tradition of context-aware computing in their work on privacy as contextual integrity is not surprising. Early in this field, contributors Ackerman et al. [2001] anticipated that context-aware computing would lead to privacy by design, arguing that technical systems and legal frameworks would be co-designed. But hints of the connection between context-aware computing and contextual integrity, which was not formulated as a framework until later, were present at least as early as Dey et al.'s anchor article:

As computational capabilities seep into our environments, there is an ever-growing and legitimate concern that technologists are not spending enough of their intellectual cycles on the social implications of their work. There is hope, however. Context-aware computing holds a promise of providing mechanisms to support some social concerns. For example, using context to tag captured information may at first seem like an intrusion to individual privacy (“I do not want to meet in a room that records what I say.”). However, that same context can be used to protect the concerns of individuals by providing computational support to establish default access control lists to any information that has been captured [Lau et al., 1999], limiting distribution to those who were in attendance. [Dey et al., 2001]

Most relevant to this article are the implicit connections these authors draw between the design of technology responding to a particular *situation* (certain people meeting in an office room) and a general expectation of privacy. The norm that literal, unfiltered information about what happens in meetings is available only to those who attended could be attributed to the abstract social *sphere* of office meetings. This prefigures a result of our study, which finds computer scientists taking up ‘context’ in ways that reflect both senses of the word, and in so doing implicitly drawing connections between them.

This early work on context-aware computing reflected the state of the art in sensors and the kind of sensing they made possible: explicitly representing context as a kind of fixed container in which people could act. This was famously critiqued in a paper by Dourish [2004], connecting how ‘context’ is approached in ubiquitous computing to broader questions in philosophy of science. Dourish argued that representing context (e.g. location or time in which an application is used), as a kind of container for activity whose boundaries are delineable draws from the *positivist* tradition in social science that sees context as stable and separable from the activity taking place within it. Dourish contrasts this understanding of context with a different one deriving from the *phenomenological* tradition of social science. According to it, context is *occasioned*, “relevant to particular settings, particular instances of action, and particular parties to that action”, not an abstract category that generalizes over settings, actions, and parties. This kind of context arises from and is maintained by its activity, sometimes dynamically adjusting along with the activities themselves. For example, a private conversation between friendly colleagues at work can shift from a formal, professional discussion into an informal, personal discussion and back again. These shifts will occur as and through changes in the conversational activity, such as changes in tones of voice or comments such as, “Well, we should really get back to work; I have to go in twenty minutes.”

Dourish investigates how this conception of context ties into the sociological mystery of how social order comes into being. There is a tension between explanations of social order that attribute it to rules, expectations, and conventions that have a broader reality beyond particular occasions of interaction (what we might call a ‘top-down’ ordering), and explanations that see all social order as arising from interaction itself as an achievement of the social actors (‘bottom-up’ ordering).

While it may be argued that top-down and bottom-up ordering are always co-occurring, often one or the other process is emphasized in scholarly work. Contextual integrity, in its original articulations [Nissenbaum, 2004, 2009], tends to emphasize the top-down pressure of contextual ends, purposes, and values shaping norms that in turn guide

information flows. In contrast, while Dourish acknowledges the role of top-down orderings, he highlights the bottom-up processes that make each context occasioned and dynamic, in the spirit of his interactional, phenomenological objection to static representations of context. We find that both ways of thinking about context are prominent in the literature that we review, even though we have limited this review only to computer science literature that refers to contextual integrity.

2.3 Context in Privacy Policy

While we were looking specifically for computer science papers that referenced contextual integrity, it was interesting to find many papers that took “privacy in context” as an idea (which also happens to be the title of Nissenbaum’s book about contextual integrity [Nissenbaum, 2009]), that do not draw from the framework of contextual integrity. If the direct origin of “privacy in context” was not contextual integrity, what was it?

Our contention is that interest was prompted by the general uptake of context and contextual integrity in the formulation of several policy documents from 2010 and later. For example, the White House Report, “Consumer data privacy in a networked world: A framework for protecting privacy and promoting innovation in the global digital economy“ [White House, 2012], lists “Respect for Context” as one of its seven principles: “Consumers have a right to expect that companies will collect, use, and disclose personal data in ways that are consistent with the context in which consumers provide the data.” Around the same time, a report issued by the U.S. Federal Trade Commission also invoked context when it stipulated that data collection by companies should be restricted to what was appropriate for the “context of interaction” or else they should make “appropriate disclosures”. In a similar vein the World Economic Forum’s 2012 report, *Rethinking Personal Data* invokes the importance of context for policy governing data in numerous places [World Economic Forum, 2012].

We have found significant variation in how computer scientists have interpreted the term “context”, often reflecting their disciplinary back-

ground and research agendas. Some follow contextual integrity quite closely. Others cite ‘contextual integrity’ and ‘privacy in context’, possibly to situate the privacy - context connection within a scholarly lineage without following CI substantively. Most of these papers were written after the policy arena acknowledged CI theory, while in parallel ubicomp researchers had already established a concept of context. For some computer scientists, context as *situation*, reminiscent of ubiquitous computing research, informs their work, as it does some of the work on privacy regulation (albeit outside the focus of this paper.) Overall, though computer scientists have, characteristically, explored the relationship between various forms of context and privacy in a rigorous and pragmatic way, they have not made the definition of context the subject of explicit theoretical commitment.

Nevertheless this work in computer science at the boundaries of contextual integrity makes important contributions to the theory itself. Inspired broadly by contextual integrity, computer scientists have explored aspects of the relationship between privacy and context in detail. Our systematic study of these works has found in the variations and commonalities within this literature insights that can inform and inspire further developments in contextual integrity.

3

The Study

The main objective of this study is to characterize the different ways CI has been interpreted and applied in computer science, reveal its technical projection, and thereby, capture gaps in CI itself. The long term objective of this study is to identify ways that CI can be made more actionable for computer scientists and systems developers. In order to do so we systematically reviewed literature coming out of different sub-fields of computer science explicitly stating the use of contextual integrity in their problem or solution definition. We made use of techniques proposed by Kitchenham et al. [2007] to make our study as comprehensive and transparent as possible.

In their projects invoking CI, computer scientists have taken on the hard task of translating an elaborate philosophical framework into computer science research practice – in which different theoretical and methodological traditions apply. This renders the translation of CI into technical contexts a non-trivial task. For these reasons alone an assessment of current uses of the theory in CS is valuable for understanding how well the theory translates, what new questions arise when applied in a technical context, and what obstacles become evident. Through this survey we evaluate its uptake in computer science and begin to

sharpen the theory to make it more actionable for researchers who want to use it in the future.

3.1 Research Questions

Driven by the motivations listed above, we decided to focus on the four research questions as we take stock of the use of CI in computer science research and assess it:

RQ1. For what kind of problems and solutions do computer scientists use CI?

As an initial question for our inquiry, we wanted to know if there were any particularly notable categories of problems being addressed by computer scientists using contextual integrity. Computer science is a broad field; researchers may have found contextual integrity useful for solving particular kinds of problems, focus on certain domains, or be more likely to invoke CI in certain subfields of computer science.

RQ2. How have the authors dealt with the conceptual aspects of CI?

Contextual integrity is partly a conceptual theory that is predictive of social concerns about privacy that originate and manifest themselves especially with technological change. The theory posits *social contexts* as evolved abstract spheres of activity characterized by *ends, purposes, and values*. Social contexts have *information norms*, parameterized by *actors (senders, recipients, and subjects)*, *information types*, and *transmission principles*. Contextual integrity identifies privacy as appropriate information flow; such flow would be characterized by contextual informational norms.

We wanted to know to what extent the computer science researchers using contextual integrity used this conceptualization of privacy. Do the researchers define context in the way contextual integrity does, or in other ways? And do they define privacy in terms of appropriate information flow according to norms?

RQ3. How have the authors dealt with the normative aspects of CI?

Contextual integrity is a normative framework of privacy. It argues that privacy is an important value because appropriate information flow promotes the data subject's interests in balance with others; it is an important *social* value because it promotes societal and ethical values, and maintains the ability of social contexts to fulfill their ends, purposes, and values.

We wanted to know if computer scientists using contextual integrity take up this normative aspect of the theory. If not, from where do they perceive the normative clout of privacy coming? How do they evaluate whether privacy is addressed effectively through their proposed mechanisms or solutions?

3.1.1 RQ4. Do the researchers expand on CI?

In developing technical systems computer scientists have to make a number of substantive and specific design decisions. This is also the point at which the rubber meets the road: how does a researcher translate a philosophical theory into a formulation useful for technical design? In executing this translation computer scientists are likely to attend to concrete questions that CI may not provide explicit guidance for. In these moments, researchers are likely to identify gaps in CI and propose techniques to make up for these gaps. What are the gaps that researchers identify, how do they expand on these, and how do they stretch the theory – explicitly or implicitly?

3.2 Study Methodology

In compiling and revising the relevant papers, we followed empirical research methodologies recommended for use in software engineering studies [Kitchenham et al., 2007]. In order to answer our research questions, we conducted the following four steps:

1. Based on our research questions (Section 3.1), we iteratively developed an initial template of analytic questions using a selection of CI articles.

2. We searched in online repositories for papers using CI as its reference theory. To ensure we have a reasonable collection, we searched digital libraries (Google Scholar, IEEE Xplore, ACM DL) for papers that appeared in CS venues that had CI in their title or main body. To cast a wider net, we included the key terms “contextual integrity” and “context AND privacy”. For those papers that explicitly invoked CI, we combed through later publications that cited them to see whether the use of CI propagated. We carefully evaluated the inclusion of papers that only reference CI without making further use of the framework. In the process, we found a number of papers that refer to context and contextual norms that do not refer to Nissenbaum’s work and removed these from the study. Evaluating whether and how CI may have proven useful in these papers is out of the scope of the current work. Some papers claimed they used CI and integrated other conceptions of “context” in CS, we kept these papers in our study. We initially categorized papers with respect to the subfields of computer science from which they originated. The represented fields of research included security engineering (including privacy engineering and access control); artificial intelligence (including papers on multi-agent systems, machine learning, semantic web, social network analysis and community detection); systems (distributed systems, pervasive and mobile computing); HCI (usable security and privacy, ubiquitous computing); and software engineering (requirements engineering and business process design).
3. Once we had completed our search, we tested the completeness and consistency of the template based on close reading of additional articles. Once the template was stable (see Appendix A), we (the authors) independently read each paper and answered each question of the analytic template for it. We did a comparative analysis of the answers in order to distill those aspects of the papers that answered our research questions. At this stage, we also concluded a quality assessment of each paper with respect to its contributions to computer science and removed those that failed

our assessment. We documented all of our analysis and discussions in an online repository.

