A Design Space for Privacy Choices: Towards Meaningful Privacy Control in the Internet of Things

Yuanyuan Feng
Carnegie Mellon University
Pittsburgh, PA, USA
yuanyuanfeng@cmu.edu

Yaxing Yao
Carnegie Mellon University
Pittsburgh, PA, USA
yaxingyao@cmu.edu

Norman Sadeh
Carnegie Mellon University
Pittsburgh, PA, USA
sadeh@cs.cmu.edu

ABSTRACT

"Notice and choice" is the predominant approach for data privacy protection today. There is considerable user-centered research on providing effective privacy notices but not enough guidance on designing privacy choices. Recent data privacy regulations worldwide established new requirements for privacy choices, but system practitioners struggle to implement legally compliant privacy choices that also provide users meaningful privacy control. We construct a design space for privacy choices based on a user-centered analysis of how people exercise privacy choices in real-world systems. This work contributes a conceptual framework that considers privacy choice as a user-centered process as well as a taxonomy for practitioners to design meaningful privacy choices in their systems. We also present a use case of how we leverage the design space to finalize the design decisions for a real-world privacy choice platform, the Internet of Things (IoT) Assistant, to provide meaningful privacy control in the IoT.

CCS CONCEPTS

• **Security and privacy → Usability in security and privacy**; • **Human-centered computing → HCI theory, concepts and models**; **Ubiquitous and mobile computing systems and tools**.

KEYWORDS

usable privacy, privacy choice, design space, Internet of Things

ACM Reference Format:
Yuanyuan Feng, Yaxing Yao, and Norman Sadeh. 2021. A Design Space for Privacy Choices: Towards Meaningful Privacy Control in the Internet of Things. In *CHI Conference on Human Factors in Computing Systems (CHI '21), May 8–13, 2021, Yokohama, Japan.* ACM, New York, NY, USA, 16 pages. https://doi.org/10.1145/3411764.3445148

1 INTRODUCTION

"Notice and Choice" as introduced in the Fair Information Practices Principles (FIPPs) [119] has become the de facto model for privacy protection around the world [120]. Privacy notices inform people about existing or potential data collection, use, and sharing practices regarding their personal data, while privacy choices provide people actual control over certain aspects of such data practices. Research also shows that privacy choices work best when people make informed privacy decisions based on effective privacy notices [63, 115]. However, privacy notice and choice in real-world systems often fall short in protecting people's data privacy [18, 105]. Many privacy notices are in the form of lengthy privacy policies, which often contain legal jargon and are hard for users to read and understand [42, 104]. To make matters worse, privacy choices are often absent [18], difficult to locate [51], and hard to understand [62] making it difficult, if not impossible, for users to effectively control the collection and use of their data. Even when privacy choices are available, the options are often limited and misaligned with people's concerns, leaving them unable to express their true privacy preferences, which in turn leads to frustration and a general sense of resignation [11, 18, 26, 75]. As a result, making privacy notices and choices more usable has emerged as an important research area in both the human-computer interaction (HCI) and privacy fields [113, 126].

Compared to the amount of user-centered research on privacy notices over the past two decades [2, 48, 49, 64, 85], existing research on designing privacy choices is scattered and less cohesive. One possible reason might have been the absence of clear legal requirements for privacy choices [18, 117, 118]. Recent data privacy regulations such as General Data Protection Regulation (GDPR) and California Consumer Privacy Act (CCPA) [15, 24, 93] have introduced new, more specific requirements for privacy choices. Yet, these new legal requirements vary greatly and their interpretations will require further clarification [39], leaving practitioners with insufficient guidance when it comes to designing privacy controls. In fact, regulations such as GDPR and CCPA directly refer to high-level usability concepts (e.g., GDPR's article 7 requirement that consent should be "freely given" and that requests for consent be given in an "intelligible and easily accessible form, using clear and plain language," or that "it shall be as easy to withdraw as to give consent"). These high-level concepts beg for HCI research and methodologies that provide more specific guidance to practitioners.

While some HCI and privacy research has been done on effective communication of privacy choices [1, 63, 89, 92, 106, 110], most of this work has been conducted in a piecemeal fashion, with limited attempts to more comprehensively map the different dimensions involved in communicating and providing access to privacy choices to users. In addition, most prior work has focused on Web and mobile privacy scenarios, with limited research conducted in the context of the Internet of Things environments, which are known to further exacerbate many of the challenges involved in allowing data subjects to effectively control their data [21, 82]. A primary focus

CHI '21, May 8–13, 2021, Yokohama, Japan

ACM ISBN 978-1-4503-8096-6/21/05.
https://doi.org/10.1145/3411764.3445148

of this work is to construct a comprehensive conceptual framework for privacy choices whose applicability extends to IoT contexts.

We argue that meaningful privacy choices should address several facets that extend beyond traditional considerations of usability. Thus, the second motivation of this work is to provide design guidelines to help researchers and practitioners to implement meaningful privacy choices in real-world systems.

To this end, we employed the design space methods to construct a design space for privacy choices from our systemization of knowledge [56]. Design space is a mapping of all dimensions of a subject matter, which is an effective approach to guide practitioners in designing new features related to the subject matter [100, 108]. We believe our design space for privacy choices not only fills the research gaps mentioned above but also helps practitioners identify privacy choice requirements and develop a comprehensive privacy choice concept to provide users meaningful privacy control.

In this paper, we first analyze how users interact with privacy notices and choices provided by real-world systems based on a technology review and iterative discussion sessions. This user-centered analysis sheds light on three possible relationships between privacy notice and privacy choice in real-world systems, situating our design space for privacy choices in relation to privacy notices. Then, we further examine the design space for privacy choices, organizing it around five dimensions, namely: type, functionality, timing, channel, and modality. We also offer rich examples and considerations for listed design options under each dimension. Finally, we present a use case of designing a privacy choice platform for IoT to demonstrate how practitioners can leverage the design space to provide system users with meaningful privacy control.

This work makes two main contributions. First, our user-centered analysis of privacy choices contributes a conceptual framework that considers privacy choice as a process and articulates its relationship with privacy notices in real-world systems. This framework connects previously scattered HCI and privacy literature on designing privacy choices under a cohesive umbrella. Second, the constructed design space supports the development of practical design guidelines organized around the different design dimensions we identify, thereby contributing to the development of meaningful privacy choices that are better aligned with people's privacy expectations, and also contributing to a framework to analyze regulatory requirements pertaining to the usability considerations.

2 DEFINE MEANINGFUL PRIVACY CHOICES

Privacy choice, also referred to as privacy control [117], is regarded as the primary goal of many data privacy regulations, which emphasizes individual privacy choice that is facilitated by adequate notice [18]. The assumption behind these regulations is that people's data privacy is protected when they have adequate individual control over their personal data in the form of privacy choices.

In practice, privacy choices are far from adequate compared to the Fair Information Practice Principles [17, 119] or the robust privacy rights defined in new privacy regulations [15, 24]. They predominantly take the form of "notice and consent". However, consent is merely one type of privacy choice, which is insufficient when people want to express more granular privacy preferences than "consent or not" could offer. Legal scholars argued that privacy choice has been overly simplified in legalistic mechanisms and thus failed its purpose to protect people's data privacy [18, 105, 118].

To construct a comprehensive conceptual framework for designing privacy choices, we avoid using varying and sometimes narrow legal definitions of privacy choice. Instead, we broadly define privacy choices in this paper as "the capabilities provided by digital systems for users to control different data practices over their personal data". Data practices here include but are not limited to the collection, process, disclosure, and retention of personal data. In this paper, we focus on privacy choices provided by digital systems, with an emphasis on ubiquitous IoT systems. This is because choices can also be exercised outside of a system or non-digitally. For example, people may choose to protect their data privacy by avoiding shopping at an Amazon Go store due to the store's overly invasive data practices [55].

We introduce the notion of **"meaningful privacy choices"** that extend beyond traditional usability considerations to include several facets that are more specifically tied to supporting users in making privacy decisions that capture their true privacy preferences. Some of these facets fall under traditional considerations of **"effectiveness"** (the ability to specify privacy choices that accurately and comprehensively align with the data collection and use practices with which a user feels comfortable) and **"efficiency"** (the ability to specify these choices with minimal time and effort). We believe that meaningful privacy choices should effectively accommodate people's diverse privacy preferences at the appropriate level granularity [11], which often go beyond the minimum legal requirements. Currently, many systems only offer "take-it-or-leave-it" choices that leave users little space to bargain about their privacy [118]. Such choices are ineffective because they restrict users from expressing their true privacy preferences. Also, meaningful privacy choices should enable users to efficiently configure available privacy options primarily by minimizing users' burden in the process. Many existing privacy choices are not efficient because they impose unrealistically high cognitive and time burden on users [26, 52, 125].

However, effectiveness and efficiency do not fully capture the complexity of supporting users in their privacy decision making. For one, users may not be aware of the choices available to them. This falls into the facet of **"user awareness"**, where meaningful privacy choices should be clearly conveyed to users in a prominent manner. More importantly, users may fail to fully understand the options they have to choose from and the potential ramifications of their decisions. This touches on the concept of "informed privacy decisions", namely that users should have a sufficient understanding of their options, how these options impact the collection and potential use of their data, including what might be inferred from this data and the possible consequences of these inferences. Research shows that disclosing detailed inferences of mobile apps using location data motivates users to re-evaluate and potentially adjust their mobile privacy settings [3]. Hence, **"comprehensiveness"** of privacy choices is another necessary facet to consider. Last but not least, how privacy choices are disclosed is often open to possible manipulation through what is now commonly referred to as "dark patterns" [92, 127]. We believe meaningful privacy choices should be free of potential manipulation or biased framing. Therefore,

the **"neutrality"** of privacy choices should also be evaluated for meaningful privacy control.

By defining meaningful privacy choices through the five facets described above, we propose new criteria for researchers and practitioners to think about the usability of privacy choices.

3 RELATED WORK

3.1 Improve the Usability of Privacy Choices

Both HCI and privacy research communities call for more usable privacy notices and choices for years [18, 26, 113, 117, 118]. Privacy notices are necessary to increase data privacy transparency but cannot guarantee privacy protection [2, 26]. The privacy choices available to people ultimately determine the level of protection to their data privacy. Also, when people cannot take actions to control their personal data privacy after being notified about certain data practices, they feel a higher level of privacy violation [98]. Their frustration on the lack of privacy choices may also lead to privacy resignation[22].

There is considerable amount of user-centered research on what makes privacy notices effective [54, 64, 85, 104] and how to design more usable privacy notices [48, 58, 114]. In comparison, the research on privacy choices does not provide the same level of design guidelines. Early research drew attention to the widespread problem of the absence of real privacy choices in cyberspace [18, 118]. In the US, online privacy choices governed by industry self-regulatory groups such as the Digital Advertising Alliance [34] fails to provide sufficient privacy protection [26], because consumers often do not know the existence of such privacy choices [47, 84, 124]. Research also found usability and noncompliance issues with many online privacy choices such as opt-outs for email communications and targeted advertising [33, 41, 65, 72]. More recently, the introduction of more stringent data privacy regulations like GDPR leads to greater availability of online privacy choices [32, 52], but many hastily implemented choices have poor usability. For example, privacy choices are often hidden in lengthy privacy policies or account settings that are difficult to find [52], and exercising privacy choices require too much user efforts [51]. So far, only a handful of research studies provide concrete design recommendations to improve the usability of privacy choices in specific areas such as mobile permissions [46] and online opt-outs [33, 51].

There are also user-centered studies examining how to effectively communicate privacy choices to users through different design elements, such as using pop-ups [92], certifications [10], labels [63], icons [89, 110], dashboards [106] and nudges [1, 4]. While generating insights on designing more usable privacy choices, these studies do not provide cohesive design guidelines for practitioners to implement meaningful privacy choices. Therefore, we aim to extend existing research on improving the usability of privacy choices by developing a comprehensive conceptual framework for meaningful privacy choices.

