Chapter 13
Context-Adaptive Privacy Mechanisms

Florian Schaub

Abstract Sensing and context awareness are integral features of mobile computing and emerging Internet of Things systems. While context-aware systems enable smarter and more adaptive applications, they also cause privacy concerns due to the extensive collection of detailed information about individuals and their behavior, as well as the difficulties for individuals to understand and manage information flows. However, context awareness also holds significant potential for supporting users in managing their privacy more effectively. Context-adaptive privacy mechanisms can inform users about how changes in context may impact their privacy, recommend privacy-preserving actions tailored to the respective situation, as well as automate certain privacy configuration changes for the user. This chapter provides an overview of research on context-adaptive privacy mechanisms, including an introduction to context-aware computing and the context dependency of personal privacy; a discussion and model for operationalizing context awareness for privacy management, including privacy-relevant context features; as well as an overview of existing context-adaptive privacy mechanisms with various applications. The chapter concludes with a discussion of research challenges for context-adaptive privacy mechanisms.

13.1 Introduction

Context awareness is an essential aspect of mobile computing and the emerging Internet of Things. Today's smartphones, vehicles and other "smart" devices include a multitude of sensors that allow devices and respective applications to determine physical location, orientation, ambient noise, light levels, and many other context features; as well as collect this information periodically or continuously. Such context information can be leveraged to infer user behavior, activities, mobility

F. Schaub (✉)
University of Michigan, Ann Arbor, MI, USA
e-mail: fschaub@umich.edu

A. Gkoulalas-Divanis, C. Bettini (eds.), *Handbook of Mobile Data Privacy*,
https://doi.org/10.1007/978-3-319-98161-1_13

patterns, emotions or mood, as well as learn a user's interests and preferences. Situational awareness gained this way facilitates the adaptation of systems and applications to align with the user's context and appropriately support the user in their activities.

While context-awareness enables smarter and more adaptive technology, the extensive collection of sensor, context and mobility data has implications for personal and information privacy as discussed in prior chapters. The increased sensing, processing, and sharing of detailed information about users and their context further make it inherently difficult for individuals to determine who has access to information, let alone effectively control information flows. This creates inherent user interaction and usability challenges for solutions that aim to provide users with effective information and controls for privacy management.

Sensing and context-aware systems are often collecting data continuously [87, 122]. At the same time, the number of situations and entities potentially requiring privacy decisions and configuration increase constantly. This creates a scaling issue for privacy self-management [141], as it becomes unrealistic to correctly specify privacy settings for each system or situation in advance. Furthermore, the shift towards recording mundane activity and mobility information rather than specific events makes it difficult to grasp potential privacy implications of information collection [87]. Yet, advances in data mining and information retrieval make formerly ephemeral activities more accessible [1, 25, 80, 87, 141] and facilitate profiling through discovery of new patterns and knowledge by combining information from multiple sources [43]. Users may have inconsistent mental models of the capabilities and data practices of systems [117], which hampers their ability to predict what information is actually collected or to whom it is disclosed [2]. Long-term privacy implications of decisions and actions are typically hard to foresee without appropriate support [120], yet, typically "decisions about privacy must be made individually, in isolation, and far in advance" [141]. Users may also not realize that data access once authorized is still active in other situations, or that information collection may occur in unanticipated contexts [10, 19].

However, context awareness and privacy do not have to be mutually exclusive. Prior chapters presented and discussed methods to mitigate privacy issues associated with context and mobility data collection in multiple domains. Furthermore, context awareness can also be leveraged to actively protect privacy and support users in managing their privacy more effectively [128]. Context-adaptive privacy mechanisms can inform users how changes in context may impact their privacy, recommend privacy-preserving actions tailored to the respective situation, as well as automate certain privacy configuration changes for the user.

In this chapter,[1] we first introduce context-aware computing in more detail (Sect. 13.2), before discussing how context information can be leveraged in privacy management (Sect. 13.3). We further provide an overview of different types of

[1]Parts of this chapter have appeared in the author's doctoral dissertation [123] and a prior article [128]. This chapter provides an expanded and revised overview of research on context-adaptive privacy mechanisms.

existing or proposed context-adaptive privacy mechanisms (Sect. 13.4) and discuss research challenges in this domain (Sect. 13.5).

13.2 Context-Aware Computing

Context awareness in technology has been studied extensively. Subsequently, we provide a brief introduction to context awareness and context-aware computing from the perspective of ubiquitous computing research.

13.2.1 Defining Context Awareness

In 1987, as part of a critique of the artificial intelligence paradigm of planning, Suchman argued that computer systems should respond to the settings in which they are used [144]. Schilit and Theimer first introduced the term *context-aware computing* in relation to human-computer interaction and ubiquitous computing [129, 130]. They named the current location, other persons in the vicinity, and nearby resources as important aspects of context.

Schmidt [132] distinguishes multiple categories of context-aware computing applications. *Context-adaptive systems* are systems that perform actions when certain context conditions are met. *Adaptive and context-aware user interfaces* dynamically adjust to services and resources available in the current context and make them available to the user. *Context-aware resource management* dynamically maps system functions onto context features and available services. Context-awareness can further facilitate the *management of interruptions based on situations*. Schmidt's *sharing context* category encompasses applications that exchange context information between different systems or users. The category *metadata generation and implicitly user-generated content* reflects the idea that context information can serve as metadata to enrich created information or even as implicitly generated content on its own. For instance, location and mobility data collected via smartphones, connected vehicles or sensing infrastructure with the goal of improving traffic prediction models, maps, or localization accuracy. In context-aware computing, multiple context factors are typically combined to increase the accuracy of collected context information [14] and infer the current situation [1].

The diversity of context-aware applications raises the question what constitutes context. Dey defines context as "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" [44]. Abowd and Mynatt identified five general dimensions to describe context [1]:

- **Who.** What persons and entities are present in the user's or system's proximity.
- **What.** What are the current activities of the user, present entities, and systems.

- **Where.** What is the current location of the user, the system, or the application.
- **When.** What is the point in time to which the context relates to.
- **Why.** What are the user's reasons and intentions behind an activity.

Context-aware applications might examine the first four dimensions (*who*, *what*, *where*, *when*) in order to determine the user's intentions (*why*) and initiate system actions that support and satisfy these intentions [45, 46]. While the employed terms suggest that context is mainly based on physical aspects, context information can encompass physical, social, emotional, as well as informational aspects [46]. For example, the *who* dimension can also extend to virtually present entities or services. Thus, an entity can be any person, application, service, or object of relevance [44]. Baldauf et al. [13] distinguish between *physical*, *virtual*, and *logical sensors*. Physical sensors measure real-world phenomena, such as temperature, pressure, light intensity, or radio signal strength. Virtual sensors provide access to digital information, such as calendar data, emails, contacts, or social media posts. Logical sensors combine output from multiple sensors and information sources, also called sensor fusion [133], to obtain higher-level abstractions of the current context.

Consequently, a context-aware application can leverage such context information to provide the user with information or services that are relevant to the user's activities in a specific context [44]. A context instance can be described as a *situation* [134], which is defined as a set of states of relevant context features or entities [44]—a snapshot of a specific context configuration. Different situations can be distinguished by their specific set of values for the relevant context features.

Context information can exist at different levels of abstraction and interpretation. Situations composed of sensor-based cues can be seen as low-level context requiring further interpretation to be useful [21]. High-level context can be obtained through context reasoning and interpretation of available sensor information and, thus, enriching situations with semantic interpretations. An advantage of this view on context is that it becomes possible to distinguish between low-level and high-level context changes. Thus, applications, including context-adaptive privacy mechanisms, can adapt to high-level context changes, while ignoring low-level context changes that do not affect the higher level interpretation.

Situations do not exist in isolation. By viewing context as a process rather than a state, an information space can be modeled as a directed state graph, in which each node represents a different situation and the edges between the nodes are annotated with the conditions of the context change [38]. Modeling such relations can reduce the search space for situation recognition [21], but also requires knowledge of the changing conditions and potential situations.

An important aspect of context-aware computing is the quality of obtained context information. Sensor information can be inaccurate, incomplete, and noisy. Therefore, context-aware applications and systems must take uncertainty into account, e.g., by assigning confidence metrics to context values that represent the estimated likelihood of the value reflecting reality [59, 133]. The assessment of context quality makes it possible to measure improvements in context quality, e.g., with multi-sensor fusion, or to infer high-level context under consideration of

potential ambiguities of the low-level context [21, 45]. Many context-aware systems reason over uncertain context by combining different reasoning approaches that support uncertainty [21].

13.2.2 Modeling Context

Context frameworks and middleware have emerged to ease development of context-aware applications and reuse of context information [74]. Most context frameworks support sensor fusion and aggregation of context information, as well as interpretation of raw context data. Many context-aware systems and context middleware follow a similar layered conceptual approach [13, 74], as depicted in Fig. 13.1. Context data is retrieved from physical and virtual sensors. Context processing components interpret the retrieved sensor data. As a first step, logical sensors may combine context information from multiple sources. Reasoning and interpretation of context information provide semantic interpretations or support the transformation of context data into different representations. A context management layer retains historical context information and provides interfaces for applications to access and process context information. Applications either pull context information from the context system or subscribe to specific context updates and events.

Context models are an integral part of context-aware systems. Context models reduce the complexity of context-aware applications by separating representation and evaluation of context information from application logic. They govern preprocessing, interpretation, and representation of context information, improve extensibility of context-aware applications [21] and facilitate exchange and reuse of context information. Formal context models further enable consistency checking

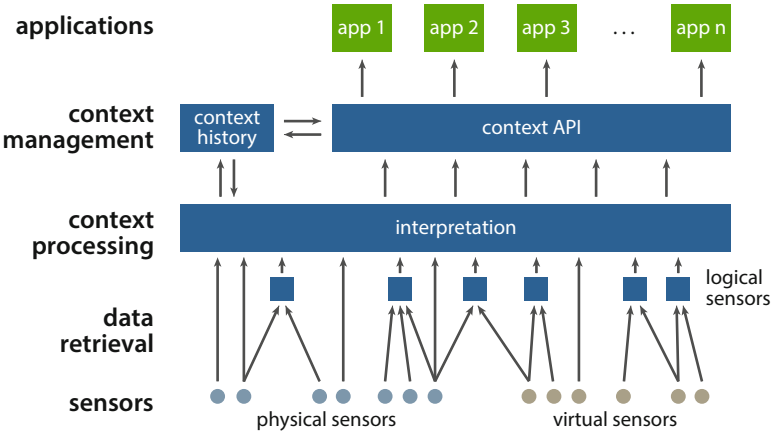

Fig. 13.1 Conceptual view of a context-aware system based on the layer models by Baldauf et al. [13] and Knappmeyer et al. [74]

and support inference of high-level context knowledge [21]. A basic context item can be modeled with five parameters [13]: its *semantic data type* (e.g., location or temperature), the current context *value*, a *time stamp* when the context item was acquired, the *source* it was acquired from (e.g., sensor id), and a *confidence* value that reflects the uncertainty of the measured context item. Bettini et al. [21] further note that a comprehensive context model should be able to represent context information on *different abstraction levels*, model *relations and dependencies between context information*, and support *efficient reasoning with uncertain information*. Furthermore, *context history* should be stored to support adaptation, *modeling formalisms* should ease modeling of real world concepts and support efficient context provisioning. Many models also distinguish between *primary and secondary context* [21]. Primary context information serves to index context situations and enables efficient access, while additional information is considered secondary context. Often location and time are used as primary context [1, 21].

While many different context models have been proposed, existing context models are characterized by a small set of major categories [21, 143]. *Key-value models* are basic context models that represent context as key-value tuples. They provide a flat view on context data. *Hierarchical markup models* can represent hierarchical data structures. Hierarchical markup models are typically represented in an XML dialect, such as RDF [143]. *Object-oriented context models* can be realized in object-oriented programming languages. Such models also enable hierarchical representation of context information, and additionally enable encapsulation [143]. They can combine representation and processing of context information, which can provide advantages for application-specific context models but reduces reusability.