4. We used the output of the templates to complete a thematic analysis of each paper. We consolidated what we had discovered into major categories of themes, one for each research question. Our work indicated the most productive way to interpret these questions. For RQ1, we found the most significant way we could characterize the variety of problems addressed in the literature was by looking at the kind of technological architecture researchers were designing. For RQ2, we focused on how researchers characterized “context” in their work. We split this concept of ‘context’ down into many finer-grained variables in order to show the variability between papers. For RQ3, we looked specifically for sources of normativity used by each paper and coded them accordingly. For RQ4, we analyzed the ways in which the papers expanded on contextual integrity. Our analysis did not reveal that the initial categorization of papers according to subfields in CS revealed further insights for our study.

In the remainder of this Section, we provide detailed accounts of select papers as illustrations of how we thematically analyzed each paper in accordance with the steps described above.

3.3 Three exemplars of analysis

In order to provide the reader with a demonstration of how we got to the different themes in our results, we pulled out three of the papers to serve as exemplars. We selected these three papers as they deeply engage CI; they stem from different subfields in CS with varying methods and techniques; and, they allow us to demonstrate the rather different ways in which the authors have elaborated on CI. The curious reader is encouraged to read these full papers which are rich in ideas and thoughtful in their use of CI. All other papers are analyzed according to respective categories and themes, extracted through the template that had guided our reading of them.

3.3.1 **Privacy and Contextual Integrity: Framework and Applications (Barth, Datta, Mitchell, and Nissenbaum)**

The first published computer science paper to reference contextual integrity was coauthored by Helen Nissenbaum and therefore can be said to be an authoritative expression of the theory. It is not, strictly speaking, a paper about the design of a technological artifact. Rather, it is an articulation of a subset of the principles and parameters of contextual integrity in a formal logic [something further discussed in Section 4.1 under RQ1]. Formalization is a prerequisite to computational implementation, and so this paper demonstrated the potential of contextual integrity as a guide to the design of computational systems. For the purposes of our study it is just as notable what it did not formalize into logic, as this has left open many challenges to computer scientists seeking to use contextual integrity.

After grounding the work in an exposition on contextual integrity theory, the first major contribution of the paper is a careful translation of principles of contextual integrity into formal logic. The particular flavor is Linear Temporal Logic (LTL), a type of logic which is capable of expressing formulae of relationships of variables arranged in time. This translation refines the ontology of contextual integrity by making explicit that information flows have a temporal order. This allows the authors to define specific transmission principles that condition appropriate flow on previous flows [further discussion under RQ2 in Section 4.2]. The logical specification allows a particular history or trace of information flows to be audited for appropriateness according to formal rules.

One of the benefits of having to make the logic of contextual integrity explicit is that it brings to light aspects of the theory that are easy to take for granted but which have far-reaching implications. The paper explicitly models both the knowledge available to each actor at different points in time as well as the ways that different attributes are related to each other via inference. This paper therefore provides an epistemic model that is only implicit in other accounts of CI. Having provided a formal language for expressing policies in the style of CI's context-specific information norms, the authors go on to prove a

number of theorems about the computational complexity of auditing traces based on these policies, testing for the possibility of complying with the policy, and comparing policies.

The authors do not tie their formalization back to the origin of norms through the evolution of social sphere and its ends, purposes, and values. Rather, after formalizing the aspects of contextual integrity that they are able to, they validate their work by showing that it is expressive of United States sectoral privacy laws: HIPAA, GLBA, and COPPA (see Datta et al. [2011] for further work along these lines). They also argue that the expressivity of their formalization compares favorably with other proposed access control policy languages such as XACML, ECAP, and P3P.

This paper is particularly notable as the first published computer science paper concerning contextual integrity. Explicitly only a formalization of *part* of CI, [Barth et al., 2006] provide a way of expressing norms as policies that can be used in computational tests for compliance. This sets a precedent for computer science papers using contextual integrity to consider ‘context’ in a structured, abstracted, and normative way [see RQ2 in Section 4.2]. It sets aside parts of contextual integrity that account for how norms form through adaptive social processes. By focusing on regulatory compliance, it brackets the social source of privacy norms [RQ3 in Section 4.3]. If there is something lost in this usage of contextual integrity in computer science, it may be recovered through other uses and understandings of social context that have influenced technical research.

3.3.2 Android Permissions Remystified: A field Study on Contextual Integrity (Wijesekera, Baokar, Hosseini, Egelman, Wagner, and Beznosov)

The potential role that permissions in mobile platforms can play in providing users with control and transparency over how their information flows to apps and mobile platforms has recently attracted much research. For a long time, Android and iOS platforms asked users for permissions at install time. Recently they have extended the framework to also make it possible to prompt users for permissions during runtime.

Prior research has shown that few people read the Android install-time permissions and even fewer comprehend the complexity that the permissions are loaded with – for example, the permission to access Internet may be bundled with the permission to load ads. Prompting users too frequently causes “habituation”. Limiting prompts, however, raises questions about which ones to select in order to effectively protect user privacy. Wijesekera et al. [2015] leverage contextual integrity to approach the usable security and privacy problems that arise when interfacing a permission request model to the users (a user interface and technical platform problem as discussed in RQ1 in Section 4.1).

In their study, the authors examine how applications are currently accessing user data and assess whether or not it corresponds to users’ expectations. To do so, they instrumented Android phones to log whenever an application accesses a permission-protected resource and distribute these to 36 participants who use the phones for a week. For each permission request, they keep a log of the app name and permission, and further “contextual factors” which include whether the requesting application was visible to the user (running with or without user interaction, notifications, in the foreground or background); screen status (whether the screen was on or off); connectivity (the wifi connection state); location (the user’s last known coordinates); the UI elements that were exposed to the user during the request; history of interactions with applications; and, the path to the specific content that was requested.¹ After the week, study subjects participated in an exit survey where they were shown a sample of screenshots to inquire about their expectations relating to requested permissions. The authors use the outcomes of the study to start specifying a classifier to automatically determine whether or not to prompt the user based on contextual factors.

During the one week study, the 36 phones logged 27 million application requests to protected resources, translating to 100,000 requests per user/day. The authors found that 75.10% of the permissions were

¹For example, if Spotify requests a wi-fi scan while the user is playing Solitaire, then visibility is set to false, the history shows that prior to the Spotify prompt, the user had viewed Solitaire, the screen status was on etc.

requested by apps that are invisible to the user (most of these were requested when the screen was turned off, which is 95% of most phones lifetime).² Using the data they collected, they analyze which permissions are requested and the different ways in which certain information can be accessed (e.g., there are multiple ways to access location). They argue that due to invisibility, frequency and exposure, what the authors' have dubbed as users' contextual integrity – meaning what they expect from apps and their permission behavior – is violated.³

The result of the exit survey shows that users' decision to block a permission request was based on a variety of contextual factors. When asked why they would want to block certain permissions, 53% of survey subjects stated that they didn't think the permission was necessary for the functionality of the application. However, users do not categorically deny permission requests based solely on the type of resources being accessed by an app. They also take into consideration how much they trust the application and whether they are actively using it. Moreover, the status of the screen and whether the app is in the foreground has an impact on whether users are more likely to appreciate the permission type in their decision.

The authors use these insights to develop a classifier that can infer when the user is likely to deny a permission request and prompt for explicit run-time permissions. Their classifier makes use of originating application, permission and visibility for prompting users as well as personalization factors to meet users' contextual expectations. They complete the study of this classifier with a short evaluation of its accuracy.

In their reading of contextual integrity, the authors abstract away the social contexts of apps [RQ2]. They are not concerned with the information norms an app may be subject to due its social context, e.g., is it appropriate for a health app to collect user location? Rather, they equate privacy violations with occurrences of the collection of

²The applications making the most permission requests are Facebook, Google Location Reporting and Facebook Messenger.

³The user study provides greater insights as to when users feel that their expectations are not met which is worth reading but too detailed for the study at hand.

personal information in ways that defy user expectations in the context of an interaction based on their list of contextual factors. Starting from this definition, they go on to study those permissions and contextual factors that are most likely to defy users' expectations and that may be good candidates and situations for prompting users at run-time. If we were we to describe the study using Dourish's vocabulary, we would say that the authors study which of these factors users consider to be "contextual" to their interaction with their apps and mobile devices. This sets this study apart from typical context-aware computing papers that have a more static view of what counts as context. A more detailed discussion on how contexts are handled in the different papers can be found in Section 4.2.

While the authors lean on CI, they do not make explicit use of the parameters part of the conceptual framework nor invoke its normative aspects. Implicitly, we can interpret the model that underlies the study to treat users as senders, apps as recipients, type of data as a kind of contextual factor. Moreover, we can regard the permission prompts as implementing transmission principles that make select information flows conditional on user's approval. However, by evaluating appropriateness of information flows with respect to an app, rather than the social context that the app serves, the study also falls short of understanding user expectations with respect to information flows that may be initiated by the organization, be it sharing user data with other companies or users [RQ3]. In general, relying on users expectation as a normative source leaves out other potential sources of information norms which may have been very useful in further pruning those prompts that request permissions for inappropriate information flows. As a result, the authors clearly deviate from the normative ambitions of the framework and hold its conceptual premises only in our interpretation.

Foregrounding apps does reveal interesting results that go beyond what is typically in the scope of a CI analysis [RQ4]. First, the authors find that users wanted to accept some permissions because they were convenient and others they wanted to reject because they requested access to privacy sensitive information (e.g., SMS messages) regard-

less of the social context. Second, users were more likely to expect and accept requests that occurred in the foreground of an application than in the background, and they were more likely to want to block a permission if it was from an app or process in the background, too frequent or when the phone screen was locked. In other words, users consider additional factors when it comes to evaluating the appropriateness of an information flow. This result stands to inform CI by pointing out the need to acknowledge technical and operational contexts, which we discuss in Section 4.2.