3.2 Towards Comprehensive Design Guidelines for Meaningful Privacy Choices

Current research on designing privacy choices is not comprehensive enough partially due to the varying and evolving regulatory landscape. Recent studies primarily evaluate privacy choice design options for specific regulations including GDPR [92, 110, 112, 125] and CCPA [28, 53], lacking cohesive design guidelines for privacy choices across different regulations. Additionally, existing work heavily focuses on web and mobile privacy choices [27, 28, 46], with limited considerations of the more challenging ubiquitous computing contexts [80, 130] such as the Internet of Things (IoT). IoT systems, often capable of location tracking, environmental sensing, and facial recognition, collect and use various potentially privacy sensitive information in the environments (e.g. people, time, location, activity) to deliver context-aware utilities [101, 102]. Implementing meaningful privacy choices in the IoT context is extremely challenging, not only because of technical challenges to secure data on various IoT devices and sensors [19, 35, 37], but also due to the ubiquity of IoT systems, the persistence of data practices, and practical constraints of limited user interface [21, 82]. Therefore, we aim to provide comprehensive design guidelines for privacy choices that address different regulatory considerations and the unique privacy challenges in IoT.

Specifically, we choose the design space approach, which is a set of well-established methods in design science and has been adopted in information systems [56, 100, 108] to guide system practitioners in designing new products or features. This approach is also used in privacy research studies to develop taxonomies to better map social network data types [107] and privacy-enhancing tools in web browsers [133]. Schaub and colleagues [114] has developed a design space for privacy notices to help system practitioners conceptualize and implement effective privacy notices for their systems. Their design space for privacy notices has a dimension of "control", where they suggested that such control should be implemented as privacy options relevant to the data practices disclosed in privacy notices. Their suggestion has its merits because integrated notice and choice helps users make informed privacy decisions [63]. However, one dimension in the design space for privacy notices fails to cover the full range of considerations that come into designing meaningful privacy choices as we discussed in Section 2. If practitioners only consider the design space for privacy notices, they may have the misconception that a system is privacy-friendly as long as the privacy notice includes certain privacy options to users, which has been the case in many privacy choices implemented by for-profit companies [118]. Therefore, we believe constructing a design space for privacy choices is a stride towards comprehensive design guidelines for meaningful privacy choices.

4 PRIVACY CHOICE AS A PROCESS

Privacy choices are predominantly implemented using simple mechanisms like "notice and consent" and "opt-in/out" to satisfy legal requirements. These mechanisms often assume that users interact with notice and choice sequentially as one action. In practice, how users interact with privacy choices is as straightforward as assumed and largely depends on various contextual attributes [91]. Therefore, we argue that a user's interaction with available privacy choices is a process consisting of complex, serendipitous, or sometimes recurring privacy decision making.

In this section, we first present a conceptual framework for the process of exercising privacy choices through a user-centered analysis of possible user interactions with privacy notice and choice

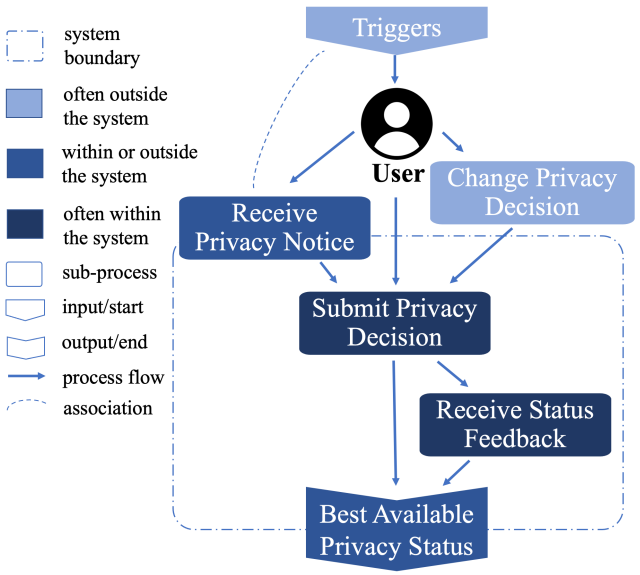

Figure 1: The Process of Exercising Privacy Choices

in real-world scenarios. Based on the process, we articulate three possible relationships between notice and choice that users may experience.

4.1 Analyze User Interactions with Privacy Choices in Real-World Scenarios

To develop a holistic design space for privacy choices, we start with a user-centered analysis to understand how users exercise available privacy choices together with their interactions with relevant privacy notices. We first conducted a technology review of commonly available privacy choices on the Web, on mobile devices, and with the Internet of Things technologies both in private and (e.g., smart home devices) and in public (e.g., Bluetooth location tracking). We further discuss the provision of the functionality of these available privacy choices. After identifying a list of existing real-world privacy choices and their available functionality, we brainstorm a wide range of real-world scenarios in which users interact with these choices. In the brainstorming sessions, we specifically considered various contextual attributes (e.g., temporal, spatial, social) that may affect user interactions with privacy choices, and whether users receive privacy notices or not before they exercise their choices. Then, we enumerated the paths that users may take to exercise their privacy choices, encompassing activities both within and outside the system that provides such privacy choices. After multiple iterations, we presented the process in a user-centered activity diagram (Figure 1) using workflow modeling [79] to and Unified Modeling Language activity diagram [9] methods.

4.2 The Process of Exercising Privacy Choices

Figure 1 outlines the complicated process when users exercise privacy choices in real-world scenarios. Our analysis shows that exercising privacy choices involves activities within and outside the system that provides the choices, so we used "sub-processes" to represent all user activities within or outside the system.

The diagram in Figure 1 includes several key elements relevant to the actor – a user who interacts with the system that provides privacy choices. The process follows the input-output model, which starts with an input that we refer to as "triggers" and ends with an output status where the user achieves the "best available privacy status" at a point in time. Below we explain all the key elements.

"Triggers", the element that initiates the process, is defined as an event or input of information experienced by the user. Triggers often lead the user to see a system's privacy notice or to make a privacy decision, which motivates the user to start the process of exercising privacy choice. Triggers vary in forms depending on real-world contextual factors. Triggers can be informational messages, such as an app permission prompt on a smartphone. They can be tangible things, such as a physical sign showing that "this area is under video surveillance." Triggers may also be social or ad-hoc events that prompt users to take further privacy actions, such as a conversation with friends about certain privacy settings that the user was unaware of. **"Best available privacy status"** is defined as the end state of a round of the process as perceived by the user. We use this element as the output of the process because the user's privacy preferences may not be fully fulfilled by the available privacy choices provided by the system at a point in time. If additional privacy choices are made available in the future, the user may start another round of the process to achieve a different best available privacy status.

The main elements of the diagram are four sub-processes between input and output of the process: **"receive privacy notice"** means the user receives the system's privacy notice, which could be provided by the system or from somewhere else; **"submit privacy decision"** means the user communicates a privacy decision to the system regarding available privacy choices; **"change privacy decision"** means the user changes their mind about previous privacy decisions, often outside the system; **"receive feedback"** means the user receives some form of feedback from the system about their privacy choice status. Because some triggers are related to the system's privacy notice, we use an association line to represent the potential link between triggers and "receive privacy notice".

Note that the user may not perform all these sub-processes in one round of the process. This can be caused by contextual factors (e.g., no time to read the privacy notice) or the design of the available privacy choices (e.g., absence of feedback of privacy choice status). All the elements are connected by process flows (i.e., arrow lines). Typically, the user transverses different combinations of sub-processes from the start point to the end point by following the arrow lines. The color shades of the elements indicate if the element takes place within the system (deep blue), outside the system (light blue), or both (medium blue). Furthermore, the process shown in Figure 1 is repeatable because it is common for people to change their minds about privacy decisions over time. Therefore, the diagram also encompasses the scenarios when the user changes their mind about previous privacy decisions.

Here is an example of a possible path. Alice (actor/user) sees a cookie banner (triggers) on a shopping website that she just opened. She clicks the button "more info" on the banner, which takes her to the privacy notice (receive privacy notice). After reading the

privacy notice, she makes a change to the default cookies settings to disable cookies for advertising purposes (submit privacy decision). However, the website does not show her any feedback on her change, so she can only assume the website has executed her new cookie settings (best available privacy status).

In summary, we depicted the complex process of exercising privacy choices in real-world scenarios based on a user-centered analysis. This process is essential to untangle how users interact with available privacy choices in relation to relevant privacy notices, which leads to our following discussion on the relationship between privacy notice and choice from a user-centered perspective.

4.3 The Relationship between Notice and Choice as Experienced by Users

"Notice and choice" are often considered together in privacy research [18, 26, 111]. However, our user-centered analysis above indicates that users may not interact with privacy notices when exercising privacy choices in many real-world scenarios, as shown in Figure 1. We further synthesize three types of relationship between notice and choice as experienced by users: decoupled, integrated, and mediated. This categorization helps distinguish our design space for privacy choices from that for privacy notices [114].

Decoupled. Privacy notice and choice are often decoupled in real-world scenarios, which means they are not necessarily communicated to users at the same time. Privacy notices can be delivered through physical signs, emails, or websites, often in the form of privacy policies [88]. However, many privacy notices do not contain sufficient or actionable information for users to exercise privacy choices available to them [26], and some merely disclose the absence of choice to be legally compliant [18]. A prevalent problem in privacy notices is the failure to convey meaningful choices that users can act upon [18, 118]. Also, notice and choice can be decoupled due to how users interact with them, particularly when users make privacy decisions without receiving effective privacy notices. As shown in Figure 1, privacy choice actions can be initiated by various triggers other than receiving privacy notices (e.g., A user changes privacy settings in a system with the help of a trusted friend). In short, the decoupled relationship between notice and choice can be caused by how notice and choice are designed in the system, or how users interact with such notice and choice.

Integrated. Notice and choice can be integrated if the system communicates them to users together or sequentially. This way, users can easily make decisions right after receiving privacy choices. The integrated model is recommended by privacy advocates [63, 115], since privacy notices are more effective when users have easy access to privacy controls that they can act upon right away [26, 63, 115]. In reality, integrated notice and choice are currently not widespread because many factors discourage system practitioners to implement integrated notice and choice in systems. First, privacy choices are not universally required by law worldwide. Even under relatively strict data privacy laws like GDPR, regulators are still refining guidelines on how data processors should integrate notice and choice with greater usability [39]. Second, users choosing privacy-preserving options may be perceived as a threat to certain business models that rely on the pervasive collection of users' personal information (e.g., data-driven advertising). As a result, there is not much incentive for many companies to provide integrated notice and choice. To promote more usable integrated notice and choice, a comprehensive solution requires joint efforts from regulators, industry, and the research community.

Mediated. Notice and choice can also be mediated by privacy-enhancing technologies such as privacy agents [13, 27, 122]. Though uncommon today, it is a promising technical approach to increase the usability of privacy choices in addition to legislative efforts. Currently, a central usability problem with the notice and choice approach to privacy is the high user burden for data privacy management. Particularly with new data privacy laws worldwide, digital systems involving certain data privacy practices are increasingly required to provide appropriate privacy choices. However, to effectively manage their data privacy, users would have to repeatedly enter and re-enter the process depicted in Figure 1, and submit numerous privacy decisions for countless systems. Such high user burden contributes to privacy fatigue [20] or resignation [22], where users often give up managing their data privacy at all. Research has shown that data-driven software privacy agents can reduce user burden by helping users understand privacy choices and assisting them in making privacy decisions that match their privacy preferences [43, 70]. For example, machine learning-based privacy assistants are shown to be effective to help users manage their Android permissions settings through personalized recommendations [77, 121]. Also, a recent research study has effectively identified and extracted opt-out links from official website privacy policies using natural language processing techniques [6]. Similar technologies can help users find available privacy choices in notices that would otherwise be too time-consuming or challenging for users themselves to read [42, 83]. By combining privacy-enhancing technologies such as privacy assistants and the automatic extraction of privacy choices from privacy notices, it is feasible to facilitate the mediated relationship between notice and choice. The mediate notice and choice facilitated by privacy-enhancing technologies are promising to help users effectively exercise privacy choices according to their diverse privacy preferences with significantly reduced user burden. Note that the mediated notice and choice requires a certain level of integration between notice and choice, such as following machine-readable privacy standards [25]. Therefore, the privacy-enhancing technology mediated notice and choice could be a more usable alternative to the integrated notice and choice currently recommended by privacy advocates.

