

Privacy-Aware IoT: State-of-the-Art and Challenges

Shukun Tokas^a, Gencer Erdogan^b and Ketil Stølen^c
Sustainable Communication Technologies, SINTEF Digital, Oslo, Norway

Keywords: Privacy, GDPR, IoT, Privacy-Enhancing Technologies, Privacy Challenges, Privacy Recommendations, Privacy Assistance, State of the Art.

Abstract: The consumer IoT is now prevalent and creates an enormous amount of fine-grained, detailed information about consumers' everyday actions, personalities, and preferences. Such detailed information brings new and unique privacy challenges. The consumers are not aware of devices that surround them. There is a lack of transparency and absence of support for consumers to control the collection and processing of their personal and sensitive data. This paper reports on a review of state-of-the-art on privacy protection in IoT, with respect to privacy enhancing technologies (PETs) and GDPR-specific privacy principles. Drawing on a thorough analysis of 36 full papers, we identify key privacy challenges in IoT that need to be addressed to provide consumers with transparency and control over their personal data. The privacy challenges we have identified are (1) the lack of technical expertise in privacy notice comprehension, (2) the lack of transparency and control of personal data, and (3) the lack of personalized privacy recommendations.

1 INTRODUCTION

Leaving a group of people to take a personal call, tilting a document to protect it from prying eyes, or lowering our voice during intimate conversations are examples of offline privacy behaviors we instinctively engage in (Acquisti et al., 2020). Violation of privacy invokes reactions such as shifting of gaze, adopting protective postures (Altman, 1977), or some form of actions to protect privacy. Altman's insights apply also to how individuals interact in an online world, engaging in everyday internet-enabled technologies such as e-mail, social media, instant messaging, and Internet of Things. The Internet of Things (IoT) devices are deployed in public and private spaces and enable use cases that enhance productivity and quality of life. For example, fine-grained information about consumers' everyday habits and actions can be used for accurate personalized recommendations to enhance consumer experience. Collectively, the information can be used in the context of smart cities to assist in data-driven strategic decisions regarding infrastructures. However, the data captured by IoT devices and services is often personal and sensitive, revealing consumer's preferences, health, finances –

all kind of information that they prefer not to reveal.

Several research studies and surveys reveal that privacy concerns are at an all-time high, as the collection and use of data in IoT are happening with very little or no control. Organizations collecting data are most of the time unknown to consumers. For example, a Norwegian population survey (Norwegian Consumer Council, 2021) reveals that more than 50% of the participants feel uncomfortable about commercial actors collecting information about them. They also state that they have no other choice but to share their data. An increased lack of transparency and absence of support for consumers to control the collection and processing of their personal data in IoT may heavily affect many areas of our lives and even constitute a long-term danger for democracy.

In 2019, Cisco conducted a consumer privacy survey (CISCO, 2019) of 2601 adults from twelve of the world's largest economies in Europe, Asia Pacific, and America and confirmed that privacy has become a critical matter for individuals worldwide. The survey reports that 84% of the participants indicated that they care about privacy and want more control over the use of their data, and 32% had already taken action to protect their privacy. Essentially, the need to protect privacy is emphasized by the fact that *privacy* is a fundamental human right. The General Data Protection Regulation (GDPR) (EP and CEU, 2016)

^a <https://orcid.org/0000-0001-9893-6613>

^b <https://orcid.org/0000-0001-9407-5748>

^c <https://orcid.org/0000-0002-8810-9902>

strengthened this further by explicitly adding key data protection principles. Several studies investigate this strengthening, e.g., Godinho and Adjerid (Godinho de Matos and Adjerid, 2021) found that *opt-in consent* to the personal data collection such as location data was only 6.2%, and Cisco (CISCO, 2019) confirms that 55% of respondents view GDPR very favorably. GDPR offers an extended definition of data protection in terms of privacy principles which are required to be implemented as a part of an overall approach to protect privacy with respect to processing of personal data. To reap the full benefits of IoT technologies and protect consumers and their *Right to Privacy*, coordinated R&D efforts in privacy are necessary. Research efforts in privacy-aware IoT, especially on consumer-centric privacy control, are in the early stages. We need new technology to restore control over personal and sensitive data in IoT systems. In this paper we review the state-of-the-art, by considering privacy enhancing technologies (PETs) and GDPR-specific privacy principles.