3.3.3 Implicit Contextual Integrity in Online Social Networks (Criado and Such)

In this fascinating paper coming from the field of Multi-Agent Systems in Artificial Intelligence applied in the context of Online Social Networks (OSNs) the authors theorize and develop an agent that responds to the problem of “Implicit Contextual Integrity”.⁴ The main motivation for the authors is to introduce mechanisms that address issues related to “inappropriate exchanges and undesired disseminations” that happen due to lack of effective privacy controls in OSNs [RQ1]. Pointing to numerous studies in computer science, the authors argue that contextual integrity, a model upon which effective computational mechanisms can be built, is the right framework for developing effective controls for OSN users. However, prior computational models have assumed the existence of well-defined contexts with predefined roles and explicit norms. These are not always available in OSNs, as context, roles and associated informational norms are “implicit, ever changing and not a-priori known”.

In order to support users with these implicit norms, roles and contexts, the authors propose an Information Assistant Agent (IA-agent) that can infer the context that the user is in and the information norms belonging to that context. In describing their solution, they first present an information model for Implicit Contextual Integrity and

⁴The authors have written two papers with the same title. Here we refer to the longer and more detailed version published in the Information Sciences Journal 325 (2015) 48–69.

then characterize the IA-agent. The agent uses the information model and further modules to learn implicit contexts, relationships and the information sharing norms. It uses this information to warn users before they make potentially inappropriate information exchanges (within a context) or engage in undesirable dissemination of information previously exchanged (across contexts).

Criado and Such leverage a plethora of techniques available to them to compose a formalization of appropriateness that can be used by the IA-Agent. First, they assume that each information exchange can be mapped to a range of finite topics, e.g., that a post about a tennis match is about sports. The frequency with which certain topics gets mentioned by members of a context is crucial – messages pertaining to topics that are rarely mentioned are considered inappropriate and vice versa. Some exception is made, however, to infrequently communicated topics: if reciprocity underlies a given communication between members of the context, then the information flow is reconsidered as appropriate. Furthermore, if information on a topic has been previously disclosed in a given context, then a repeat disclosure in that context is not seen as inappropriate, and hence is not regarded as entailing new privacy risks. Appropriateness of a topic may increase if members of a context start exchanging messages on the subject. It may also decay, as information flows pertaining to the topic decrease or disappear.

A message may flow to people in multiple contexts, in which case it is assumed to be flowing to the context with most recipients. For example: if Mary is Alice’s friend and workmate, and Alice sends a message to Mary and three other people from her work context, then it is assumed to be a message flowing in the work context. The agent also takes into consideration whether the information in a message is known in the different contexts shared by the recipients of that message. Since the IA-agent needs to keep track of frequency, past mentions and reciprocity, the corresponding design requires keeping track of all past communications.

In summary, the total appropriateness of a given information flow is based on three different metrics: appropriateness of topic to individuals, appropriateness of an information flow in a context, and appropriateness of a message across contexts. Four modules allow the IA-Agent to

complete its tasks:

1. Community finding algorithm: identifies new contexts made up of densely connected members.
2. Passing time function: updates appropriateness of information flows over time also based on the knowledge about different topics in a context.
3. Message sending function: uses received messages to update the different appropriateness and knowledge functions.
4. Message reception function: processes messages before they are sent to either flag them to the user as inappropriate, avoid undesirable dissemination of previously exchanged information, and update appropriateness and knowledge functions.

The authors conclude the paper with experiments based on simulations of exchanges among multiple IA-agents. The results show that the agents are able to infer information sharing norms even if a small proportion of the users follow the norms; agents help reduce the exchange of inappropriate information and the dissemination of sensitive information; and, they minimize the burden on the users by avoiding raising unnecessary alerts.

This paper mostly remains faithful both to the definition of CI as well as its parameters [RQ2]. The model includes sender, receiver, messages, topics and context. The authors make no explicit comments about the transmission principle, however, one could argue that the agent implements transmission principles: information may flow as long as it passes the contextual norms of a context, or norms of dissemination across contexts. Otherwise, the user is presented with an alert which gives her an opportunity to double-check on the appropriateness of an information flow.

The authors assume that contexts emerge in interaction, an approach very much aligned with Dourish [2004]. Contexts are not pre-defined, but as communities of users establish connections and communications they are detected by the “community finding module”. Hence,

users' communication patterns, networking patterns as well as the IA-agent become sources of normativity [RQ3]. This implies a division of labor between the user and the agent: the agent plays an active role in maintaining informational norms and the user is still able to practice discretion when it comes to determining what is considered an appropriate information flow in a context.

In addition, the authors introduce a number of distinctions and parameters that go beyond those of CI [RQ4]. The distinction between inappropriate exchange and undesirable dissemination allows the authors to express norms with respect to information flows within and *across multiple contexts*. This is a pragmatic extension of CI to also cover information that cross contexts.

The different functions for frequency, reciprocity and prior knowledge give the authors the tools to explore adaptivity of informational norms throughout time and in multiple contexts. This allows the authors to capture norm development and also make explicit the role that users play in maintaining norms. In many ways, the "implicit CI model" the authors introduce is complementary to CI, in that it provides means to extend social norms in a context with changes to those norms through interactions over time. Adaptivity, multiple-contexts, temporality and user engagement in contextual norms are further discussed in Section 4.4.

The model underlying the IA-Agent also exhibits some differences in interpretation of aspects of CI. The agent relies on frequency of exchanges on topics as a means to infer norms. Norms are not the same as the most frequent information flows, nor would such a definition do justice to topics that are pertinent but infrequently exchanged.

Finally, the proposed IA-agent helps maintain contextual integrity but is outside of the scope of CI analysis. The appropriateness of information flows to the OSN provider, the provider of the IA-Agents as well as other third parties is not discussed. It is as if CI only applies to social relations but the service providers are outside of the scope of CI. This leaves out questions like whether an IA-Agent should compile and keep all past communications of all members of a social network, and if so, who can have access to the Agent's memory? This aligns

with industrial practices where OSN companies claim that they are only facilitating information flows deemed appropriate by their users. It is possible to argue that what norms should apply to an IA-Agent is too much to ask of a single CS paper. However, this type of scoping is not exceptional among the papers we found and worthy of a lengthier discussion which we come back to in Section 5.

4

Results

Through our study of 20 computer science papers invoking contextual integrity we discovered a variety of themes and innovations in privacy engineering that also reflect on improvements to the privacy framework. After parsing each paper into our review template (see Section 3), we coded our results and surfaced a number of recurring themes. We then consolidated these themes into answers to our research questions. We discuss those answers in this section.

The papers included in the survey were: Barth et al. [2006, 2007], Criado and Such [2015], Datta et al. [2011], Jia et al. [2017], Kayes and Iamnitchi [2013a,b], Krupa and Vercoouter [2012], Netter et al. [2011], Omoronyia et al. [2012, 2013], Salehie et al. [2012], Samavi and Consens [2012], Sayaf et al. [2014], Shih and Zhang [2010], Shih et al. [2015], Shvartzshnaider et al. [2016], Tierney and Subramanian [2014], Wijesekera et al. [2015], Zhang et al. [2013]. The three categories we derived to answer RQ1, RQ3, and RQ4, and the themes within each category as they apply to each paper, are in table 4.2 at the end of Section 4.2. A separate table (Table 4.1) is provided for our results for RQ2 in Section 4.2.2. In Sections 4.2.1–4.2.4, we detail each set of results and its relation to our research questions. Blank fields in the tables stand

Table 4.1

Paper	Substantiality	Domain	Stability	Valence	Epistemology
Barth et al. [2006]	ABSTRACT	SOCIAL	REPRESENTATIONAL	NORMATIVE	MODEL
Barth et al. [2007]	ABSTRACT	BOTH	REPRESENTATIONAL	NORMATIVE	MODEL
Criado and Such [2015]	CONCRETE	SOCIAL	INTERACTIONAL	NORMATIVE	EMPIRICAL
Datta et al. [2011]	ABSTRACT	SOCIAL	REPRESENTATIONAL	NORMATIVE	MODEL
Jia et al. [2017]	CONCRETE	TECHNICAL	REPRESENTATIONAL	DESCRIPTIVE	EMPIRICAL
Kayes and Iamnitchi [2013a]	CONCRETE	SOCIAL	REPRESENTATIONAL	DESCRIPTIVE	EMPIRICAL
Kayes and Iamnitchi [2013b]	CONCRETE	SOCIAL	REPRESENTATIONAL	DESCRIPTIVE	MODEL
Krupa and Vercouter [2012]	ABSTRACT	SOCIAL	REPRESENTATIONAL	NORMATIVE	MODEL
Netter et al. [2011]	CONCRETE	SOCIAL	REPRESENTATIONAL	DESCRIPTIVE	MODEL
Omoronyia et al. [2012]	ABSTRACT	BOTH	REPRESENTATIONAL	NORMATIVE	EMPIRICAL
Omoronyia et al. [2013]	ABSTRACT	BOTH	REPRESENTATIONAL	DESCRIPTIVE	MODEL
Salehie et al. [2012]	CONCRETE	BOTH	REPRESENTATIONAL	DESCRIPTIVE	MODEL
Samavi and Consens [2012]	ABSTRACT	BOTH	INTERACTIONAL	DESCRIPTIVE	EMPIRICAL
Sayaf et al. [2014]	CONCRETE	SOCIAL	INTERACTIONAL	NORMATIVE	MODEL
Shih and Zhang [2010]	BOTH	BOTH	REPRESENTATIONAL	DESCRIPTIVE	EMPIRICAL
Shih et al. [2015]	CONCRETE	TECHNICAL	REPRESENTATIONAL	DESCRIPTIVE	EMPIRICAL
Shvartzshnaider et al. [2016]	ABSTRACT	SOCIAL	INTERACTIONAL	DESCRIPTIVE	EMPIRICAL
Tierney and Subramanian [2014]	BOTH	SOCIAL	REPRESENTATIONAL	NORMATIVE	BOTH
Wijesekera et al. [2015]	CONCRETE	TECHNICAL	REPRESENTATIONAL	NORMATIVE	MODEL
Zhang et al. [2013]	CONCRETE	TECHNICAL	INTERACTIONAL	DESCRIPTIVE	EMPIRICAL

for cases where none of the themes in our taxonomy were applicable to respective papers in question.