5 A DESIGN SPACE FOR PRIVACY CHOICES

The purpose of the design space described in this section is to help researchers and practitioners better understand the key dimensions to be considered when designing privacy choices for their systems. The design space also provides a taxonomy to categorize, evaluate, and communicate different privacy choice design options with all involved stakeholders, including users and legal professionals.

To construct the design space, we followed the design space analysis methods [81, 86] and conducted iterative discussion sessions. In these sessions, we revisited our technology review on a range of web, mobile, and IoT systems to enumerate existing and proposed design options for privacy choices. Then we organized these design options under appropriate dimensions. We also incorporated our

user-centered analysis (Figure 1) to articulate the possible relationships between notice and choice in our design space.

We have identified five dimensions in the design space: two dimensions that are unique to privacy choices include **type** (what kinds of choices are offered) and **functionality** (what capabilities are offered to support the privacy choice process), and three similar dimensions shared with the previously proposed design space of privacy notices [114], which are **timing** (when the choice is provided), **channel** (how the choice and user's privacy decision of the choice are communicated), and **modality** (what interaction modes are used to deliver the choice and record user's privacy decision of the choice). We significantly extend Schaub et al.'s work [114] by updating three established dimensions in their design space and articulating these dimensions for privacy choices.

Figure 2 is a visual representation of our design space for privacy choices, which also outlines the three possible relationships between notice and choice discussed in Section 4. Note that the five dimensions in the design space should be considered in parallel rather than sequentially because different dimensions may impact one another. Under each dimension is a list of possible design options synthesized from our technology review. Note that these lists are not exhaustive and can be expanded to accommodate future novel systems or interactions and new data privacy regulations.

5.1 Type

Current privacy choices provided by systems are often limited in type largely due to regulations or existing industry practices [26]. However, many different types of privacy choices could be made available to users by systems. We argue that type is a unique innate dimension in the design space for privacy choices. We describe four major types of privacy choices, demonstrating that privacy choice in its broad definition can accommodate users' diverse privacy preferences beyond just being legally compliant. Note that the different types of privacy choices in this dimension are not mutually exclusive but often build upon one another. Practitioners may choose one or more types according to the system requirements and applicable legal requirements.

5.1.1 Binary Choice. Many existing privacy choices provided by systems are essentially binary, meaning users have to choose one out of two options. **"Notice and consent"** is a common type of binary choice. For example, population-specific privacy regulations in the US require explicit consent from patients [50] and guardians of children [59].for collecting and process data from these populations. However, "notice and consent" is a simplified version of notice and choice, where users are given a binary choice to the disclosed data practices. It is a convenient design option to achieve legal compliance but often does not provide users with real privacy choices [118]. Many digital systems adopt the same approach by combining privacy notices with relevant agreements to obtain user consent. Such user consent is often tied with the eligibility to use the system, and many users become habituated to ignore these agreements and provide consent anyway [12]. In digital systems, this type of binary consent design limits users' ability to express their privacy preferences. Since users are often motivated to proceed and use the system, this design often pressures users into consent and thus deprives them of richer privacy choices that they may otherwise be entitled to. Therefore, the binary "notice and consent" design is only appropriate when there are straightforward legal requirements.

"Opt-in/out" is another common type of binary choice. Different from "notice and consent", opt-in/out choices are not necessarily tied to the eligibility to use the system. Users can still use the system but have more options in deciding if and how their personal data is collected and used.However, the opt-in/out choices made by the user may affect their ability to use the full functionality of the system, such as penalization services made available by processing user data. When designing opt-in/out choices, the **default value** is a critical design decision. Opt-in as the default means certain data practices are allowed until the user indicates otherwise. This usually means more data practices could happen in the background as users seldom change their privacy defaults [127, 129]. Opt-in as the default is often in the interest of system providers and may enable certain personalized services for users. Contrarily, opt-out as the default means certain data practices cannot happen until the user allows them. This is more privacy-preserving and backed by many privacy advocates and legal researchers [118]. The regulatory requirements on the default value differ worldwide. For example, GDPR requires an opt-out default for certain data practices [24] while CCPA allows the default to be opt-in for sales of California residents' personal data [93]. Binary opt-in/out choices are usually sufficient to achieve legal compliance and relatively easy to implement in a system but fall short to accommodate users' diverse privacy preferences.

5.1.2 Multiple choices. Multiple choices provide users more than one privacy options to choose from. They can be implemented as the sum of multiple binary choices. A come example is the GDPR-compliant cookie banners that have been increasingly adopted by websites using cookies to track user data. These cookie banners provide users several binary opt-in/out choices to allow or disable cookies for different purposes (e.g., strictly necessary, performance, advertising) [23]. Multiple choices can also be non-binary privacy choices. For example, mobile platforms (e.g., iOS and Android) provide users several options for an app's access to location data collected by the device, including "always", "while using the app", "never", and more recently on iOS "just once" [131].

Generally speaking, multiple choices provide users with more options to better capture and accommodate their diverse privacy preferences, which is increasingly required by new data privacy regulations and recommended by privacy advocates. However, implementing them often requires more effort in system design and development compared to binary choices.

5.1.3 Contextualized choices. We define contextualized choices as context-specific privacy choices that are often implemented as a combination of binary and multiple choices in context. Contextualized choices stem from the contextual integrity framework to understand privacy [8, 91]. In complicated, inter-related systems, data privacy must be considered in context (e.g., time, location, purpose) against data collection and use practices [7, 91]. For example, contextualized choices can be fine-grained privacy settings that allow inhabitants of a smart building to share the occupancy status of their offices during working hours but not in after-hours. [97]. According to CI, A privacy violation can be identified when the

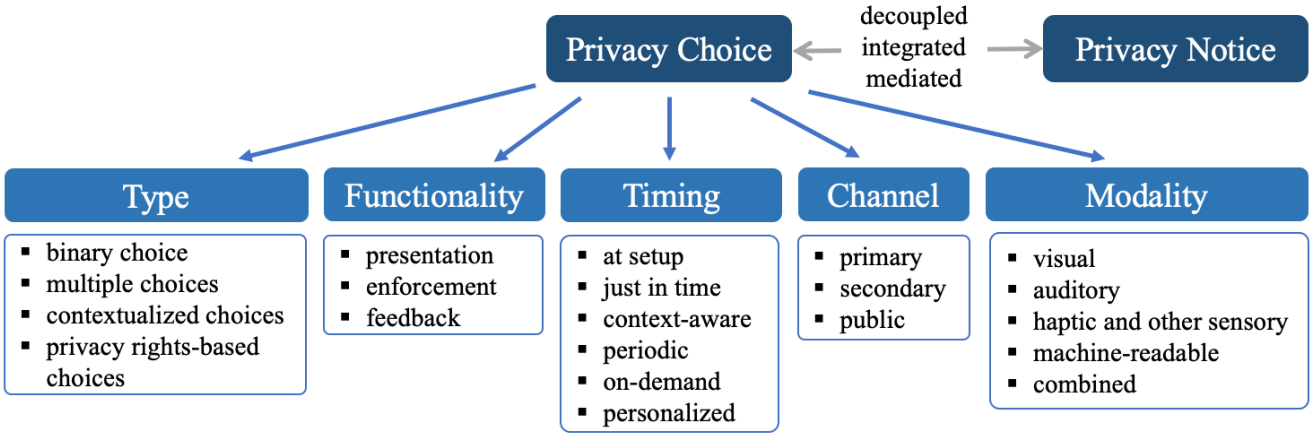

Figure 2: A Design Space for Privacy Choices

data transmission principles diverge from certain social norms [91]. Ideally, contextualized choices should be provided to users when potential privacy violations are likely to happen. Contextualized choices can better align people's diverse privacy preferences, but implementing them faces several challenges. First, the system needs to support fine-grained contextual attributes to contextualized choices (e.g., the ability to allow/deny data collection for a period of time at a specific location). Second, the system needs to collect additional data to distinguish different contextual attributes, introducing the trade-off between sacrificing certain data privacy and gaining more granular privacy control. Third, contextualized choices are likely to overwhelm users with countless privacy decisions, particularly considering the exponentially increasing number of deployed IoT systems and their ubiquitous data practices. Fortunately, software privacy agents could mitigate such user burden by selectively delivering contextualized choices according to users' privacy preferences [4]. Nevertheless, to achieve the vision of smart buildings and smart cities, contextualized privacy choices are necessary to protect people's data privacy in IoT.

5.1.4 Privacy rights-based choices (access, rectification, erasure, portability, etc. There is an emerging type of privacy choices that support the robust privacy rights recognized by new data privacy regulations, which we refer to as privacy rights-based choices. These are relatively complicated choices beyond the capability of binary or multiple choices, which often require additional communication between users and systems in the form of various user requests. Take GDPR [24] as an example, data subjects have the right of access, which allows them to request a copy of their personal data undergoing processing (article 15). Enabling access can be done at different cost levels: systems can provide an interface for users to download their own data (e.g., Facebook), which has a high upfront development cost; systems can build simple request mechanism and only respond to users who send in requests, which costs less if the request number is low. GDPR also acknowledge the right to rectification (article 16), where data subjects can request data controllers to correct inaccurate or incomplete personal data about them, the right to erasure, also known as "the right to be forgotten" [109](article 17), where data subjects can request data controllers to delete their personal data being collected, as well as the right to data portability(article 20), where (if requested) data subjects should receive their personal data in a commonly used machine-readable format that can be easily transferred into other systems or applications.

Privacy choices supporting these robust privacy rights often cannot be fulfilled by systems immediately due to the complexity of user requests. Currently, only large corporations with resources can implement full-fledged privacy rights-based choices[134]. For organizations with limited resources, an alternative is to rely on privacy choice platforms (e.g., OneTrust) to handle users' privacy rights-based requests in order to provide this type of choices.

5.2 Functionality

Our user-centered analysis on the process of exercising privacy choices calls for a dimension that captures the functionalities needed to support different aspects of the process. The functionality dimension distinguishes the design space for privacy choices from that for privacy notices because most notices can be considered as a single function (i.e., presentation of privacy-relevant information). We have initially identified three items under the functionality dimension that meaningful privacy choices should offer. This list is not exhaustive and more items (e.g., future novel privacy features) can be added to this dimension as long as they support meaningful privacy choices as defined in Section 2.

5.2.1 Presentation (of privacy choices). A system must present available privacy choices to users. This is the indispensable functionality shared by privacy notices and privacy choices. Ideally, the presentation of privacy choices should be easy to understand. Users should be clear about what data practices could happen, what options they have, and how to tell the system about their privacy decisions. As a result, the presentation of privacy choices may contain multiple components, which are often integrated with related privacy notices so that users can fully understand the choices they have. To decide how to effectively present privacy choices to users,

other dimensions of the design spaces (e.g. timing, channel, modality) should be thoroughly considered against opportunities and constraints of the system.

5.2.2 Enforcement (of users' privacy decisions). Since the process of exercising privacy choices involves users' privacy decisions, a system must be equipped with the functionality to enforce the different privacy decisions submitted by users regarding the available privacy choices. The enforcement often includes several sub-functions. First, to enforce a user's privacy decisions, appropriate authentication with the user is often required. This ensures that a system correctly enforces the privacy decisions of a specific user and minimizes security risks like identity theft. Second, the actual enforcement actions can be fully automated or mediated by humans (e.g. customer service personnel). Fully automated enforcement requires significant up-front system development efforts but costs less if the system scales up. Human-mediated enforcement can be more attentive to users' individual needs but its human resource cost can be high. Third, since users' privacy decisions are not one-time actions, the system should be able to record any changes to users' privacy decisions and enforce them in a timely manner. Note that the functionality to enforce users' privacy decisions is technically more challenging for complicated types of privacy choices, such as contextualized choices.