These basic modeling approaches have a number of limitations according to Bettini et al. [21]. Due to pre-defined schemas, such models are limited in the variety of context information they can capture and have limited capabilities for expressing relationships and dependencies between context items. Quality of context information can be included but is often not expressive. Typically, basic models are also limited in terms of consistency checking, reasoning, and inference of high-level context abstractions.

Graphical modeling approaches, such as the Unified Modeling Language (UML) or the more specialized Context Modeling Language (CML) [58, 59], combine expressiveness and formality with ease of specification [143], especially for the modeling of relationships between context items. *Ontological models* leverage knowledge representation methods to describe relationships between context items, as well as the semantics of those relationships and context items. The formalization of context semantics enables consistency checking, the inference of context abstractions, and the derivation of new knowledge from asserted facts with semantic reasoning tools [21]. Ontological models facilitate knowledge sharing and interoperability between different context systems, because semantics of context items are represented in the models. Ontological context models can be described in OWL DL [21].

While ontological context models are most expressive [143], they may not be well suited for the representation of dynamic context information and corresponding adaptation preferences, due to performance limitations of online ontological reasoning and inadequate support for uncertainty [21]. As a consequence, Bettini et al. [21] advocate a *hierarchical hybrid model* that utilizes different context representations with varying levels of formal specification on different layers. The first layer performs low-level sensor fusion of raw sensor data. The second layer provides a basic context representation, e.g., based on markup or RDF, and supports efficient context reasoning to infer context abstractions. A third layer defines context semantics in an ontological model. Applications can choose the context representation layer that is most suitable to their requirements. For instance, Henricksen et al. combine CML with OWL to enable reasoning on uncertain information as supported by CML with OWL's semantic reasoning capabilities [58].

In addition to a context model, context-aware adaptation typically also requires the modeling of user characteristics and preferences to enable personalized adaptation. User characteristics can be integrated into a context model [67] or be maintained in a separate user model [31, 111]. Depending on the application, relevant user characteristics may include personal characteristics, the user's role, user preferences, user tasks and social relations [67, 111].

13.2.3 Privacy Protection in Context-Aware Systems

Context information in itself is often privacy sensitive, because it not only supports context-aware adaptation but can also facilitate undesired user profiling [2, 4, 64, 87]. Privacy risks of mobility data have been discussed in detail in the chapters "Privacy risks and inferences with mobility data" by Gambs and "Privacy in location-sensing technologies" by Solti et al. To address these risks of mobility and context data collection, privacy protections for context information have been studied extensively [22, 62]. Especially privacy of location information has received considerable attention as discussed in multiple chapters in Part II of this book. Here, we provide a short introduction on mechanisms and research directions for privacy protection in context-aware systems.

General privacy protection approaches can be applied to and adapted for context-aware systems. This includes strict access control, obfuscation of data through generalization or addition of noise, anonymization and de-identification methods, as well as private information retrieval and privacy-preserving data mining [22]. Privacy engineering and privacy by design principles facilitate the design of context-aware applications and systems that can meet both data quality and privacy requirements [40]. Heiber and Marrón [56] propose a privacy threat modeling framework for context-aware systems, which consists of a *data model*, an *adversary model*, and *inference rules*. The data model describes what context information is

available and the adversary model defines what information the adversary could gain access to. The inference rules are a set of rules that the adversary can apply to obtain data, for example, linking or matching of context items. The framework enables evaluation of the amount of information that could be potentially gained by adversaries with varying capabilities against previously defined privacy requirements for the context-aware system or application.

Privacy extensions for context models and systems have been proposed to protect privacy-sensitive context information [13], typically centering on access control. For instance, the CML privacy extensions [60] enable expression of ownership of context facts, object types, fact types, and situations, as well as corresponding usage preferences. Rei is a privacy policy language for the CoBra context middleware [69]. Rei privacy policies govern actions by defining *rights*, *prohibitions*, *dispensations*, and *domain-dependent policies*. Available actions are pre-defined in an ontology. For example, location sharing is defined by an action that describes what to share (the location) with whom (a set of recipients). Corresponding privacy policies govern what entity can perform this action. The *info spaces* approach [68] supports access control and privacy management. An info space has a *user*, or a group of users, a *user agent* that handles privacy enforcement for the user, an *owner* that defines permissions for the info space, and a set of *information objects*, which are subject to authorization. Privacy policies are enforced in the info space system when accordingly tagged information crosses info space borders.

The info space approach further introduces support for adapting the granularity of context information. Similarly, Wishart et al. [154] extend CML with granularity support by representing granularity for specific data types as a hierarchical ontology. In their approach, privacy preferences are evaluated first, then granularity preferences are applied to the context information instance before disclosure. They later added dynamic discovery and processing of context sources with declarative rules [155] in order to form a complete privacy-aware context management system based on *context ownership* [60], *privacy and granularity preferences*, and *dynamic handling of disclosure requests*. Pareschi et al. [113] propose semantic aggregation based on local context in order to provide high quality of service while preserving privacy. Information from individual users is aggregated into *stereotypes* in order to enhance privacy by generalizing quasi-identifiers in order to reduce information that could identify the individual. Sheikh et al. [138] draw a connection between the quality of context and its privacy sensitivity. They propose that applications should only receive context information with a granularity that corresponds to the required quality of context.

Bettini and Riboni [22] caution that while hierarchical context models inherently support generalization of context facts—for instance generalizing location data to the city or region level—hierarchical levels of specificity are not necessarily expressions of sensitivity and that continuous data streams pose further challenges for maintaining privacy. Potential solutions are discussed in multiple chapters in this book.

13.3 Leveraging Context in Privacy Management

Privacy expectations and privacy behavior in social interactions have been shown to be subject to dynamic adaptation processes [6, 11, 101]. This dynamism of privacy is often not sufficiently supported in computing systems. Many systems and applications allow a priori configuration of static privacy settings, but do not support dynamic adaptations of those settings to meet the user's privacy expectations in different situations. Context-aware systems, however, have the ability to dynamically adapt to changes in the user's context, environment, and activities. Such context awareness also holds significant potential for dynamically supporting users in managing their privacy [123, 128].

13.3.1 Privacy is Contextual and Dynamic

Throughout their days, individuals constantly adjust their privacy expectations and their sharing behavior based on their activities and surroundings [149]. For example, the amount of information revealed in a conversation depends on who one is talking to, the topic of the conversation, and who else is around. Individual privacy expectations and perceptions of privacy infringement are highly contextualized and shaped by individual, social, and cultural expectations and norms [6, 103, 109].

13.3.1.1 Contextual Integrity

Marx introduced the notion of *personal border crossings* to characterize privacy violations [103]. He argues that privacy expectations are shaped by cultural and individual boundaries. *Natural borders* (e.g., walls and clothes) limit what can be perceived by others. *Social borders* reflect expectations in the roles of persons, e.g., lawyers and doctors keeping client and patient information confidential. *Spatial and temporal borders* separate disjoint events and episodes of life. They reflect the expectation that such events are not linked. *Ephemeral and transitory borders* reflect the expectation that fleeting moments are not recorded. If such borders are breached privacy expectations are being violated and the action that caused the breach is perceived as privacy infringing. For instance, when a user's location traces are used to infer socio-economic status or behavior patterns for targeted advertising.

Nissenbaum expands this perspective by framing privacy as *contextual integrity* [108, 109]. Privacy expectations are shaped by context-relative norms of information flow. The context considered in contextual integrity is elaborate and nuanced, going beyond the primarily sensor-oriented context common in context-aware systems. Context-relative norms of information flow are characterized by contexts, actors, attributes, and transmission principles [109].

Contexts—in the framework of contextual integrity—encompass the general institutional and social circumstances of a situation (e.g., healthcare, education, family, religion, etc.), the activities in which actors engage, as well as the purposes, goals, and *values* associated with those activities. Nissenbaum notes that individuals often engage in multiple such contexts at the same time which can be associated with different, potentially conflicting informational norms. For instance, talking about private matters at work in a specific society and culture [109].

Actors are senders, receivers, and information subjects who participate in activities and contexts. Actors have specific roles and capacities depending on the context. Roles define relationships between various actors, which express themselves through the level of intimacy, expectations of confidentiality, and power dynamics between actors [109]. Informational norms regulate information flow between actors.

Attributes describe the type and nature of the information being collected, transmitted, and processed. Informational norms render certain attributes appropriate or inappropriate in certain contexts. The concept of appropriateness in Nissenbaum's framework serves to describe what are acceptable actions and information practices.

Transmission principles constrain the flow of information between entities. They are associated with specific expectations. Typical transmission principles are confidentiality, reciprocity or fair exchange of information, and whether an actor deserves or is entitled to receive information.

Context-relative norms may be explicitly codified or only implicitly established. Common types of norms are morals, conventions of etiquette, rules, and procedures. Information flows that violate respective norms are perceived as privacy violations by individuals. Furthermore, technology may affect moral and political factors, e.g., power structures, fairness, or social hierarchies; as well as impact goals and values in a specific context.

The aspect of informational norms is also apparent in the concept of *collective information practices* proposed by Dourish and Anderson [48]. In their view, information flows not only transmit information but also serve as social boundaries, which help to define identity, membership, and affiliation in social groups. The acceptance and utilization of the same information practices shapes a group's identity.

13.3.1.2 Privacy Regulation Theory

Contextual integrity provides a framework for understanding privacy expectations in social contexts, and identifying privacy issues of information technology. Because privacy expectations vary with context, privacy regulation in social interactions occurs in a continuous adaptation process in which individuals balance their personal privacy needs with their desire for disclosure [149]. Understanding this process is essential for designing context-adaptive privacy mechanisms to effectively support it.

Altman's *privacy regulation theory* [11] describes privacy as a dynamic, dialectic, and non-monotonic process. In this process, individuals regulate what they disclose (*outputs*) and what level of potential intrusion they are willing to accept (*inputs*) based on internal changes (e.g., changes in personal preference, past experiences, or new knowledge), as well as external changes in the environment and current context. In social interaction, adjustments rely on *verbal*, *paraverbal*, and *nonverbal behavioral mechanisms*, such as revealing or omitting information (verbal), changing intonation and speaking volume (paraverbal), or using posture and gestures to non-verbally express and control personal space and territory.

A critical part of Altman's theory is the distinction between *desired privacy* and *achieved privacy*. Individual privacy preferences and privacy expectations may differ from the level of privacy obtainable in a given situation with the available privacy control means. If achieved privacy is lower than desired privacy, privacy expectations are violated and the individual feels exposed. Achieving more privacy than desired causes *social isolation*. Thus, the privacy regulation process aims for an optimal privacy level in which desired and achieved privacy are aligned.

Validation studies have shown that Altman's theory can be considered a realistic model of individual privacy adaptation behavior [101]. Despite its focus on privacy regulation in social interactions, Altman's theory suggests itself for application to privacy regulation in interactions with information technology, primarily to identify tensions affecting individual dynamic privacy regulation in the presence of technological systems [26, 92, 112].

Lehikoinen et al. extend Altman's theory for privacy in ubiquitous computing [92]. Focused on the bidirectional, dialectic nature of the privacy regulation process, they map Altman's inputs and outputs to different interaction patterns in ubiquitous computing environments. When interacting with an interactive environment or others' personal devices, the inputs are determined by those technical components, own outputs are partially dependent on the sensing capabilities of those components. When the user's personal device interacts with other devices or the environment, inputs and outputs are digital information. Lehikoinen et al. further introduce the concept of *leaking* to describe situations where the actual outputs exceed desired privacy [92]—a case of importance in information systems where individuals may not be fully aware of their outputs, i.e., what information about them is being sensed or communicated. Romero et al. focus on the dialectic aspect of the privacy regulation process [118]. They propose additional phases (*collaboration*, *signaling*, *joint understanding*) before the actual boundary regulation in order to better capture the influence of technology support in mediated communication in contrast to Altman's verbal, paraverbal, and nonverbal regulation mechanisms.