The main contributions of this study are (i) a review of literature on privacy protection in IoT, and (ii) identification of research challenges that need to be addressed to empower consumers in making informed privacy choices. The rest of the paper is structured as follows: Section 2 presents background information about the underlying privacy principles and privacy-enhancing technologies. Section 3 presents the research method used to conduct our study. In Section 4 we present and discuss the results of the review, while in Section 5 we present the identified privacy research challenges. Section 6 gives an overview of related work. Section 7 concludes the paper.

2 BACKGROUND

IoT devices are now ubiquitous in the very spaces the consumers move through, e.g., cars, wearables, smart buildings. These connected devices generate unprecedented quantities of data about consumer behavior. All this information undoubtedly enhances productivity and quality of life. For example, activity trackers allow consumers to track their daily physical activity, including the number of steps, quality of sleep, heart rate, oxygen saturation etc. Wearables for head, e.g., Muse headband can measure brain activity, brain health and performance, in real world environments to track the user's ability to focus. Information from these connected devices can in combination measure consumers' body movement, behavioral patterns, and productivity. But these very useful IoT devices also raise unique privacy challenges, mainly due

to current malpractices such as unknown data collection, complicated privacy notices, dark patterns, notification overload, or low privacy notice comprehension, leading to provocation of paradoxical behavior and privacy resignation (Feng et al., 2021). With the GDPR now fully applicable, it has become a legal obligation for all data controllers to take account of the state of the art privacy solutions. Next, we present an overview of the data protection principles and data protection friendly technology.

2.1 GDPR'S Privacy Principles

The GDPR's processing principles (EP and CEU, 2016) are set out in Article 5(1) and required to be followed by entities responsible for processing personal data. Data controllers are prescribed with the duty to demonstrate compliance (in Article 5(2)) with the privacy principles. In the following we present a brief description of the privacy principles.

Lawfulness, Fairness and Transparency: Under the regulation, personal data shall be processed lawfully, fairly, and in a transparent manner. In short, this principle requires honest usage and communication with the data subject about their personal data. *Purpose Limitation:* It restricts the collection and processing of personal data for specific, explicit and legitimate purposes only, and requires that the data is not processed beyond such purposes. *Data Minimization:* It means that the data controllers must only collect and process personal data that is relevant, necessary, and adequate to accomplish the purposes for which it is processed. *Accuracy:* It means that the data controllers must take reasonable measures to prevent inaccuracies and ensure that the data is accurate and up to date. It also includes taking necessary measures to respond to data subjects' request to correct inaccurate or incomplete information. *Storage Limitation:* It means that personal data must not be kept for longer than necessary for the purposes for which it was collected for. Once the data is no longer needed, it must be securely deleted. *Confidentiality and Integrity:* It means that the controllers must take appropriate security measures to protect the data against unauthorized and unlawful processing, accidental loss, and so on. *Accountability:* The GDPR strengthens the six privacy principles by explicitly adding the accountability requirement (in Article 5(2)). It means that the data controller is responsible for complying with the aforementioned six principles, and they must be able to evidence their compliance.

Article 25 introduces data protection by design and data protection by default obligations on the data

controllers. The *data protection by design* principle requires the controller to implement appropriate technical and organizational measures (EP and CEU, 2016), but it is not specified how such measures can be embedded into the design. In addition, it introduces a specific *data protection by default* obligation. This principle necessitates that privacy is built into the system by default, i.e., no measures to protect privacy are required by the data subject, when he/she acquires a new product or service.

The regulation prescribes that the data subject has *Right of Access* (Article 15), which requires the data controllers to provide a data subject that requests to know, with his or her personal data. For example, the controller needs to inform the data subject (upon request) about the purposes of processing, the legal bases for processing, recipients of data when personal information has been or will be disclosed, and the right to lodge a complaint with a supervisory authority. In this paper, we focus on clause *I(a)-I(e)*, which are particularly concerned with information related to access control on personal data.