5.2.3 Feedback (of privacy choice status). Since privacy choice is a process, it is crucial for the system to provide accurate feedback in accordance with user actions. Some privacy decisions made by users can be executed by the system immediately, so it is crucial for the system to provide timely feedback indicating users' privacy settings have been adjusted based on their most recent actions. Such feedback of privacy choice status is not only a heuristic in user interface design [90], but also an important sub-process in Figure 1. Some privacy decisions from users require additional processing time, such as a user's request for a copy of their personal data collected by the system. Some systems may not be able to immediately provide such data to the user and need more time to fulfill the request. In this case, the best practice for the system is to provide the current status of the request and update the user with new statuses (e.g. when the data is ready and how the user can obtain the data). Overall, providing clear and timely feedback about the privacy decisions made by users is a crucial usability functionality of privacy choices.

5.3 Timing

Timing impacts the effectiveness of privacy notice [38] and subsequently affects how users engage in privacy decision making [49, 99]. Decoupled notice and choice affect people's perception of potential privacy risks [2, 98], causing mismatches between their privacy decisions and their privacy preferences. To choose appropriate timing to deliver privacy choices, a range of factors should be considered including the type of privacy choices required, users' primary task at hand, the availability of notice, and other contextual factors. Since timing is a shared dimension of both notice and choice, we explain the potential relationship between notice and choice (i.e., decoupled, integrated, or mediated) for each timing design option listed below.

5.3.1 At setup (often integrated). Privacy choices can be provided when the user interacts with the system for the first time, often along with relevant privacy notices. This timing design is often associated with the consent type of privacy choices, such as software license agreements during installation or terms of use during account signup. Delivering privacy choices at setup has two potential benefits: (1) Users can have the opportunity to review privacy notices and make privacy choices regarding the system before using it; and (2) Systems can front-load obtaining user consent to meet applicable legal requirements. However, privacy choices at setup share similar drawbacks of "notice and consent", which push users to choose the option that allows them to use the system. A study shows that many users regret their choice later after agreeing to the software license agreement displayed at setup [49].

5.3.2 Just-in-time (often integrated). Privacy choices can be presented to users when the specific data practice is about to happen, which is often integrated with relevant privacy notices. A common example is the "ask on first use" (AOFU) model of app permissions on mobile platforms(e.g., Android, iOS). Under AOFU, when a mobile app requests to access certain data (e.g., location, microphone) on a mobile device for the first time, the mobile platform will ask the user for permission via a pop-up dialog box. The just-in-time design is advantageous because it allows users to make privacy decisions in the actual context that better matches their privacy preferences [44]. The disadvantages are the interruption to the user's task at hand [98] and potential privacy fatigue if there are too many just-in-time decisions to be made [20].

5.3.3 Context-aware (often integrated). The specific temporal, spatial, or social context that the user is in can be leveraged to determine the timing of privacy choices. This timing design is particularly suitable for contextualized privacy choices described in the type dimension. Context-aware timing is more relevant to the user's situation and makes privacy choices more meaningful if integrated with a tailored privacy notice in that specific situation. For example, in a smart building equipped with a multi-sensor indoor location tracking system that can track visitors' real-time whereabouts in the building [87], it is ideal to deliver privacy choices with a tailored privacy notice of the building's tracking practices before prospective visitors enter the building [97]. However, the key challenge for context-aware timing design is detecting relevant contexts. Accurate detection of relevant context requires not only technical solutions to identify various contextual attributes, but also data about users' privacy preferences to determine what contextual attributes are most relevant to individual users.

5.3.4 Periodic (integrated or decoupled). Privacy choices can be shown to users multiple times. Particularly, when a system's data practices change, it is important and sometimes legally required to provide periodic privacy notice and choice [29]. Delivering privacy choices periodically also accommodates cases where users change minds about previous privacy decisions. Periodic privacy choices can be integrated with or decoupled from privacy notices, while the former is recommended by letting users re-evaluate their privacy decisions. Privacy choices can also be delivered periodically according to other appropriate criteria. For example, iOS 13 reminds iPhone users of apps having background location access

and prompts users with the location permission choices again for them to reconsider [57]. However, periodic privacy choices may increase user burden, resulting in privacy fatigue [20, 61] or habituation [14, 60]. Ideally, users should have a say in the frequency of these periodic privacy choices.

5.3.5 On Demand (integrated or decoupled). The timing design options above pertain to systems that actively deliver privacy choices to users. Users can also actively seek available privacy choices to submit their privacy decisions on demand. Figure 1 shows that people's privacy decisions can be influenced by triggers outside the system, sometimes without receiving privacy notices. As a result, on-demand privacy choices can be integrated with or decoupled from privacy notice depending on the user's actual interactions with notice and choice. For on-demand privacy choices to be meaningful, users should be able to easily locate available privacy choices, instead of spending unrealistic efforts on finding them [51].

5.3.6 Personalized (integrated, decoupled, or mediated). When to present privacy choices to users can be personalized through a combination of the timing mechanisms tailored to individual user's preferences. This could be additional settings that allow users to choose the preferred timing for privacy choices or more sophisticated context-aware solutions. Personalized solutions are advantageous because one-size-fits-all solutions towards privacy poorly align with people's diverse privacy preferences [76]. However, asking for users' privacy decisions every time or periodically introduces a higher user burden. One promising approach is to provide data-driven personalized nudges [4] or recommendations [77] to assist users in making privacy decisions. Personalized privacy choices can be integrated with or decoupled from relevant privacy notices depending on user preferences. This design option also supports the mediated relationship between notice and choice, where personalized privacy agents can decide when to present users with the most relevant privacy choices.

5.4 Channel

Different channels can be used to present privacy choices to users and communicate users' privacy decisions back to systems. We adopt the same categorization of channels in [114]: primary, secondary, and public. Note that privacy choices must consider channels for two-way communication between systems and users, compared to the one-direction communication of privacy notices. Systems may choose the most appropriate channel or leverage multiple channels to support meaningful privacy choices.

5.4.1 Primary. Primary channel refers to the same platform or device the user interacts with the system, such as a website's cookie settings are presented on the website. Using a primary channel means that privacy choices are embedded in the user's interaction with the system, enabling users to make privacy decisions within the context of the system [115]. Hence, primary channels are usually preferred for privacy choices between systems and users. However, primary channels are limited in the IoT context due to lack of user interfaces [21] and technical challenges. For example, IoT sensor-based smart building systems capable of collecting presence and environmental data lack explicit user interfaces on their privacy channel – IoT sensors [87]. Also, the primary channel for a smart home speaker is the voice user interface but recording users' privacy decisions via voice commands is error-prone due to challenges in automatic speech recognition [94].

5.4.2 Secondary. When systems are limited in their primary channels to deliver privacy choices or receiving users' privacy decisions, providing privacy choices via a secondary channel is recommended. An example of leveraging secondary channels is Amazon Echo smart speakers. Although users of Echo devices can delete their recordings by speaking to the voice assistant, the full set of privacy choices are still delivered on the Alexa privacy settings page via its website or mobile app [71]. Secondary channels that are already widely adopted by users, such as websites and mobile apps, are particularly suitable for privacy choices in the IoT context.

5.4.3 Public. Public channels have long been used to deliver privacy notices, such as physical signage in public places. This physical signage can be leveraged to point users to available privacy choices delivered through other channels, such as the video surveillance signage guideline for GDPR compliance [39]. Public channels can also fill the privacy gap when the data subjects are not the users of the system (e.g. passers-by, incidental users), which are increasingly required by new data privacy regulations as part of data subjects' privacy rights. However, public channels are often limited in the amount of information being communicated. Obviously, the physical size of a sign determines how much content can be shown, meaning not all privacy-related information can be fully conveyed. Therefore, privacy choices through public channels require support from other channels, such as a layered approach, to deliver privacy-related information [40], where a concise summary is offered via a public channel with clear information to access the actual choices via a different channel.

Public channels can also efficiently communicate users' privacy decisions to systems with low user burden, under the premise that standardized privacy formats are widely adopted. In the web context, the Platform for Privacy Preferences (P3P) [25] enables users to specify their privacy decisions in P3P's machine-readable format to be communicated back to web systems that support P3P [27]. In the IoT context, privacy choices can be conveyed via public markers [103] or beacons [66]. Privacy beacons[67] is a particularly promising approach because users can broadcast their predefined, standardized privacy preferences via public IoT channels (e.g., Wi-Fi, Bluetooth). Systems supporting this approach can automatically record and execute users' privacy preferences without bothering users to make repetitive privacy decisions for different systems.

5.5 Modality

Similar to the channel dimension, different modalities can be leveraged to facilitate the two-way communication of privacy choices between systems and users. The type of privacy choices, the user's likely attention level, as well as the opportunities and constraints of the system should be considered to determine the most appropriate modality. Accessibility issues and possible user distraction are other important aspects to consider in privacy choice modality design. Not all modalities are available to systems at all times, it might be ideal to use multiple modalities to support different functionalities during the process of exercising privacy choices.

5.5.1 Visual. Privacy choices are commonly delivered visually, often in the form of text, images, icons, signage, or a combination thereof. Visual privacy choices can be text-based, which allows systems to provide clear descriptions of the choices available to users and obtain affirmative privacy decisions from users when required. However, linguistic properties including the framing, the length, and the use of jargon, affect people's comprehension of the choices and their ability to make appropriate privacy decisions [2, 72]. For text-based privacy choices to be meaningful, the specific wording of the choice requires design attention and user testing [5]. There are also image-based or icon-based visual privacy choices that are relatively intuitive if users are familiar with them [45]. However, using images or icons to represent complex privacy concepts may cause user confusion if not well-designed [73]. Visual methods can also be used to communicate users' privacy decisions back to the system. Take digital video surveillance cameras as an example, it is technically feasible for people to use visual markers (e.g. special colored apparels) to convey their privacy preferences to the Respectful Camera System [116]. Although most visual privacy choices are communicated digitally today (e.g., emails, websites, mobile apps), they can also be delivered via a wide range of tangible media (e.g., paper consent forms and physical signage). Tangible interactive kiosks may also be used for users to submit their privacy decisions in certain IoT smart building scenarios.

5.5.2 Auditory. Auditory privacy choices are primarily delivered via spoken words or various sounds. Delivering privacy choices via spoke words is important because it provides an accessible option for the blind and visually impaired community [36, 74]. Sound by itself cannot fully communicate complex privacy concepts, but can serve as alerts or reminders for otherwise invisible privacy choices [132]. The limitation of sounds is that their meanings need to be learned, so it is recommended to use sounds that are familiar to users or widely accepted in a specific culture. Auditory privacy choices are promising in the IoT context as voice assistant-controlled smart home devices and appliances gain popularity. Given accurate speech recognition and appropriate user authentication, auditory methods can communicate users' privacy decisions back to IoT systems [71].

5.5.3 Haptic and Other Sensory. Although rarely available, haptic and other sensory methods could be potentially useful to communicate privacy choices between systems and users. Since haptic and sensory methods are highly abstract and difficult to convey a large volume of information efficiently, they should be used in addition to other modalities for best outcomes. For example, using haptic signals is a non-intrusive way to draw users' attention to important privacy choices that need their attention right away. Other sensory methods, such as specialized gestures or body motions that are agreed upon, can be leveraged to convey users' privacy decisions to the systems [96] if the system can detect specialized gestures from users. However, a textual user guide is often needed to define what a particular gesture or body motion means.

5.5.4 Machine-readable. The above modalities engage different senses of the user and require user attention. An alternative supported by digital systems is the machine-readable modality. This means a system's data practices and available privacy choices are encoded in a machine-readable format, which can be communicated to other systems automatically. Machine-readable privacy choices are also the foundation for software agents that help reduce users' privacy management burden [43, 70, 77]. Privacy agents or assistants can communicate users' privacy preferences to systems and negotiate privacy choices on users' behalf at different automation levels [22]. The machine-readable format is also the premise for the mediated relationship between notice and choice. P3P is an early machine-readable standard in web context [25] that enabled a privacy agent prototype Privacy Bird [27]. Although P3P failed to achieve wide adoption, machine-readable modality coupled with privacy agents may be appealing in IoT context where modalities based on human senses are limited [115].