13.3.2 Operationalizing Context Awareness for Privacy

Altman's privacy regulation theory significantly influenced reasoning about privacy and has been found to be a realistic model to describe privacy regulation from an

individual's perspective [101]. Altman's theory, as well as Nissenbaum's framework of contextual integrity, recognize that privacy regulation is a dynamic and dialectic process. Perception and awareness of the given situation shape a person's privacy concerns and expectations, together with personal privacy preferences, individual knowledge and experiences, as well as cultural and social background and constraints. This privacy decision making process results in a consciously or subconsciously desired privacy level, which is put into practice with available means of privacy control. Subsequently, the individual may receive feedback on the effectiveness of the exercised control, i.e., what level of privacy was actually achieved. Such feedback in turn leads to internal adjustments of individual privacy concerns, expectations, and preferences.

Context-adaptive privacy mechanisms can mirror aspects of a user's cognitive privacy regulation processes in order to provide privacy adaptation and privacy decision making support specific to the user's situation. For the sake of operationalizing privacy regulation theory for context-adaptive privacy mechanisms, the dynamic regulation process can be coarsely divided into three inter-related phases [128]:

- **Awareness.** Awareness of privacy-relevant processes and information flow shape individual privacy concerns [16]. An individual becomes aware of a contextual aspect or change in her environment that potentially necessitates regulative action to maintain a desired level of privacy (e.g., another person appears that could overhear a private conversation). The recognition of context changes as potential privacy risks depends on the individual's perception. However, with modern sensing technologies, a user's awareness and privacy perception is likely incomplete [87, 92], because sensors and information flows may not be apparent. Potential consequences are wrongly formed mental models and misconceptions about afforded privacy in a given context.
- **Decision.** Based on contextual and situational awareness, personal preferences and experiences, as well as cultural background and social motivations, an individual decides whether to decrease or increase exposure in the changed situation (e.g., including or excluding the new person from the conversation). Privacy decisions often need to be made based on incomplete information and are subject to cognitive biases and decision heuristics [6, 7], as well as susceptible to framing and manipulation [5, 6].
- **Control.** Once the individual formed a privacy decision, the decision needs to be mapped onto controls available in the current context. Available controls are determined by the means at disposal for asserting control (e.g., a door that can be closed to prevent eavesdropping) as well as the prevalent socio-cultural expectations and norms, which may restrict available controls (e.g., closing a door may be considered inappropriate in some cultures [135]). Although deciding on a regulation action and acting upon it are closely related, we argue that decision and control should be considered separately. Consciously or subconsciously forming an intention for a desired level of privacy (*desired privacy*) is an internal process, while the ability to implement the desired privacy is subject to external contextual constraints in a given situation (*achieved*

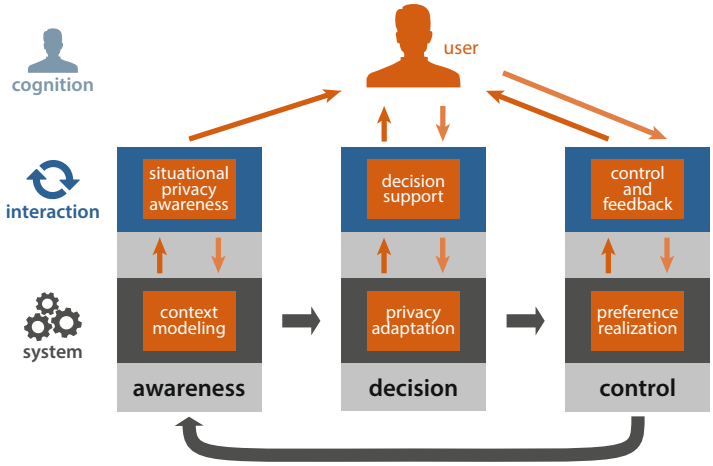

Fig. 13.2 Context-adaptive privacy mechanisms align context-aware system capabilities with the cognitive privacy regulation process in order to support privacy decision making and regulation

privacy). The results of performed control actions can be potentially perceived and verified by the user or the system and thus influence future awareness and subsequent regulation decisions [11].

In social interactions, these phases may overlap and influence each other. For instance, control actions may lead to awareness about their effectiveness, which may potentially require re-evaluation and re-adaptation. The regulation process runs continuously, resulting in micro-level privacy adaptations, such as adjusting what degree of information is revealed in a conversation, as well as macro-level adaptations, such as moving to a different, more private location.

Context-adaptive privacy mechanisms can support privacy regulation by supporting the different phases of this cognitive process on the system level, as shown in Fig. 13.2. Context-aware systems for privacy align well with the cognitive privacy regulation process. An individual combines situational awareness with individual preferences and experiences to make informed decisions on how to regulate privacy boundaries. Such decisions are implemented through actions, and personal preferences are adapted by learning from positive and negative experiences. Similarly, a context-adaptive privacy mechanism can leverage context awareness, elicited privacy preferences, and previous decisions to predict the user's privacy preferences and desired level of privacy for a given situation. Context-triggered changes in privacy expectation can either be addressed by automatically reconfiguring privacy controls and settings to the changed needs, or by suggesting privacy actions suitable for the current context to the user. This may include suggestions for more restrictive privacy configuration, but could also lead to a more permissive configuration [128]. Oppermann and Zimmermann [111] similarly distinguish three components of context-adaptive systems: the *sensory function* for obtaining relevant information,

the *inference function* guiding adaptation, and the *effector function* implementing adaptations in the system.

Context-adaptive privacy mechanisms can further tailor what information and recommendations they provide to the user's current situation and information needs in order to optimally support privacy decisions without getting in the way of the user's activity or overburdening them with irrelevant privacy information or settings [124]. Interaction strategies may be constrained by application requirements that determine the appropriateness and opportunities for user interactions in order to avoid disrupting the user's primary activities. Awareness of the user's context and preferences also offers the potential for integrating privacy regulation tasks into the user's primary activities. Output modalities and the presentation of recommendations can be tailored to the user's primary context and activity. Thus privacy management has the potential to become a natural by-product of using a system rather than a burdensome configuration task [91].

The level of automation should further align with the expectations of individual users. For instance, some users may be content with largely automated adaptations of their privacy settings, while others may prefer explicit awareness and control. Automation preferences may also be different for different situations. For instance, most people would likely not object to their location being automatically shared with emergency services in the case of a severe car crash, but may have diverse preferences for everyday situations.

In behavioral privacy regulation, individuals leverage combinations of different privacy mechanisms to achieve a desired privacy level, depending on the environment and context [11]. Similarly, technical privacy control mechanisms often function as systems, which are combined and configured according to application needs and privacy requirements. Rather than merely preventing information exchange, control mechanisms should enable users to form and maintain realms of exclusion within a certain socio-technological context [142], which would facilitate interaction with specific desired people, devices, or systems without interference.

13.3.3 Privacy-Relevant Context Features

Context-adaptive privacy mechanisms aim to identify and adapt to privacy-relevant changes in context. This requires maintaining a context model composed of privacy-relevant context features. Much research has been conducted to gain a deeper understanding of what contextual factors affect privacy perceptions, concerns, and behavior.

In contrast to general-purpose context models and context systems, a privacy context model constitutes a high-level abstraction of context that focuses on privacy-relevant context features only, in order to most effectively support dynamic adaptation of privacy mechanisms. In practice, context-adaptive privacy mechanisms can act as a consumer of more detailed context information provided by a context middleware.

Adams and Sasse identified *information sensitivity*, *information receivers*, and *information usage* as key factors for privacy perceptions [8]. The framework of contextual integrity provides a more generalized perspective on those factors as *actors*, *attributes*, and *transmission principles* [109]. In order to support privacy decision making of individual users, a privacy context model should provide a user-centric perspective on context [126]. Therefore, we can distinguish between the user, the environment (including other entities), as well as activities that link the user to other entities.

13.3.3.1 User Features

The *user* is characterized by *privacy-sensitive items*, which are potentially exposed in the current situation. Privacy-sensitive items loosely correspond to attributes defined by the contextual integrity framework [109]. Privacy-sensitive items can either be information sources or disturbance endpoints, corresponding to Altman's outputs and inputs [11].

Information sources potentially reveal information about the user, this could be the user's behavior, presence or activity that can be observed, e.g., by sensors, or digital information created by or about the user. Information sources can reveal time-variant or static information [116]. For time-variant information, an observer's scope is limited by the observation window. While an observer may be able to predict past or future values of an information source, the prediction scope is bounded in relation to the observation scope. Other information is static, such as a name, social security number or fingerprint. Once disclosed, they are known to the observer, and can only be changed with substantial effort.

Some information sources can further provide information at different levels of *granularity* and abstraction [8, 57, 68, 154]. For example, location can be provided as exact geo coordinates, as an address, on a street level, city level, or region level. Changes in information granularity are privacy relevant, because coarser information increases the difficulty for an observer to derive the exact information and thus potentially affords higher anonymity or privacy. In the contextual integrity framework, granularity adjustments are considered as a transmission principle for restricting information flow [109]. How granularity is specified depends on the semantic type of an information source. If granularity is expressed numerically, the scale must be mapped to a generalization function for the specific information source, which then transforms the original information into a version with the respective specificity. If specificity is expressed by class identifiers (e.g., *street* or *city* for location), a partial order of semantics between classes, e.g., provided by an ontology, is required in order to be able to determine if granularity is increased or decreased.

Disturbance endpoints are an individual's physical aspects that constitute potential targets for physical disturbances [28, 75, 126]. Any action occurring in the user's physical proximity can be seen as a potential disturbance. For instance, a household robot may have no direct means to observe the user, but the device's

activity or presence could still be perceived as an intrusion by the individual. Similarly, notifications in the user's environment can disrupt the user's solitude. The endpoint of a disturbance can be seen as a privacy-sensitive item, as it constitutes the connection point where the actual intrusion on an individual's privacy occurs. Therefore, disturbance endpoints should be included in privacy context models for smart environments and physical spaces [126]. Intuitively, the user's body is a potential endpoint for physical disturbances, as the user's senses perceive physical privacy intrusions. The user can be disturbed by touch, sound, smell, taste, and visual aspects [75]. In addition to the user's body, devices that are closely associated with the user may also be considered disturbance endpoints of the user. This may include the user's smartphone and its notifications but also wearable or implanted devices. Including such additional disturbance endpoints in privacy context models when relevant, allows to model changes of such endpoints and resulting privacy implications. In contrast to information sources, disturbance endpoints have typically no granularity. A disturbance endpoint is either exposed or not. Yet, certain disturbance endpoints can be perceived more invasive than others. For example, a vibrating phone is generally perceived less disturbing than a ringing phone.

13.3.3.2 Environment Features

The user's physical and virtual surroundings can be described as the user's *environment*. The environment contains context features that have an extrinsic effect on the user's privacy, compared to the intrinsic effects of user features. If one takes a user-centric perspective on modeling privacy context [126], the environment and its context features change based on the user's actions, e.g., when the user changes location. A different perspective can be to model physical and virtual aspects in multiple environments of which the user can be a part [34]. In this approach a user can participate in multiple environments at once. Chang et al. give the example of a video conference at the office, in which the user is in the virtual environment video conference and the physical environment office [34]. With a user-centric modeling perspective, the environment is implicitly defined by the user's location (the office), but includes any virtual entities that are able to participate in the user's physical environment, including any communication partners in a video call [126].

While environment models for context awareness can be highly detailed, a small number of environmental context features is most relevant for privacy. Primary privacy-relevant context features are other *entities* that participate physically or virtually in the user's environment, as well as the *observation or disturbance channels* [77], with which they are connected to the user's information sources and disturbance endpoints [126]. This corresponds to findings identifying receivers as a salient factor for privacy decisions [8, 109, 114]. Modeling of other environmental aspects, such as room layout or inanimate objects is typically not required, because any privacy-relevant change to such environment features would be reflected by

changes to the set of present entities and whether and how they can access information sources and disturbance endpoints. For example, closing an office door should remove people in the hallway from the model.