4.1 (RQ1) Architecture

Our first research question was “What kind of problems and solutions do computer scientists use CI for?” CI is a philosophical theory of privacy with social and legal implications that are designed to apply across a wide range of technologies. Computer scientists do not have the luxury of this insensitivity to technical detail. Their work reveals how specific classes of technical architecture have different socially meaningful implications for privacy.

There was variation in the kind of system described in each paper. Far from being neutral with respect to the way CI was used by the papers, focus on different technical architectures resulted in different manifestations of the privacy theory. Some themes within this category were **user interfaces and experience** (4.1.1), **infrastructure** (4.1.2), and **decentralized** architectures (4.1.3).

4.1.1 User Interfaces and Experiences

Four papers surveyed Shih and Zhang [2010], Shih et al. [2015], Zhang et al. [2013], Wijesekera et al. [2015] studied the experience of users with applications or with designing user facing interfaces or applications. Since contextual integrity theory operates at the level of social norms and says little about user interfaces, and user experience in different situations, these papers raise the question of how user-facing apps and their interfaces are related to broader normative questions about what is appropriate information flow in differentiated socio-technical situations. One paper [Zhang et al., 2013] explicitly drew its motivation from the FTC’s view of the importance of “context of interaction” rather than a broader social or normative view of privacy. Nevertheless, these cited contextual integrity as part of its motivation and study set up. This prompts contextual integrity theorists to address the theoretical connection between ‘context of interaction’ and social spheres.

In general, these papers were not concerned with modeling social norms of a large population of users. Rather, they were more concerned

with individual user’s activity and their interaction with a device in different situations. Situations could be environmental conditions (e.g., where the user is located, night or day); social situations (e.g., work, home, among friends); or, technical situations (e.g., whether an app is in use when it asks for permissions), comparable to conceptions of context a la Dey et al. [2001].

Two themes were common to these papers: they implicitly highlighted that in addition to what may be appropriate flow of information in different *social spheres*, users may have further criteria for what information flows they expect or prefer in different *situations* (See Section 4.2 for further discussion on this topic). Second, these papers aspired to generalize their results in order to provide recommendations concerning infrastructural design that can be implemented to respect contextual integrity. For example, Wijesekera et al. [2015] propose introducing techniques to improve Android permission models to better cater to user preferences and expectations in different interactional situations.

4.1.2 Infrastructure

Many of the papers in our sample were about formal models for or techniques specific to systems that serve as *infrastructure*. By infrastructure, we mean technology that is designed to cater to a large set of users and diversity of applications. We distinguish between social technical platforms since they raise different kinds of challenges to applying CI in practice.

Social platform: A social platform is a technology that mediates social interaction as an affordance or service. In the papers we surveyed, it was used synonymously with online social networks (OSN) and social ecosystems [Kayes and Iamnitchi, 2013b], examples varying from Facebook, Snapchat, SMS, to Amazon and Google Play store reviews, and email.

From the perspective of contextual integrity and privacy, what is most pressing about social platforms is how they can potentially mediate activity pertaining to multiple social spheres. Friends, family, classmates, work associates, and so on may all interact using the same social platform.

This poses a challenge to contextual integrity because while the framework is well tailored to evaluating the ethics and impact of a particular technology by identifying the (singular) context it is in, social platforms are designed to mediate more than one social context and perhaps to create entirely new social spheres through their use.

In order to accommodate the uses of the technology in multiple spheres simultaneously, computer scientists are challenged with modeling not just the norms within a single social sphere, but contexts in general and how they interact. Contexts may be very fluid in social platforms. Papers we reviewed looked at scenarios where contexts may collapse; multiple contexts may produce conflicting norms [Tierney and Subramanian, 2014]; contexts and social norms may change over time [Netter et al., 2011, Sayaf et al., 2014]; and, as in the case of [Criado and Such, 2015], how contexts may emerge in interaction. Contextual integrity scholars have not yet provided much guidance on how to deal with the fluidity of social contexts and its impact on how to interpret informational norms, leaving computer scientists to come up with creating solutions themselves.

Note that the definition of a social platform is agnostic about the particular implementation or location of the technology that undergirds a social service. The technology may be distributed, federated or centralized; include apps on a smartphone; web pages in a browser; servers hosted in a “cloud”, and telecommunications infrastructure supporting information flow via Internet protocols. This technical complexity is addressed in what we call *technical platforms*.

Technical platform: A technical platform is a technology that mediates the interactions between other heterogeneous technologies connecting multiple users. Examples include Android smartphones [Wijesekera et al., 2015], Platforms for Smart Things [Jia et al., 2017], the Web (and web browsers) and Smart Grids [Salehie et al., 2012].

A difficulty in defining ‘technical platform’ is that the technology in question is often designed as a ‘stack’ with multiple layers, each layer being a ‘platform’ on which the next one operates. Hence there are many technologies, such as Facebook, that are both in a sense “applications” that stand on technical platforms and are also technical platforms in their own right as they mediate other applications through

a developer API. This is in part due to a design principle that has influence on the Internet [Clark et al., 2002] that recommends having as few controls as possible introduced on each layer to allow for a wider range of possibilities at the higher layers.

From the perspective of contextual integrity, the challenge with analyzing technical platforms is that they necessarily involve the participation of many social actors who may access and process data (and especially personal information) flowing through them. In contemporary applications the actors involved with operating the technical platform are subject to a number of technical, legal and social norms, some of which are substantial to the social contexts their users see themselves as operating in. We tentatively propose the concept of “operator context” that defines the roles and norms of the operator of communications infrastructure that acts between users.

Formal models: By formal models we refer to papers that conceptualize frameworks that can be part of an infrastructure that serves many different social contexts or technologies, but the implementation details of which are either irrelevant or considered only at an abstract level. Such papers come with verification of the consistency and completeness of the formal model as well as a prototype to show the feasibility of actually implementing the system. These papers provide useful insight into how CI can be operationalized, raising issues at the logical level that are difficult to surface in more empirical work.

Examples include papers on access control models that preserve contextual integrity in an enterprise, like Barth et al. [2006] and Barth et al. [2007]; frameworks that describe and evaluate ontologies to audit privacy relevant processes in a linked data environment [Samavi and Consens, 2012]; or adaptive systems that monitor when new threats arise, reconfiguring information flows to continue matching user privacy requirements [Omoronyia et al., 2013], or that identify when information norms themselves change [Shvartzshnaider et al., 2016].

4.1.3 Decentralization

The rare paper in our sample dealt with a specifically challenging technical feature: decentralized architectures. We highlight this theme,

however rare, because of the way it positions technology relative to social spheres and interactions. While user interfaces and experiences are connected to individual users (and their expectations), social platforms are central and common to a large number of diverse social contexts. Decentralized architectures have an interactivity and topological complexity that mirrors that of society itself, and trust and reputation mechanisms come to play a greater role in the absence of a centralized entity that can arbitrate information norms. We look forward to more papers on this theme and CI.

4.2 (RQ2) Capturing Context

Our second research question was “How have the authors dealt with the conceptual aspects of CI?” Contextual integrity theory has a specific understanding of context as social sphere, parameterized by roles, norms, purposes, and values. The norms are parameterized by their actors (senders, receivers, and subjects in contextually defined roles), information topics, and transmission principles. We wanted to know whether and how computer science papers used this conceptualization of privacy. We found that while several papers drew closely from the concepts in CI, others represented context very differently. As we have discussed, many computer scientists interpreted ‘context’ in a way that draws from the research field of ubiquitous computing (See Section 2.2.2). Because of these discrepancies, we have chosen to focus on the nuanced differences in how context is represented rather than on which of the parameters are used.

We have coded the way each paper has defined and used context across five binary dimensions, which we have named: substantiality, domain, stability, valence, and epistemology. Within each dimension there are two opposed poles.

1. **Substantiality.** Some papers discuss contexts as an **abstract** type or ideal of a situation. Others discussed contexts as **concrete** happenings and situations. *Example: hospitals in general are an abstract context. Mount Sinai Beth Israel hospital in Manhattan is a concrete context.*

2. **Domain.** Some papers discuss **social** contexts, defined by configurations of people to each other. Others discuss **technical** contexts, defined by objective properties of mechanical devices and the environment they were in. Some papers understood contexts as combining **both** social and technical factors. *Example: a classroom with a teacher and students is a social context. A language education mobile app that prompts the user with questions and sends results back to a server for analysis is a technical context.*
3. **Stability.** We draw on Dourish [2004] for this distinction. Some papers treat context as a **representational** problem, as if they were stable, delineable, and distinct from the activity that contained them. Others treat them more as an **interactional** problem, as arising from interactions between people and things, defined by specific circumstances. *Example: The Oval Office in the White House is a stable context. A flash mob is an interactional context.*
4. **Valence.** Some papers see the **normative** aspects of privacy as being inherent in context. Others treat contexts merely **descriptively**, without normative force. *Example: A conference Code of Conduct is an account of norms inherent in a context. A list of attendees, keynote speakers, and program committee members is a description of the context.*
5. **Epistemology.** Some papers adopt a **model-building** approach to defining contexts. They posit a schema or model of context and derived conclusions from it. Other papers take a more **empirical** approach, deriving context definitions from data. A parameterized definition of a context, e.g., context is location, time, and activity, is an example of a model based approach, whereas applying traffic and topic analysis to communications in order to surface contexts is an example of an empirical approach that can be used to characterize different contexts.

We note that as far as CI is concerned, it is essential that contexts be understood as **normative**, as one important trait of contexts is that

they have ends, purposes, and values. They are **social** contexts, pertaining to relationships between people in defined roles, but they are oriented around functions, purposes, aims, goals, activities, values, etc. As these social norms evolve in society in general and then are applied to particular cases of information flow, contextual integrity conceptualizes contexts **abstractly**. “Context” interpreted to mean *sphere*, as discussed above, has these three properties (i.e. they are normative, social, and abstract). To the extent the papers draw on different meanings of context, they diverge from CI. For example, when the literature interprets context as *situations*, as discussed in Section 2.2.2, it conceptualizes contexts as **concrete** and at least partly **technical**. Our study has surfaced that computer scientists, in trying to make CI actionable, have encountered the problem of applying abstract social norms to concrete socio-technical situations.