5.5.5 Combined. Systems can leverage multiple modalities to provide meaningful privacy choices. We envision a semi-automated software privacy agent that can help its user configure available privacy choices offered by a system in a machine-readable format based on the user's predefined privacy preferences. In cases where the privacy agent cannot configure certain privacy choices, it sends a visual or haptic notification to the user to require additional input (e.g., privacy decisions). In summary, a combination of multiple modalities can provide users with more smooth, less burdensome experiences when they interact with privacy choices.

6 USE CASE: A PRIVACY CHOICE PLATFORM FOR THE INTERNET OF THINGS

In this section, we present a use case on how we leverage the design space to design a privacy choice platform for IoT – the **IoT Assistant (IoTA)** app[1]. We describe and evaluate our design decisions for IoTA to support meaningful privacy control in IoT.

6.1 Towards Meaningful Privacy Control in IoT

It is challenging to implement privacy notice and choice for IoT systems [16, 82]. Unlike large technology companies, smaller providers of IoT applications and services often have limited resources to develop full-fledged privacy choice components for their IoT systems to be compliant with new data regulations like GDPR [134]. Also, organizations or individuals who purchase and deploy these IoT technologies are data controllers or data processors in certain circumstances under GDRP [68]. However, these small stakeholders have few tools to provide data subjects with appropriate privacy notice and choice regarding the IoT technologies they own. Therefore, we aim to build a privacy choice platform for deployed IoT systems to support meaningful privacy control in IoT. This platform is particularly valuable for small stakeholders involved in IoT data collection and processing, where they can take advantage of our platform to achieve legal compliance around IoT data privacy.

The platform will deliver relevant privacy notices together with available privacy choices to users, following the best practices [69, 114]. The information included in these notices and choices comes from a privacy infrastructure implemented by our research team [30]. This privacy infrastructure provides a mechanism for stakeholders of IoT systems (manufacturers, service providers,

[1] At the time of publication, a stable version (1.2.3) is available to download in more than 30 countries at iOS App Store and Google Play Store (search term: "IoT Assistant").

and owners, etc.) to register deployed IoT systems. Using the privacy infrastructure, stakeholders can easily describe an IoT system's deployed location, the approximate range of data collection, and privacy-related information in a machine-readable format. IoTA communicates with the privacy infrastructure and delivers available notices and choices of these registered IoT systems to users. Note that, to leverage our privacy choice platform, stakeholders of IoT systems need to achieve a certain level of integration with the privacy option management component of the privacy infrastructure.

6.2 Requirements, Advantages, and Constraints

6.2.1 Requirements. We identified three high-level system requirements to be achieved by the privacy choice platform for IoT, which ultimately guide our design decisions.

- To help stakeholders navigate and potentially achieve compliance with new data privacy regulations.
- To support users in exercising their privacy choices through the different paths shown in Figure 1.
- To integrate with various IoT systems to deliver their available privacy choices to the platform users and communicate users' privacy decisions back to the IoT systems.

6.2.2 Advantages and Constraints. We have technical advantages to build the privacy choice platform thanks to the underlying privacy infrastructure [30]. First, we have access to the growing amount of structured information about deployed IoT systems (e.g., geographic location, data practices, the availability of privacy choice). Second, the privacy infrastructure helps the platform handle the two-way communication of privacy choices and users' privacy decisions with integrated IoT systems. We also face several constraints. First, the platform relies on the information provided by stakeholders of IoT systems. The completeness and the accuracy of such information cannot be guaranteed, so a verification mechanism is needed. Second, the privacy choices available to users on the platform are managed and enforced by stakeholders of the IoT systems. The platform has no control over these IoT systems' actual data practices, so an auditing mechanism is needed. Finally, as a small research team, we are constrained by our resources to implement advanced functionality in the early versions of IoTA.

6.3 Design Decisions for the IoT Assistant

Given the requirements, advantages, and constraints above, we considered the overall feasibility of all design options under each dimension of the design space and arrived at the following high-level and detailed design decisions.

6.3.1 High-level Design Decisions. We first decide that the privacy choice platform should be in the form of an application on mobile devices. With 3.5 billion people and 44.8% of the world population are smartphone users [123], a mobile app is suitable for a privacy choice platform that aims to serve a large number of people to take control of their data privacy in IoT. Also, various capabilities (e.g., GPS location, motion sensor, Bluetooth, notification) embedded in mobile devices provide us the flexibility to consider more design options under each dimension. Therefore, we implement the platform as a location-aware mobile app with a map-based main interface, as shown in Figure 3 (a). Second, we decide to rely on the privacy infrastructure for IoT because the advantages offered by the privacy infrastructure significantly outweigh the constraints. Also, the constraints can be mitigated by adopting content moderation practices and integrating with other IoT auditing technologies [95] in the future. Finally, we decide to implement integrated notice and choice. IoTA offers a concise privacy notice based on the machine-readable information from the privacy infrastructure, along with any available privacy choices of the IoT system, as shown in Figure 3 (b) (c). This decision is crucial because effective privacy notices help users make informed privacy decisions that match their preferences.

6.3.2 Detailed Design Decisions by Dimension. We also considered all possible design options under each dimension in the design space and arrived at the detailed design decisions below.

Type: multiple choices and privacy rights-based choices. To help stakeholders of IoT systems navigate and potentially achieve compliance with new data privacy regulations, IoTA offers four privacy options with flexibility to accommodate different regulatory requirements, namely *data collection*, *data sharing*, *request copy of data*, and *request data deletion*, as shown in Figure 3(d). The *data collection* and *data sharing* choices are binary in nature but together they are multiple choices for users to control potential IoT data practices. We use allow/deny as options for these two choices without dictating a default value, where stakeholders of IoT systems can set their default value (opt-in or opt-out) according to the applicable regulation. Note that the *data sharing* choice can be adjusted to serve the "do not sell my information" opt-out required by CCPA [93]. In addition, we provide two privacy rights-based choices – *request copy of data* and *request data deletion* – to initially support access, portability, and erasure rights in GDPR [24]. The responsibility to enforce these privacy rights-based choices and to provide feedback on privacy choice status falls on the individual IoT systems. IoTA only displays the privacy choice status provided by the system to users. In summary, we do not claim that IoTA supports all types of privacy choices required by different regulations. Instead, we consider IoTA as a valuable start to include multiple types of privacy choices with the possibility to expand further.

Functionality: presentation and feedback. To support different paths shown in Figure 1, IoTA implements the functionality of presentation and feedback under this dimension. First, IoTA presents available privacy choices of an IoT system to users in a machine-readable format along with a concise privacy notice. Currently, registered users can click the "manage" button to access a web page (managed by specific IoT stakeholders) to submit their privacy decisions. Second, for each type of privacy choice, IoTA provides a list of standardized statuses to indicate the current status of this privacy choice. Whenever a user makes choices and clicks the "refresh" button, IoTA fetches the latest status from the IoT system to provide users timely feedback. Note that the enforcement of users' privacy decisions is not supported by IoTA due to our constraints, and the responsibility of enforcement lies in the stakeholders who deploy IoT systems. In summary, IoTA offers a simple interface for users to interact with available privacy choices of various IoT systems within one mobile app.

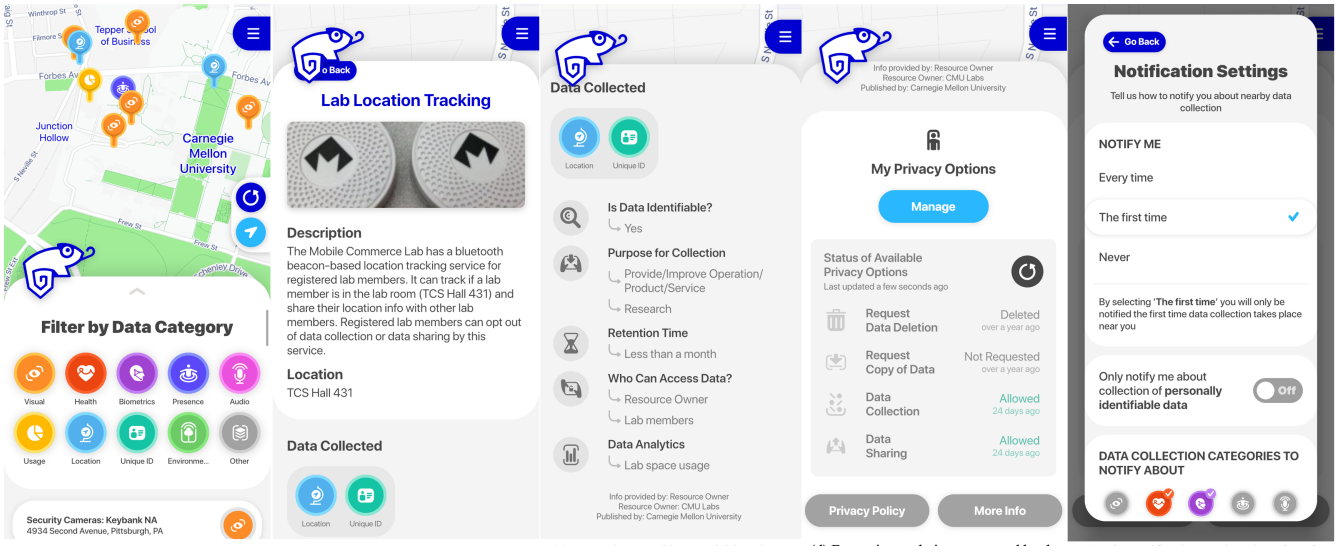

(a) The map view of the main interface for users to explore IoT systems by location

(b) An IoT system called "lab location tracking" and a brief description of it

(c) A concise, machine-readable privacy notice for "lab location tracking"

(d) Four privacy choices supported by the IoT Assistant (v1.2.3) with status updates

(e) The notification settings interface for users to adjust notification preferences

Figure 3: The Main Interfaces of the IoT Assistant (IoTA) App

Timing: context-aware, personalized, and on demand. IoTA uses several timing designs to deliver integrated privacy notice and choice to users. Users can choose to receive context-aware (i.e., location-based) privacy notice and choice via notifications on their mobile devices. Notification is done by fetching information of nearby IoT systems based on the location of their mobile device. IoTA also supports personalized timing to some extent. Using the notification settings in Figure 3(e), users can choose a notification frequency and customize what to be notified about by data category. Note that our team is actively researching improved methods to further reduce user burden to configure their notification settings, including leveraging data-driven privacy profiles [76, 78] to provide recommended settings that match users' diverse privacy preferences. Finally, the map-based interface in IoTA as shown in Figure 3(a) allows users to explore potentially privacy-sensitive IoT systems in other places by moving the map. Users can configure any available privacy options on demand for IoT systems afar.

Channel: Secondary and Public. For deployed IoT systems, IoTA is a secondary channel to communicate privacy choices between the systems and their users. One requirement for any platform that offers privacy choice is the ability to integrate with a wide range of IoT systems that may use a diversity of primary channels. As such, it is more feasible to choose a widely used secondary channel (i.e., a mobile app) that can be easily integrated with other technologies and is capable to leverage other channels such as Bluetooth and Wi-Fi. Additionally, we have recently implemented the initial QR code functionality for easy discovery of individual IoT systems, which leverages the public channel. Stakeholders can request a unique QR code for their deployed IoT system and publicize the presence of the IoT system and its available privacy choices via the QR code. The general public can scan the QR code using their phones: IoTA users will be directed to the integrated privacy notice and choice of the specific system in IoTA, while non-IoTA users will be prompted to download IoTA on their phones to continue.

Modality: Combined: The IoTA currently uses a combination of modalities to deliver IoT systems' available privacy choices to users and communicate users' privacy decisions back to these IoT systems. At the app interface level, privacy choices are primarily delivered visually via text and custom icons as shown in Figure 3(d). At the notification level, users can customize the notifications on their mobile devices to be haptic (e.g., vibration) or auditory (e.g., sound). At the technical level, the privacy choices provided by the IoTA are machine-readable format in nature, which opens the door for wider integration with other IoT systems or other modalities. For example, our research team is also working on an improved privacy-preserving facial obfuscation technology [31, 128] that enables people to opt out of IoT video analytics captured by cameras.