Entities determine the relevant actors [109] in relation to privacy. For privacy preferences pertaining to the sharing of information, the receiving entity has been shown to be a key factor [8, 41, 63, 66, 95, 110, 115, 139]. Patil et al. identified it as the highest ranked factor for location sharing preferences [114]. Thus, present entities are often considered in the modeling of privacy-relevant context [53, 126]. An entity could be a *person*, *device*, *software agent*, or *service* that participates in the user's environment. Both *physical and virtual entities* can be represented in the same model [75, 126]. Entities in the user's physical environment potentially forward information to virtual entities. Virtual entities rely on physical entities in order to participate in the user's physical environment. For example, the user's location can only be observed by users of a location sharing service if some physical entity senses the information and relays it, e.g., a smartphone with a GPS sensor.

Channels model how entities participate in the user's environment. They define the underlying transmission principles considered in the contextual integrity framework [109]. *Observation channels* originate at an information source and are connected to one or more observing entities. *Disturbance channels* originate from one or more entities and end at one of the user's disturbance endpoints. This perspective corresponds to Altman's inputs and outputs in the sense that inputs, i.e., disturbances, depend on the capabilities of entities in the environments and that the user's own outputs also partially depend on the sensing capabilities of physically present entities [92]. A channel can consist of multiple *links* between a set of entities. Such a multi-link channel defines a directed graph between the user and multiple hierarchically organized entities [75].

In addition to entities and channels, *location* [53, 66, 77, 136] and *time* [77, 121] are often considered as environmental context factors when modeling privacy-relevant context. Benisch et al. find that location, the time of the day, and the day of the week influence privacy preferences of users [20]. Tsai et al. also note the importances of time [146], while Massimi et al. find that location is an important aspect in determining privacy sensitivity [104]. Location should be considered on different levels of abstraction, including the user's physical location (e.g. geo coordinates), a semantic interpretation of the location (e.g., a specific room in a building), as well as the type of the location or environment. Semantic location information can be derived from the user's position or other location-specific environment cues, such as nearby WiFi access points. Kargl et al. propose a semantic retrieval process for geographic locations [71]. The environment type gives the location further semantic meaning [106] and describes the prevailing social context [109]. Human association of location is based on actions rather than coordinates [29]. Massimi et al. found that the type of environment strongly influenced the expectations and perceptions of being recorded [104]. *Home*, *work*, and *other* are common types to categorize environments [55, 95]. The environment type can be derived to some extent from the semantic location and the user's

activities. Krumm et al. propose a method for deriving general semantic labels for environments from geographic positions [81].

13.3.3.3 Activities

Dourish frames context as an interactional rather than a representational problem [47]. Grounded in an analysis of the sociological origins of context, he argues that context arises from the user's activity rather than being purely representational. Thus, the relevance of context features is dynamically defined for individual users by their activities. The same intuitively holds true for privacy in context. Depending on activity, certain personal aspects that are shared as part of an activity, may be considered sensitive in other activities or situations. Therefore, user's *activities* have been considered as a privacy-relevant context feature [77, 106, 121, 126]. For privacy in ambient assisted living systems, Shankar et al. note that activity can be a crucial discriminant [137], because such systems typically pertain to the same location—the home—and different activities at that location are likely associated with different preferences.

Purpose plays an important role in privacy decision making [19, 37, 65, 104]. The notion of activity can describe or be associated with purposes. An activity is an abstract description of what the user is doing. To a certain extent, the user's activities reflect the user's intentions and goals in a specific situation [109].

For context-adaptive privacy mechanisms, activities can further describe which entities must have access to certain privacy-sensitive items of the user so that the user can actually pursue the activity. However, activities should not be confused with privacy preferences or privacy settings. An activity describes what the user is doing in a situation and with whom, and is therefore part of context information; privacy preferences describe which entities are allowed to observe or disturb that situation.

Activity recognition is a well-researched topic of ubiquitous and pervasive computing research. Recognition of complex human activities is challenging because individuals may engage in concurrent or interleaved activities, which results in ambiguity for interpretation [72]. Common approaches for inferring activity from sensor data rely on machine learning or rule-based inference from different context cues and sensors [99], including wearable sensors [30, 89]. Activity recognition typically requires the collection of training data, which is used for relevant feature extraction and training and validation of a recognition model for specific activities or classes of activities [89]. Activity recognition in context-aware systems requires the real-time processing and analysis of continuous sensor data streams [80], such as location, body motion, or interaction with a system. Activities can be organized hierarchically, by decomposing them into smaller actions that are easier to detect than complex activities [66, 100]. Higher-level activities composed from a set of actions can then be used in reasoning [100]. Knowledge about previously recognized activity and past activities can improve activity recognition [80].

13.4 Context-Adaptive Privacy Mechanisms

Privacy-relevant context factors can be utilized by context-adaptive privacy mechanisms to actively support privacy decision making or automatically configure privacy settings in reaction to context changes. Context-adaptive privacy mechanisms take context information and knowledge about the user's privacy preferences, expectations or concerns as inputs and determine whether privacy adaptation is required, as well as how privacy aspects should be adapted to align with the user's privacy needs or which actions should be recommended to the user. While privacy preferences may be used in the reasoning process, reasoning results need to be translated into privacy specifications that map privacy decisions—made by the user or the system on behalf of the user—onto configurations for privacy controls and tools available in the current situation. The differentiation between higher level privacy preferences for reasoning and lower level policies for enacting preferences allows to decouple privacy reasoning from technical realization and enforcement aspects, in order to better align with the user's cognitive privacy regulation process [123].

In this section, we first discuss a number of requirements for context-adaptive privacy mechanisms before providing an overview of common approaches and applications for context-adaptive privacy mechanisms (see Table 13.1).

13.4.1 Requirements

The process of privacy reasoning is not only subject to the *internal constraints* posed by the preferences of the individual user and the *external constraints* posed

Table 13.1 Overview of context-adaptive privacy mechanisms

Category	Approaches
Privacy context acquisition and dissemination	Privacy labeling of context features
	Automated analysis of privacy implications
	Collaborative identification of sensors
	Machine-readable privacy specification
	Proactive communication of privacy information
Context-aware authentication	Adapt required level of authentication to context
	Context as authentication secret
Context-adaptive information disclosure	Context-based disclosure policies
	Limit access to context of origin
	Create and enforce contextual privacy borders
Context-aware content adaptation	Content hiding based on context
	Content adaptation based on present entities
	Privacy-friendly output modality selection
Context-adaptive privacy automation	Predict desired level of user involvement
	Automated blocking of disclosure and access requests

by the given context, but must also respect additional *systemic constraints* posed by system characteristics and available technology. Each of these aspects may introduce uncertainty into the privacy reasoning process that must be taken into account by context-adaptive privacy mechanisms. Context features represented in the privacy context model are derived from sensor data, which can be noisy, inaccurate, or incomplete [12]. This uncertainty needs to be reflected in privacy context models [126]. The user model, which captures the privacy preferences of an individual user in reference to contextual factors, is also subject to uncertainty. Any elicited or inferred privacy preferences can only be a discrete approximation of the user's true privacy preferences. Even if the user explicitly states a preference, uncertainty remains, because the user may not have been able to properly express the desired preference with the available interaction methods, or may not even be able to articulate a privacy preference consistently [6]. Furthermore, users' preferences may change over time or may only apply to specific situations. Hence, context-adaptive privacy mechanisms need to consider uncertainty in their reasoning processes. The confidence in the outcome of the reasoning process should reflect the uncertainty of the considered inputs.

In order to be trusted by users, context-aware systems must be perceived as reliable, which can be achieved with predictable and consistent behavior [35]. Following the principle of least astonishment, context-aware systems should aim to consistently match the user's expectations and preferences. At the same time, reasoning results should be explainable to the user [12]. The reasoning process must be intelligible and understandable to ensure that reasoning results are perceived as credible by the user [49]. Therefore, context-adaptive privacy mechanisms should align with the user's privacy decision making, for instance by modeling it after Altman's privacy regulation process [11] and its three phases [128], as described in Sect. 13.3. Furthermore, reasoning processes and their outcomes should be accompanied by intelligible explanations that can help the user understand actions taken by the system or recommendations provided to them [90, 93], e.g., by reducing the complexity of rules constituting the user model [35].

In order to be able to provide meaningful decision support in previously unknown situations and dynamically adapt to such new situations, context-adaptive privacy mechanisms should operate under the open world assumption. User models should be extensible in order to accommodate new situations and adapt to the individual user's privacy preferences over time by integrating explicit and implicit user feedback. Users should also be enabled to inspect and adjust inferred privacy preferences. Learning of privacy preferences and new contexts needs to occur online during normal operation. Considering the subjective and fluid nature of privacy preferences [6], the user model requires continuous maintenance [12] to account for changes in preferences, to add new preferences, and to correct erroneously learned preferences.

Furthermore, context-adaptive systems typically have to operate continuously which has interesting consequences for context-adaptive privacy mechanisms. For instance given a context change, a context-adaptive privacy mechanism may have to anticipate future context changes in order to prevent infringing situations before they

occur, as well as perform adaptations immediately when certain context features change, e.g., remove sensitive information from a display before someone entering a space can see it [127]. Similarly, context-adaptive privacy mechanisms need to consider carefully when to prompt users for input. On the one hand, privacy mechanisms should not alienate the user with unexpected autonomous actions; on the other hand, the system should not annoy the user with copious messages and notifications. The crucial issue is to develop a system that provides user support for privacy management without being obtrusive or overwhelming.

13.4.2 Privacy Context Acquisition and Dissemination

A crucial aspect of any context system is the acquisition and dissemination of context information. In the case of privacy mechanisms, sensing of context information has been complemented with proactive approaches that communicate privacy-relevant information. For instance, sending out wireless privacy beacons [76, 87] enables devices to communicate their sensing and actuation capabilities, as well as their data practices in machine-readable formats.

Multiple machine-readable privacy specification formats and policy languages have been proposed [42, 82]. A well-known example is P3P [148]—the platform for privacy preferences—which was designed to enable website operators to express a legal privacy policy in machine-readable form. When users also specify their privacy preferences in a machine-readable format (e.g., with APPEL [88] or XPref [9]), a website's data practices can be matched against the user's personal privacy preferences with privacy practices of visited websites and detect conflicts. Machine-readable privacy policies, such as P3P, can be seen as *labeling protocols* [3] that enhance context features, namely present services and entities, with privacy-relevant information. Such machine-readable privacy specifications can be either integrated with service discovery protocols or actively announced—either by the respective entity [76, 85] (e.g., a localization system, a surveillance camera, a smart thermostat, or a vacuuming robot) or a separate infrastructure component or third party that gathers, aggregates and disseminates machine-readable privacy information about devices and systems in an environment to user devices (e.g., the user's smartphone). *Privacy proxies* [85]—dedicated entities trusted by the user and other stakeholders—can manage privacy and match user preferences with a system's data practices in the case of sensing systems and sensors with limited resources.

Langheinrich implemented these approaches in pawS [85]. Privacy beacons communicate data collection and processing practices of nearby systems to the user's trusted device. The user's privacy proxy obtains the respective machine-readable privacy policies and matches them with the user's pre-specified privacy preferences. Yee proposes a similar architecture [156], in which a smart environment has a dedicated privacy controller that performs policy matching between a user's privacy preferences and present system's privacy policies. In case of mismatch, the privacy controller initiates policy negotiation among the involved parties. Similarly,

P4P [83] is a context-based negotiation framework that arbitrates between a service provider's P3P privacy policy and the user's privacy preferences selected according to the user's context.

A challenge for approaches that rely on machine-readable privacy specifications is that they require cooperation by the entity collecting information and translating their natural language privacy policies into machine-readable formats. This process is often met by resistance due to fears and uncertainty regarding the legal nature of machine-readable formats and associated liability. As a result, even widely popularized and standardized formats, such as P3P, lacked adoption and have been largely abandoned [39]. However, recent advances in natural language processing, machine learning, crowdsourcing, and static code analysis, provide the opportunity for service operators, third parties as well as regulators to automate the analysis of systems' natural language privacy policies and program code to infer their data practices [23, 27, 119, 125, 151, 152, 157]. Such analysis results could then be encoded in machine-readable formats to support context-adaptive privacy mechanisms.

