


RESEARCH

Open Access

Making context the central concept in privacy engineering



Tore Hoel^{1*} , Weiqin Chen¹ and Jan M. Pawlowski²

* Correspondence: tore.hoel@oslomet.no

¹Oslo Metropolitan University, Oslo, Norway
Full list of author information is available at the end of the article

Abstract

There is a gap between people's online sharing of personal data and their concerns about privacy. Till now, this gap is addressed by attempting to match individual privacy preferences with service providers' options for data handling. This approach has ignored the role different contexts play in data sharing. This paper aims at giving privacy engineering a new direction putting context centre stage and exploiting the affordances of machine learning in handling contexts and negotiating data sharing policies. This research is explorative and conceptual, representing the first development cycle of a design science research project in privacy engineering. The paper offers a concise understanding of data privacy as a foundation for design extending the seminal contextual integrity theory of Helen Nissenbaum. This theory started out as a normative theory describing the moral appropriateness of data transfers. In our work, the contextual integrity model is extended to a socio-technical theory that could have practical impact in the era of artificial intelligence. New conceptual constructs such as 'context trigger', 'data sharing policy' and 'data sharing smart contract' are defined, and their application is discussed from an organisational and technical level. The constructs and design are validated through expert interviews; contributions to design science research are discussed, and the paper concludes with presenting a framework for further privacy engineering development cycles.

Keywords: Privacy engineering, Contextual integrity, Context, Context trigger, Personal data, Online data sharing

Introduction

People who are concerned about privacy do not necessarily make choices about data sharing reflecting the gravity of their concerns. This gap defines the 'privacy paradox', observed in a number of studies (Baruh, Secinti, & Cemalcilar, 2017; Norberg, Horne, & Horne, 2007; Taddei & Contena, 2013). In real life, intentions only explain part of our behaviour (Sheeran, 2002). In online practices, this is demonstrated by the fact that most of us use popular online search engines, well knowing the 'free services' are paid for by sharing our personal data. The gap to be concerned about, however, is not that our actions do not follow our intentions, but the fact that available privacy solutions are so far behind our online practices. We share an unprecedented amount of personal data aligning our lives to data-driven smart cities, smart public services, intelligent

campuses and other online practices utilising artificial intelligence (AI). We know little about how this data is used. When pushing back, for example using the European General Data Protection Regulation (GDPR) to stop blanket acceptance of opaque privacy policies, we only seem to get more obfuscation, having to fight pop-up windows asking for permission to use our private data for every new site to be accessed. To close the gap and prevent ‘privacy fatigue’ (Choi, Park, & Jung, 2017), we need better privacy solutions, but both the research and design community are struggling to see where the solutions should come from.

The inventor of the world wide web, Tim Berners-Lee, admitted in 2017 that ‘we’ve lost control of our personal data’ (Berners-Lee, 2017). This paper is premised on what some may characterise as a defeatist position on data sharing: We will not be able to scale back sharing of personal data, no matter how much we appeal to the GDPR principles of purpose limitations and data minimization (EU, 2012). The craving for data exposing our behaviour as consumers, citizens and persons caring for our health and cognitive development is already strong (Mansour, 2016). And it is being strengthened by the AI arms race, where the fierce competition lessen the appetite to address contentious AI issues, such as data privacy, public trust and human rights related to these new technologies (Nature, 2019). The challenge needs to be addressed by stepping up efforts in privacy engineering searching for more adequate solutions to manage personal data sharing in a world of digital transformation.

This paper aims at advancing privacy engineering through contributions addressing semantic, organisational and technical aspects of future solutions. In the ‘Background’ section, we pinpoint the weaknesses of the current discourse on privacy and point to a better understanding of context as a fruitful direction of development. In the following sections, we construct conceptual artefacts and draw up designs that may support digital practices in a society embracing big data and more and more use of artificial intelligence.

Background

In this paper, we want to advance the field of privacy engineering, defined by Kenny and Borking as ‘a systematic effort to embed privacy relevant legal primitives [concepts] into technical and governance design’ (Kenny & Borking, 2002, p. 2). We would argue that not only legal primitives need to be embedded, but realise that adding philosophical, social, pedagogical and other perspectives make privacy engineering utterly complex. No wonder Lahlou, Langheinrich, and Rucker (2005) found that engineers were very reluctant to embrace privacy: Privacy ‘was either an abstract problem [to them], not a problem yet (they are ‘only prototypes’), not a problem at all (firewalls and cryptography would take care of it), not their problem (but one for politicians, lawmakers or, more vaguely, society) or simply not part of the project deliverables’ (Lahlou et al., 2005, p. 60). When the term “privacy” is so often misunderstood and misused in human-computer interaction (HCI) (Barkhuus, 2012), there is a need to converge on a subset of core privacy theories and frameworks to guide privacy research and design (Badillo-Urquiola et al., 2018), taking into account the new requirements of data-driven society (Belanger & Crossler, 2011).

Figure 1 gives an overview of how privacy theories have developed from mainly focusing on the individual handling ‘small data’ to dealing with data sharing in group and

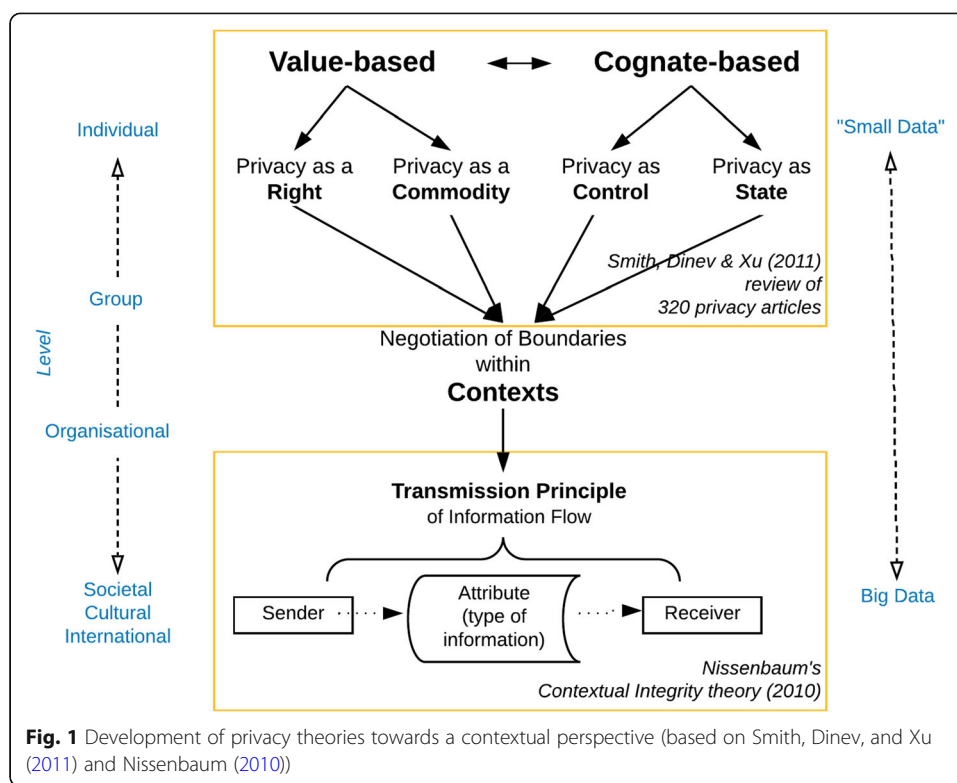


Fig. 1 Development of privacy theories towards a contextual perspective (based on Smith, Dinev, and Xu (2011) and Nissenbaum (2010))

societal settings, where new technologies using big data set the scene. In this paper, we see the development of a contextual approach to privacy as necessary to make progress within privacy engineering.

In their interdisciplinary review of privacy-related research, Smith et al. (2011) found that definitional approaches to general privacy could be broadly classified either as value-based or cognate-based (the latter related to the individual’s mind, perceptions and cognition). Sub-classes of these definitions saw privacy as a right or a commodity (that could be traded for perceived benefits), or privacy as individual control of information, or as a state of limited access to information. The problem with these theories is their lack of explanatory power when it comes to shed light on the boundaries drawn between public and private information in actual online practices in our digital age (Belanger & Crossler, 2011; Smith et al., 2011). In Fig. 1, we have indicated that when met with challenges from group-level interactions in data-rich networked environments, both value-based and cognate-based theories will be drawn towards contextual perspectives on how information flows. We would claim that when boundaries between private and shared information are negotiated—often mediated by some ICT tool—the perspectives from the individual privacy theories may play an active role. There will still be arguments referring to individual data ownership and control, data sharing with cost-benefit considerations and trade-offs or the ability to uphold solitude, intimacy, anonymity or reserve (the four states of individual privacy identified by Westin (2003)). However, these perspectives will serve as a backdrop of deliberations that require another set of privacy constructs, for which context will serve as the key concept.

It may be objected that to highlight context may just be to exchange one elusive concept (privacy) with another borderless concept (context). Smith et al. (2011) were not

at all sure that context-driven studies may produce an overall contribution to knowledge, ‘unless researchers are constantly aware of an over-arching model’ (ibid., p. 1005). To contribute to an understanding of privacy context, we have pointed to the theory of contextual integrity (CI) as a candidate for further development (see Fig. 1). In the following, we introduce the CI theory, focussing on how this theory’s concept of context is to be understood.

The contextual integrity theory

Over the last 15 years, CI has been one of the most influential theories explaining the often conflicting privacy practices we have observed along with the development of ubiquitous computing. When Helen Nissenbaum first launched CI, she used philosophical arguments to establish ‘[c]ontexts, or spheres, [as a] a platform for a normative account of privacy in terms of contextual integrity’ (Nissenbaum, 2004, p. 120). The two informational norms she focussed on were norms of appropriateness and norms of information flow or distribution. ‘Contextual integrity is maintained when both types of norms are upheld, and it is violated when either of the norms is violated’ (ibid., p. 120).