6.4 Evaluation

The design decisions above helped us implement IoTA to achieve the system requirements while balancing advantages and constraints. In this subsection, we briefly review the main features of IoTA and evaluate these features against the notion of "meaningful privacy choices" from five facets described in Section 2 (i.e., effectiveness, efficiency, user awareness, comprehensiveness, and neutrality). Given that the app is still at its early stage entering the mainstream market, we also describe our future evaluation plan.

IoTA currently offers four privacy options, providing a range of choices that better accommodate users' diverse privacy preferences and support certain privacy rights in GDPR. This improves the **effectiveness** of privacy choices in the IoT context. Through the implementation of these privacy options, IoTA serves as a centralized privacy choice platform that enables users to more **efficiently** manage their data privacy in IoT. IoTA also supports three discovery mechanisms of IoT systems (i.e., location-based map interface, push notifications, and QR codes). These discovery mechanisms are likely to increase **user awareness** of deployed IoT systems and their available privacy choices. In addition, the unified view

of the integrated notice and choice in IoTA provides concise yet **comprehensive** information about IoT data practices to help users understand the ramifications of their privacy decisions. Further, we strive to implement the integrated notice and choice in IoTA without framing or bias, trying to provide a **neutral** ground for users to enact their privacy choices. As a result, we believe IoTA is a significant step towards "meaningful privacy choices" in IoT.

Given the fact that IoTA is publicly released very recently, real-world evaluations are yet to be conducted. In the future, we plan to systematically evaluate the effectiveness, efficiency, comprehensiveness, and neutrality of the privacy choices supported by IoTA through multiple quantitative and qualitative user studies. First, we would like to evaluate the usability of the existing discovery mechanisms in IoTA. In particular, since IoTA supports customization of notifications, it is necessary to understand users' notification preferences in IoT. Second, we seek to understand how users interact with the integrated notice and choice and what information they pay more attention to. This will inform future iterations of IoTA. Finally, we intend to measure whether and how these supported privacy options influence people's privacy-related behaviors and expectations in IoT.

6.5 Summary and Implementation Status

In this use case, we presented our design decisions for IoTA by fully considering the design space for privacy choices and initially evaluated of the design of IoTA against the notion of meaningful privacy choices. This use case provides practitioners an example on how to leverage the design space to implement meaningful privacy choices in their systems. Currently, a stable version of IoTA (1.2.3) is available on both iOS and Android platforms in over 30 countries around the world. We will continue to improve the usability of IoTA and revisit the design decisions when necessary.

7 CONCLUSION

Under the backdrop of lacking cohesive design guidelines for privacy choices, we constructed a comprehensive design space for privacy choices, which is flexible enough for different regulatory requirements and particularly applicable to in IoT context. We also presented a use case of how we navigate the design space to design a privacy choice platform for IoT, showcasing the applicability of the design space in building real-world systems. Overall, our constructed design space contributes a user-centered conceptual framework that considers privacy choice as a process, offering a taxonomy to understand meaningful privacy control, particularly in the IoT context. Also, the design space for privacy choices provides practitioners comprehensive design guidelines on evaluating both system and legal requirements before implementing privacy choices, as well as considering different design options under five dimensions to design meaningful privacy choices.

ACKNOWLEDGMENTS

We thank our colleague Justin Donnell for his work on the IoT privacy infrastructure and all the team members who have contributed to the IoT Assistant app: Akshath Jain, Yoshua Torralva, Salil Deshpande, Maahin Beri, Elizabeth Louie, and Gaurav Misra. This research was supported in part by grants from the Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Laboratory (AFRL) under the Brandeis program (FA8750-15-2-0277) and in part by grants from the National Science Foundation (NSF) Secure and Trustworthy Computing program (CNS-1801316, CNS-1914486). The US Government is authorized to reproduce and distribute reprints for Governmental purposes not withstanding any copyright notice thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as representing the official policies or endorsements, either expressed or implied, of DARPA, AFRL, NSF, or the US Government.

REFERENCES

[1] Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, et al. 2017. Nudges for privacy and security: Understanding and assisting users' choices online. *ACM Computing Surveys (CSUR)* 50, 3 (2017), 1–41.

[2] Idris Adjerid, Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2013. Sleights of Privacy: Framing, Disclosures, and the Limits of Transparency. In *Proceedings of the Ninth Symposium on Usable Privacy and Security* (Newcastle, United Kingdom) *(SOUPS '13)*. Association for Computing Machinery, New York, NY, USA, Article 9, 11 pages. https://doi.org/10.1145/2501604.2501613

[3] Hazim Almuhimedi. 2017. *Helping Smartphone Users Manage their Privacy through Nudges.* Ph.D. Dissertation. Carnegie Mellon University, Pittsburgh, PA, USA.

[4] Hazim Almuhimedi, Florian Schaub, Norman Sadeh, Idris Adjerid, Alessandro Acquisti, Joshua Gluck, Lorrie Faith Cranor, and Yuvraj Agarwal. 2015. Your Location Has Been Shared 5,398 Times! A Field Study on Mobile App Privacy Nudging. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) *(CHI '15)*. Association for Computing Machinery, New York, NY, USA, 787–796. https://doi.org/10.1145/2702123.2702210

[5] Rebecca Balebako, Richard Shay, and Lorrie Faith Cranor. 2014. Is your inseam a biometric? a case study on the role of usability studies in developing public policy. In *Workshop on Usable Security*, Vol. 23.

[6] Vinayshekhar Bannihatti Kumar, Roger Iyengar, Namita Nisal, Yuanyuan Feng, Hana Habib, Peter Story, Sushain Cherivirala, Margaret Hagan, Lorrie Cranor, Shomir Wilson, Florian Schaub, and Norman Sadeh. 2020. Finding a Choice in a Haystack: Automatic Extraction of Opt-Out Statements from Privacy Policy Text. In *Proceedings of The Web Conference 2020* (Taipei, Taiwan) *(WWW '20)*. Association for Computing Machinery, New York, NY, USA, 1943–1954. https://doi.org/10.1145/3366423.3380262

[7] Louise Barkhuus. 2012. The Mismeasurement of Privacy: Using Contextual Integrity to Reconsider Privacy in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas, USA) *(CHI '12)*. Association for Computing Machinery, New York, NY, USA, 367–376. https://doi.org/10.1145/2207676.2207727

[8] Adam Barth, Anupam Datta, John C Mitchell, and Helen Nissenbaum. 2006. Privacy and contextual integrity: Framework and applications. In *2006 IEEE symposium on security and privacy (S&P'06)*. IEEE, 15–pp.

[9] Ricardo Melo Bastos and Duncan Dubugras A Ruiz. 2002. Extending UML activity diagram for workflow modeling in production systems. In *Proceedings of the 35th Annual Hawaii International Conference on System Sciences*. IEEE, 3786–3795.

[10] Paola Benassi. 1999. TRUSTe: an online privacy seal program. *Commun. ACM* 42, 2 (1999), 56–59.

[11] Michael Benisch, Patrick Gage Kelley, Norman Sadeh, and Lorrie Faith Cranor. 2011. Capturing location-privacy preferences: quantifying accuracy and user-burden tradeoffs. *Personal and Ubiquitous Computing* 15, 7 (2011), 679–694.

[12] Rainer Böhme and Stefan Köpsell. 2010. Trained to accept? A field experiment on consent dialogs. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 2403–2406.

[13] John J Borking, BMA Van Eck, and P Siepel. 1999. *Intelligent software agents and privacy.* Registratiekamer The Hague.

[14] Cristian Bravo-Lillo, Saranga Komanduri, Lorrie Faith Cranor, Robert W Reeder, Manya Sleeper, Julie Downs, and Stuart Schechter. 2013. Your attention please: designing security-decision UIs to make genuine risks harder to ignore. In *Proceedings of the Ninth Symposium on Usable Privacy and Security*. 1–12.

[15] Brazil National Congress. 2018. General Data Protection Law (English translation). https://iapp.org/resources/article/brazils-general-data-protection-law-english-translation.

[16] Ramon Caceres and Adrian Friday. 2011. Ubicomp systems at 20: Progress, opportunities, and challenges. *IEEE Pervasive Computing* 11, 1 (2011), 14–21.

[17] Fred H Cate. 2006. The failure of fair information practice principles. *Consumer protection in the age of the information economy* (2006).

[18] Fred H Cate. 2010. The limits of notice and choice. *IEEE Security & Privacy* 8, 2 (2010), 59–62.

[19] Andrew Chio, Georgios Bouloukakis, Cheng-Hsin Hsu, Sharad Mehrotra, and Nalini Venkatasubramanian. 2019. Adaptive Mediation for Data Exchange in IoT Systems. In *Proceedings of the 18th Workshop on Adaptive and Reflexive Middleware*. 1–6.

[20] Hanbyul Choi, Jonghwa Park, and Yoonhyuk Jung. 2018. The role of privacy fatigue in online privacy behavior. *Computers in Human Behavior* 81 (2018), 42–51.

[21] Richard Chow. 2017. The last mile for IoT privacy. *IEEE Security & Privacy* 15, 6 (2017), 73–76.

[22] Jessica Colnago, Yuanyuan Feng, Tharangini Palanivel, Sarah Pearman, Megan Ung, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. 2020. Informing the design of a personalized privacy assistant for the internet of things. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.

[23] CookieYes. 2019. *GDPR cookie consent banner examples*. Retrieved September 12th, 2020 from https://www.cookieyes.com/gdpr-cookie-consent-banner-examples/

[24] Council of European Union. 2016. General Data Protection Regulation. https://gdpr-infor.eu.

[25] Lorrie Faith Cranor. 2003. P3P: Making privacy policies more useful. *IEEE Security & Privacy* 1, 6 (2003), 50–55.

[26] Lorrie Faith Cranor. 2012. Necessary but not sufficient: Standardized mechanisms for privacy notice and choice. *J. on Telecomm. & High Tech. L.* 10 (2012), 273.

[27] Lorrie Faith Cranor, Praveen Guduru, and Manjula Arjula. 2006. User interfaces for privacy agents. *ACM Transactions on Computer-Human Interaction (TOCHI)* 13, 2 (2006), 135–178.

[28] Lorrie Faith Cranor, Hana Habib, Yixin Zou, Alessandro Acquisti, Joel Reidenberg, Norman Sadeh, and Florian Schaub. 2020. *Design and Evaluation of a Usable Icon and Tagline to Signal an Opt-Out of the Sale of Personal Information as Required by CCPA*. Retrieved September 13th, 2020 from https://oag.ca.gov/sites/all/files/agweb/pdfs/privacy/cranor-design-eval-usable-icon.pdf

[29] Lorrie Faith Cranor, Kelly Idouchi, Pedro Giovanni Leon, Manya Sleeper, and Blase Ur. 2013. Are they actually any different? Comparing thousands of financial institutions' privacy practices. In *Proceedings of the Twelfth Workshop on the Economics of Information Security*, Vol. 13.

[30] Anupam Das, Martin Degeling, Daniel Smullen, and Norman Sadeh. 2018. Personalized privacy assistants for the internet of things: providing users with notice and choice. *IEEE Pervasive Computing* 17, 3 (2018), 35–46.

[31] Anupam Das, Martin Degeling, Xiaoyou Wang, Junjue Wang, Norman Sadeh, and Mahadev Satyanarayanan. 2017. Assisting users in a world full of cameras: A privacy-aware infrastructure for computer vision applications. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. IEEE, 1387–1396.

[32] Martin Degeling, Christine Utz, Christopher Lentzsch, Henry Hosseini, Florian Schaub, and Thorsten Holz. 2019. We value your privacy... now take some cookies: Measuring the GDPR's impact on web privacy. In *Proceedings of Network and Distributed System Security Symposium(NDSS '19)*.

[33] Jayati Dev, Emilee Rader, and Sameer Patil. 2020. Why Johnny Can't Unsubscribe: Barriers to Stopping Unwanted Email. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12.