The first paper of our study in publication order is [Barth et al., 2006], which we have detailed in Section 3.3.1 of this paper. Helen Nissenbaum is a coauthor and the paper includes a summary of contextual integrity theory. The technical contribution of the paper focuses on a “fragment” of contextual integrity. It is this technical contribution that we have assessed according to the criteria above. The [Barth et al., 2006] technical presentation of context is as one that is **abstract, social, representational, normative, and modeled**. Their work models the specific normative logic of contextual integrity. It shows how norms and laws can be represented as abstract policies amenable to automated enforcement.

This paper is one end of a spectrum. Other papers in our sample drew their understandings of context from other traditions, including ubiquitous computing (discussed in Section 2.2.2 of this paper). Following Dourish [2004], some papers eschewed explicit abstract representational modeling of context for what resembles interactional views of context derived from empirical data about user behavior or human-computer interaction. Several papers considered the narrow context of a user and their device, as opposed to social relations more generally. Most papers did not see norms as inherent to the contexts they studied, but rather saw contexts descriptively. (Some of these papers sourced their normativity from other factors, see Section 4.3). Our paper

exemplars (3.3.1–3.3.3) provide deeper explanations of the dimensions used to classify contexts here.

What we have discovered in answer to RQ2 is the distribution of papers across these dimensions. This tells us how well contextual integrity as a conceptual theory of privacy has made it into computer science. CI conceptualizes contexts as **normative** and **social**. Papers that have modeled context as either purely technical or purely descriptive have missed some of the core intent of CI.

To the extent that it sees the formation and maintenance of a social context as an adaptive social process, we argue that contextual integrity is consistent with the **interactional** view of contexts from Dourish [2004], though in its concrete application it has a tendency to work from a **representation** of context. We believe this leads to deep sociological questions about how social norms and purposes, which can seem **abstract** and theoretical, can form from **concrete** human interactions.

We note with special interest Criado and Such [2015], detailed in 4.1.3, which stands out as a paper that addresses a particularly difficult challenge. It is the only paper in our sample that manages to be both concrete, interactional, and empirical as well as socially normative. We see this as an important innovation in the use of contextual integrity in computer science.

4.3 (RQ3) Source of Normativity

Our third research question was, “How have the authors dealt with the normative aspects of CI?” In contextual integrity, the normative (in the sense of prescriptive or ethical) force of information norms comes from the purposes, ends, and values associated with each social sphere. This complex metaethical theory rarely finds its full expression in the computer science literature. Instead, the papers in our sample take a variety of narrower positions, implicitly or explicitly, on the source of normative values that motivate the importance of privacy.

The subsections here explain the themes we found in this category. **Compliance and policy** refers to when normativity was taken from

legal authority or some other unquestioned source of policy. **Threats** refers to the computer security research practice of positing a threat model to motivate research. **User preferences and expectations** locates the source of normativity in the subjective perspective of individual users. **Engagement** refers to designs that allow users to dynamically engage with each other to determine norms.

4.3.1 Compliance and Policy

Some of the papers in our sample took their motivation from the practical problem of compliance with legal regulation, such as HIPAA. These papers effectively outsource their normative questions to the legal system. They at times argue as if compliance is relevant because it is internalized as a business interest [Barth et al., 2007]. One line of this compliance-based research is contiguous with other work on formalizing privacy regulations in ways that are less closely tied to contextual integrity [DeYoung et al., 2010]. Datta et al. [2011] synthesize the contributions of this research trajectory.

Other papers are less specific about source of the specific form of their restrictions, but nevertheless have an explicit mechanism for stating *policy*. Some computer research in this field culminates in the articulation of a policy language, which is valid for its expressivity, not for the specific character of the content of any particular expression it allows.

In both the cases of compliance and policy, normativity is exogenous to the technical design.

4.3.2 Threats

Some of the papers motivated their research goals in terms of privacy *threats*. These presumably adopted this stance as a continuation of practices from security research, which typically posits a threat model of potential attacks and adversarial capabilities before detailing how a technical improvement can mitigate these threats.

Taking this position alleviates the researcher from having an overarching theory of privacy; they can instead work from specific cases that are plausible or undisputed cases of privacy violation.

4.3.3 User Preferences and Expectations

Some papers motivated their research either explicitly or implicitly in terms of whether a technical design was able to meet user preferences or expectations of privacy. Preferences and expectations are not the same thing, but they are related in that they depend primarily on the individual subjectivity of the user. A user's expectation is the outcome they desire or is in their acknowledged interest, and a number of papers explore the expectations users have in different social or interactional context. User preferences on the other hand were often used to study what kind of controls users may prefer to have or exercise when using systems. User perceptions also played a role in papers where researchers explored what information flows users noticed or how they perceived them [Wijesekera et al., 2015, Zhang et al., 2013].

Measuring user expectations and preferences as a way of assessing the appropriateness of information flow is consistent with contextual integrity. This can be done explicitly through survey methods, as is done by Shvartzshnaider et al. [2016]. In CI, appropriateness is a function of social norms, and these norms do codify social expectations and values. Certainly in some cases user expectations will track social expectations. But though they are related, we caution researchers against conflating social norms with user expectations and preferences. This is because individual users are more prone to becoming unreflectively habituated to a new technology than society as a whole. Also, individual user preferences may at times be opposed to the interests of society. We have identified elaborating on the relationship between individual preferences and social norms as a way to improve CI.

4.3.4 Engagement

Some papers explicitly articulated mechanisms through which users could engage with a system to define what's normative for the system. Rather than accept a policy or threat model exogenously or see an individual's opinions and satisfaction as the ends of design, these papers allowed for the establishment of norms to be a dynamic social process accomplished through use of the technology itself. For a more in depth discussion of how this can work, see the more detailed discussion of

Criado and Such [2015] in Section 3.3.3. Another example is Tierney and Subramanian [2014] who describe a marketplace or library of abstract context definitions, complete with roles and access controls corresponding to transmission principles, that are developed by a community of context designers. Users can then instantiate the context template that best fits their social needs.

4.4 (RQ4) Expanding Contextual Integrity

Our fourth research question was “Do the researchers expand on contextual integrity?” The rigors of computer science led many paper authors to innovate and improvise as they used contextual integrity in their designs. We grouped these innovations into the category **Expanding Contextual Integrity**. We found many papers were engaged in developing mechanisms for technological **adaptation** to changing social conditions (4.4.1). Some addressed the challenges associated with technologies that operated within **multiple contexts** at once (4.4.2). Some developed ideas concerning the **temporality and duration** of information and how this affects privacy (4.4.3). Others were particularly concerned with **user decision making** (4.4.4) with respect to privacy and information controls. While all these innovations are compatible with contextual integrity as outlined in [Nissenbaum, 2009], we found the detail with which the paper authors engaged these topics showed ways to expand contextual integrity.

We note that many of these themes echo discoveries made with respect to our other three research questions. For example, those papers that addressed the design of social infrastructure (see Section 4.1) had to address the problem of how to handle multiple contexts in the same technology, and as they did so they had to make decisions about how to represent context that did not necessarily accord with CI’s concept of context as social sphere (see section 4.2). Of the four research questions, this one reflects the technical accomplishments discussed in sections 4.1–4.3 back on CI in order to identify the limits of the framework itself. Table 4.2 shows how themes from different research questions were distributed across the papers in the survey.

Table 4.2: Contextual Integrity Through the Lens of Computer Science

Paper	RQ1. Architecture	RQ3. Source of Normativity	RQ4. Expanding Contextual Integrity
[Barth et al., 2006]	FORMAL MODEL	COMPLIANCE	TEMPORALITY
[Barth et al., 2007]	FORMAL MODEL	COMPLIANCE	TEMPORALITY
[Criado and Such, 2015]	INFRASTRUCTURE: SOCIAL	ENGAGEMENT	ADAPTATION, MULTIPLE CONTEXT
[Datta et al., 2011]	FORMAL MODEL	COMPLIANCE	TEMPORALITY
[Jia et al., 2017]	INFRASTRUCTURE: TECHNICAL	THREATS	
[Kayes and Iammitchi, 2013a]	INFRASTRUCTURE: SOCIAL	ENGAGEMENT, POLICY, USER PREFERENCES	MULTIPLE CONTEXTS, ADAPTATION, DURATION
[Kayes and Iammitchi, 2013b]	INFRASTRUCTURE: SOCIAL	ENGAGEMENT, USER PREFERENCES	ADAPTATION
[Krupa and Vercouter, 2012]	DECENTRALIZED	THREATS, USER PREFERENCES	
[Netter et al., 2011]	INFRASTRUCTURE: SOCIAL	THREATS, USER PREFERENCES	ADAPTATION
[Omoronyia et al., 2012]	INFRASTRUCTURE: SOCIAL	THREATS, USER PREFERENCES	ADAPTATION
[Omoronyia et al., 2013]	INFRASTRUCTURE: FORMAL MODEL	THREATS, USER PREFERENCES	ADAPTATION
[Salehie et al., 2012]	INFRASTRUCTURE: TECHNICAL	THREATS	ADAPTATION
[Samavi and Consens, 2012]	INFRASTRUCTURE: TECHNICAL	COMPLIANCE, USER PREFERENCES	TEMPORALITY
[Sayaf et al., 2014]	INFRASTRUCTURE: SOCIAL	USER PREFERENCES	CONTEXT CLASH, TEMPORALITY
[Shih and Zhang, 2010]	USER INTERFACE	USER PREFERENCES	MULTIPLE CONTEXTS
[Shih et al., 2015]	USER INTERFACE	USER PREFERENCES	USER DECISION MAKING
[Shvartzshnaider et al., 2016]	FORMAL MODEL	USER EXPECTATIONS	ADAPTATION
[Tierney and Subramanian, 2014]	INFRASTRUCTURE: SOCIAL	ENGAGEMENT	
[Wijesekera et al., 2015]	USER INTERFACE	USER EXPECTATIONS, USER PREFERENCES	APPLICATIONS AS ACTORS
[Zhang et al., 2013]	USER INTERFACE	THREATS	MULTIPLE CONTEXTS

4.4.1 Adaptation

The most common way in which computer science papers expanded on contextual integrity was to address questions of social adaptation.