Privacy norms are not universal, ‘contextual integrity couches its prescriptions always within the bounds of a given context’ (Nissenbaum, 2004, p. 136). Non-universal norms may seem like an oxymoron, as a norm is supposed to cover more than one case. What role does CI give to context is a key question we see the originator of the theory grapples with in the 2004 article. ‘One of the key ways contextual integrity differs from other theoretical approaches to privacy is that it recognises a richer, more comprehensive set of relevant parameters’, Nissenbaum (2004, p. 133) states, reflecting on her application of CI on three cases (government intrusion, access to personal information and to personal space) that has dominated privacy law and privacy policies in the USA. However, access to more detail—a richer set of parameters—does not alter the way traditional privacy approaches have worked, implying violation of privacy or not from the characteristics of the setting matched against individual preferences. Barkhuus observes as follows:

It (..) appears rather narrow to attempt to generate generalized, rule-based principles about personal privacy preferences. Understanding personal privacy concern requires a contextually grounded awareness of the situation and culture, not merely a known set of characteristics of the context. (Barkhuus, 2012, p. 3)

This questions on how context should be understood in relation to preference—as *something more* than characteristics of individual preferences—represents a research gap that will be addressed in this paper as it goes to the heart of the privacy discourse: Where are norms of the appropriateness of the exchange anchored—internally or externally—in the value system of the individual or in the negotiated relationships to others in situations?

First, we will explore how context is to be understood, before we return to the question of how preference and context are related.

Understanding context

Context is the set of circumstances that frames an event or an object (Bazire & Brézillon, 2005). This generally accepted meaning of the term is not very helpful when wanting to

use it in a specific discipline where we need a clear definition. There are, however, many definitions of context to choose from. Bazire and Brézillon (2005) collected a corpus of more than 150 definitions, most of them belonging to cognitive sciences (artificial intelligence being the most represented area). In human cognition, they note, there are two opposite views about the role of context. The first view considers cognition as a set of general processes that modulate the instantiation of general pieces of knowledge by facilitation or inhibition. In the second view (in the area of situated cognition), the context has a more central role as a component of cognition by determining the conditions of knowledge activation as well as the limit of knowledge validity.

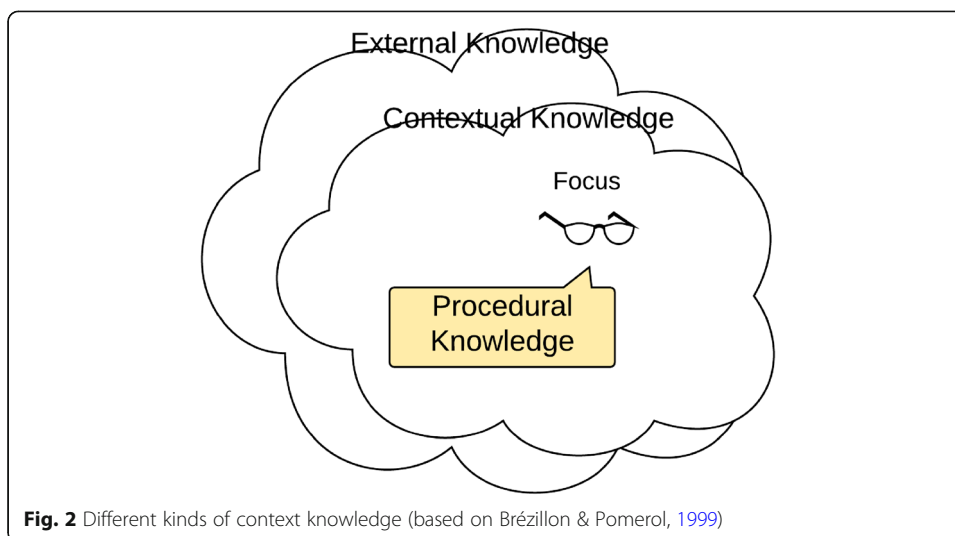
These two opposite views have parallels with our question above about the basis of a decision on the appropriateness of data exchange. The context may have an internal nature or an external one. ‘On the one hand, context is an external object relative to a given object, when, on the other hand, context belongs to an individual and is an integral part of the representation that this individual is building of the situation where he is involved. According to this second viewpoint, ‘context cannot be separated from the knowledge it organises, the triggering role context plays and the field of validity it defines’. (Bazire & Brézillon, 2005, citing Bastien, 1999).

In the context of privacy deliberations, in this paper, we would follow the second viewpoint focussing on knowledge organisation. Context is itself contextual; context is always relative to something—described as the focus of attention. For a given focus, Brézillon and Pomerol (1999) consider context as the sum of three types of knowledge. *Contextual knowledge* is the part of the context that is relevant for making decisions, and which depends on the actor and on the decision at hand. By definition, this creates a type of knowledge that is not relevant, which Brézillon and Pomerol (1999) call *external knowledge*. However, what is relevant or not evolves with the progress of the focus, so the boundaries between external and contextual knowledge are porous. A subset of the contextual knowledge is *proceduralised* for addressing the current focus. ‘The proceduralized context is a part of the contextual knowledge that is invoked, assembled, organised, structured and situated according to the given focus and is common to the various people involved in decision making’ (Brézillon, 2005, p. 3), see Fig. 2.

The first point of view on contexts discussed above—modulation of an external object by facilitation or inhibition—has similarities with what we could call a preference approach (matching individual preference characteristics with alternative actions). In the next sections, we will explore limitations to this approach reviewing literature from the field of personalised learning and privacy standardisation.

Limitations of a preference approach

The individual privacy preference approach has similarities with the approach of personalised learning (Campbell, Robinson, Neelands, Hewston, & Massoli, 2007), which has been critiqued for lack of nuanced understanding for how the needs of different learners can be understood and catered for in school. Prain et al. (2012) argues that the critical element in enacting personalised learning is the ‘relational agency’ operating within a ‘nested agency’ in the development of differentiated curricula and learners’ self-regulatory capacities.



The construct of ‘nested agency’ recognises that the agency of both groups [teachers and learners] as they interact is constrained by structural, cultural and pedagogical assumptions, regulations, and practices, including prescriptive curricula, and actual and potential roles and responsibilities of teachers and students in school settings. (Prain et al., 2012, p. 661)

The main lesson learnt from the well-researched field of personalised learning is the need for a better understanding of the contexts, in which the learning occurs.

A special group of learners are persons with disabilities. ISO/IEC published in 2008 the Access for All standard aiming ‘to meet the needs of learners with disabilities and anyone in a disabling context’ (ISO, 2008). The standard provided ‘a common framework to describe and specify learner needs and preferences on the one hand and the corresponding description of the digital learning resources on the other hand, so that individual learner preferences and needs can be matched with the appropriate user interface tools and digital learning resources’ (ISO, 2008). The Canadian Fluid Project has proposed to use the same framework to define *Personal Privacy Preferences*, working as ‘a single, personalised interface to understand and determine a privacy agreement that suits the function, risk level and personal preferences’, so that, ‘private sector companies would have a standardised process for communicating or translating privacy options to a diversity of consumers’ (Fluid Project, n.d.).

Using the ISO 24751:2008 framework to define personal privacy preferences implies acceptance of the standard’s definition of context as ‘the situation or circumstances relevant to one or more preferences (used to determine whether a preference is applicable)’. Then privacy is seen as a characteristic of the individual, rather than a relationship between different actors mediated by contexts. The Canadian project proposes to ‘leverage ISO 24751* (Access for All) to discover, assert, match and evaluate personal privacy and identity management preferences’ (Fluid Project, n.d.). However, the challenge is not to facilitate matching between predefined preferences and alternative representations of web content (which was the focus of the Access for All standard); the

challenge is to orchestrate dynamic privacy policy negotiations in the particular contexts of a great number of online activities. If one sees only individuals with needs, one tends to overlook other factors, like culture, social norms and activity patterns embedded in complex settings.

To make context ‘a first-class citizen’ (Scott, 2006) in privacy engineering CI needs to be developed from a normative ex post theory to a theory positioned more in the middle of privacy negotiations supported by information technology. In the next subsection, we will see how the theoretical base of CI has been broadened by Helen Nissenbaum and different research groups.

Formalising CI

Barth, Datta, Mitchell, & Nissenbaum, 2006 made a first attempt to make a formal model of a fragment of CI, focussing on “communicating agents who take on various roles in contexts and send each other messages containing attributes of other agents” (Barth et al. (2006), p. 4). In 2010, Nissenbaum provided a nine-step decision heuristic to analyse new information flows, thus determining if new practice represents a potential violation of privacy. In this heuristic, she for the first time specified precisely which concepts should be defined to fulfil a CI evaluation, see Fig. 3 (Nissenbaum, 2010, pp. 182-183).

From this heuristic, authors in collaboration with Nissenbaum have developed templates for tagging privacy policy descriptions (e.g. Facebook’s or Google’s privacy policy statements) to be able to analyse how these documents hold up to the CI theory (Shvartzshnaider, Apthorpe, Feamster, & Nissenbaum, 2018).

CI describes information flows using 5-parameter tuples, which include actors (data subjects, senders and receivers) involved in the information flow, the type (attribute) of the information, and the condition (transmission principle) under which the information flow occurs. This combination of five factors defines contexts, which determine privacy norms.