2.2 Privacy Enhancing Technology

Privacy Enhancing Technologies (PETs) is an umbrella term covering a broad range of emerging technologies and approaches that protects privacy by eliminating or preventing unnecessary processing of personal data. PETs allow both data protection and privacy-preserving data analytics. PETs can range from a piece of tape masking a webcam to an advanced blockchain based technique. A report on PETs from The Royal Society (The Royal Society, 2019), explores promising PETs that have the potential to enable privacy-aware data collection, analysis and dissemination of results: homomorphic encryption, trusted execution environments, secure multi-party computation, differential privacy, and personal data stores. In this section, we briefly describe above-mentioned PETs, along with PETs we found during our analysis of the state-of-the-art. *Encryption* (E), in general, it is the process of encoding information. Encryption converts the original representation of the information, using cryptographic keys, into a cipher. *Differential privacy* (DP) protects from sharing private information about individuals. It is a property of the cryptographic algorithm that performs analysis on a dataset. DP adds sufficient noise to the dataset or aggregates to hide the impact of any one individual, as it enables to describe patterns of groups within the dataset while maintaining the privacy of individuals. *Trusted Execution Environment* (TEE), also known as secure enclaves, is a secure area inside the main pro-

cessor that allows the isolation of secret code from the rest of the software running on a system. It prevents the operating system or the hypervisor from reading code in the TEE. *Obfuscation* (O) is a technique for protecting personal and sensitive data through computer algorithms and masking techniques, e.g., adding misleading data. *Anonymization* (A) is the least restrictive way to utilize data because anonymized data is not regulated. Anonymization techniques attempt to either obscure personally identifiable data or remove any attributes associated with an individual that could make the individual identifiable in a dataset. Suppression, generalization, and noise addition are well known anonymization approaches. *Framework* (F) refers to privacy frameworks, such as a framework to ensure consumer trust, a framework to reason about keeping humans in the loop to reduce discrimination, or a framework to fulfill privacy compliance obligations, etc. *Architecture* (Arch.) refers to privacy architectures that at a high-level envision several complementary requirements, such as, empower data subjects to effectively manage privacy settings, privacy-aware data collection and processing, reduce disclosure of personal data, etc. *Ontology* (Ont.) refers to privacy-aware metamodels for capturing privacy requirements, e.g., notice, preferences, purpose, privacy profiles, interactions between data collectors and processors across the value chain, etc. *Privacy assistant* (Asst.) refers to a privacy management interface or an app, endowed with dialog and machine learning to assist users in making informed privacy decision and managing IoT privacy settings. Users can interact with such assistants to review their privacy settings of IoT devices, get recommendations in case of difficult privacy decision-making situations, and also semi-automate the process of setting privacy preferences to reduce the notification overload. *Guidelines* (G) refers to recommendations and procedures, e.g., privacy-by-design guidelines intended for developers to assist them in creation of privacy-aware systems, actionable recommendations on how to effectively present privacy attributes on an IoT label to better communicate privacy risk to consumers. *Privacy predictive model* (PPM) refers to a specific model which is used to predict future privacy decisions or privacy preferences of users. *Model* (M) refers to privacy-aware models, e.g., a privacy-preserving data collection and access control model. *Voice and gaze communication cues* (VGCC) refers to a unique blend of voice volume level and gaze direction to minimize data collection by smart speakers. *Blockchain* (B) refers to a blockchain based technique to build trusted relationship between data sharing parties. *Privacy Announcement Mechanisms* (PAM) refers to a pri-

vacy notice announcement mechanism in context of body-worn cameras, to indicate whether the camera in question is recording or idle. *Denaturing Mechanisms* (DM) refers to the technique of scrubbing a video or picture of personal identifiers, such as face and landmark locations. *IoT Locator* (Loc.) refers to visual, auditory, or contextual pictures as locators to help users physically find a nearby IoT device. Locators enhance user's awareness of the presence of nearby IoT devices. *Design space for privacy choices* (DSPC) refers to a design space for privacy choices based on a user-centered analysis of what makes privacy notices effective. In addition to privacy notices, DSPC can be used to implement more effective privacy controls in privacy assistants.