As noted in Section 2.1 above, CI theorizes that norms are the result of a process of social adaptation. Social spheres have ends, purposes, and values robustly as a function of their evolution. Norms within these spheres are legitimate to the extent that they serve their contextual purposes, but environment changes (such as the prevalence of new digital technologies) are the stimulus for further adaptation of norms. To the extent that CI has a conservative impulse, it is to warn against the agitation caused by disruptive technologies that change the environment too quickly for social evolution to adapt.

This grand theory of privacy is not actionable for computer scientists. In the papers we found that dealt with *adaptation*, the researchers were interested in designing technology that is responsive to social change at a much smaller scale in both space and time. Criado and Such [2015] discuss the adaptation of an informal sports discussion group emerging out of a collegial working forum. If large-scale evolution of social spheres and privacy norms depends on variation on the level of social interaction, it is challenging to design technology that keeps up with this variation. If large scale agitation about threats to privacy happens when technology disrupts a shared social expectation, then small scale agitation can occur when technology fails to address emerging norms. For computer scientists to deal with these challenges, they have to be more specific about these processes of adaptation than CI currently is.

Many of the papers we reviewed concerned themselves with the problem of maintaining contextual privacy under conditions of social change. Few adopted the theory proposed by Nissenbaum [2009]; instead these papers proposed their own mechanisms to account for and capture changes in context and norms. Most often these did not take into account the stability of contextual ends, purposes, and values. Rather, they generally took on the problem of having technology react appropriately to exogenous social change. Criado and Such [2015] design agents that guess rules for appropriate information flow from

regularities in user behavior. Shvartzshnaider et al. [2016] experiment with a method for empirically surveying for opinions on social norms and translating results into a logical specification. Such mechanisms could be used to build a system that is robust to changes in social opinion.

Papers that addressed social adaptation were likely to also use concrete, interactional and empirical concepts of context (see Section 4.2). Some designed methods to have users engage in the process of determining norms (see 4.4). In general, technical systems that are adaptive to changes in social behavior can be prone to the failure of maladaptation. To be actionable for these designs, CI would benefit from more specificity regarding the process of social evolution that legitimates the norms of social spheres.

4.4.2 Multiple Contexts, and Context Clash

Another common way in which computer science papers expanded on contextual integrity is that many discussed technologies that recognized the existence of multiple contexts at once. This was common for those papers that addressed the design of social infrastructure (see Section 4.1), for example. Contextual integrity as a privacy framework posits many different social spheres with different norms of information flow. But as it is currently resourced, CI provides little conceptual clarity as to how different contexts relate to each other, and no normative clarity as to how this multiplicity of contexts affects the appropriateness of information flow.

As a result, many of the paper in our study improvised solutions to the problems associated with representing multiple social contexts. In some, system users were registered or detected as being in one or another context, with shifting access control policies in a context-appropriate way, something the agent in the Criado and Such [2015] paper is tasked with reasoning about. Some papers accommodated the relationship between contexts through a mechanism of context adaptation (see above). Others addressed the specific problem of what happens when information *flows between contexts*. For example, Sayaf et al. [2014] raised the privacy concern that a photograph might move from a

context where it was interpreted as a swimsuit advertisement into one where it was sexually objectified.

All the papers that dealt explicitly with the problem of using CI when multiple contexts affected a situation used a concrete and empirical concept of context (see Section 4.2). This points to an insight about CI that we see as a research finding: a more actionable CI would address how situations (concrete context) can be empirically analyzed to determine which sphere or spheres (abstract, normative, social contexts) apply. For example, could a system that monitors communication within a university in general classify a particular message as belonging to a classroom, employment, or social sphere? It may be possible to formulate this as a machine learning problem.

4.4.3 Temporality and Duration (Read/Write)

Several of the papers in our sample extended contextual integrity by explicitly addressing restrictions or allowances on information flow based on the timing of flows. For example, a flow might be allowed after the sender has received permission, but not before, or until certain actions are completed in the future. These extensions are not a challenge to contextual integrity as a theory; they are fully within the scope of what is possible as a *transmission principle*. However, the specific elaborations of the relationship between timing and information flow policies were notable.

A related theme which does more conceptual work within contextual integrity is that of data's *duration*. In technical terms, this was expressed in our sample as restrictions of reading, writing, and deleting data, as found in [Kayes and Iamnitchi, 2013a]. These operations stretch the idea of information "flow" so much that they perhaps require an entirely different notion, that of information "stock".

Another line of research discusses the relationship between temporality and the possibility of privacy policy enforcement. Datta et al. [2011] note that some aspects of privacy policies cannot be completely enforced at the time when information flows because the policy mandates what happens *after* the flow. For example, some policies impose restrictions on how information is used.

4.4.4 User decision making

Contextual integrity as a theory of privacy abstracts away from individuals in order to provide a normative framework that is independent of specific actors and their interests. It is this stability that gives it much of its normative power. Nevertheless, many computer science papers that used contextual integrity were concerned with user's individual decision making.

While voluntariness is one factor that can affect the transmission principles of information norms, contextual integrity has little to say about the role of the individual in shaping norms and social contexts more generally. These computer science papers put emphasis back on the individual and her decisions in context.

5

Findings and Discussion

We have summarized the achievements of computer scientists in developing contextual integrity. The research we have reviewed has variously used parts of contextual integrity and innovated on the relationship between privacy and context. Through our analysis, we have identified new research questions and opportunities at the intersection of CI and computer science.

In the time since contextual integrity first emerged, it has attracted useful insights from legal and ethical theorists as well as social scientists. Some of the toughest challenges have come from those seeking to apply CI to problems in their home fields, whether law and public policy or computer science, design, and engineering – the focus of this paper. Like most efforts to apply theory and other abstractions to concrete or real-world challenges, these, too, require that a distance be traveled to leverage the theoretical constructs of contextual integrity, to concrete privacy challenges of computer science, design, and engineering. In traveling this distance, the efforts we have surveyed reveal unanswered questions, conceptual gaps, and realities that do not align fully with the CI model. These findings call attention to several specific ways to expand, explain, and adjust CI in order to make it more

responsive to the needs of computer science and engineering researchers seeking to inform their work with a meaningful account of privacy.

In this final section, we present *theoretical gaps in CI* that our literature survey has exposed, systematically organizing them into four subsections, each associated with our four research questions: RQ1 - Architecture; RQ2 - Character of Contexts; RQ3 - Privacy as a Moral Imperative; and RQ4 - Expanding CI. In each subsection, we describe the nature of the theoretical gaps, i.e between theory and practical application, followed by a discussion of lessons learned that could translate into guidance for those embarking on new technical privacy research and design projects. The task is challenging because although the parameterized informational norms of contextual integrity offer greater specificity than other normative privacy theories, there remains significant room for interpretation. This room for interpretation, on the one hand, is what distinguishes contextual integrity from accounts of privacy that are not adaptive in the face of historical, cultural, social, and even personal variations, but it can be frustrating for those looking for precise, literal rules that are both correct and directly implementable.

For each research question we also have sections that we have titled, “call to action,” in which we discuss the lessons learned from past applications that can positively inform further forays into using CI in privacy research. We encourage the creative spirit we observed in our survey and recommend lessons learned and open questions to inspire future researchers in the field of context integrity.

5.1 Architecture: Towards a Modular Contextual Integrity

Corresponding to our **RQ1**, we have discovered that the way contextual integrity is used in technical design depends on the architectural properties of the technology being designed. This presents an opportunity for faceting CI into more specialized programs that are targeted at specific classes of technical problems. At the same time, our study revealed that the technical designs of computer science researchers often bracketed the social roles of those operating technical and social platforms, despite these being central to public discussion and scholarship on privacy and technology.

5.1.1 Theoretical gaps

We see the demand for “modular contextual integrity”, faceting CI and giving guidelines for design and research at specific levels of the technical stack, be it in designing user interfaces or experiences, technical or social platforms, or devising formal conceptualizations. Providing these guidelines may require that we derive frameworks of heuristics and principles from CI’s conceptual and normative facets. We expect to do this in tandem with further elaboration of the fundamental concept of “context” (see Section 5.2) and the concept of Transmission Principle, both distinctive of CI and often not well understood, despite their importance for CI’s power as a normative, or ethical theory.

An example of a promising strategy to address this problem is to identify and describe social spheres specific to the design, provision, operation, and use of technology. This is especially relevant in those papers where the designers explicitly delegated responsibilities for enabling contextual integrity to technical elements. In the case of Criado and Such [2015], the agent co-regulates norms. In Wijesekera et al. [2015] apps actively take part in asking for information flows and the authors consider a classifier that would reason as to when information flows may breach contextual integrity. In Samavi and Consens [2012], the authors produce an auditing mechanism that checks logs for potential breaches to contextual integrity. These mechanisms are very different with respect to the degree of autonomy they provide to technical agents and those who are operating them. However they all invoke the question: to what extent these mechanisms are subject to the norms of the context they are co-regulating, acting in or auditing? Should these mechanisms be subject to other contextual norms (pertaining to intelligent agents and their administrators)? In the practical world, this is comparable to the question of whether operators of systems and the technical infrastructures they deploy can simply posit themselves as (providers of) communication channels that are not bound by the social context of their users. The papers we surveyed consistently treat them as a product of but not as subject to the application of the CI. We have raised the possibility of an “operational” context, with an ‘operator’ role empowered with certain privileges and responsibilities

with respect to information flowing on the platform. Further guidance on this matter will be pertinent to enabling a holistic application of CI to technical designs.

On a related matter, further guidance is also necessary with respect to systems that provide infrastructure to multiple contexts: Such systems are expected to reflect on the normative aspects of CI while promoting a logic that can provide the flexibility for multiple social or technical contexts with potentially diverging informational norms to co-exist. What role the normative and conceptual aspects of CI can and should play in infrastructure underlying multiple contexts is an open research question.