Contextual integrity describes a desirable state that people strive towards by keeping perceived-private information private according to the context. For example,

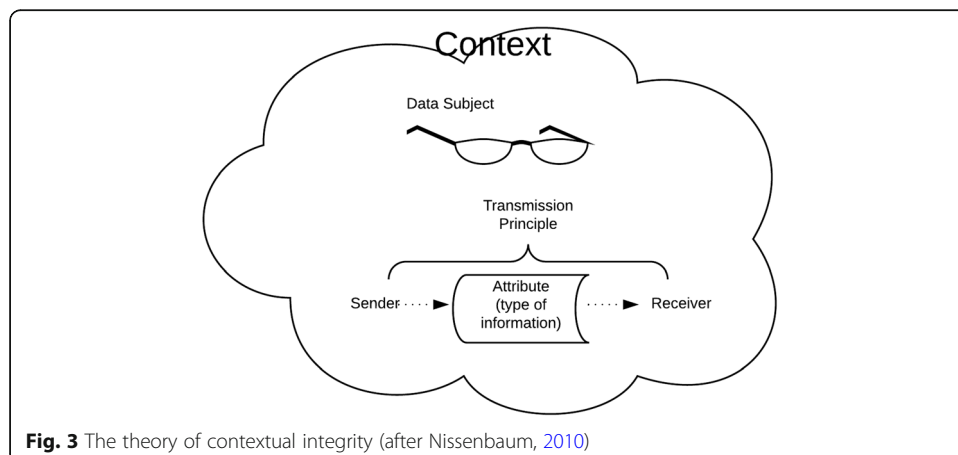


Fig. 3 The theory of contextual integrity (after Nissenbaum, 2010)

people expect to share medical information with doctors but not with employers. Where it in some cultures yearly salary is perceived as private, within others it is normal to share this information. Contextual integrity thereby explains how privacy is grounded in each context, governed by pre-existing norms and values. (Barkhuus, 2012, p. 3).

However, mapping pre-existing norms and values is not easy to achieve in situations where there is no pre-defined understanding of privacy. The use of social media is a case in point. 'Understanding personal privacy concern requires a contextually grounded awareness of the situation and culture, not merely a known set of characteristics of the context' (Barkhuus, 2012, p. 3).

CI bridges two worlds—a world inhabited by humanists, social scientists, lawyers and regulators and another world inhabited by mathematicians, computer scientists and engineers, according to Benthall, Gürses, and Nissenbaum (2017). For the latter world, '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.)' (Benthall et al., 2017, p. 12). The CI theory gives the framework and the key concepts for a contextual understanding of information privacy. However, by positioning contexts as 'well-defined settings', these authors limit the potential of CI in an era when sensors and computational power allow a more dynamic reasoning about contexts. As Barkhuus has stated, the theory needs to be further developed, and new research should be informed with privacy.

Appropriation of behavior in the situation (..) and not a behavioristic belief that people's actions are based on a structured set of privacy concerns. Instead of focusing on the how and what in terms of people's preferences for personal data sharing, we need a foundation of research that looks at why.' (Barkhuus, 2012, p. 8)

However, computational reasoning about *why* goes beyond the affordance of current information technologies as practical benefits of *general AI* still needs to materialise (Tuomi, 2018). Data-driven machine learning (ML), on the other side, is available, and it should be explored how this field could contribute to a contextual approach to privacy engineering.

This review of how context has become a more central concept in privacy engineering has identified a number of research gaps. We have pointed to the new era of Big Data, AI and unparalleled access to processing power, all factors that open up for dynamic and synchronous reasoning about privacy decisions. For this to happen, we need to further develop 'context' as a concept that reaches beyond just framing pre-defined preference settings. This means to advance CI in the direction of a prescriptive theory, giving context a pivotal role in IT systems that answer user requirements. Therefore, we have identified design challenges both on conceptual, process and technical levels of design.

Methodology

This work follows a conceptual and exploratory approach; however, it is situated in the tradition of Design Science Research (DSR) (Gregor & Hevner, 2013) where the

developed constructs are tested against empirical cases derived from literature and in interaction with the practice field. The methodology implies several cycles of design; however, in this paper, we present the results of initial development, focussing on conceptual analysis. The main objective of this paper is to come up with novel design artefacts (Baskerville, Baiyere, Gergor, Hevner, & Rossi, 2018) that a selected group of experts will see as valuable in future privacy and information sharing scenarios. We have chosen a design approach addressing the design task as different layers of interoperability. Design theory implications of this approach play, however, a minor role in this paper.

We will use the lens of the European Interoperability Framework (EIF) to structure our approach. EIF (Fig. 4) has four levels, the first covering legal interoperability. For now, we leave this aside as we might say that GDPR and other legal frameworks have levelled the legal ground for privacy engineering. The different actors know how to interoperate to exchange and handle information legally, and developers are committed to the principle of ‘privacy by design’ (PbD) (Cavoukian, 2009), i.e. the obligation, from the very beginning to build privacy into their solutions. However, in this field, clarity is lacking in the other interoperability levels of EIF—at the organisational, semantic and technical levels.

In order to make privacy more than an afterthought, after solutions are designed and implemented, we must know what PbD means at different levels. We will explore privacy design at three levels, in this order:

- **Semantic:** how is privacy conceptualised, and how could privacy concepts be formalised to be used in technical design to achieve precise format and meaning of exchanged data, so that information is preserved and understood throughout exchanges between parties;
- **Organisational:** how institutions align their business processes, responsibilities and expectations in relation to privacy to achieve commonly agreed and mutually beneficial goals;

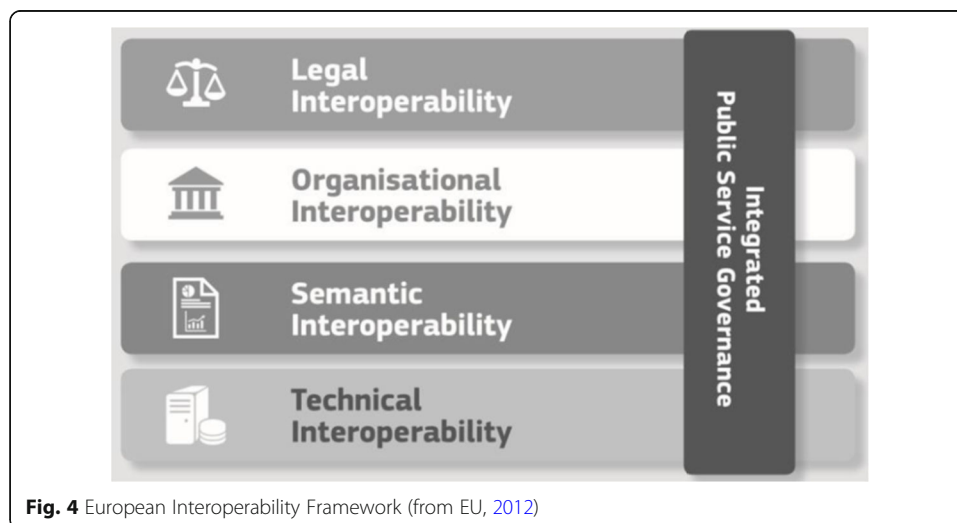


Fig. 4 European Interoperability Framework (from EU, 2012)

- **Technical:** what system architecture would ensure that data could be shared between system and services without violating the privacy of users.

A first validation of design is done through semi-structured interviews with a group of 11 experts from the standard community, the learning analytics community and the information science community, whom the authors knew have touched privacy issues in their academic work. In a web form that was used as a scaffold for the interviews, the experts were given a link to an early draft of this paper and presented with the design artefacts, i.e. Figs. 1, 2, 6 and 10 in this paper; and a figure of a context graph template; and the guiding definition of data privacy. The recorded online Skype interviews were staged as a stepwise discussion, moving through 8 pages from a general discussion on privacy (Fig. 1), to discussing the data flow of the application scenario of the proposed solution (Fig. 10). Substantial parts of the interviews, each lasting from 35 to 60 min, were transcribed and analysed to probe acceptance of the proposed privacy engineering approach and to capture suggestions for improvement of constructs. In addition, the constructs were discussed in the context of an educational scenario to see how the suggested approach holds up to well-known use cases in learning, education and training (LET).

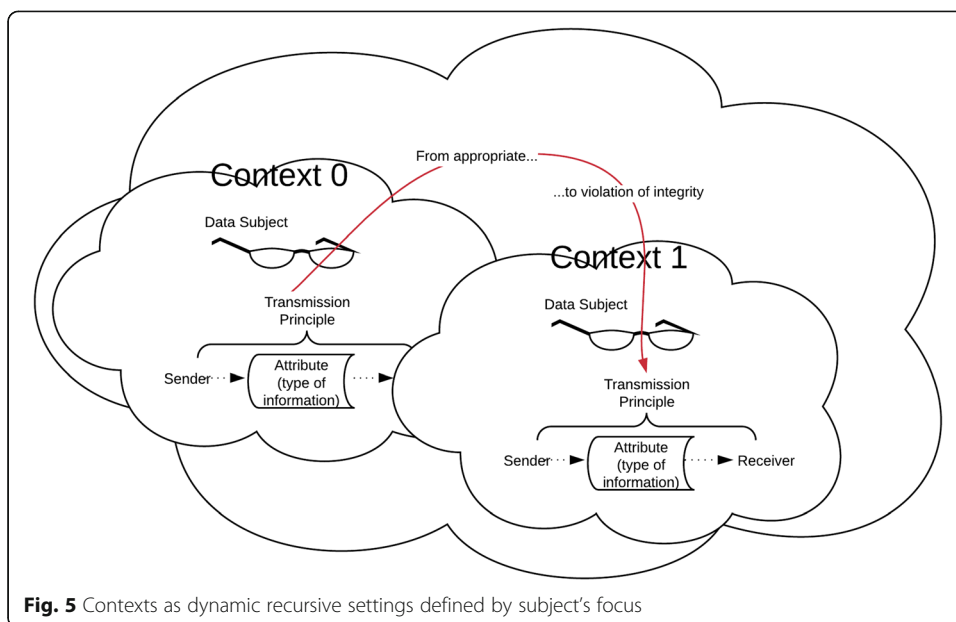
Conceptual design—privacy in context

‘Taking context seriously means finding oneself in the thick of the complexities of particular situations at particular times with particular individuals’, as what Nardi (1992, p. 35) observed when considering HCI design challenges. Applying a top-down approach, finding commonalities across situations is difficult ‘because studies may go off in so many different directions, making it problematic to provide the comparative understanding across domains’, she concluded. Nissenbaum’s CI model, we have pointed out, needs to be extended with a better understanding of context. Nowadays, we do not need to handle context as a container; we can sense context more in real time as an occasioned event—relevant to particular settings, particular instances of action and particular parties to that action—not an abstract category that generalises over settings, actions, and parties.