3 RESEARCH METHOD

Our research method consisted of three main steps. In Step 1, we identified scientific papers addressing privacy in IoT. We used the following online databases to search for relevant papers: IEEE Xplore, ACM Digital Library, SpringerLink, and Science Direct. These were selected because they are central sources for scientific literature related to privacy and computer science. For each of the databases, we conducted manual searches within the database, using the following keywords: privacy, data protection, privacy enhancing technology (PET), privacy engineering, GDPR, ubiquitous computing, and internet of things (IoT). To limit the result set and support the screening process, we defined a set of inclusion and exclusion criteria. Our *inclusion criteria* state that the studies should be: Related to privacy engineering, privacy enhancing technology within IoT domain; Written in English; Peer reviewed; Published between 2013-2021. As per our *exclusion criteria* we excluded: Repeated studies found in different search engines; Inaccessible papers; Studies in the form of patents, general web pages, books, thesis, tutorials, reports or white papers. In Step 2, for each paper included in our analysis, we extracted information related to privacy enhancing technologies and GDPR-specific privacy principles. We classified the papers with respect to the aforementioned topics, and created a mapping of privacy solutions in IoT that supports GDPR's core privacy principles. Finally, in Step 3, we analyzed the extracted information from the papers to identify the state-of-the-art and privacy challenges in IoT.

Through our research, we identified 36 studies. The studies were then assessed and information was extracted as per the classification scheme by two independent researchers. There was no duplication of ef-

forts, in terms of assessment and information extraction. Each researcher got a unique set of 18 papers. A calibration exercise, on four randomly selected studies, was carried out to address the variances in interpretations of classification parameters. In cases of uncertainties in interpretations, the researchers concluded and collated the results in a joint session.

4 FINDINGS

In the following, we introduce the classification scheme, present an overall summary of information extracted from selected papers with respect to the classification scheme and discuss the results. With respect to the headings of the columns in Figure 1, *Paper* refers to the selected papers, *PETs* refers to the PET abbreviations (Section 2.2), columns *P1-P7* refer to the seven core privacy principles (Section 2.1), column *P8* refers to the data protection by design and by default, and column *P9* refers to the data subject access right. Essentially, *P1-P9* covers Article 5, Article 25, and Article 15 [*clause 1(a)-1(e)*] of the GDPR. Note that P1 consists of three different yet connected components. For simplicity reasons, if a study supports one or more components of P1, then we consider that the study supports the principle as a whole. Similarly, P8 includes the terms *by-design* and *by-default*, if either or both are supported then we consider that the study supports P8 as a whole.

4.1 Privacy Enhancing Technologies

Figure 2 shows the privacy enhancing technologies (PETs) that are addressed by two or more primary studies. One primary study may address more than one PET. In total, we have identified 19 PETs (see Section 2.2). From Figure 2 we see that there are nine PETs that are addressed by two or more papers. The remaining PETs that are not listed in Figure 2 were each addressed by only one primary study. The PETs addressed by two or more papers are privacy guidelines, privacy assistants, Ontology, privacy models, privacy architectures, privacy frameworks, privacy predictive models, encryption-related solutions to support privacy, and anonymization.

There are several PETs that can provide users with tools to enhance transparency and control in IoT. For instance, Wang et al. (Wang et al., 2017) present a privacy-aware IoT architecture comprising several components, including a software module privacy mediator. The mediator receives the user's privacy preferences through a smartphone app called the IoT Assistant (IoTA). It applies the privacy preferences to

Paper	PETs	P1	P2	P3	P4	P5	P6	P7	P8	P9
(Benhamida et al., 2021)	A, E	✓		✓			✓		✓	
(Arachchige et al., 2020)	E, DP	✓					✓		✓	
(Koelle et al., 2018b)	PAM, G	✓							✓	
(Feng et al., 2021)	DSPC	✓							✓	
(Fernandez et al., 2021a)	Asst.	✓								
(Emami-Naeini et al., 2017)	PPM	✓							✓	
(Wang et al., 2017)	DM	✓	✓	✓					✓	
(Pappachan et al., 2017)	F	✓	✓	✓			✓			
(Emami-Naeini et al., 2021)	G	✓								
(Lee and Kobsa, 2019)	Arch., PPM	✓	✓	✓		✓	✓		✓	
(Huang et al., 2021)	F								✓	
(Fernandez et al., 2021b)	A, Arch.	✓		✓			✓		✓	
(Bose et al., 2015)	E						✓		✓	
(Koelle et al., 2018a)	A	✓					✓		✓	
(Antignac and Le Métayer, 2014)	Arch.			✓					✓	
(Ren et al., 2021)	A						✓			
(Mehrotra et al., 2016)	A								✓	
(Escher et al., 2020)	F	✓								
(Fernández et al., 2020)	E, M	✓					✓		✓	✓
(Conzon et al., 2019)	Arch.								✓	
(Toumia et al., 2020)	Ont.	✓	✓			✓			✓	✓
(Song et al., 2020)	Loc.	✓								
(Mhaidli et al., 2020)	VGCC			✓					✓	
(Foukia et al., 2016)	Arch.	✓					✓	✓	✓	
(O'Connor et al., 2017)	G	✓					✓		✓	
(Fernandez et al., 2019)	Arch.	✓	✓			✓	✓	✓	✓	
(Zhang et al., 2020)	PPM	✓	✓						✓	
(Morel et al., 2019)	F	✓				✓	✓	✓		✓
(Sadique et al., 2020)	M, G					✓	✓			
(Kammüller, 2018)	F	✓	✓			✓			✓	
(Das et al., 2018)	Asst.	✓	✓			✓			✓	
(Alkhariji et al., 2021)	Asst.	✓	✓	✓	✓	✓	✓	✓	✓	✓
(Elkhodr et al., 2013)	Asst., O	✓							✓	
(Shi et al., 2021)	B, TEE	✓					✓	✓	✓	
(Agarwal et al., 2020)	Ont.	✓	✓		✓				✓	
(Sanchez et al., 2020)	Ont.	✓	✓			✓	✓		✓	