Similar approaches can be employed to detect and identify sensors in physical environments. Winkler and Rinner propose collaborative tagging of cameras to gain privacy awareness with respect to video surveillance [153], i.e., individuals mark locations and characteristics of spotted cameras in a mapping system. Korayem et al. propose an automated method to identify computer screens from camera images [78]. Information about identified sensors or devices and their locations can be leveraged by context-adaptive privacy mechanisms to determine privacy risks.

How and what privacy-relevant context information is gathered ultimately depends on the purpose and requirements of the respective context-adaptive privacy mechanism. Furthermore, environment constraints, such as the level of trust in the present infrastructure or the willingness and capability of other entities to cooperate with privacy mechanisms. Next, we discuss different kinds of context-adaptive privacy mechanisms.

13.4.3 Context-Aware Authentication

Context awareness has been proposed as an improvement for user authentication. Sigg notes that context awareness could enable password-less authentication and unobtrusive adaptive security [140]. He also suggests the use of ambient audio as a location- or situation-based secret key. Bardram et al. leveraged the user's location in an environment as an additional factor in user authentication [15]. Langheinrich proposed a password-free authentication scheme for RFID tags [86]. Instead of revealing the complete tag ID, an RFID tag releases secret shares of the ID over time, thus requiring an attacker to be in close proximity for a longer time, while legitimate users can identify their tags efficiently by relying on simple caching strategies. Mayrhofer and Gellersen leverage parallel motion as a context factor

in device paring [105]. They propose a method that relies on shaking two mobile devices together in order to derive a shared key from accelerometer data.

Gupta et al. propose what they call intuitive and sensible access control (ISAC) as an approach to autonomously select adequate authentication methods for smartphones depending on the user's context [52, 53]. For instance, password or PIN entry may be required in public but not when the user is alone. Their approach uses location, present Bluetooth devices, and WiFi access points as features to define *contexts of interest*. How frequently users encounter a context of interest is used to determine *familiarity of context* in order to identify likely "safe" locations.

Hayashi et al. propose context-aware scalable authentication (CASA) [55], which employs a similar approach. They uses passively collected context features to determine whether a user is at home, work or another place (other). Depending on the user's context, the user is either authenticated implicitly without interaction based on the passive context features, or active authentication is required (e.g., with PIN or password) when not at home or work.

Context-aware authentication approaches can be leveraged in context-aware privacy mechanisms to identify and authenticate the primary user as well as other entities detected in the environment.

13.4.4 Context-Adaptive Information Disclosure

The integration of context awareness into information disclosure and access control mechanisms has received considerable attention. Particularly location is a frequently used context feature to manage information sharing and access control. For instance, Behrooz and Devlic propose the context-aware privacy policy language (CPPL) [17]. With CPPL, machine-readable privacy specifications are scoped to certain contexts, e.g., based on a where a specific system is active. CPPL facilitates filtering of relevant privacy policies based on current context.

Jagtap et al. propose a privacy system that constrains information flow from mobile devices with dynamic semantic reasoning over context and pre-specified privacy preferences [66]. Their context features include user location, surroundings, other present entities, and inferred user activity, which are associated with entity roles. The Android-based implementation of their approach [50] supports context-aware privacy preferences encoded as privacy policies. These policies can be used for instance to specify under what circumstances smartphone apps should receive correct or fake location information and other sensor data.

Context and proximity have further been used to limit access to information to the context in which it has been collected originally [70, 73, 84]. For example, Kriplean et al. developed an approach that makes RFID readouts only accessible to devices that were physically present when the readout occurred [79].

The *info spaces* concept [61, 68] is a technical realization of Marx's personal border crossings theory [103]. Information spaces are defined by physical, social or activity-based borders, which are supported by location, entity and activity

detection as context features. Privacy preferences (encoded as privacy policies) are enforced when accordingly-tagged information crosses such borders. This may include granting or denying read or write access to the information, adjusting the granularity and accuracy with which information is sensed or shared, as well as aggregating information to obtain higher-level interpretations.

Moncrieff et al. use a context model to manage privacy in a smart home equipped for ambient assisted living [106]. Their context features include a user's location (room in the house), social interactions (other present persons), hazardous activity (e.g., leaving the kitchen stove on), and unusual periods of inactivity. Context are matched against pre-defined privacy disclosure rules for care givers. Their policies support different granularity levels to regulate sensor access. An active feedback display provides occupants with information about current and past observers on demand.

13.4.5 Context-Aware Content Adaptation

Approaches in the previous category primarily provide context-aware adaptation of what information is disclosed or forwarded to outside parties. However, particularly in the context of smart environments and media spaces, context awareness has further been used to dynamically adapt content and information within the environment in order to provide privacy protections.

An early example of context-aware content adaptation is Schmidt's Context NotePad, which is a PDA application that dynamically hides its content when the user is not alone in the room and is not actively using the NotePad [131]. The TreasurePhone system [136] uses location to specify and activate preference spheres. The user-specified preference spheres limit which information on a mobile phone can be accessed in a certain location.

Presence and position of individual users in smart spaces can been used to implicitly regulate privacy and content visibility. Neustaedter and Greenberg propose a system in which the video stream in a media space automatically stops when the user leaves the chair or other persons enter the room [107]. The ProD system uses access control lists to define privacy preferences for content adaptation [36]. Similarly, the Angel system [51] poses privacy restrictions on displayed content based on the user's activity, which is associated with user-defined privacy rules. Marquardt and Greenberg suggest the use of proxemic interaction, which leverages context information about present persons and devices to guide device adaptation, for privacy management [102]. The PriCal system [127] enables context-aware privacy adaptation on calendar wall displays. More specifically, displayed calendar views are dynamically adapted to present persons and their privacy preferences for each other. In addition to known persons, users can also specify privacy preferences for unknown persons, i.e., people who are not registered system users. The system uses case-based reasoning to learn a user's nuanced privacy preferences based on

users adjusting the visibility of individual calendar entries (full, busy, hidden) in deviation from their pre-specified rules.

The ATRACO system [77] considers both physically present entities as well as virtual entities connected via communication channels to the user's environment in context-adaptive privacy management. Privacy preferences describe what information items or spaces can be accessed by whom and how, in relation to a specific context. Changes in context trigger dynamic privacy evaluation. Privacy preferences are pre-specified as ontology instances, which are part of a larger user profile for dynamic adaptation of the smart environment.

When multiple output modalities are available in the user's physical environment, information display and means of interaction can be adapted to block observations by undesired entities. For example, private notifications could be displayed only on the user's personal device, such as a smartphone, rather than on a wall display to reduce the opportunity of visual observation channels by other entities [32]. Similarly, auditive output can be moved from speakers to earphones or translated into a visual representation. Furthermore, observation granularity can be tied to the physical arrangement of entities. For instance, Vogel and Balakrishnan propose a calendar application for public displays that only displays a user's calendar entries if the user is close enough to the display so that the display's content is shielded by the user's body [147] .

13.4.6 Context-Adaptive Privacy Automation

A major promise of context-adaptive privacy mechanisms is that they can reduce privacy management and configuration effort for users. However, not all privacy or information sharing decisions can or should be automated—automation and user autonomy need to be balanced carefully. Furthermore, systems need to enable users to correct decisions [54]. Bellotti and Edwards discuss how intelligibility of context-aware systems can be enhanced through user involvement [18]. Depending on the level of uncertainty about an inferred decision, a system should either provide means for correction, require confirmation from the user, or offer the user the available choices for selection. Multiple approaches have been proposed to leverage context awareness in determining if and when users should be involved in privacy decisions and when certain decisions can be automated.

The Super-Ego framework dynamically determines whether the user should be involved in decisions about location disclosure requests [145]. A decision engine uses a set of previous disclosure decisions from the larger user base to decide about user involvement based on two thresholds for manual decision and automatic decision. Location requests where the average of previous disclosure decisions is below the manual decision threshold are denied, requests between the manual and automatic decision threshold require user intervention, and requests above the automatic decision threshold are granted automatically. By adjusting these thresholds different decision strategies can be supported. Toch finds that

mixed strategies provide the best tradeoff between accuracy and automation [145]. Similarly, SPISM [24] semi-automatically determines if for a given location request the user's location should be shared and at what granularity. Context information is used in classifying the request and assigning a sharing class (no; yes with low/medium/high granularity). Depending on the level of confidence in the classification result, the system prompts the user to make the decision.

Wijesekera et al. leverage context awareness to explore opportunities for automatically blocking permission requests from smartphone apps [150]. They conducted a field study to determine how context of an app's permission access affects users' privacy concerns. They identified the app's visibility (is the display on, is the user interacting with the app or not) and the request frequency as important factors in user's privacy preferences. In further experiments, they find that automatically blocking requests when the screen is off is unlikely to interfere with the user experience but enhances privacy. If an app that the user is currently not using requests resource access a prompt should be shown to obtain the user's consent. Apps that are running in the background should use passive indicators, such a GPS icon, when resources are being accessed.

13.5 Research Challenges

The variety of context-adaptive privacy mechanisms in different domains demonstrates the benefits of leveraging context awareness to actively support and partially automate privacy management. Context-adaptive privacy mechanisms are subject of active research, with multiple research challenges requiring further investigation:

- **Secure and trustworthy context acquisition.** Context-adaptive privacy mechanisms and other context systems rely on the integrity and trustworthiness of context information. Some sensors and context information can potentially be spoofed, which could trick context-adaptive privacy mechanisms into revealing personal information in the wrong situations. Therefore, context-adaptive privacy mechanisms should be designed to be resilient against spoofing attacks. For instance, by triangulating context information with different types of sensors and context acquisition methods, sensor fingerprinting, or distributed trust management in sensing and context infrastructures. In reasoning, confidence in context information and uncertainty about undetected context features should be considered.
- **Protection of contextual privacy information.** As discussed in Sect. 13.2.3, context information itself can often be privacy-sensitive and needs to be adequately secured, especially if systems retain context and privacy decision history. Similarly, context-specific privacy preferences and rules should be treated as privacy-sensitive information in their own right, as they detail what entities are given access to what kind of information in different situations. Some approaches rely on a trusted privacy assistant or privacy proxy [85] to aggregate

and protect a user's privacy information. Many of the approaches for protecting context information, such as authentication, obfuscation and other cryptographic techniques, can also provide opportunities for protecting privacy preference information against information leakage.

- **Privacy adaptation in heterogeneous environments.** A particular challenge for privacy mechanisms is to provide individual users with control over the collection, dissemination, and use of their information in collaborative and shared information ecosystems [1]. Different systems, infrastructure, sensors, services, or devices may be controlled by different stakeholders [2]. Thus, how a specific privacy decision can be translated into a privacy adaptation largely depends on the control capabilities available in the given context, i.e., the level of control and trust concerning other devices, infrastructure, and entities [75]. Thus, context-adaptive privacy mechanisms may have to consider different adaptation strategies within the same context as well as across contexts and systems. Furthermore, privacy controls available within a specific environment should be discoverable in order to facilitate privacy management and adaptation in previously unknown environments and contexts.
- **Privacy adaptation in physical environments.** Context-aware adaptation of content and information flows faces a particular challenge in physical spaces: other persons may be able to observe content and information adaptation, which could potentially be interpreted negatively and reveal the user's privacy preferences, potentially resulting in awkward situations. Thus, context-adaptive privacy mechanisms should be designed with plausible deniability in mind. Resulting adaptations should either be difficult to observe by others or others should not be able to determine whether the adaptation occurred because of them. For instance, the PriCal system [127] implements a hide-then-reveal paradigm: Whenever a person enters a room, the calendar display is cleared instantly when the new entity is detected and subsequently populated with content that has been adapted to the user's privacy preferences for the changed context. Because the display updates every time someone enters the room, study participants did not perceive those adaptations as specifically related to them [127]. Similarly, ATRACO [77] adapts a photo slideshow dynamically to the user's privacy preferences for present persons by seamlessly filtering out photos they should not see.
- **Privacy adaptation in multi-user environments.** Depending on the application and system context, context-adaptive privacy mechanisms may have to take privacy preferences from multiple users into account. In such situations, diverging privacy preferences have to be resolved by the system while respecting individuals' privacy preferences. Privacy preference negotiation and resolution could occur automatically or delegate the final decision whether some information should be disclosed to each user.
- **User trust in adaptation capabilities.** Context-adaptive privacy mechanisms must be perceived as trustworthy and reliable by users in order for users to trust the mechanisms with dynamically regulating privacy for them. This is particularly relevant for autonomous privacy adaptation on the user's behalf.