In this section, we will put the CI model in context of Patrick Brézillion’s theory (Brézillion & Pomerol, 1999; Brézillion, 2003, 2005), extending CI with an event dimension. We then suggest a representation format to describe reasoning about privacy events. The semantic design concludes with provisional definition of a data privacy process to guide organisational and technical design.

To allow sensors and computers to work on privacy data, we need a formal representation of the CI model. The core attributes of CI are formally represented as $C(ds, s, r, a, t)$, where C = context, ds = data subject, s = sender, r = receiver, a = attribute (type), and t = transmission principle. However, C is not understood as a top-down container but as dynamic recursive settings where C_1 is understood in relation to C_0 and the shift of focus determined by other contexts (Fig. 5).

Figure 5 describes a situation where sharing of data becomes a violation of integrity as the focus changes when moving from one context to another. Both contexts have a common backdrop context that could be activated. The calculation of the



appropriateness of the transmission is based on the same principles in C_0 and C_1 , but the change of perspective due to a new context gives different results.

We are interested in describing the knowledge production related to different situations regulating the appropriateness of information flows. What happens when the shift of focus changes the situation from one of being appropriate to one violating the integrity of the data subject? The change is caused by moving from C_0 and C_1 in Fig. 5. However, the outcome may be influenced by events which activate knowledge held in a supra context to these two contexts. We see this process as a negotiation where the data subject interacts with different contexts drawing on different types of knowledge, see the three categories of knowledge described in Fig. 2: procedural, contextual and external (Brézillon & Pomerol, 1999).

A subset of the contextual knowledge is foregrounded to address the current focus. In our model, this is the ‘calculation’ of contextual integrity (the appropriateness) of the data sharing under the contextual circumstances. In terms of CI, we have the syntax for building the proceduralised context (see above); we have just to find data for who are data subjects and who are senders and receivers of what information to make a decision about the appropriateness of the applied transmission principle.

Events that trigger contextual knowledge creation

Once this proceduralised context has satisfied the focus, this piece of knowledge goes back to the contextual knowledge. The context will remain active depending on the decision. If the decision does not raise further questions, or other events do not occur, the context is dissolved and the knowledge stored as external knowledge. However, just a small incident is enough to trigger the CI dynamic of establishing a context and integrating external knowledge in building proceduralised contexts (to test transmission principle appropriateness) and further movements between the body of contextual knowledge and proceduralised contexts.

From a privacy engineering perspective, the context triggers—the events that challenge the data sharing—are of special interest (Fig. 6).

We assume that these events vary with contextual factors like culture, legal domain, trust, institutional actors, tools to be used, etc. In the model in Fig. 6, it is the context triggers that activate the reflection on data sharing contexts, which in turn leads to confirmation of revision of data sharing policies for practice in pervasive online environments.

A context trigger is defined as an event, which implies notification and different sources triggering the event like user input, interaction with other data subjects or systems, environmental conditions, exposure to information flows (e.g. news), etc. These events may be internal or external to the current activity context. In the end, context triggers can be understood as something similar to messages between different contexts, being able to spark knowledge processes related to information flows.

Contextual graphs

The concept of context trigger extends Nissenbaum's CI model. Since events add a new dimension to contexts, we need a different way to represent the process other than just adding a new attribute to the $C(ds, s, r, a, t)$ -formula. We suggest to use contextual graphs, a notation system developed by Brézillon (2003). This is a scheme that makes it possible to 'represent and clearly organise operators' activities and all their variants (procedures and practices), include automatic learning and adaptation capabilities in a system and make context explicit in the representation of operators' reasoning' (Brézillon, 2003, p. 21-22).

Conceptual graphs have been used in knowledge management projects as a tool for incremental learning (e.g. for incident solving on a subway line) (Brézillon, 2003). New practices have been compared to the existing knowledge graph and added as new nodes in the conceptual graph if they were valuable for future events. Conceptual graphs do not deal with probabilities, and there is no decision node, only 'chance' nodes where the contextual element is analysed to select the corresponding path (Mostéfaoui & Brézillon, 2004).

Schematically, a contextual graph is an acyclic graph with a unique input, a unique output, and a serial-parallel organisation of nodes connected by oriented arcs. A node can be an action (square box), a contextual node (circular box) or a recombination

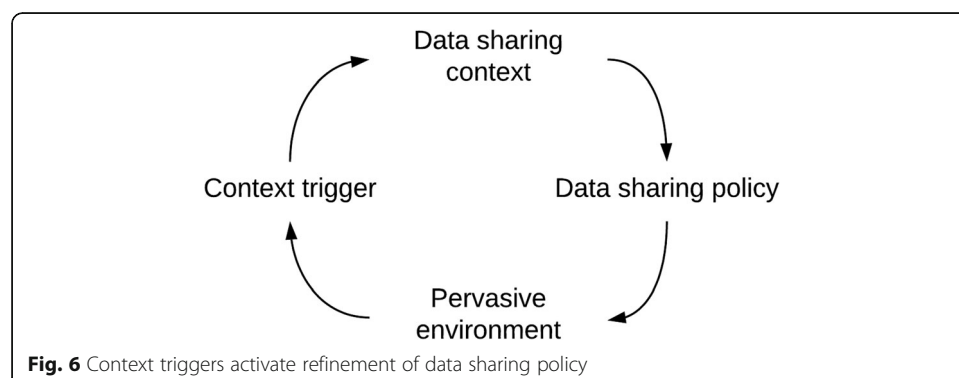


Fig. 6 Context triggers activate refinement of data sharing policy

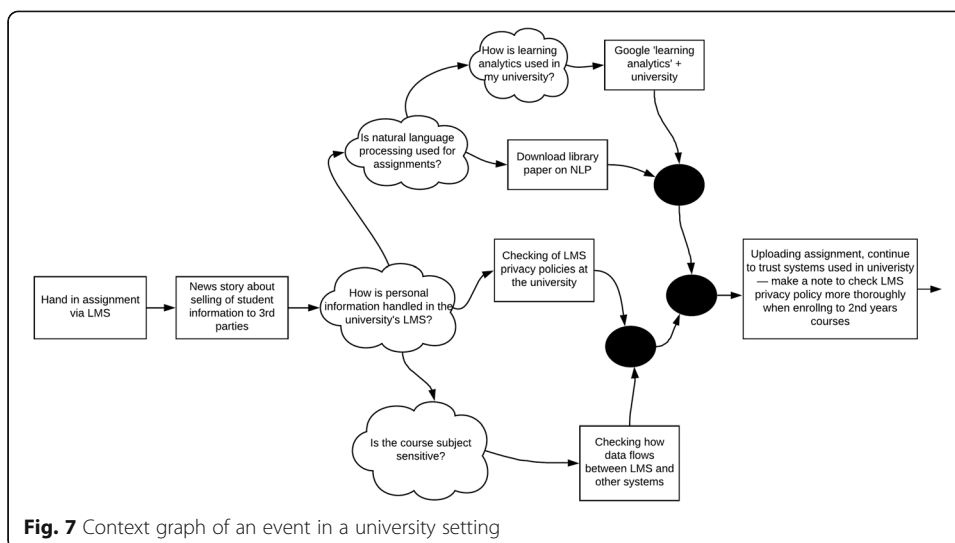


Fig. 7 Context graph of an event in a university setting

node (black box) (Mostéfaoui & Brézillon, 2004) (see Fig. 7 for an example of a context graph describing an event in a university setting).

A context trigger is registered as an action in the context graph notation. In the example in Fig. 7, the triggering news story mentioned in the second box from the left could have been the news of the August 12, 2019, letter to all major edtech companies from three US senators, expressing concerns that educational technologies ‘may put students, parents and educational institutions at risk of having massive amounts of personal information stolen, collected or sold without [the students’] permission or knowledge’ (Durbin, Markey, & Blumenthal, 2019). Other actions in the graph may be related to negotiations about data transmission principles. The output action is opening up or closing down the data sharing.

Understanding data privacy

Based on the conceptual development in this section and a validation in the first round of expert interviews, we will suggest a provisional understanding of data privacy to be used in privacy engineering. We see work towards improving data privacy as follows:

The process to be used in privacy engineering. We see work towards improving data privacy as to data transmission principles. The output action is openinfluence how personal data will be handled by service providers.

This definition gives agency to the user, which may be a natural person or an IT agent. Execution power is delegated to a policy that is understood as relatively stable principles derived through negotiations with actors present in different contexts of interaction. Even if it is the user that defines its data-sharing policies, it is the operational context that constrains the interpretation of the defined scheme of actions. The user may have a principled point of view, and a situational point of view—both views allowing practices that may seem to be in conflict with each other. However, from a

contextual perspective, expressed in the same concerted policy, a range of apparently conflicting positions could be played out.