[34] Digital Advertising Alliance. 2009. *Self-regulatory principles for online behavioral advertising,*. Retrieved September 13th, 2020 from https://digitaladvertisingalliance.org/principles

[35] Benchaa Djellali, Kheira Belarbi, Abdallah Chouarfia, and Pascal Lorenz. 2015. User authentication scheme preserving anonymity for ubiquitous devices. *Security and Communication Networks* 8, 17 (2015), 3131–3141.

[36] Hilko Donker, Palle Klante, and Peter Gorny. 2002. The design of auditory user interfaces for blind users. In *Proceedings of the second Nordic conference on Human-computer interaction*. 149–156.

[37] Yitao Duan and John Canny. 2004. Protecting user data in ubiquitous computing: Towards trustworthy environments. In *International Workshop on Privacy Enhancing Technologies*. Springer, 167–185.

[38] Serge Egelman, Janice Tsai, Lorrie Faith Cranor, and Alessandro Acquisti. 2009. Timing is everything? The effects of timing and placement of online privacy indicators. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 319–328.

[39] Marc Elshof. 2019. *GDPR Update - EDPB video surveillance guidelines*. Retrieved September 7th, 2020 from https://www.jdsupra.com/legalnews/gdpr-update-edpb-video-surveillance-94566/

[40] Pardis Emami-Naeini, Yuvraj Agarwal, Lorrie Faith Cranor, and Hanan Hibshi. 2020. Ask the Experts: What Should Be on an IoT Privacy and Security Label? (2020). arXiv:2002.04631

[41] José Estrada-Jiménez, Javier Parra-Arnau, Ana Rodríguez-Hoyos, and Jordi Forné. 2017. Online advertising: Analysis of privacy threats and protection approaches. *Computer Communications* 100 (2017), 32–51.

[42] Benjamin Fabian, Tatiana Ermakova, and Tino Lentz. 2017. Large-Scale Readability Analysis of Privacy Policies. In *Proceedings of the International Conference on Web Intelligence* (Leipzig, Germany) *(WI '17)*. Association for Computing Machinery, New York, NY, USA, 18–25. https://doi.org/10.1145/3106426.3106427

[43] Kassem Fawaz, Thomas Linden, and Hamza Harkous. 2019. The Applications of Machine Learning in Privacy Notice and Choice. In *2019 11th International Conference on Communication Systems & Networks (COMSNETS)*. IEEE, 118–124.

[44] Federal Trade Commission. 2013. Mobile Privacy Disclosures: Building Trust Through Transparency (FTC Staff Report). https://www.ftc.gov/reports/mobile-privacy-disclosures-building-trust-through-transparency-federal-trade-commission.

[45] Federal Trade Commission. 2015. Internet of Things: Privacy & Security in a Connected World (FTC Staff Report). https://www.ftc.gov/reports/federal-trade-commission-staff-report-november-2013-workshop-entitled-internet-things.

[46] Adrienne Porter Felt, Serge Egelman, Matthew Finifter, Devdatta Akhawe, and David Wagner. 2012. How to ask for permission. In *Proceedings of the 7th USENIX conference on Hot Topics in Security*. 7–7.

[47] Stacia Garlach and Daniel Suthers. 2018. I'm supposed to see that?'AdChoices Usability in the Mobile Environment. In *Proceedings of the 51st Hawaii International Conference on System Sciences*.

[48] Joshua Gluck, Florian Schaub, Amy Friedman, Hana Habib, Norman Sadeh, Lorrie Faith Cranor, and Yuvraj Agarwal. 2016. How short is too short? Implications of length and framing on the effectiveness of privacy notices. In *Twelfth Symposium on Usable Privacy and Security ({SOUPS} 2016)*. 321–340.

[49] Nathaniel S Good, Jens Grossklags, Deirdre K Mulligan, and Joseph A Konstan. 2007. Noticing notice: a large-scale experiment on the timing of software license agreements. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 607–616.

[50] Lawrence O Gostin. 2001. National health information privacy: regulations under the Health Insurance Portability and Accountability Act. *Jama* 285, 23 (2001), 3015–3021.

[51] Hana Habib, Sarah Pearman, Jiamin Wang, Yixin Zou, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. 2020. " It's a scavenger hunt": Usability of Websites' Opt-Out and Data Deletion Choices. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12.

[52] Hana Habib, Yixin Zou, Aditi Jannu, Neha Sridhar, Chelse Swoopes, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. 2019. An empirical analysis of data deletion and opt-out choices on 150 websites. In *Fifteenth Symposium on Usable Privacy and Security ({SOUPS} 2019)*.

[53] Hana Habib, Yixin Zou, Yaxing Yao, Alessandro Acquisti, Lorrie Faith Cranor, Joel Reidenberg, Norman Sadeh, and Florian Schaub. 2021. Toggles, Dollar Signs, and Triangles: How to (In)EffectivelyConvey Privacy Choices with Icons and Link Texts. In *Conference on Human Factors in Computing Systems (CHI)*. ACM. https://doi.org/10.1145/3411764.3445387

[54] Marian Harbach, Markus Hettig, Susanne Weber, and Matthew Smith. 2014. Using personal examples to improve risk communication for security & privacy decisions. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 2647–2656.

[55] Drew Harwell and Abha Bhattarai. 2018. *Inside Amazon Go: The camera-filled convenience store that watches you back*. Retrieved September 13th, 2020 from https://www.washingtonpost.com/news/business/wp/2018/01/22/inside-amazon-go-the-camera-filled-convenience-store-that-watches-you-back/

[56] Alan R Hevner, Salvatore T March, Jinsoo Park, and Sudha Ram. 2004. Design science in information systems research. *MIS quarterly* (2004), 75–105.

[57] Chris Hoffman. 2020. *Why Your iPhone Keeps Asking About Background Location Use*. Retrieved September 15th, 2020 from https://www.howtogeek.com/563557/why-your-iphone-keeps-asking-you-about-background-location-use/

[58] Giovanni Iachello and Jason Hong. 2007. *End-user privacy in human-computer interaction*. Vol. 1. Now Publishers Inc.

[59] Laurel Jamtgaard. 2000. Big Bird Meets Big Brother: A Look at the Children's Online Privacy Protection Act. *Santa Clara High Technology Law Journal* 16, 2 (2000), 385.

[60] Farzaneh Karegar, John Sören Pettersson, and Simone Fischer-Hübner. 2020. The dilemma of user engagement in privacy notices: Effects of interaction modes and habituation on user attention. *ACM Transactions on Privacy and Security (TOPS)* 23, 1 (2020), 1–38.

[61] Mark J Keith, Courtenay Maynes, Paul Benjamin Lowry, and Jeffry Babb. 2014. Privacy fatigue: The effect of privacy control complexity on consumer electronic information disclosure. In *International Conference on Information Systems (ICIS 2014), Auckland, New Zealand, December*. 14–17.

[62] Patrick Gage Kelley, Sunny Consolvo, Lorrie Faith Cranor, Jaeyeon Jung, Norman Sadeh, and David Wetherall. 2012. A Conundrum of Permissions: Installing Applications on an Android Smartphone. In *Financial Cryptography and Data Security*, Jim Blyth, Sven Dietrich, and L. Jean Camp (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 68–79.

[63] Patrick Gage Kelley, Lorrie Faith Cranor, and Norman Sadeh. 2013. Privacy as part of the app decision-making process. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 3393–3402.

[64] Agnieszka Kitkowska, Mark Warner, Yefim Shulman, Erik Wästlund, and Leonardo A Martucci. 2020. Enhancing Privacy through the Visual Design of Privacy Notices: Exploring the Interplay of Curiosity, Control and Affect. In *Sixteenth Symposium on Usable Privacy and Security ({SOUPS} 2020)*. 437–456.

[65] Saranga Komanduri, Richard Shay, Greg Norcie, and Blase Ur. 2011. Adchoices-compliance with online behavioral advertising notice and choice requirements. *ISJLP* 7 (2011), 603.

[66] Bastian Könings, Florian Schaub, and Michael Weber. 2013. PriFi beacons: piggybacking privacy implications on wifi beacons. In *Proceedings of the 2013 ACM conference on Pervasive and ubiquitous computing adjunct publication*. 83–86.

[67] Bastian Könings, Sebastian Thoma, Florian Schaub, and Michael Weber. 2014. Pripref broadcaster: Enabling users to broadcast privacy preferences in their physical proximity. In *Proceedings of the 13th International Conference on Mobile and Ubiquitous Multimedia*. 133–142.

[68] Costas Lambrinoudakis. 2018. The general data protection regulation (GDPR) era: ten steps for compliance of data processors and data controllers. In *International Conference on Trust and Privacy in Digital Business*. Springer, 3–8.

[69] Marc Langheinrich. 2001. Privacy by design—principles of privacy-aware ubiquitous systems. In *International conference on Ubiquitous Computing*. Springer, 273–291.

[70] Daniel Le Métayer and Shara Monteleone. 2009. Automated consent through privacy agents: Legal requirements and technical architecture. *Computer law & Security review* 25, 2 (2009), 136–144.

[71] Nicole Lee. 2019. *You can now ask Alexa to delete your voice history.* Retrieved September 16th, 2020 from https://www.engadget.com/2019-05-29-amazon-alexa-voice-deletion.html

[72] Pedro Leon, Blase Ur, Richard Shay, Yang Wang, Rebecca Balebako, and Lorrie Cranor. 2012. Why Johnny can't opt out: a usability evaluation of tools to limit online behavioral advertising. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 589–598.

[73] Pedro Giovanni Leon, Justin Cranshaw, Lorrie Faith Cranor, Jim Graves, Manoj Hastak, Blase Ur, and Guzi Xu. 2012. What Do Online Behavioral Advertising Privacy Disclosures Communicate to Users?. In *Proceedings of the 2012 ACM Workshop on Privacy in the Electronic Society* (Raleigh, North Carolina, USA) *(WPES '12)*. Association for Computing Machinery, New York, NY, USA, 19–30. https://doi.org/10.1145/2381966.2381970

[74] Stefan Leuthold, Javier A Bargas-Avila, and Klaus Opwis. 2008. Beyond web content accessibility guidelines: Design of enhanced text user interfaces for blind internet users. *International Journal of Human-Computer Studies* 66, 4 (2008), 257–270.

[75] Jialiu Lin, Shahriyar Amini, Jason I. Hong, Norman Sadeh, Janne Lindqvist, and Joy Zhang. 2012. Expectation and Purpose: Understanding Users' Mental Models of Mobile App Privacy through Crowdsourcing. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing* (Pittsburgh, Pennsylvania) *(UbiComp '12)*. Association for Computing Machinery, New York, NY, USA, 501–510. https://doi.org/10.1145/2370216.2370290

[76] Jialiu Lin, Michael Benisch, Norman Sadeh, Jianwei Niu, Jason Hong, Banghui Lu, and Shaohui Guo. 2013. A Comparative Study of Location-sharing Privacy Preferences in the United States and China. *Personal Ubiquitous Comput.* 17, 4 (April 2013), 697–711.

[77] Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhimedi, Shikun Zhang, Norman Sadeh, Alessandro Acquisti, and Yuvraj Agarwal. 2016. Follow My Recommendations: A Personalized Privacy Assistant for Mobile App Permissions. In *Proceedings of the Twelfth USENIX Conference on Usable Privacy and Security* (Denver, CO, USA) *(SOUPS '16)*. USENIX Association, USA, 27–41.

[78] Bin Liu, Jialiu Lin, and Norman Sadeh. 2014. Reconciling Mobile App Privacy and Usability on Smartphones: Could User Privacy Profiles Help?. In *Proceedings of the 23rd International Conference on World Wide Web* (Seoul, Korea) *(WWW '14)*. ACM, New York, NY, USA, 201–212. https://doi.org/10.1145/2566486.2568035

[79] Duen-Ren Liu and Minxin Shen. 2003. Workflow modeling for virtual processes: an order-preserving process-view approach. *Information Systems* 28, 6 (2003), 505–532.

[80] Kalle Lyytinen and Youngjin Yoo. 2002. Ubiquitous computing. *Commun. ACM* 45, 12 (2002), 63–96.

[81] Allan MacLean, Richard M Young, Victoria ME Bellotti, and Thomas P Moran. 1991. Questions, options, and criteria: Elements of design space analysis. *Human–computer interaction* 6, 3-4 (1991), 201–250.