Trust evaluation in inter-personal relations as well as in technical systems has to rely on external trust signals in order to obtain information about internal trust facets [33], thus context-adaptive privacy mechanisms need to provide indications to users that they are functioning as configured or expected.

- **Personalized Privacy Adaptation.** Context-adaptive privacy mechanisms leverage context information to support users with privacy management. Privacy mechanisms can further learn a user's privacy preferences over time to not only contextualize but also personalize privacy management. Existing personalized privacy approaches learn from user feedback [63, 120] or derive privacy preference profiles from many users [94, 96, 98]. These approaches can further be combined to bootstrap privacy preferences by matching a user to one of a small set of privacy profiles [97, 98] and then leveraging user feedback and behavior to further refine and extend the user model to account for individual nuances in privacy preferences [123].

13.6 Summary

While context aware systems and the collection of context information pose challenges for personal privacy, we outlined the potential for supporting privacy management with context awareness in this chapter. Interpersonal privacy regulation has been shown to be highly dynamic and dependent on context. We presented a model for operationalizing context awareness by developing privacy mechanisms that align in their way of operation with individuals' privacy decision making processes. Context-adaptive privacy mechanisms can leverage context awareness to detect and determine privacy-relevant context changes in a user's environment and either provide context-specific privacy recommendations to the user or automatically adjust and adapt privacy configurations to ensure that the user's privacy preferences are respected in the changed context. We further provided an overview of privacy-relevant context features that have been shown to play a role in individual privacy decision making, as well as an overview of existing context-adaptive privacy mechanisms in various domains and associated research challenges. Context-adaptive privacy mechanisms are a promising approach for reducing the user effort in privacy management, in particular in sensor-rich environments.

References

1. G. D. Abowd and E. D. Mynatt. Charting past, present, and future research in ubiquitous computing. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 7(1):29 –58, 2000.
2. M. Ackerman and T. Darrell. Privacy in context. *Human-Computer Interaction*, 16(2–4):167–176, 2001.
3. M. S. Ackerman. Privacy in pervasive environments: next generation labeling protocols. *Personal and Ubiquitous Computing*, 8(6):430–439, Sept. 2004.

4. M. S. Ackerman and S. D. Mainwaring. Privacy Issues and Human-Computer Interaction. In L. F. Cranor and S. Garfinkel, editors, *Security and Usability*, chapter 19, pages 381–400. O'Reilly, 2005.
5. A. Acquisti, I. Adjerid, R. Balebako, L. Brandimarte, L. F. Cranor, S. Komanduri, P. G. Leon, N. Sadeh, F. Schaub, M. Sleeper, Y. Wang, and S. Wilson. Nudges for privacy and security: Understanding and assisting users choices online. *ACM Computing Surveys*, 50(3), 2017.
6. A. Acquisti, L. Brandimarte, and G. Loewenstein. Privacy and human behavior in the age of information. *Science*, 347(6221):509–514, 2015.
7. A. Acquisti and J. Grossklags. What Can Behavioral Economics Teach Us About Privacy? In *Digital Privacy: Theory, Technologies, and Practices*, chapter 18, pages 363–377. Auerbach Publications, 2008.
8. A. Adams and M. Sasse. Privacy in multimedia communications: Protecting users, not just data. In *Human-Computer Interaction/Interaction d'Homme-Machine (IMH-HCI '01)*, pages 49–64, 2001.
9. R. Agrawal, J. Kiernan, R. Srikant, and Y. Xu. An XPath-based preference language for P3P. In *12th international conference on World Wide Web (WWW '03)*. ACM, 2003.
10. H. Almuhimedi, F. Schaub, N. Sadeh, I. Adjerid, A. Acquisti, J. Gluck, L. F. Cranor, and Y. Agarwal. Your location has been shared 5,398 times! a field study on mobile app privacy nudging. In *Proc. CHI '15*. ACM, 2015.
11. I. Altman. *The Environment and Social Behavior: Privacy, Personal Space, Territory, Crowding*. Brooks/Cole Publishing company, Monterey, California, 1975.
12. A. Aztiria, A. Izaguirre, and J. C. Augusto. Learning patterns in ambient intelligence environments: a survey. *Artificial Intelligence Review*, 34(1):35–51, May 2010.
13. M. Baldauf, S. Dustdar, and F. Rosenberg. A survey on context-aware systems. *International Journal of Ad Hoc and Ubiquitous Computing*, 2(4):263–277, 2007.
14. J. Bardram and A. Friday. Ubiquitous Computing Systems. In J. Krumm, editor, *Ubiquitous Computing Fundamentals*, chapter 2, pages 37–94. CRC Press, 2009.
15. J. E. Bardram, R. E. Kjær, and M. Ø. Pedersen. *Context-Aware User Authentication – Supporting Proximity-Based Login in Pervasive Computing*, pages 107–123. Springer Berlin Heidelberg, Berlin, Heidelberg, 2003.
16. L. Barkhuus. The Mismeasurement of Privacy: Using Contextual Integrity to Reconsider Privacy in HCI. In *ACM annual conference on Human Factors in Computing Systems (CHI '12)*, page 367, New York, New York, USA, 2012. ACM.
17. A. Behrooz and A. Devlic. A Context-aware Privacy Policy Language for controlling access to context information of mobile users. In *Conference on Security and Privacy in Mobile Information and Communication Systems (MobiSec '11)*. Springer, 2011.
18. V. Bellotti and K. Edwards. Intelligibility and Accountability: Human Considerations in Context-Aware Systems. *Human-Computer Interaction*, 16(2):193–212, Dec. 2001.
19. V. Bellotti and A. Sellen. Design for Privacy in Ubiquitous Computing Environments. In *Third European Conference on Computer-Supported Cooperative Work (ECSCW '93)*. Springer, 1993.
20. M. Benisch, P. G. Kelley, N. Sadeh, and L. F. Cranor. Capturing location-privacy preferences: quantifying accuracy and user-burden tradeoffs. *Personal and Ubiquitous Computing*, 15(7):679–694, Dec. 2010.
21. C. Bettini, O. Brdiczka, K. Henricksen, J. Indulska, D. Nicklas, A. Ranganathan, and D. Riboni. A survey of context modelling and reasoning techniques. *Pervasive and Mobile Computing*, 6(2):161–180, Apr. 2010.
22. C. Bettini and D. Riboni. Privacy protection in pervasive systems: State of the art and technical challenges. *Pervasive and Mobile Computing*, 17:159–174, 2015.
23. J. Bhatia, T. D. Breaux, and F. Schaub. Mining privacy goals from privacy policies using hybridized task recomposition. *ACM Trans. Softw. Eng. Methodol.*, 25(3):22:1–22:24, May 2016.

24. I. Bilogrevic, K. Huguenin, B. Agir, M. Jadliwala, and J.-P. Hubaux. Adaptive information-sharing for privacy-aware mobile social networks. In *ACM international joint conference on Pervasive and ubiquitous computing (UbiComp '13)*, page 657, New York, New York, USA, 2013. ACM.
25. J. Bohn, V. Coroam, M. Langheinrich, F. Mattern, and M. Rohs. Social, Economic, and Ethical Implications of Ambient Intelligence and Ubiquitous Computing. In W. Weber, J. M. Rabaey, and E. Aarts, editors, *Ambient Intelligence*, chapter 1, pages 5–29. Springer, 2005.
26. M. Boyle and S. Greenberg. The language of privacy: Learning from video media space analysis and design. *ACM Transactions on Computer-Human Interaction*, 12(2):328–370, June 2005.
27. T. D. Breaux and F. Schaub. Scaling requirements extraction to the crowd: Experiments with privacy policies. In *2014 IEEE 22nd International Requirements Engineering Conference (RE)*, pages 163–172, Aug 2014.
28. P. Brey. The Importance of Privacy in the Workplace. In S. O. Hansson and E. Palm, editors, *Privacy in the Workplace*, chapter 5, pages 97–118. Fritz Lang, 2005.
29. B. Brown, A. S. Taylor, S. Izadi, A. Sellen, J. J. Kaye, and R. Eardley. Locating Family Values: A Field Trial of the Whereabouts Clock. In *9th international conference on Ubiquitous computing (UbiComp '07)*, pages 354–371. Springer, 2007.
30. A. Bulling, U. Blanke, and B. Schiele. A tutorial on human activity recognition using body-worn inertial sensors. *ACM Comput. Surv.*, 46(3):33:1–33:33, Jan. 2014.
31. H. Byun and K. Cheverst. Exploiting user models and context-awareness to support personal daily activities. In *Personal Daily Activities, Workshop in UM2001 on User Modelling for Context-Aware Applications*, 2001.
32. H. Cao, P. Olivier, and D. Jackson. Enhancing Privacy in Public Spaces Through Crossmodal Displays. *Social Science Computer Review*, 26(1):87–102, Feb. 2008.
33. C. Castelfranchi and R. Falcone. *Trust Theory: A Socio-Cognitive and Computational Model*. John Wiley & Sons, 1 edition, 2010.
34. Y.-L. Chang, E. Barrenechea, and P. Alencar. Dynamic user-centric mobile context model. In *Fifth International Conference on Digital Information Management (ICDIM '10)*, pages 442–447. IEEE, July 2010.
35. K. Cheverst, N. Davies, K. Mitchell, and C. Efstratiou. Using Context as a Crystal Ball: Rewards and Pitfalls. *Personal and Ubiquitous Computing*, 5(1):8–11, Feb. 2001.
36. B. Congleton, M. S. Ackerman, and M. W. Newman. The ProD framework for proactive displays. In *21st annual ACM symposium on User interface software and technology (UIST '08)*, page 221, New York, New York, USA, 2008. ACM.
37. S. Consolvo, I. Smith, T. Matthews, A. LaMarca, J. Tabert, and P. Powledge. Location disclosure to social relations: why, when, & what people want to share. In *SIGCHI conference on Human factors in computing systems (CHI '05)*, pages 81–90. ACM, 2005.
38. J. Coutaz, J. L. Crowley, S. Dobson, and D. Garlan. Context is key. *Communications of the ACM*, 48(3):49–53, Mar. 2005.
39. L. F. Cranor. Necessary but not sufficient: Standardized mechanisms for privacy notice and choice. *Journal on Telecommunications and High Technology Law*, 10:273, 2012.
40. G. Danezis, J. Domingo-Ferrer, M. Hansen, J.-H. Hoepman, D. L. M tayer, R. Tirtea, and S. Schiffner. Privacy and data protection by design – from policy to engineering. resreport, European Union Agency for Network and Information Security (ENiSA), 2015.
41. S. Davis. Using relationship to control disclosure in Awareness servers. In *Graphics Interface (GI '05)*, pages 145–152, 2005.
42. J. De Coi and D. Olmedilla. A review of trust management, security and privacy policy languages. In *International Conference on Security and Cryptography (SECRYPT '08)*. INSTICC Press, 2008.
43. P. De Hert, S. Gutwirth, A. Moscibroda, D. Wright, and G. González Fuster. Legal safeguards for privacy and data protection in ambient intelligence. *Personal and Ubiquitous Computing*, 13(6):435–444, Oct. 2009.