In this section on conceptual development, we have extended the five-tuple model of CI with a concept of context triggers to allow for a more dynamic reasoning about knowledge management in different contexts of online pervasive environments. To prepare the ground for organisational and technical design, we have indicated a direction for formalising these concepts using context graph notation. The next sections in this construction part of the paper will address the two remaining levels of the EIF framework we use to structure the design process.

Organisational design—defining data sharing policies

The ultimate goal of this DSR process is an implementable and context-aware privacy-preserving system, which is many design cycles away. However, even in the initial design phase, it is important to develop artefacts that give an idea of the direction of project. Business processes at the organisational level are part of this picture and are the design object in this section.

In the first stage of this DSR project, we will focus on *data sharing policies*, which is a key element in our guiding definition of data privacy (see section above). Data sharing policies act on behalf of the user by allowing exchange of data without intervening in or distracting from the primary activities of the user. Furthermore, these policies are statements that are directed towards receivers of personal data, describing the user's expectations and restrictions related to use of the data. Data controllers, e.g. universities running a learning management system, are a target group of data sharing policies, and they may use these policies in setting up their own systems and interacting with their users.

For the users, defining data sharing policies is part of their personal data management. This process should be non-intrusive, i.e. it should work behind the scenes; only to be activated as a negotiation process in two cases: when data transmission is about to violate the appropriateness defined by the user or when there is an event that triggers revisiting of the previously defined data sharing policy.

We have defined the following requirements for the process of initiating data sharing policies in a system:

- The system learns from data sharing practices related to the tools used;
- The system learns from event handling and is able to alert of potential threats based on prior actions or practices of the user or the community of practice she participates in related to different data sharing policies;
- The system can be tuned, i.e. the sensibility of alerts can be adjusted; and
- The data sharing policies are expressed in smart contracts that IT systems act upon on behalf of the user.

For the proposed solution to be adopted, a number of use cases need to be satisfied. The privacy model must be as meaningful for the users as for the institutions, the tool providers and other stakeholders. All these stakeholders should have agency and be motivated to take part in negotiation of the transmission principles, which will be written into the privacy policies. The business process should work in diverse cultural

settings; and furthermore, new technologies based on AI—in particular machine learning (ML)—should play a role. As stated in the ‘Introduction’ section, we do not see an option to rule out AI technologies and big data from future privacy solutions. We want a privacy process to be non-intrusive without blocking the opportunity to intervene when integrity is violated or threatened. This may be achieved by delegating support for event handling to ML and execution of the privacy policies to smart contracts.

Training the system to know what is appropriate—the role of ML

We have abandoned an approach following pre-programmed rules in favour of a contextual and dynamic approach using ML. This approach enables computers to perform specific tasks intelligently by learning from data, and the system continues to improve accuracy of results after it has been deployed (Shalev-Shwartz & Ben-David, 2014). However, a weakness of ML is that it is difficult to develop systems with contextual understanding of a problem, or ‘common sense’. ‘When our expertise fails, humans fall back on common sense and will often take actions, which while not optimal, are unlikely to cause significant damage. Current machine learning systems do not define or encode this behaviour, meaning that when they fail, they may fail in a serious or brittle manner’ (Royal Society, 2017, p.30). In our proposal, we suggest to make contexts the very object of ML, highlighting context triggers as the key concept for supervised machine learning. This implies we will need a certain amount of labelled data to train the system, and adaptiveness—further training—built into the system when going live. Ideally, the system should be able to know when to foreground a context trigger, based on the online activities of the user. So first, it must build a repertoire of events; next, it must learn what causes these events and what sensibility each user has towards these events.

Figure 8 describes how training data is used for machine learning, deployed in the proposed system, which in turn generates more data used to continuously improve the system.

ML will also play a role in managing the data sharing policies. It is out of scope of this paper to explore how all the diverse data sharing policies a user will meet could be reduced to a structured set of policies that could be added to and updated by the user in order to facilitate appropriate data sharing streams. However, we see that ML will be

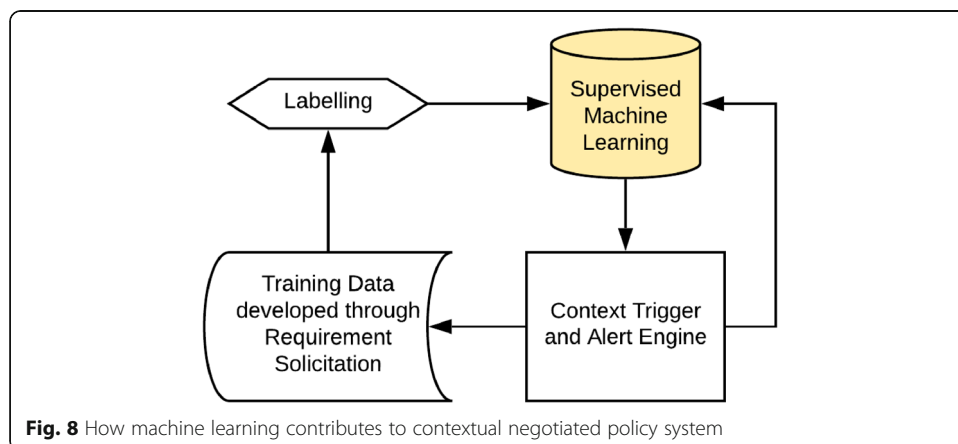


Fig. 8 How machine learning contributes to contextual negotiated policy system

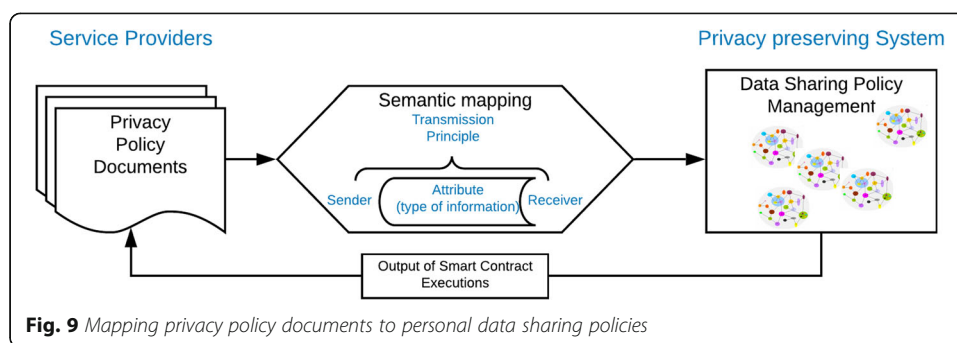


Fig. 9 Mapping privacy policy documents to personal data sharing policies

used to describe how service providers' privacy policies map to the extended CI model we propose. Shvartzshnaider et al. (2018) described how such a mapping could be done. In Fig. 9, we describe how policy documents are mapped to a semantic model, which can be modified to represent each user's preferred data sharing configuration. In turn, this structure is used to generate a smart contract (Lyons, Courcelas, & Timsit, 2018), which will guide data sharing and in the end influence the policies of service providers.

Smart contracts execute the data sharing policies the user subscribes to. As an example, these contracts may allow sharing of one's data with 3rd party companies that might be doing special analysis to be used by the service provider. However, if an alert goes off regarding one of these companies being part of a data breach scandal, a revision of the data sharing policy may lead to a change in the smart contract blocking further data transfer.

Process summary

It is a challenge for privacy engineering to design a process that gradually will move practice to a safer ground. All stakeholders need to see the benefits of new solutions and be able to influence them. This is why we have made negotiation of data sharing policies the starting point for organisational design.

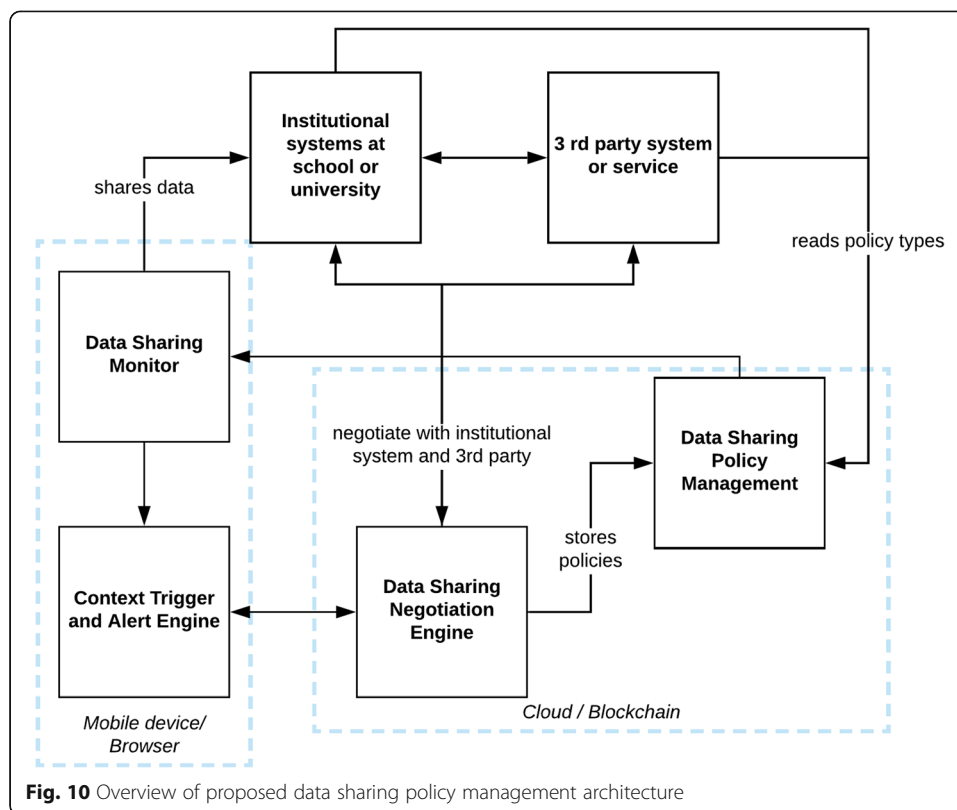
First, we have outlined a process of using the CI concepts to formalise the unstructured and long-winded privacy policy documents presented by service providers to construct an ontology that gives users a foundation to define their own data sharing policies.