5.1.2 Calls to Action

We call computer science researchers to be as explicit as possible about how the technologies they design are situated in the broader complex of platforms, operators, users, and moderators. If there is an implicit hierarchy (such as users whose activities are logged, agents that track all conversations, and auditors who use these records, or an operating system with many dependent applications), computer scientists can be explicit about this and address the privacy and information flow issues resulting from this differentiation of roles. If there are critical roles in the operation of the system (such as an auditor or operator), can privacy tools inspired by contextual integrity be built for them?

Most papers we selected did not focus on social spheres but on situations, proposing techniques that surface or implement informational norms that arise in a representational or interactional context. This focus often meant that in their models the authors did not consider normative rules applying to a specific context, abstracting the social sphere away. Some of this is justified: the intention is to develop designs that are flexible enough to function in different social spheres. It is also possible that the authors are more comfortable making normative judgements about what constitutes a relevant situation, e.g., some combination of location and activity, these also being things that the researchers can measure using sensor data. However, the numerous research papers showing user concerns due to context clash in online social

networks, as well as ever increasing public calls for curation of user generated content suggests that lack of attention to informational norms in social spheres may have negative consequences and should not be an afterthought. One way to guarantee this in abstract formal models as well as in infrastructure, is to at least provide a placeholder in associated conceptual frameworks that can be used to express and enforce normative rules when these systems are implemented. Better would be to also consider how well a proposed technique can sustain divergent informational norms pertaining to different social spheres that an infrastructure comes to play a role in. Developing and evaluating systems that enable a fluid interaction between informational norms in a social sphere and user preferences presents itself here as an interesting research question.

5.2 Diverse concepts of context

Our investigation into how computer scientists conceptualize contexts when they employ CI revealed diverse and divergent theoretical assumptions. Some researchers were well aligned with CI's concept of contexts as abstract, normative, social spheres; others drew on other traditions such as ubiquitous computing's concept of context as situations including users and technology. Still others supported users co-creating contexts through their engagement with each other and the technology. Some drew inspiration from multiple sources in order to provide a new technical solution to privacy.

This variety of work demonstrates that privacy and context are closely linked. It also demonstrates that *context* is a polysemous (many-meaning) term. The different senses of context have different implications for privacy by design. Our survey suggests that no one sense of context supports either a complete normative theory or technical design, and that there is a rich design space at the interplay between diverse specific meanings.

5.2.1 Theoretical gaps

Our investigation revealed inductively that computer scientists use diverse meanings of context that vary across many dimensions. "Context"

can refer to something abstract or concrete, social or technical, representational or interactional, normative or descriptive, and a priori modeled or empirically discovered. Only a subset of this space of meanings are addressed by CI in its current form, specifically, in the way it conceptualizes contexts as abstract, normative, social spheres continuously evolving within differentiated society.

It is not surprising that technical designs are concerned with the concrete circumstances of both users and technical applications. In order for CI to be actionable in this sense, what is needed is a theoretical account of how social spheres relate to sociotechnical situations. This account may well address other tensions between the many senses of “context”. For example, an advanced CI would be able to guide how to infer from the observed, descriptive details of a situation a model of the norms appropriate to guide behavior within it. This is a philosophical problem, but one that is made urgent by the demands of existing research on privacy by design.

Another theoretical challenge to CI is raised by Dourish’s [2004] critique of ubiquitous computing. CI’s model of contexts as social spheres parameterized by roles, information types, and transmission principles does suggest what Dourish describes as a positivist model of social contexts: contexts as containers of social action with specific expectations and prescriptions associated with them. To the theorist, we raise the question: what is the relationship between the situated, interactional account of context in Dourish and the social spheres in CI? The theory of “Implicit Contextual Integrity” invented by Criado and Such [2015] has suggested that the spirit of CI can be extended to social situations that evolve on a much smaller and more specific scale than is currently suggested by CI. Philosophical theorists can work to make this claim more precise.

5.2.2 Calls to Action

Computer scientists need not wait for theoretical prescriptions to continue to do good work at the intersection of CI and privacy by design. There is much to be done designing systems that address the reality that supporting users privacy requires skillful creation and moderation

of context. We anticipate that the best work will be explicitly aware of the challenges of matching concrete situations with the abstract spheres from which CI posits users get their normative expectations. Applications of CI are especially likely to be relevant in the smart environment applications, where sensors and actuators will interact with many users at once, making it hard to rely on individual preferences and expectations.

Indeed, any one of the dimensions of variation in the meaning of context (abstract or concrete, social or technical, representational or interactional, normative or descriptive, and a priori modeled or empirically discovered) presents a technical problem to computer scientists wishing to implement privacy by design based on CI. One concrete issue that persisted throughout many papers is the scoping of context. Papers, for example, that focused only on the social context, be it either due to their focus on user interfaces and experiences or social platforms, neither considered what we call the operational context, nor did they pay attention to how social informational norms may be impacted by flows of personal information to third parties. It would be very valuable to have studies that not only consider norms of information flow among users or towards an app provider, but also flows to other third parties, like other services, companies or governments. If a study intends to focus only on a subset of the information flows, than the limitations of the results due to this decision should be made explicit. We call upon computer scientists to work on pragmatic solutions to the problems these conceptual discrepancies pose to designers and users.

5.3 Privacy as a Moral Imperative: Between bits and norms

One major finding from our investigation of **RQ3** is that *none* of the papers in our review used the normative aspect of contextual integrity as a basis for their technical contributions. In contextual integrity, the normativity of privacy comes from the ends, purposes, and values of social contexts (spheres) as they have evolved over time. These ends, purposes, and values legitimize the norms that determine the appropriateness of information flow, even as technology changes what those

norms should be. Computer scientists sometimes acknowledge this aspect of contextual integrity, but they do not ground their technical contributions in it. Rather they draw on other sources of normativity, such as threats, user expectations, or legal compliance, to motivate their work.

For a number of reasons, these moves are understandable. Computer science has not traditionally equipped itself to deal with the hard problems surrounding the origins of ethics and morals. Threat modeling is narrowly pragmatic and has proven to be suitable for engineering purposes. User expectations are measurable and so attractive to those concerned with empirical validity. Considerations of legal compliance are part of the real business logic of functioning organizations. By focusing on these sources of normativity, computer scientists make their research more actionable. But these methods also carry the risk of falling short of socially maintained norms of privacy. Threat modeling may miss the mark; user expectations can be habituated by technology that works at odds with social principle; laws may be unjust. The burden is on Contextual Integrity theorists to show how its social and philosophical theory of social norms relates to these more concrete factors. In turn, we call computer scientists to stretch towards the social and philosophical sources of normativity. Our survey has shown that such ambition can lead to new technical innovation.

5.3.1 Theoretical gaps

Contextual integrity theorists need to address how their metaethical theory, whereby norms arise from the evolution of social spheres, ties in with the concrete sources of normativity used by computer scientists. We have identified three areas that need elaboration.

Contextual Integrity must provide a way of translating from the information norms of social spheres into a characterization of enumerated and discrete privacy *threats*. This is connected to the task of deriving mid-level theories of CI for modules of the technical stack (see Section 5.1.1).

Contextual integrity must also articulate the special place for user expectations, preferences, and control within the general framing of

appropriate flow. This would require greater fleshing out of situations where user control is legitimate, given its importance in the sphere of technical device usage. It would also address questions of how to resolve conflicts between user preferences and social norms, and between users with different preferences and expectations.

Finally, CI theorists must clarify the relationship between social spheres and the law. While there is in the United States an attractive synergy between the structure of sectoral privacy laws (like HIPAA, GLBA, and the like) with the view of society differentiated into social spheres, the relationship between social spheres and the law is less clear in jurisdictions of omnibus data protection laws such as the EU's GDPR [Herrmann et al., 2016]. The CI theorist must address what circumstances social norms provide important guidelines to appropriate information flow that are not covered by law, and what advantages they provide to technology designers who heed them.

5.3.2 Calls to Action

Computer scientists need not wait for passive instruction on the normative goals of their work. Rather, the problem of measuring *social norms*, in contrast to user's expectations, is one that requires technical sophistication. Shvartzshnaider et al. [2016] is one example of a paper that takes this task on explicitly in service of contextual integrity.

Computer scientists are particularly well situated to study users' perception and expectation of informational norms in different social spheres (and not independent of them). Developing and evaluating techniques to do so remains an open question. Coming back to Dou-rish, exploring how far users and different actors can be brought into engage in the evolution of informational norms is another avenue of exploration that has not been exhausted by researchers in our survey. Computer scientists may also consider designing systems that support communities of users' to determine their own norms.

Many of the studies did not consider conflicts among actors: these could be conflicts in informational norms across contexts, in different situations, even for individual users due to how their expectations evolve in relation to norms throughout time. Discrepancies between ideal vs.

actual norms may also lead to conflicts. Looking at these conflicts not as something to be designed away but as productive points of departure for engagement presents itself as another interesting research question.

5.4 Expanding and Sharpening Contextual Integrity

This leads to our findings from **RQ4**, where we look for aspects of CI that computer science researchers expanded on through their work. Computer scientists sometimes worked through mechanisms of technical adaptation to social change as they tried to respect privacy norms that were grounded in descriptions of concrete social and technical interaction.

5.4.1 Theoretical gaps

CI theorists must develop the framework's account of normative change and adaptation. The work we surveyed suggests a need for technical systems that automatically recognize contexts and that are sensitive to the norms of their users, even though social contexts and norms change. What principles can CI offer to adaptive system designers to ensure that these coevolving sociotechnical systems maintain their legitimacy with respect to the purposes of some more abstract social sphere? Do such systems challenge the theory that social contexts are robust in their ability to maintain their purpose? On what grounds would such a system be considered maladaptive? Is there any danger that technology will derail the social processes that reproduce contexts, or can society always be trusted to correct its own use of technology over time? What if powerful actors leverage existing systems with appropriate flows for ends, purposes and values that lack legitimacy? These thorny theoretical questions are both profound and of practical import for system design.

CI must also address the critical "sore" point in the present-day when many systems and devices span multiple contexts. Our inquiry here into the many relevant senses of 'context' sheds light on this phenomenon. Contexts can clash when the norms of multiple social spheres applicable to the same situation conflict with each other. Information

can also flow inappropriately between different situations. A more actionable version of CI will address these complex privacy challenges specifically.