44. A. K. Dey. Understanding and Using Context. *Personal and Ubiquitous Computing*, 5(1):4–7, Feb. 2001.
45. A. K. Dey. Context-Aware Computing. In J. Krumm, editor, *Ubiquitous Computing Fundamentals*, chapter 8, pages 321–352. CRC Press, 2009.
46. A. K. Dey and G. D. Abowd. Towards a Better Understanding of Context and Context-Awareness. In *CHI 2000 Workshop on The What, Who, Where, When, and How of Context-Awareness*, 2000.
47. P. Dourish. What we talk about when we talk about context. *Personal and Ubiquitous Computing*, 8(1):19–30, Feb. 2004.
48. P. Dourish and K. Anderson. Collective Information Practice: Exploring Privacy and Security as Social and Cultural Phenomena. *Human-Computer Interaction*, 21(3):319–342, Sept. 2006.
49. B. J. Fogg and H. Tseng. The elements of computer credibility. In *SIGCHI Conference on Human Factors in Computing Systems (CHI '99)*. ACM, 1999.
50. D. Ghosh, A. Joshi, T. Finin, and P. Jagtap. Privacy Control in Smart Phones Using Semantically Rich Reasoning and Context Modeling. In *IEEE Symposium on Security and Privacy Workshops*, pages 82–85. IEEE, May 2012.
51. G. Go aszewski and J. Górski. Context sensitive privacy management in a distributed environment. In *On the Move to Meaningful Internet Systems (OTM '10)*, pages 639–655, 2010.
52. A. Gupta, M. Miettinen, and N. Asokan. Using context-profiling to aid access control decisions in mobile devices. In *International Conference on Pervasive Computing and Communications (PerCom '11) Workshops*, pages 310–312. IEEE, Mar. 2011.
53. A. Gupta, M. Miettinen, N. Asokan, and M. Nagy. Intuitive security policy configuration in mobile devices using context profiling. In *2012 ASE/IEEE International Conference on Privacy, Security, Risk and Trust (PASSAT '12)*. IEEE, 2012.
54. B. Hardian, J. Indulska, and K. Henricksen. Balancing Autonomy and User Control in Context-Aware Systems - a Survey. In *3rd Workshop on Context Modeling and Reasoning (CoMoRea'06)*, pages 51–56. IEEE, 2006.
55. E. Hayashi, S. Das, S. Amini, J. Hong, and I. Oakley. CASA: context-aware scalable authentication. In *Ninth Symposium on Usable Privacy and Security (SOUPS '13)*, New York, New York, USA, 2013. ACM.
56. T. Heiber and P. Marrón. Exploring the Relationship between Context and Privacy. In P. Robinson, H. Vogt, and W. Wagealla, editors, *Privacy, Security and Trust within the Context of Pervasive Computing*. Springer, 2005.
57. K. Henricksen and J. Indulska. A software engineering framework for context-aware pervasive computing. In *Second IEEE Annual Conference on Pervasive Computing and Communications (PerCom '04)*, pages 77–86. IEEE, 2004.
58. K. Henricksen and J. Indulska. Developing context-aware pervasive computing applications: Models and approach. *Pervasive and Mobile Computing*, 2(1):37–64, Feb. 2006.
59. K. Henricksen, J. Indulska, and A. Rakotonirainy. Modeling context information in pervasive computing systems. In *1st International Conference on Pervasive Computing (Pervasive '02)*, pages 167–180, 2002.
60. K. Henricksen, R. Wishart, T. McFadden, and J. Indulska. Extending Context Models for Privacy in Pervasive Computing Environments. In *2nd Workshop on Context Modeling and Reasoning (CoMoRea '05), PerCom '05 workshops*, pages 20–24. IEEE, 2005.
61. J. I. Hong and J. A. Landay. An architecture for privacy-sensitive ubiquitous computing. In *2nd International Conference on Mobile systems, applications, and services (MobiSys '04)*. ACM, 2004.
62. J.-Y. Hong, E.-H. Suh, and S.-J. Kim. Context-aware systems: A literature review and classification. *Expert Systems with Applications*, 36(4):8509–8522, May 2009.
63. G. Hsieh, K. P. Tang, W. Y. Low, and J. I. Hong. Field Deployment of IMBuddy: A Study of Privacy Control and Feedback Mechanisms for Contextual IM. In *9th International Conference on Ubiquitous Computing (UbiComp '07)*, pages 91–108. Springer, 2007.

64. S. E. Hudson and I. Smith. Techniques for addressing fundamental privacy and disruption tradeoffs in awareness support systems. In *ACM conference on Computer supported cooperative work (CSCW '96)*, pages 248–257, New York, New York, USA, 1996. ACM.
65. G. Iachello, K. N. Truong, G. D. Abowd, G. R. Hayes, and M. Stevens. Prototyping and sampling experience to evaluate ubiquitous computing privacy in the real world. In *SIGCHI conference on Human Factors in computing systems (CHI '06)*, page 1009, New York, New York, USA, 2006. ACM.
66. P. Jagtap, A. Joshi, T. Finin, and L. Zavala. Preserving Privacy in Context-Aware Systems. In *Fifth International Conference on Semantic Computing*, pages 149–153. IEEE, Sept. 2011.
67. A. Jameson. Modelling both the Context and the User. *Personal and Ubiquitous Computing*, 5(1):29–33, Feb. 2001.
68. X. Jiang and J. A. Landay. Modeling privacy control in context-aware systems. *IEEE Pervasive Computing*, 1(3):59–63, July 2002.
69. L. Kagal, T. Finin, and A. Joshi. A policy language for a pervasive computing environment. In *4th International Workshop on Policies for Distributed Systems and Networks (POLICY '03)*, pages 63–74. IEEE, 2003.
70. A. Kapadia, T. Henderson, J. J. Fielding, and D. Kotz. Virtual Walls: Protecting Digital Privacy in Pervasive Environments. In *5th International Conference on Pervasive Computing (PERVASIVE '07)*, pages 162–179. Springer, 2007.
71. F. Kargl, G. Dannhäuser, S. Schlott, and J. Nagler-Ihlein. Semantic information retrieval in the COMPASS location system. In *Third International Symposium on Ubiquitous Computing Systems (UCS '06)*. Springer, Oct. 2006.
72. E. Kim, S. Helal, and D. Cook. Human activity recognition and pattern discovery. *IEEE Pervasive Computing*, 9(1):48–53, Jan 2010.
73. T. Kindberg and A. Fox. System software for ubiquitous computing. *IEEE Pervasive Computing*, 1(1):70–81, Jan. 2002.
74. M. Knappmeyer, S. L. Kiani, E. S. Reetz, N. Baker, and R. Tonjes. Survey of Context Provisioning Middleware. *IEEE Communications Surveys & Tutorials*, 15(3):1492–1519, 2013.
75. B. Könings and F. Schaub. Territorial Privacy in Ubiquitous Computing. In *Eighth International Conference on Wireless On-Demand Network Systems and Services (WONS '11)*, pages 104–108, Bardoneccia, Jan. 2011. IEEE.
76. B. Könings, F. Schaub, and M. Weber. PriFi Beacons: Piggybacking Privacy Implications on WiFi Beacons. In *UbiComp '13 Adjunct Proceedings*. ACM, 2013.
77. B. Könings, F. Schaub, and M. Weber. Privacy and Trust in Ambient Intelligence Environments. In *Next Generation Intelligent Environments*, chapter 7, pages 133–164. Springer, second edition, 2016.
78. M. Korayem, R. Templeman, D. Chen, D. Crandall, and A. Kapadia. Enhancing lifelogging privacy by detecting screens. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, pages 4309–4314, New York, NY, USA, 2016. ACM.
79. T. Kriplean, E. Welbourne, N. Khoussainova, V. Rastogi, M. Balazinska, G. Borriello, T. Kohno, and D. Suciu. Physical Access Control for Captured RFID Data. *IEEE Pervasive Computing*, 6(4):48–55, 2007.
80. N. C. Krishnan and D. J. Cook. Activity recognition on streaming sensor data. *Pervasive and Mobile Computing*, 10(Part B):138–154, 2014.
81. J. Krumm and D. Rouhana. Placer: semantic place labels from diary data. In *ACM international joint conference on Pervasive and ubiquitous computing (UbiComp '13)*, page 163, New York, New York, USA, 2013. ACM.
82. P. Kumaraguru, L. Cranor, J. Lobo, and S. Calo. A survey of privacy policy languages. In *Third Symposium on Usable Privacy and Security (SOUPS '07) Workshops*, 2007.
83. O. Kwon. A pervasive P3P-based negotiation mechanism for privacy-aware pervasive e-commerce. *Decision Support Systems*, 50(1):213–221, Dec. 2010.
84. M. Langheinrich. *Privacy by Design — Principles of Privacy-Aware Ubiquitous Systems*, pages 273–291. Springer Berlin Heidelberg, Berlin, Heidelberg, 2001.

85. M. Langheinrich. A Privacy Awareness System for Ubiquitous Computing Environments. In *4th International Conference on Ubiquitous Computing (UbiComp '02)*, pages 237–245. Springer, 2002.
86. M. Langheinrich. RFID privacy using spatially distributed shared secrets. In *4th International Symposium on Ubiquitous Computing Systems (UCS '07)*, pages 1–16. Springer, 2007.
87. M. Langheinrich. Privacy in Ubiquitous Computing. In J. Krumm, editor, *Ubiqutious Computing Fundamentals*, chapter 3, pages 95–160. CRC Press, 2009.
88. M. Langheinrich, L. F. Cranor, and M. Marchiori. A P3P Preference Exchange Language 1.0 (APPEL1.0). W3c working draft, W3C, 2002.
89. O. D. Lara and M. A. Labrador. A survey on human activity recognition using wearable sensors. *IEEE Communications Surveys and Tutorials*, 15(3):1192–1209, 2013.
90. D. Leake, A. Maguitman, and T. Reichherzer. Cases, Context, and Comfort: Opportunities for Case-Based Reasoning in Smart Homes. In *Designing Smart Homes: The Role of Artificial Intelligence*, pages 109–131. Springer, 2006.
91. S. Lederer, J. I. Hong, A. K. Dey, and J. A. Landay. Five Pitfalls in the Design of Privacy. In L. F. Cranor and S. Garfinkel, editors, *Security and Usability*, chapter 21, pages 421–446. O'Reilly, 2005.
92. J. T. Lehikoinen, J. Lehikoinen, and P. Huuskonen. Understanding privacy regulation in ubicomp interactions. *Personal and Ubiquitous Computing*, 12(8):543–553, Mar. 2008.
93. B. Y. Lim, A. K. Dey, and D. Avrahami. Why and why not explanations improve the intelligibility of context-aware intelligent systems. In *27th international conference on Human factors in computing systems (CHI '09)*, page 2119, New York, New York, USA, 2009. ACM.
94. J. Lin, S. Amini, J. I. Hong, N. Sadeh, J. Lindqvist, and J. Zhang. Expectation and Purpose: Understanding Users Mental Models of Mobile App Privacy through Crowdsourcing. In *ACM Conference on Ubiquitous Computing (Ubicomp '12)*. ACM, 2012.
95. J. Lin, M. Benisch, N. Sadeh, J. Niu, J. Hong, B. Lu, and S. Guo. A comparative study of location-sharing privacy preferences in the United States and China. *Personal and Ubiquitous Computing*, 17(4):697–711, 2013.
96. J. Lin, B. Liu, N. Sadeh, and J. I. Hong. Modeling users' mobile app privacy preferences: Restoring usability in a sea of permission settings. In *Symposium On Usable Privacy and Security (SOUPS 2014)*, pages 199–212, Menlo Park, CA, 2014. USENIX Association.
97. B. Liu, M. S. Andersen, F. Schaub, H. Almuhimedi, S. A. Zhang, N. Sadeh, Y. Agarwal, and A. Acquisti. Follow my recommendations: A personalized privacy assistant for mobile app permissions. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, pages 27–41, Denver, CO, 2016. USENIX Association.
98. B. Liu, J. Lin, and N. Sadeh. Reconciling mobile app privacy and usability on smartphones: Could user privacy profiles help? In *Proceedings of the 23rd International Conference on World Wide Web*, WWW '14, pages 201–212, New York, NY, USA, 2014. ACM.
99. S. Loke. On representing situations for context-aware pervasive computing: six ways to tell if you are in a meeting. In *3rd Workshop on Context Modeling and Reasoning (CoMoRea '06)*, pages 35–39. IEEE, 2006.
100. P. Lukowicz, S. Pentland, and A. Ferscha. From Context Awareness to Socially Aware Computing. *IEEE Pervasive Computing*, 11(1):32–41, 2012.
101. S. T. Margulis. On the Status and Contribution of Westin's and Altman's Theories of Privacy. *Journal of Social Issues*, 59(2):411–429, June 2003.
102. N. Marquardt and S. Greenberg. Informing the Design of Proxemic Interactions. *IEEE Pervasive Computing*, 11(2):14–23, Feb. 2012.
103. G. Marx. Murky conceptual waters: The public and the private. *Ethics and Information technology*, pages 157–169, 2001.
104. M. Massimi, K. N. Truong, D. Dearman, and G. R. Hayes. Understanding Recording Technologies in Everyday Life. *IEEE Pervasive Computing*, 9(3):64–71, July 2010.
105. R. Mayrhofer and H. Gellersen. Shake Well Before Use: Intuitive and Secure Pairing of Mobile Devices. *IEEE Transactions on Mobile Computing*, 8(6):792–806, June 2009.