Second, we have pointed towards a process of describing personal data sharing policy structures that are represented in executable scripts that handles sharing of personal data according to negotiated policies.

In this paper, we have only introduced the concept of privacy policies; we realise that there is more work to be done to outline in detail how these policies will be defined and formalised. In the next section, we will give a brief description of how a technical architecture can realise the proposed contextual model and high-level business principles.

Technical design—an application scenario

The technical design task at hand has similarities with coming up with an email spam detection system, but the scope is broader and more complex. In spam detection, ML



is used to train systems to recognise spam, and when deployed in live systems, the user is given a role to further train the system identifying incorrect classifications. One could imagine that today’s spam detection systems were extended to include spam policies, which were given roles of its own, living outside of the user’s e-mail system. This is the direction of the application scenario we present below.

Even if the main thrust of this paper is not on technical architecture, the overview of proposed system modules presented in Fig. 10 demonstrates the direction of development and will help in soliciting expert feedback to validate our proposal.

We suggest a system that consists of following parts:

- Data sharing monitoring
- Context trigger and alert system
- Data sharing policy negotiations
- Data sharing policy management
- Institutional systems
- 3rd party systems or services

In Fig. 10, these parts are integrated into an ecosystem for data sharing for a particular organisation (in the figure, an educational institution is used as an example).

Data sharing monitoring takes place on the user’s devices, and this engine shares the data with the institutional systems in the school or university, which in turn may share data with 3rd parties. The sharing happens according to policies stored outside of the device. On the device, the context trigger and alert engine runs, getting information

from the data sharing monitor and passing on information to the data sharing negotiation engine if there is a triggered event. Then the negotiation engine interacts with institutional systems and 3rd parties to set up or revise personal data sharing policies, which in turn are represented as smart contracts. Both artefacts are stored in the data sharing policy management system, which may be hosted on the blockchain or in the cloud. The institutional and 3rd party systems have access to the description of the data-sharing policies, but no personal information is transferred about who subscribes to particular policies.

Figure 10 describes an educational scenario, where, as an example, a student using a learning management system (LMS) in a university disagrees to share data with 3rd party companies that may use her data for profiling not relevant to her learning. Denying the institution to share data with the analytics company, she knows she will meet with an impaired version of the LMS next time she uses it. However, she feels it might be worth the change of policy, perhaps for a time, till the university comes up with a new practice of data sharing with 3rd parties. Another student doing the same deliberations (see use case exemplified in a context graph, Fig. 7) may decide to trust the institution and dissolve contextual knowledge about profiling threats into the external knowledge, for the time being.

This application scenario gives just a minimum of prescriptions about technologies. Ubiquitous online presence implies mobile devices and cloud services. We have hinted to use of smart ledger technology as smart contracts play an important role in use cases for blockchains. Smart contracts are capable of facilitating, executing and enforcing the negotiation or performance of an agreement (i.e. contract) without the use of intermediaries (Lyons et al., 2018). In this paper, we do not prescribe the use of blockchain technologies per se, but point to the fact that this direction would simplify trust and management issues related to the execution of the data sharing policies. In any case, there is a need to look into the steps necessary to turn policies into code that can be executed to do the switching of a person's data streams.

We see the following steps for the technical design of the solution:

- Developing a high-level view of the system architecture
- Designing requirement solicitation process for collecting training data
- Designing and deploying an app for collecting users' CI data
- Training context trigger and alert engine
- Training data sharing policy negotiation engine
- Designing a data sharing policy system

In summary, in this section, we have presented ideas for a high-level technical design that uses ML and other cutting edge technologies to develop a system that should work behind the scenes and only be called upon when privacy incidents happen. The main objective is to design a system that does not clutter one's screens with privacy alerts, but does its work in the background by learning the user's contextual preferences without leading the user down a tunnel of ignorance of the adverse consequences of data sharing choices.

Validation

The constructs proposed in this paper are the result of the first design cycle (Gregor & Hevner, 2013). There are many more cycles to go to develop the solution and to be able to prove that our proposal could change the direction of privacy engineering from the traditional individual preference based solutions. The aim of the first round of validation, interviewing 11 experts, is therefore to assure that we have a sound direction of the design and to solicit ideas for improvement of designed constructs using the lens of the different levels of interoperability defined in EIF (EU, 2012), focussing on semantic, organisational and technical aspects of design.

In analysing the recordings and transcripts of the interviews, the first question we asked ourselves was whether the interviewees bought into the idea of making context more prominent in privacy engineering? In a dialogue where the interviewer uses the artefacts (figures, concepts and definitions) to build up a storyline (Strauss & Corbin, 1990), interviewees react through speech acts of different kinds (affirmations, nods, pauses, interjections, etc). Noting these acts of approval, hesitance and disagreement, we conclude that the interviewed experts did find the direction of this research valuable. In this context, it is especially interesting when several interviewees used some of our proposed concepts to elaborate their position and to explain back how the presented constructs should be understood. In the group of 11, we had three experts who have been actively developing specifications for the preference approach that our research criticises. All three expressed support for what they were presented. 'Happy to see sound judgmental advice about 24751 [Access for All project] and I think it should be shared with...' 'I think your model would work well (implemented around objects) to deal with accessibility needs, their implementation and contexts...(..) I like your ideas for privacy very much'.

A theme-based analysis of the transcripts collected feedback on key concepts in this research. The paper posits contexts as dynamic and active, but what does it mean? 'Context is itself contextual, in some way, the principle of fluidity is at play', one interviewee said before mentioning complex adaptive systems and pointing to the Cynefin framework (Snowden, 2005) to describe the *habitat* of privacy engineering, stating that we were in the complex, if not the chaotic quadrant. The Cynefin framework was also mentioned by another interviewee who added: 'The difficulty is that what context you're in is a highly individual thing... you've bitten off one of the possibly biggest problems you can bite off, because it entrains all of human psychology.'

Even if the interviewees supported the proposed solutions to the semantic task of grasping the complexity of privacy through the concept of context they offered critical comments, the individual focus in the presented e proposed solutions to the semantic tascy' was questioned. 'Could you replace that process of an individual with an entity, being a group, or society, or an institution, or a team? Given that it is not only individuals that have to maintain privacy or confidentiality ' As a response, now in the aUnderstanding data privacy' section, the concept of a 'user' is introduced opening up for a more active role for also IT systems. The interviews made it clear that the support of machine learning and other intelligent technologies for privacy management was not controversial, at least among our experts.

At an organisational level, the interviews showed that there is a need for further work. The idea of delegating decisions and negotiations to data sharing policies was

supported as nobody wanted to be bothered with these issues on a daily basis; however, how is this system supposed to work? ‘What I hear you are saying—and I agree—is that privacy is not the property of the user or their preferences, it is actually something that is transactional, and maybe even conversational. (...) actually, the privacy concerns I have will vary substantially according to the topic, what we are talking about, what conversation we are having...’. This expert foresaw a system with a clear taxonomy of relationships, ‘which allows us to put pressure on people and legal systems to give us a vocabulary for ‘we want more of this one and less of that one’...’. Other interviewees drew a parallel with Creative Commons licencing of online content. ‘So, is what you are talking about, slapping a CC licence on every piece of information that goes out that tells you under what circumstances it could be reproduced?’

These comments highlighted a gap in the current research, the need to explore how policies could regulate data sharing. Are these policies organised as an ontology? How are adjustments based on individual preferences registered and enacted? These questions are out of scope for this paper, but need to be addressed to prove the viability of the proposal.

When walking through the technical architecture, another gap in the presented work became evident. ‘I think the architecture to some extent makes sense. The question is—there is an element of trust missing. How can I know that I can trust the system? Where is the trust block? How do I know this is to my benefit?’ In terms of the European Interoperability Framework, we have left out the legal level in our design. Some interviewees missed an outline of the fundamentals that restricted the processes described in the high-level architecture.

Overall, the interviewees saw the technical architecture as important to understand the proposed design. Even if the model indicated the use of cutting-edge technologies as blockchain, no expert objected to the idea. On the contrary, one interviewee even extended the design: ‘In your model – you have blockchain – that means the data are not really stored any more at the university - then the institution only have a token to your data - and you can revoke that token. As opposed to now, where I have to make a request to be erased’.

Table 1 summarises the implications and changes based on the interviews, and in the rest of this section, we will discuss directions for further development cycles as a result of the first validation.

The experts confirmed the complexity of the problem, reinforcing the need to be very clear about the ideas premising our proposals and to choose a design strategy with care. In the ‘Introduction’ section, we confessed a somewhat defeatist position to whether privacy can be vigorously pursued. In our research, we need to turn this into a requirement that directs the design of solutions that gradually work towards improved privacy, without alienating users nor service providers. In technical terms that means to specify how the user gradually may become more able to manage and control her own data sharing, which is not necessarily in the interest of companies. In information science, this is a classical design challenge balancing two problems, bootstrapping (meeting users’ need now) and adaptiveness (adapting to unforeseen and new needs) (Hanseth & Lyytinen, 2010).