Figure 1: Identified primary studies, PETs, and principles (Abbr. for principles: P1: Lawfulness, fairness, transparency; P2: Purpose limitation; P3: Data minimization; P4: Accuracy; P5: Storage Limitation; P6: Confidentiality & integrity; P7: Accountability; P8: Privacy-by-design and by-default; P9: Data subject access request).

a real-time video before the video is stored or made available for analytics. The IoTA is used to communicate the available privacy notices and choices of registered IoT devices to its users. The privacy mediator runs on a cloudlet (edge device) located near an IoT device. When a user comes near a video camera, IoTA identifies the camera, retrieves camera-specific privacy policies, and alerts the user if the user has not set privacy preferences for this camera. Feng et al. (Feng et al., 2021) address the issue of providing meaningful privacy choices to users. In particular, they propose a design space for privacy choices based on a user-centered analysis of what makes pri-

vacuity notices effective. A newer version of the IoTA app leverages the design space (Feng et al., 2021) to implement meaningful privacy choices. Morel et al. (Morel et al., 2019) propose a generic framework for IoT to facilitate the management of information and consent in a way that reduces information asymmetry (or power imbalances) between data controllers and data subjects. They provide a Personal Data Custodian, a software tool installed on gateway devices to inform data subjects about privacy notices which are specified in a policy language endowed with formal semantics. Fernandez et al. (Fernandez et al., 2021a) provide a novel augmented reality privacy manage-




Figure 2: The number of papers addressing the various Privacy Enhancing Technologies (PETs).

ment interface, i.e., the Privacy Augmented Reality Assistant (PARA), for smart home devices. When a smartphone points to a smart IoT device, the PARA interface shows the data collected and allows the users to switch on or off the data collection, offering real-time privacy control. The evaluation results show that PARA users become more aware of IoT devices and their disclosed data, improving their privacy perception and control.

In addition, several research studies (Lee and Kobsa, 2019; Zhang et al., 2020; Emami-Naeini et al., 2017) consider the ability to predict consumers' preferences or decisions as useful. For our purposes, we refer to PETs proposed in these studies (Lee and Kobsa, 2019; Zhang et al., 2020; Emami-Naeini et al., 2017) as privacy predictive models (PPM). The studies designed and performed surveys about different IoT scenarios to capture and analyze consumers' privacy expectations, preferences, and various other factors impacting privacy decisions. Consumers' responses are then used for building machine learning models to predict privacy decisions in IoT. Naeini et al. (Emami-Naeini et al., 2017) presents a prediction model to predict user's comfort level (of data collection) and consent (to allow or deny data collection and use). Lee et al. (Lee and Kobsa, 2019) performed ML experiments to check the feasibility of privacy decision support for user's in IoT and deduce that it is feasible to generate meaningful privacy recommendations for users, but it is crucial to train the model on confident privacy decisions. Zhang et al. (Zhang et al., 2020) demonstrates that by using clustering techniques it is possible to often accurately predict user's consent in facial recognition scenarios. There are several PETs that can increase user's awareness of presence of nearby IoT devices. For instance, Song et al. (Song et al., 2020) designed, prototyped, and evaluated several designs for locators to assist the user to physically locate nearby IoT devices. The locators were, e.g., LED on devices for visual cues, beep sound for auditory cues, or contextualized pictures on a smartphone app. Koelle et al. (Koelle et al., 2018b)

investigated announcement mechanisms in context of body-worn cameras to address the problems of lack of noticeability, understandability, security and trustworthiness. They provide design recommendations for privacy notices for body-worn cameras.