106. S. Moncrieff, S. Venkatesh, and G. West. Dynamic privacy assessment in a smart house environment using multimodal sensing. *ACM Transactions on Multimedia Computing, Communications, and Applications*, 5(2), 2008.
107. C. Neustaedter and S. Greenberg. The design of a context-aware home media space for balancing privacy and awareness. In *International Conference on Ubiquitous Computing (UbiComp '03)*. Springer, Mar. 2003.
108. H. Nissenbaum. Privacy as Contextual Integrity. *Washington Law Review*, 79(1):119–159, 2004.
109. H. Nissenbaum. *Privacy in Context - Technology, Policy, and the Integrity of Social Life*. Stanford University Press, 2009.
110. J. S. Olson, J. Grudin, and E. Horvitz. A study of preferences for sharing and privacy. In *CHI '05 extended abstracts on Human factors in computing systems*, New York, New York, USA, 2005. ACM.
111. R. Oppermann and A. Zimmermann. Context Adaptive Systems. *i-com*, 10(1):18–25, May 2011.
112. L. Palen and P. Dourish. Unpacking "privacy" for a networked world. In *Conference on Human factors in computing systems (CHI '03)*, pages 129–136, New York, New York, USA, 2003. ACM.
113. L. Pareschi, D. Riboni, A. Agostini, and C. Bettini. Composition and Generalization of Context Data for Privacy Preservation. In *Sixth International Conference on Pervasive Computing and Communications (PerCom '08)*, pages 429–433. IEEE, Mar. 2008.
114. S. Patil, Y. L. Gall, and A. Lee. My Privacy Policy: Exploring End-user Specification of Free-form Location Access Rules. In *Workshop on Usable Security (USEC '12)*, 2012.
115. S. Patil and J. Lai. Who gets to know what when: configuring privacy permissions in an awareness application. In *SIGCHI conference on Human factors in computing systems (CHI '05)*, page 101, New York, New York, USA, 2005. ACM.
116. B. A. Price, K. Adam, and B. Nuseibeh. Keeping ubiquitous computing to yourself: A practical model for user control of privacy. *International Journal of Human-Computer Studies*, 63(1-2):228–253, July 2005.
117. A. Rao, F. Schaub, N. Sadeh, A. Acquisti, and R. Kang. Expecting the unexpected: Understanding mismatched privacy expectations online. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, pages 77–96, Denver, CO, June 2016. USENIX Association.
118. N. A. Romero, P. Markopoulos, and S. Greenberg. Grounding Privacy in Mediated Communication. *Computer Supported Cooperative Work (CSCW)*, 22(1):1–32, 2013.
119. N. Sadeh, A. Acquisti, T. D. Breaux, L. F. Cranor, A. M. McDonald, J. R. Reidenberg, N. A. Smith, F. Liu, N. C. Russell, F. Schaub, and S. Wilson. The usable privacy policy project: Combining crowdsourcing, machine learning and natural language processing to semi-automatically answer those privacy questions users care about. techreport CMU-ISR-13-119, Carnegie Mellon University, 2013.
120. N. Sadeh, J. Hong, L. Cranor, I. Fette, P. Kelley, M. Prabaker, and J. Rao. Understanding and capturing people's privacy policies in a mobile social networking application. *Personal and Ubiquitous Computing*, 13(6):401–412, Aug. 2009.
121. R. Saleh, D. Jutla, and P. Bodorik. Management of Users' Privacy Preferences in Context. In *International Conference on Information Reuse and Integration*, pages 91–97. IEEE, Aug. 2007.
122. M. Satyanarayanan. Pervasive computing: vision and challenges. *IEEE Personal Communications*, 8(4):10–17, 2001.
123. F. Schaub. *Dynamic Privacy Adaptation in Ubiquitous Computing*. Phd dissertation, University of ulm, 2014.
124. F. Schaub, R. Balebako, A. L. Durity, and L. F. Cranor. A design space for effective privacy notices. In *Proc. SOUPS'15*, 2015.
125. F. Schaub, T. D. Breaux, and N. Sadeh. Crowdsourcing privacy policy analysis: Potential, challenges and best practices. *it-Information Technology*, 58(5):229–236, 2016.

126. F. Schaub, B. Könings, S. Dietzel, M. Weber, and F. Kargl. Privacy Context Model for Dynamic Privacy Adaptation in Ubiquitous Computing. In *6th International Workshop on Context-Awareness for Self-Managing Systems (CASEMANS '12), ACM UbiComp 2012 workshops*, pages 752–757, New York, New York, USA, 2012. ACM.
127. F. Schaub, B. Könings, P. Lang, B. Wiedersheim, C. Winkler, and M. Weber. Prical: Context-adaptive privacy in ambient calendar displays. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '14, pages 499–510, New York, NY, USA, 2014. ACM.
128. F. Schaub, B. Könings, and M. Weber. Context-adaptive privacy: Leveraging context awareness to support privacy decision making. *IEEE Pervasive Computing*, 14(1):34–43, 2015.
129. B. Schilit, N. Adams, and R. Want. Context-aware computing applications. In *Workshop on Mobile Computing Systems and Applications*, pages 85–90. IEEE, 1994.
130. B. Schilit and M. Theimer. Disseminating active map information to mobile hosts. *IEEE Network*, 8(5):22–32, Sept. 1994.
131. A. Schmidt. Implicit human computer interaction through context. *Personal Technologies*, 4(2-3):191–199, June 2000.
132. A. Schmidt. Context-Aware Computing: Context-Awareness, Context-Aware User Interfaces, and Implicit Interaction. In M. Soegaard and R. F. Dam, editors, *Encyclopedia of Human-Computer Interaction*, chapter 14, pages 1–28. The Interaction-Design.org Foundation, Aarhus, Denmark, 2012.
133. A. Schmidt, K. Aidoo, A. Takaluoma, U. Tuomela, K. Van Laerhoven, and W. Van de Velde. Advanced interaction in context. In *First International Symposium on Handheld and ubiquitous computing (HUC '99)*, pages 89–101. Springer, 1999.
134. A. Schmidt, M. Beigl, and H.-W. Gellersen. There is more to context than location. *Computers & Graphics*, 23(6):893–901, Dec. 1999.
135. B. Schwartz. The Social Psychology of Privacy. *American Journal of Sociology*, 73(6):741–752, May 1968.
136. J. Seifert, A. De Luca, B. Conradi, and H. Hussmann. TreasurePhone : Context-Sensitive User Data Protection on Mobile Phones. In *8th International Conference on Pervasive Computing (Pervasive '10)*, pages 130–137. Springer, 2010.
137. K. Shankar, L. J. Camp, K. Connelly, and L. Huber. Aging, Privacy, and Home-Based Computing: Developing a Design Framework. *IEEE Pervasive Computing*, 11(4):46–54, Oct. 2012.
138. K. Sheikh, M. Wegdam, and M. V. Sinderen. Quality-of-Context and its use for Protecting Privacy in Context Aware Systems. *Journal of Software*, 3(3):83–93, Mar. 2008.
139. P. Shi, H. Xu, and Y. Chen. Using contextual integrity to examine interpersonal information boundary on social network sites. In *SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*, page 35, New York, New York, USA, 2013. ACM.
140. S. Sigg. Context-based security: state of the art, open research topics and a case study. In *5th ACM International Workshop on Context-Awareness for Self-Managing Systems (CASEMANS '11)*, pages 17–23, New York, New York, USA, 2011. ACM.
141. D. Solove. Privacy Self-Management and the Consent Dilemma. *Harvard Law Review*, 126:1880–1903, 2013.
142. D. J. Solove. *Understanding Privacy*. Harvard University Press, 2008.
143. T. Strang and C. Linnhoff-Popien. A context modeling survey. In *First International Workshop on Advanced Context Modelling, Reasoning And Management, UbiComp '04*, 2004.
144. L. A. Suchman. *Plans and situated actions: the problem of human-machine communication*. Cambridge University Press, 1987.
145. E. Toch. Super-Ego: a framework for privacy-sensitive bounded context-awareness. In *5th ACM International Workshop on Context-Awareness for Self-Managing Systems (CASEMANS '11)*, pages 24–32, New York, New York, USA, 2011. ACM.

146. J. Y. Tsai, P. Kelley, P. Drielsma, L. F. Cranor, J. Hong, and N. Sadeh. Who's viewed you? the impact of feedback in a mobile location-sharing application. In *27th international conference on Human factors in computing systems (CHI '09)*, New York, New York, USA, 2009. ACM.
147. D. Vogel and R. Balakrishnan. Interactive public ambient displays: transitioning from implicit to explicit, public to personal, interaction with multiple users. In *17th annual ACM symposium on User interface software and technology (UIST '04)*, pages 137–146, New York, New York, USA, 2004. ACM.
148. R. Wenning, M. Schunter, L. Cranor, B. Dobbs, S. Egelman, G. Hogben, J. Humphrey, M. Langheinrich, M. Marchiori, M. Presler-Marshall, J. Reagle, and D. A. Stampley. The Platform for Privacy Preferences 1.1 (P3P1.1) Specification. W3c working group note, W3C, 2006.
149. A. F. Westin. *Privacy and Freedom*. Atheneum, New York, 1967.
150. P. Wijesekera, A. Baokar, A. Hosseini, S. Egelman, D. Wagner, and K. Beznosov. Android permissions remystified: A field study on contextual integrity. In *24th USENIX Security Symposium (USENIX Security 15)*, pages 499–514, Washington, D.C., 2015. USENIX Association.
151. S. Wilson, F. Schaub, A. A. Dara, F. Liu, S. Cherivirala, P. G. Leon, M. S. Andersen, S. Zimmeck, K. M. Sathyendra, N. C. Russell, et al. The creation and analysis of a website privacy policy corpus. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2016.
152. S. Wilson, F. Schaub, R. Ramanath, N. Sadeh, F. Liu, N. A. Smith, and F. Liu. Crowdsourcing annotations for websites' privacy policies: Can it really work? In *Proceedings of the 25th International Conference on World Wide Web*, WWW '16, pages 133–143, Republic and Canton of Geneva, Switzerland, 2016. International World Wide Web Conferences Steering Committee.
153. T. Winkler and B. Rinner. User-centric privacy awareness in video surveillance. *Multimedia Systems*, 18(2):99–121, July 2012.
154. R. Wishart and K. Henricksen. Context obfuscation for privacy via ontological descriptions. In *First International Workshop on Location-and Context-Awareness (LoCa '05)*. Springer, 2005.
155. R. Wishart, K. Henricksen, and J. Indulska. Context Privacy and Obfuscation Supported by Dynamic Context Source Discovery and Processing in a Context Management System. In *4th International Conference on Ubiquitous Intelligence and Computing (UIC)*, number 1, pages 929–940, 2007.
156. G. Yee. A privacy-preserving UBICOMP architecture. In *International Conference on Privacy, Security and Trust (PST '06)*, New York, New York, USA, 2006. ACM.
157. S. Zimmeck, Z. Wang, L. Zou, R. Iyengar, B. Liu, F. Schaub, S. Wilson, N. Sadeh, S. Bellovin, and J. Reidenberg. Automated analysis of privacy requirements for mobile apps. In *NDSS'17: Network and Distributed System Security Symposium*, 2017.