The interviews also showed the need to be clear about trust. This is an issue that cuts through all interoperability layers. Many of the provisions in GDPR are so

Table 1 Summary of first validation by aspect

<i>Aspects raised in the interviews</i>	<i>Initial position</i>	<i>Implications of first validation</i>
Understanding of context	Dynamic entity defined by knowledge focus, not a container described by a set of characteristics (e.g. individual privacy preference statements)	Dynamic definition fruitful basis for the design
Need to extend CI theory	Context part of the theory underdeveloped	Development gap recognised
Knowledge aspects of context	Focus should be on the three types of knowledge (external, contextual and procedural)	Different aspects are understood applied to privacy use cases
Context triggers	Event-driven approach to handling privacy in context	Concept useful starting point for privacy engineering
Data sharing policy	Concept encapsulating preference handling on behalf of the user	Considered useful as an overall idea; however, many questions about structure and management not dealt with in the first development cycle
Contextual graph formalism	Graph presented as an abstract example (template)	Should be introduced in a pedagogical example related to privacy (see updated Fig. 7); the graph should be explained in relation to other graph types if it should be used in applications
Provisional definition of data privacy	Guiding definition for use in design was provided	The definition was improved (see the Understanding data privacy section) for clarity and scope
Organisational design	Focus on role of data sharing policies	Questions to the envisioned business process motivated extending the Organisational design—defining data sharing policies section, explaining more in depth the role of ML and the personal data sharing policies' relationship to institutional privacy policies
Technical design	An application scenario was presented	The role of high-level technical architecture and scenarios is highlighted; design at this level could potentially drive future design cycles
Use of smart contracts	These artefacts are part of cutting edge technologies, and in our proposal given, the role of executing data sharing decisions	According to interviews blockchain and smart contracts should be explored
Use of ML	ML is positioned as a key instrument in delegating the execution of policies to the IT system, allowing users to focus on their main activities	Interviews showed support for making ML an important part of the design

clear cut that they could be proved automatically by intelligent systems. Students seem to trust educational institutions strongly when it comes to handling of personal data (Komljenovic, 2019; Slade, Prinsloo, & Khalil, 2019), and institutional policies can be formalised so that more strict regimes for adherence could be established.

Discussion

Reflecting on this DSR, it is natural that the first phase of design is focussed on the development of artefacts and less on contributions to theories. We will discuss both aspects now looking at where to go next, starting with our engineering challenge.

Multi-level development of artefacts

The first round of validation showed that the issues of trust should be addressed explicitly in this research. Trust is an overarching concept that spans all levels of analysis we have used in our research, but also includes the legal or policy level we have left out. Even if the data sharing policies we have given a central role in our design hinge on the personally felt appropriateness of the transmission principle and its execution could be secured by smart contracts hosted on the blockchain, there is a need to also anchor trust at the societal and legal level.

At the organisational level, we assume that all relevant privacy and data sharing policies are available for semantic matching and that we could use ML to distinguish patterns in the data sharing policies. It remains to be tested how CI negotiations could be harnessed in data sharing policies and how well such policies could accommodate data sharing from different tools used by learners and institutions. We have foreseen a typology of principles that can be used to define smart and actionable contracts. It should also be tested whether the constructs that come with the extended CI theory could easily be turned into technical solutions.

Furthermore, at the technical level, blockchain technologies are being developed with the promise to eliminate some of the sources of ambiguity and conflict in domains where trust is essential (Lyons et al., 2018). We realise the need for further research to come up with a model of data sharing expressed in smart contracts that are based on laws and policies and described in a way that makes it possible for IT systems to decide whether a data stream from a user should be open or not.

The technical architecture presented in this paper proved important to explain the direction of this research. However, we realise that the high-level model does not answer a number of important questions related to the privacy of the user. How is the user identity managed by the institution? And how is the connection between user identity and data sharing policies observed by the institution? There is a scope for designing a number of more detailed models to see if it is possible to build a technical system that gradually can give the user more control over their personal data. It is also the scope for outlining the role of ML in the solution. User agency and transparency are key values to this project, and if the result is a 'black box' that is inscrutable to their users and developers, we have failed. Therefore, it should be tested if ML could be implemented in a way that fosters users' data literacy and understanding of contexts of data sharing.

We summarise the first cycle of design and the discussion of further directions in Table 2, presenting a first attempt to construct a conceptual development framework for privacy engineering making context the key concept of design.

While we describe in the table what we have done in the first development cycle, we indicate key aspects and research questions for the two next cycles. This is a dynamic framework as there will be rapid and minor design cycles and far more than three cycles before we have a working solution.

Design process

Both literature (Belanger & Crossler, 2011; Smith et al., 2011; Westin, 2003) and our interviews suggest that privacy is a complex and fuzzy field of research, something that

Table 2 Framework for privacy design, development cycles and levels

<i>Development cycle</i>	1st	2nd	3rd
<i>Key aspect</i>	Context (semantic development)	Trusted processes (organisational development)	Proof of concept (technical development)
<i>Policy/legal level</i>	(Not included)	What trust regime would integrate all interoperability levels?	How to engage policy level in development?
<i>Organisational level</i>	Process idea: Privacy as negotiation expressed in data sharing policies and executed by smart contracts. ML plays role in relieving the user of privacy tasks.	What process integrates institutional and company privacy policies, data sharing policies, and executable scripts regulating data streams? What role will ML play?	What application scenarios could change current practice without jeopardising the CI approach?
<i>Semantic level</i>	Privacy decisions triggered by events activating contextual knowledge. Data sharing policy.	Any new concepts needed?	Any new concepts needed?
<i>Technical level</i>	Modular application scenario	How to orchestrate a suite of semantic technologies that are able to transform privacy knowledge between levels (national judiciary domain, institutional domain, personal domain, tools, contracts)?	What series of self-contained and useful apps could be developed that proves key ideas of overall solution? Alert app, monitoring activity and triggering reflection on privacy? Self-storage solutions, moving data sharing control more towards user? Negotiation simulation app, using context trigger data and privacy policy ontologies.

asks for design principles and guidelines when doing privacy engineering. Reflecting on our own research, we see some ideas forming that could contribute to design theories.

First, we find the EIF (EU, 2012) used in Table 2 above useful for high-level structuring of development. Even if EIF is developed for another purpose, i.e. specifying how administrations, businesses and citizens communicate with each other within the European Union and across member state borders, the framework raises questions that are relevant also on an application level. The framework forces the developer to clarify political and legal context, specify concepts in use and processes, before embarking upon a technical design.

Second, our first validation has made us aware of the benefits of doing synchronous development at all four levels. Even if the initial development is very explorative and conceptual, focussing on business process ideas and basic constructs, we have seen the value of representing design ideas in technical application scenarios to be able to communicate ideas with the developer community. Externalisation of ideas in technical diagrams reveals design flaws when discussed with fellow experts. It would have been interesting to have a set of design templates to choose from in the more conceptual phase of design as these illustrations are more conversational artefacts than implementable technical drawings at this stage.

Third, complex issues are easier to grasp through examples. In discussing the idea of contextual graphs with our group of experts we learnt that examples, use cases, scenarios, etc., communicate much better than abstract concepts. And privacy engineering is all about communication.

Conclusions

Privacy engineering can be seen as the deliberate approach of interjecting data protection requirements into complex system development based on ethical national, institutional and corporate strategies (Kenny & Borking, 2002). The time of fast development of global and data-hungry solutions based on machine learning and analytics privacy is under pressure. As we have demonstrated in this paper, solutions based on the matching of ill-specified individual preferences with privacy-sensitive services of a myriad of data-driven companies are highly unrealistic. A new approach is needed, and we firmly believe that context negotiation is part of that approach. In this paper, we have contributed to a new understanding of privacy context, extending the theory of contextual integrity and pointing to a direction of development that uses machine learning as a technology to design solutions that work continuously and non-intrusively for the users. We have presented a condensed understanding of data privacy to give direction to the design of solutions that give context negotiations priority, but store decisions in data sharing policies for processing in the background.

The aim of this paper is to give privacy engineering a new direction. It has long been stuck in a quagmire of politicised discourse, dominated by Western centric privacy theories (Hoel & Chen, 2019). To support global system development, there is a need to realise that online practices are surprisingly similar around the world, but our understanding of the room for manoeuvring may be different. To see how actual negotiations of data sharing practices would take place in different contexts, we need to establish a semantic, organisational and technical framework that allows comparisons between cultures to happen. To make context a first-class citizen in privacy engineering is essential to move forward.

As work situated in DSR, we acknowledge that the first results give moderate contributions to design theory, having focussed on the development of artefacts that in the end will have practical application. As Peffers, Tuunanen, and Niehaves (2018) have pointed out, this is nothing new in DSR, where artefacts with value in a system or system component often are the main, or at least, the initial aim of researchers. We have, however, in this paper kept an eye on the design process itself in order to make observations that could be useful to inform design guidelines for the nascent field of privacy engineering.

Abbreviations

CI: Contextual integrity; DSR: Design Science Research; EIF: European Interoperability Framework; GDPR: General Data Protection Regulation; HCI: Human–computer interaction; LA: Learning analytics; ML: Machine learning; PbD: Privacy by design

Acknowledgements

The authors want to thank the participants that took part in this research.

Authors' contributions

Tore Hoel: conceptualization, investigation, methodology, validation, visualisation and roles/writing—original draft. Weiqin Chen: conceptualization, methodology and supervision. Jan M. Pawlowski: Methodology and supervision. The author(s) read and approved the final manuscript.

Funding

The authors have not received any funding for this work.

Availability of data and materials

Not applicable

Competing interests

None of the authors have competing interests related to this paper.

Author details

¹Oslo Metropolitan University, Oslo, Norway. ²HRW University of Applied Sciences, Mülheim, Germany.