CI also needs to better account for the relationship between privacy and time. Some papers in our survey tried to adapt CI to systems in which data did not only flow, but also was stored, processed, and deleted. Current versions of CI do not recognize that sometimes merely holding data (sometimes for great durations) can pose privacy threats. We are considering developing a concept of exposure to characterize this threat. Relatedly, there is a nuance discovered by Datta et al. [2011] that is not observed within CI: that it may not be apparent whether a case of information flow is inappropriate at the time that it flows because prescriptions refer to actions in the future. A more mature version of CI would account for the conditions under which parties can be aware of their privacy violations, and how ambiguities can be resolved.

Related to the questions resulting from the ambiguity or incompleteness of privacy norms are questions concerning the relationship between user choice and privacy. CI can in principle accommodate a wide range of preferences and a pluralistic society despite presupposing robust social agreement on information norms and the nature of social spheres. Technology is often designed to maximize adoption to diverse users and consequently can give (or restrict) users' control over how their data flows. A refinement of CI would address how user control and user diversity relate to social norms.

5.4.2 Calls to Action

Computer scientists have already made significant contributions to CI by providing valuable exemplars of research on adaptation, multiple contexts, temporality and duration, and user decision making. This work is invaluable for the evolution of CI.

We see further potential at the intersection of information theoretic approaches to privacy and contextual integrity. Many of the papers made use of techniques coming from machine learning, access control, formal methods, and user surveys, however, while inferences from

information flows were a concern in some papers [Omoronyia et al., 2012, 2013, Datta et al., 2011] we were missing works that looked at evaluating or enforcing desired norms using information theoretic models. One thing is to have a policy that expresses a norm that limits the flow of information about race, gender, class, religion and other sensitive attributes, it is another to guarantee that this information cannot be inferred otherwise. One could also imagine novel protections like differential privacy being used to develop elaborate transmissions principles. The numerous papers we surveyed demonstrate that computer scientists have actively applied and contributed to the evolution of contextual integrity using novel techniques. We hope these results serve to provide inspiration and guidance to all those researchers who are committed to leveraging or further developing CI in theory and practice.

Acknowledgements

We gratefully acknowledge the comments of Nick Doty and Blase Ur, and all other participants of our workshop session at the Privacy Law Scholars Conference (PLSC) 2017, where we were honored with Junior Scholars award. We thank Deirdre Mulligan and the anonymous reviewers for their detailed comments, which greatly improved the paper. We thank James Finlay for shepherding this paper through the publication process.

We gratefully acknowledge the support of imec-COSIC KU Leuven and FWO.

This material is based on research sponsored by DARPA under agreement number FA8750-16-2-0287. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA or the U.S. Government.

Appendices

A

Research Template

1. Provide a short summary of the objectives of the paper?
2. What subfield are the authors situated in?
3. What are the technical elements of the framework the authors are proposing? (technique, system , model, mechanism, tool, platform)
4. What problem are they solving?
5. Do they explicitly address context?
6. What parameters are recognized?
7. Are further parameters introduced?
8. How CI is challenged or extended?

References

- M. Ackerman, T. Darrell, and D. J. Weitzner. Privacy in context. *Human-Computer Interaction*, 16(2-4):167-176, 2001.
- A. Barth, A. Datta, J. C. Mitchell, and H. Nissenbaum. Privacy and contextual integrity: Framework and applications. 2006. In *Security and Privacy, 2006 IEEE Symposium on* (pp. 15-pp). IEEE.
- A. Barth, J. Mitchell, A. Datta, and S. Sundaram. Privacy and utility in business processes. 2007. In *Computer Security Foundations Symposium, 2007. CSF'07. 20th IEEE* (pp. 279-294). IEEE.
- Ann. Cavoukian. *Whole body imaging in airport scanners: building in privacy by design*. 2009.
- D. D. Clark, J. Wroclawski, K. R. Sollins, and R. Braden. Tussle in cyberspace: defining tomorrow's internet. *ACM SIGCOMM Computer Communication Review*, 32(4):347-356, 2002.
- N. Criado and J. M. Such. Implicit contextual integrity in online social networks. *Information Sciences*, 325:48-69, 2015.
- George Danezis and et al. Privacy and data protection by design-from policy to engineering. 2015. arXiv preprint arXiv:1501.03726.
- A. Datta, J. Blocki, N. Christin, H. DeYoung, D. Garg, L. Jia, D. Kaynar, and A. Sinha. Understanding and protecting privacy: Formal semantics and principled audit mechanisms. In *International Conference on Information Systems Security*, pages 1-27. Springer: Berlin, Heidelberg, 2011.

- A. Dey, G. Abowd, and D. Salber. A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications. *Human-Computer Interaction*, 16(2-4), 2001.
- H. DeYoung, D. Garg, L. Jia, D. Kaynar, and A. Datta. Experiences in the logical specification of the hipaa and glba privacy laws. In *Proceedings of the 9th annual ACM workshop on Privacy in the electronic society*, pages 73–82. ACM, 2010.
- Paul Dourish. What we talk about when we talk about context. *Personal and ubiquitous computing*, 8(1):19–30, 2004.
- Federal Trade Commission. Protecting consumer privacy in an era of rapid change. 2012.
- Seda Gürses and Jose M. del Alamo. Privacy engineering: Shaping an emerging field of research and practice. *IEEE Security & Privacy*, 14(2):40–46, 2016.
- Michael Herrmann, Mireille Hildebrandt, Laura Tielemans, and Claudia Diaz. Privacy in location-based services: An interdisciplinary approach. *SCRIP-Ted*, 13:144, 2016.
- Y. J. Jia, Q. A. Chen, S. Wang, A. Rahmati, E. Fernandes, Z. M. Mao, A. Prakash, and S. J. University. Contextiot: Towards providing contextual integrity to appified iot platforms, 2017.
- I. Kayes and A. Iamnitchi. Aegis: A semantic implementation of privacy as contextual integrity in social ecosystems. In *Privacy, Security and Trust (PST), 2013 Eleventh Annual International Conference on* (pp. 88–97), IEEE, 2013a.
- I. Kayes and A. Iamnitchi. Out of the wild: On generating default policies in social ecosystems. In *In Communications Workshops (ICC), 2013 IEEE International Conference on* (pp. 204–208). IEEE, 2013b.
- B. Kitchenham, S. Charters, D. Budgen, P. Brereton, M. Turner, S. Linkman, M. Jorgensen, E. Mendes, and G. Visaggio. Guidelines for performing systematic literature reviews in software engineering. In *Version 2.3, Keele University and University of Durham, EBSE Technical Report*, 2007. URL <https://pdfs.semanticscholar.org/e62d/bbbb70cabcd3335765009e94ed2b9883d5.pdf>.
- Y. Krupa and L. Vercouter. Handling privacy as contextual integrity in decentralized virtual communities: The privacias framework. *Web Intelligence and Agent Systems: An International Journal*, 10(1):105–116, 2012.
- T. Lau, O. Etzioni, and D. S. Weld. Privacy interfaces for information management. *Communications of the ACM*, 42(10):88–94, 1999.

- M. Netter, M. Riesner, and G. Pernul. Assisted social identity management-enhancing privacy in the social web. In *10th International Conference on Wirtschaftsinformatik*, 2011.
- H. Nissenbaum. Protecting privacy in an information age: The problem of privacy in public. *Law and philosophy*, 17(5):559–596, 1998.
- H. Nissenbaum. Privacy as contextual integrity. *Wash. L. Rev.*, 79:119, 2004.
- H. Nissenbaum. *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press, 2009.
- I. Omoronyia, L. Pasquale, M. Salehie, L. Cavallaro, G. Doherty, and B. Nuseibeh. Caprice: a tool for engineering adaptive privacy. In *Proceedings of the 27th IEEE/ACM International Conference on Automated Software Engineering* (pp. 354–357). ACM, 2012.
- I. Omoronyia, L. Cavallaro, M. Salehie, L. Pasquale, and B. Nuseibeh. Engineering adaptive privacy: on the role of privacy awareness requirements. In *Proceedings of the 2013 International Conference on Software Engineering* (pp. 632–641). IEEE Press, 2013.
- Regulation (EU). 2016/679 of the european parliament and of the council. In *On the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)*, 2016.
- I. Rubinstein. Privacy and regulatory innovation: Moving beyond voluntary codes. 2010.
- M. Salehie, L. Pasquale, I. Omoronyia, and B. Nuseibeh. Adaptive security and privacy in smart grids: A software engineering vision. In *Proceedings of the First International Workshop on Software Engineering Challenges for the Smart Grid* (pp. 46–49). IEEE Press, 2012.
- R. Samavi and M. P. Consens. L2tap + scip: An audit-based privacy framework leveraging linked data. In *Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom), 2012 8th International Conference on* (pp. 719–726). IEEE, 2012.
- R. Sayaf, D. Clarke, and R. Harper. 2014. SECURECOMM 2014.
- F. Shih and M. Zhang. Towards supporting contextual privacy in body sensor networks for health monitoring service. *W3C Workshop on Privacy and data usage control*, 4(5), 2010.
- F. Shih, I. Liccardi, and D. Weitzner. Privacy tipping points in smartphones privacy preferences. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (pp. 807–816). ACM, 2015.

- Y. Shvartzshnaider, S. Tong, T. Wies, P. Kift, H. Nissenbaum, L. Subramanian, and Prateek Mittal. Learning privacy expectations by crowdsourcing contextual informational norms. In *Fourth AAAI Conference on Human Computation and Crowdsourcing (HCOMP 2016)*, 2016.
- M. Tierney and L. Subramanian. Realizing privacy by definition in social networks. In *Proceedings of 5th Asia-Pacific Workshop on Systems*. ACM, 2014.
- White House. Consumer data privacy in a networked world: A framework for protecting privacy and promoting innovation in the global digital economy, 2012.
- P. Wijesekera, A. Baokar, A. Hosseini, S. Egelman, D. Wagner, and K. Beznosov. Android permissions remystified: A field study on contextual integrity. in *usenix security* (vol. 15), 2015.
- World Economic Forum. Rethinking personal data. 2012.
- F. Zhang, F. Shih, and D. Weitzner. No surprises: measuring intrusiveness of smartphone applications by detecting objective context deviations. In *Proceedings of the 12th ACM workshop on Workshop on privacy in the electronic society* (pp. 291–296). ACM, 2013.