[82] Tobias Matzner. 2014. Why privacy is not enough privacy in the context of "ubiquitous computing" and "big data". *Journal of Information, Communication and Ethics in Society* 12, 2 (2014), 93–106.

[83] Aleecia M McDonald and Lorrie Faith Cranor. 2008. The cost of reading privacy policies. *I/S: A Journal of Law and Policy for the Information Society* 4 (2008), 540–565.

[84] Aleecia M McDonald and Lorrie Faith Cranor. 2010. Americans' attitudes about internet behavioral advertising practices. In *Proceedings of the 9th annual ACM workshop on Privacy in the electronic society*. 63–72.

[85] Aleecia M. McDonald, Robert W. Reeder, Patrick Gage Kelley, and Lorrie Faith Cranor. 2009. A Comparative Study of Online Privacy Policies and Formats. In *Privacy Enhancing Technologies*, Ian Goldberg and Mikhail J. Atallah (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 37–55.

[86] Diane McKerlie and Allan MacLean. 1994. Reasoning with design rationale: practical experience with design space analysis. *Design Studies* 15, 2 (1994), 214–226.

[87] Sharad Mehrotra, Alfred Kobsa, Nalini Venkatasubramanian, and Siva Raj Rajagopalan. 2016. TIPPERS: A privacy cognizant IoT environment. In *2016 IEEE International Conference on Pervasive Computing and Communication Workshops (PerCom Workshops)*. IEEE, 1–6.

[88] George R Milne, Mary J Culnan, and Henry Greene. 2006. A longitudinal assessment of online privacy notice readability. *Journal of Public Policy & Marketing* 25, 2 (2006), 238–249.

[89] Mozilla Wiki. 2011. *Privacy icons*. Mozilla. Retrieved September 15th, 2020 from https://wiki.mozilla.org/Privacy_Icons

[90] Jakob Nielsen and Rolf Molich. 1990. Heuristic Evaluation of User Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Seattle, Washington, USA) *(CHI '90)*. Association for Computing Machinery, New York, NY, USA, 249–256. https://doi.org/10.1145/97243.97281

[91] Helen Nissenbaum. 2009. *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press.

[92] Midas Nouwens, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. 2020. Dark Patterns after the GDPR: Scraping Consent Pop-Ups and Demonstrating Their Influence. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3313831.3376321

[93] Office of the California Attorney General. 2020. California Consumer Privacy Act (CCPA): First Modified Regulations. https://oag.ca.gov/sites/all/files/agweb/pdfs/privacy/ccpa-text-of-mod-clean-020720.pdf.

[94] Douglas O'Shaughnessy. 2008. Automatic speech recognition: History, methods and challenges. *Pattern Recognition* 41, 10 (2008), 2965–2979.

[95] Nisha Panwar, Shantanu Sharma, Guoxi Wang, Sharad Mehrotra, Nalini Venkatasubramanian, Mamadou H Diallo, and Ardalan Amiri Sani. 2019. IoT Notary: Sensor data attestation in smart environment. In *2019 IEEE 18th International Symposium on Network Computing and Applications (NCA)*. IEEE, 1–9.

[96] Matthew Nicholas Papakipos and David Harry Garcia. 2017. Initializing camera subsystem for face detection based on sensor inputs. US Patent 9,596,084.

[97] Primal Pappachan, Martin Degeling, Roberto Yus, Anupam Das, Sruti Bhagavatula, William Melicher, Pardis Emami Naeini, Shikun Zhang, Lujo Bauer, Alfred Kobsa, et al. 2017. Towards privacy-aware smart buildings: Capturing, communicating, and enforcing privacy policies and preferences. In *2017 IEEE 37th International Conference on Distributed Computing Systems Workshops (ICDCSW)*. IEEE, 193–198.

[98] Sameer Patil, Roberto Hoyle, Roman Schlegel, Apu Kapadia, and Adam J Lee. 2015. Interrupt now or inform later? Comparing immediate and delayed privacy feedback. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 1415–1418.

[99] Sameer Patil, Roman Schlegel, Apu Kapadia, and Adam J Lee. 2014. Reflection or action? how feedback and control affect location sharing decisions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 101–110.

[100] Ken Peffers, Tuure Tuunanen, Marcus A Rothenberger, and Samir Chatterjee. 2007. A design science research methodology for information systems research. *Journal of management information systems* 24, 3 (2007), 45–77.

[101] Charith Perera, Rajiv Ranjan, Lizhe Wang, Samee U Khan, and Albert Y Zomaya. 2015. Big data privacy in the internet of things era. *IT Professional* 17, 3 (2015), 32–39.

[102] Charith Perera, Arkady Zaslavsky, Peter Christen, and Dimitrios Georgakopoulos. 2013. Context aware computing for the internet of things: A survey. *IEEE communications surveys & tutorials* 16, 1 (2013), 414–454.

[103] Nisarg Raval, Animesh Srivastava, Kiron Lebeck, Landon Cox, and Ashwin Machanavajjhala. 2014. Markit: Privacy markers for protecting visual secrets. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*. 1289–1295.

[104] Joel R Reidenberg, Travis Breaux, Lorrie Faith Cranor, Brian French, Amanda Grannis, James T Graves, Fei Liu, Aleecia McDonald, Thomas B Norton, Rohan Ramanath, Cameron Russell, Norman Sadeh, and Florian Schaub. 2015. Disagreeable privacy policies: Mismatches between meaning and users' understanding. *Berkeley Technology Law Journal* 30 (2015), 39.

[105] Joel R Reidenberg, N Cameron Russell, Alexander J Callen, Sophia Qasir, and Thomas B Norton. 2015. Privacy harms and the effectiveness of the notice and choice framework. *ISJLP* 11 (2015), 485.

[106] Joel R Reidenberg, N Cameron Russell, Vlad Herta, William Sierra-Rocafort, and Thomas B Norton. 2018. Trustworthy Privacy Indicators: Grades, Labels, Certifications, and Dashboards. *Washington University Law Review* 96 (2018),

1409.

[107] Christian Richthammer, Michael Netter, Moritz Riesner, and Günther Pernul. 2013. Taxonomy for social network data types from the viewpoint of privacy and user control. In *2013 International Conference on Availability, Reliability and Security*. IEEE, 141–150.

[108] Kay Romer and Friedemann Mattern. 2004. The design space of wireless sensor networks. *IEEE wireless communications* 11, 6 (2004), 54–61.

[109] Jeffrey Rosen. 2012. The Right to Be Forgotten. *Stanford Law Review* 64 (2012).

[110] Arianna Rossi and Monica Palmirani. 2019. DaPIS: a Data Protection Icon Set to Improve Information Transparency under the GDPR. *Knowledge of the Law in the Big Data Age. Frontiers* 252 (2019), 181–195.

[111] John A Rothchild. 2018. Against Notice and Choice: The Manifest Failure of the Proceduralist Paradigm to Protect Privacy Online (or Anywhere Else). *Cleveland State Law Review* 66, 3 (2018), 559.

[112] Iskander Sanchez-Rola, Matteo Dell'Amico, Platon Kotzias, Davide Balzarotti, Leyla Bilge, Pierre-Antoine Vervier, and Igor Santos. 2019. Can I Opt Out Yet? GDPR and the Global Illusion of Cookie Control. In *Proceedings of the 2019 ACM Asia Conference on Computer and Communications Security* (Auckland, New Zealand) *(Asia CCS '19)*. Association for Computing Machinery, New York, NY, USA, 340–351. https://doi.org/10.1145/3321705.3329806

[113] Florian Schaub, Rebecca Balebako, and Lorrie Faith Cranor. 2017. Designing effective privacy notices and controls. *IEEE Internet Computing* (2017).

[114] Florian Schaub, Rebecca Balebako, Adam L Durity, and Lorrie Faith Cranor. 2015. A design space for effective privacy notices. In *Eleventh Symposium On Usable Privacy and Security ({SOUPS} 2015)*. 1–17.

[115] Florian Schaub, Bastian Könings, and Michael Weber. 2015. Context-adaptive privacy: Leveraging context awareness to support privacy decision making. *IEEE Pervasive Computing* 14, 1 (2015), 34–43.

[116] Jeremy Schiff, Marci Meingast, Deirdre K Mulligan, Shankar Sastry, and Ken Goldberg. 2009. Respectful cameras: Detecting visual markers in real-time to address privacy concerns. In *Protecting Privacy in Video Surveillance*. Springer, 65–89.

[117] Paul M Schwartz. 1999. Privacy and Democracy in Cyberspace. *Vanderbilt Law Review* 52, 6 (1999), 1607.

[118] Paul M Schwartz and Daniel Solove. 2009. Notice & Choice. In *The Second NPLAN/BMSG Meeting on Digital Media and Marketing to Children*.

[119] Secretary's Advisory Committee on Automated Personal Data Systems. 1973. *Records, Computers, and the Rights of Citizens: Report.* US Department of Health, Education & Welfare.

[120] Robert H Sloan and Richard Warner. 2014. Beyond notice and choice: Privacy, norms, and consent. *The Journal of High Technology Law* 14 (2014), 370.

[121] Daniel Smullen, Yuanyuan Feng, Shikun Aerin Zhang, and Norman Sadeh. 2020. The Best of Both Worlds: Mitigating Trade-offs Between Accuracy and User Burden in Capturing Mobile App Privacy Preferences. *Proceedings on Privacy Enhancing Technologies* 2020, 1 (2020), 195–215.

[122] Monica Tentori, Jesus Favela, and Marcela D Rodriguez. 2006. Privacy-aware autonomous agents for pervasive healthcare. *IEEE Intelligent Systems* 21, 6 (2006), 55–62.

[123] Ash Turner. 2020. *How many smartphones are in the world?* Retrieved September 10th, 2020 from https://www.bankmycell.com/blog/how-many-phones-are-in-the-world

[124] Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012. Smart, useful, scary, creepy: perceptions of online behavioral advertising. In *proceedings of the eighth symposium on usable privacy and security*. 1–15.

[125] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. 2019. (Un)Informed Consent: Studying GDPR Consent Notices in the Field. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security* (London, United Kingdom) *(CCS '19)*. Association for Computing Machinery, New York, NY, USA, 973–990. https://doi.org/10.1145/3319535.3354212

[126] Ari Ezra Waldman. 2018. Privacy, notice, and design. *Stanford Technology Law Review* 21, 1 (2018).

[127] Ari Ezra Waldman. 2020. Cognitive biases, dark patterns, and the 'privacy paradox'. *Current opinion in psychology* 31 (2020), 105–109.

[128] Junjue Wang, Brandon Amos, Anupam Das, Padmanabhan Pillai, Norman Sadeh, and Mahadev Satyanarayanan. 2018. Enabling live video analytics with a scalable and privacy-aware framework. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* 14, 3s (2018), 1–24.

[129] Jason Watson, Heather Richter Lipford, and Andrew Besmer. 2015. Mapping user preference to privacy default settings. *ACM Transactions on Computer-Human Interaction (TOCHI)* 22, 6 (2015), 1–20.

[130] Mark Weiser. 1993. Some computer science issues in ubiquitous computing. *Commun. ACM* 36, 7 (1993), 75–84.

[131] Zack Whittaker. 2019. *iOS 13 will let you limit app location access to "just once"*. Retrieved September 12th, 2020 from https://techcrunch.com/2019/06/03/apple-ios-13-location-privacy/

[132] Michael S Wogalter, Vincent C Conzola, and Tonya L Smith-Jackson. 2002. Research-based guidelines for warning design and evaluation. *Applied Ergonomics* 33, 3 (2002), 219–230.

[133] Heng Xu, Robert E Crossler, and France BéLanger. 2012. A value sensitive design investigation of privacy enhancing tools in web browsers. *Decision support systems* 54, 1 (2012), 424–433.

[134] Jedidiah Yueh. 2018. *GDPR will make big tech even bigger.* Retrieved September 15th, 2020 from https://www.forbes.com/sites/forbestechcouncil/2018/06/26/gdpr-will-make-big-tech-even-bigger/#77a24cf82592