Received: 12 May 2020 Accepted: 17 September 2020

Published online: 19 October 2020

References

- Badillo-Urquiola, K., Yao, Y., Ayalon, O., Knijnenburg, B., PAGE, X., Toch, E., et al. (2018). *Privacy in context* (pp. 425–431). Presented at the Companion of the 2018 ACM Conference. New York: ACM Press. <https://doi.org/10.1145/3272973.3273012>.
- Barkhuus, L. (2012). *The mismeasurement of privacy: using contextual integrity to reconsider privacy in HCI. the 2012 ACM annual conference*, (pp. 367–376). New York: ACM. <https://doi.org/10.1145/2207676.2207727>.
- Barth, A., Datta, A., Mitchell, J. C., & Nissenbaum, H. (2006). *Privacy and contextual integrity: framework and applications*, (pp. 15–198). Presented at the 2006 IEEE Symposium on Security and Privacy, IEEE. <https://doi.org/10.1109/SP.2006.32>.
- Baruh, L., Secinti, E., & Cemelcilar, Z. (2017). Online privacy concerns and privacy management: a meta-analytical review. *Journal of Communication*, 67(1), 26–53. <https://doi.org/10.1111/jcom.12276>.
- Baskerville, R., Baiyere, A., Gergor, S., Hevner, A., & Rossi, M. (2018). Design science research contributions: finding a balance between artifact and theory. *Journal of the Association for Information Systems*, 19(5), 358–376. <https://doi.org/10.17705/1jais.00495>.
- Bastien, C. (1999). Does context modulate or underlie human knowledge? In A. C. Quelhas, & F. Péreira (Eds.), *Cognition and Context*. Lisbonnes: Análise psicológica.
- Bazire, M., & Brézillon, P. (2005). Understanding context before using it. In *Business Process Models. Change Management*, (vol. 3554, pp. 29–40). Berlin, Heidelberg: Springer Berlin Heidelberg. https://doi.org/10.1007/11508373_3.
- Belanger, F., & Crossler, R. E. (2011). Privacy in the digital age: a review of information privacy research in information systems. *MIS Quarterly*, 35, 1–63.
- Benthall, S., Gürses, S., & Nissenbaum, H. (2017). Contextual integrity through the lens of computer science. *Foundations and Trends® in Privacy and Security*, 2(1), 1–69. <https://doi.org/10.1561/3300000016>.
- Berners-Lee, T. (2017). Three challenges for the web, according to its inventor. Online: <https://webfoundation.org/2017/03/web-turns-28-letter/>
- Brézillon, P. (2003). Representation of procedures and practices in contextual graphs. *The Knowledge Engineering Review*, 18(2), 147–174. <https://doi.org/10.1017/S0269888903000675>.
- Brézillon, P. (2005). Task-realization models in contextual graphs. in business process models. *Change Management* (Vol. 3554, pp. 55–68). Berlin, Heidelberg: Springer Berlin Heidelberg. doi:https://doi.org/10.1007/11508373_5
- Brézillon, P., & Pomerol, J.-C. (1999). Contextual knowledge sharing and cooperation in intelligent assistant systems. *Le Travail Humain*, 62(3) Paris: PUF, 223–246.
- Campbell, R., Robinson, W., Neelands, J., Hewston, R., & Massoli, L. (2007). Personalised learning: ambiguities in theory and practice. *British Journal of Educational Studies*, 55(2), 135–154.
- Cavoukian, A. (2009). Privacy by design—the 7 Foundation Principles. Online: <https://www.ipc.on.ca/wp-content/uploads/Resources/7foundationalprinciples.pdf>. Accessed: 2019-12-09
- Choi, H., Park, J., & Jung, Y. (2017). *The role of privacy fatigue in online privacy behavior, computers in human behavior*. <https://doi.org/10.1016/j.chb.2017.12.001>.
- Durbin, R. J., Markey, E.M., & Blumenthal, R. (2019). Letter to Mr. Sundaar Pichai, August, 12, 2019. Online: <https://iblnews.org/senators-go-after-edtech-top-players-on-student-data-collection-practices/>.
- EU (2012). *Regulation of the European Parliament and of the council on the protection of individuals with regard to the processing of personal data and on the free movement of such data, COM(2012) 11 final*. Brussels: European Commission.
- Fluid Project. (n.d.). Project description: Understanding, Discovering and Asserting Personal Privacy Preferences (UDAPPP). [https://wiki.fluidproject.org/display/fluid/\(Floe\)+Privacy+Needs+and+Preferences](https://wiki.fluidproject.org/display/fluid/(Floe)+Privacy+Needs+and+Preferences)
- Gregor, S., & Hevner, A. R. (2013). Positioning and presenting design science research for maximum impact. *MIS Quarterly*, 37(2), 337–355.
- Hanseth, O., & Lyytinen (2010). Design theory for dynamic complexity in information infrastructures: the case of building internet. *Journal of Information Technology*, 25(1), 1–19. <https://doi.org/10.1057/jit.2009.19>.
- Hoel, T., & Chen, W. (2019). Privacy engineering for learning analytics in a global market: defining a point of reference. *The International Journal of Information and Learning Technology*. <https://doi.org/10.1108/IJILT-02-2019-0025>.
- ISO. (2008). Information technology -- individualized adaptability and accessibility in e-learning, education and training -- part 1: framework and reference model (ISO/IEC 24751-1:2008). International Organization for Standardization. Retrieved from <https://www.iso.org/standard/41521.html>
- Kenny, S. & Borking J. (2002). The value of privacy engineering. *The Journal of Information, Law and Technology* (IJLT). Online: <http://elj.warwick.ac.uk/ijlt/02-1/kenny.html>
- Komljenovic, J. (2019). Making higher education markets: trust-building strategies of private companies to enter the public sector. *Higher Education*, 78(1), 51–66. <https://doi.org/10.1007/s10734-018-0330-6>.
- Lahlou, S., Langheinrich, M., & Rucker, C. (2005). Privacy and trust issues with invisible computers. *Communications of the ACM*, 48(3).
- Lyons, T., Courcelas, L., & Timsit, K. (2018). Blockchain and the GDPR. Report produced by ConsensSys AG on behalf of the European Union Blockchain Observatory and Forum.
- Mansour, R. F. (2016). Understanding how big data leads to social networking vulnerability. *Computers in Human Behavior*, 57(C), 348–351. <https://doi.org/10.1016/j.chb.2015.12.055>.
- Mostéfaoui, G. K. & Brézillon, P. (2004). Modeling context-based security policies with contextual graphs. Proceedings of the Second IEEE Annual Conference on Pervasive Computing and Communications Workshops (PERCOMW'04) 0-7695-2106-1/04
- Nature (2019). Protect AI panel from interference. *Nature*, 572, 415 Editorial published 22 August 2019.
- Nissenbaum, H. (2004). Privacy as contextual integrity. *Washington Law Review*, 79, 119–157.
- Nissenbaum, H. (2010). *Privacy in context: technology, policy, and the integrity of social life*. Stanford: Stanford University Press.

- Norberg, P. A., Horne, D. R., & Horne, D. A. (2007). The privacy paradox: personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs*, 41(1), 100–126. <https://doi.org/10.1111/j.1745-6606.2006.00070.x>.
- Peffers, K., Tuunanen, T., & Niehaves, B. (2018). Design science research genres: introduction to the special issue on exemplars and criteria for applicable design science research. *European Journal of Information Systems*, 27(2), 129–139. <https://doi.org/10.1080/0960085X.2018.1458066>.
- Prain, V., Cox, P., Deed, C., Dorman, J., Edwards, D., Farrelly, C., et al. (2012). Personalised learning: lessons to be learnt. *British Educational Research Journal*, 65(3), 1–23. <https://doi.org/10.1080/01411926.2012.669747>.
- Royal Society (2017). *Machine learning: the power and promise of computers that learn by example*. ISBN: 978-1-78252-259-1. United Kingdom: The Royal Society.
- Scott, M. (2006). *Programming language pragmatics*. San Francisco: Morgan Kaufmann Publishers.
- Shalev-Shwartz, S., & Ben-David, S. (2014). *Understanding machine learning: from theory to algorithms*. Cambridge: Cambridge University Press.
- Sheeran, P. (2002). Intention behavior relations: a conceptual and empirical review. *European Review of Social Psychology*, 12, 1–36.
- Shvartzshnaider, Y., Apthorpe, N., Feamster, N., & Nissenbaum, H. (2018). Analyzing privacy policies using contextual integrity annotations. Online: <https://arxiv.org/pdf/1809.02236>. Accessed: 2019-12-11.
- Slade, S., Prinsloo, P., & Khalil, M. (2019). *Learning analytics at the intersections of student trust, disclosure and benefit*, 1–1. Tempe: Proceedings of the 9th Learning analytics and Knowledge Conference 2019 (LAK 19).
- Smith, H. J., Dinev, T., & Xu, H. (2011). Information privacy research: an interdisciplinary review. *MIS Quarterly*, 35(4), 989–1015.
- Snowden, D. J. (2005). Complex acts of knowing: paradox and descriptive self-awareness. *Bulletin of the American Society for Information Science and Technology*, 29(4), 23–28. <https://doi.org/10.1002/bult.284>.
- Strauss, A., & Corbin, J. (1990). *Basics of qualitative research*. Sage publications.
- Taddei, S., & Contena, B. (2013). Privacy, trust and control: which relationships with online self-disclosure? *Computers in Human Behavior*, 29, 821–826.
- Tuomi, I. (2018). The impact of artificial intelligence on learning, teaching, and education. Policies for the future, Eds. Cabrera, M., Vuorikari, R & Punie, Y., EUR 29442 EN, Publications Office of the European Union, Luxembourg, ISBN 978-92-79-97257-7, doi:10.2760/12297, JRC113226.
- Westin, A. F. (2003). Social and political dimensions of privacy. *Journal of Social Issues*, 59(2), 431–453. <https://doi.org/10.1111/1540-4560.00072>.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► [springeropen.com](https://www.springeropen.com)