Our analysis of existing PETs for IoT finds that the solutions are privacy-preserving to some extent, but there are certain limitations. For instance, even though the policy language in (Morel et al., 2019) focuses on enhancing informed consent, it is more relevant for privacy-sensitive data subjects who are rare and willing to spend quite some time and go through all the options to set privacy preferences. And it does not significantly reduce the cognitive burden of notice comprehension in IoT systems. Locators in (Song et al., 2020) enhances awareness of presence of nearby IoT devices, but does not offer any control over the devices or data usage. IoTA (Wang et al., 2017) and PARA (Fernandez et al., 2021a) offer simple interfaces enhancing the user awareness of deployed IoT systems and their data practices, but they do not directly work with privacy notices. IoTA requires IoTA owners to register privacy settings in a predefined template, and PARA uses hypothetical privacy settings. Several studies (Lee and Kobsa, 2019; Zhang et al., 2020; Emami-Naeini et al., 2017) have demonstrated that PPM can predict users' privacy decisions with fairly high accuracy, but a weak point of machine learning employed for privacy prediction (and recommendation) in these approaches is the need to process data at the cloud. Meta data, such as privacy profiles, privacy preferences, privacy self-efficacy, needed to train the machine learning models is personal data in itself. Processing the metadata in the cloud may leak user-specific privacy decision making information to adversaries.

Based on our analysis from primary studies, we observe that more research is needed on devising PETs with new features and capabilities that present privacy notices in a way that enhances data subjects' privacy awareness, understandability, and control without overburdening them with notice comprehension and subsequent privacy configurations. In addition, more research is needed to build privacy decision support (like PPM) that can generate privacy recommendations in a privacy-preserving manner, e.g., by using robust decentralized model training.

4.2 Privacy Principles

Figure 3 shows the number of primary studies that address the privacy principles described in Section 2.1. One primary study may address more than one privacy principle. We classify the privacy principles in

three groups based on the number of papers addressing them. The majority of the primary studies mainly address the privacy principles *lawfulness, fairness, transparency* (28 out of 36) and *data protection-by-design* (28 out of 36). This is followed by the privacy principles *integrity and confidentiality* and *purpose limitation* on a shared second place. *Storage limitation, data minimization, accountability, data subject access right*, and *accuracy* are the least supported principles. Nine studies support only one of the nine principles, while only one study supports all the nine principles. Also, we did not find a PET that follows all the principles, mainly due to very diverse privacy concerns addressed by these principles.

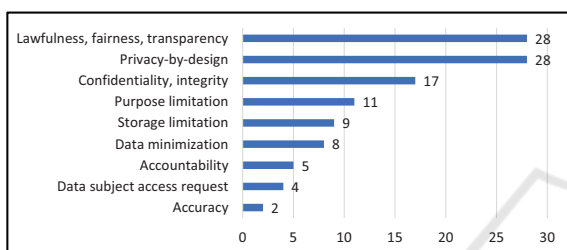


Figure 3: Privacy principles addressed by number of papers.

28 out of 36 studies addressed the lawfulness, fairness and transparency principle. As mentioned earlier, this principle consists of three different yet connected components, and for simplicity reasons, if a study supports one or more components of the principle then we consider that the study supports the principle as a whole. This is the reason that the number of studies addressing this principle is higher. Several studies are focused on enhancing individual privacy awareness through ambient notices (Koelle et al., 2018b; Mhaidli et al., 2020), privacy labels (Emami-Naeini et al., 2021; Emami-Naeini et al., 2017), and also, through privacy assistants/dashboards (Feng et al., 2021; Das et al., 2018; Zhang et al., 2020; Fernandez et al., 2021a) for awareness as well as control. The purpose limitation principle is addressed in (Lee and Kobsa, 2019; Pappachan et al., 2017; Toumia et al., 2020; Wang et al., 2017; Agarwal et al., 2020; Sanchez et al., 2020) by using ontology, privacy assistant, framework, and architecture. Data minimization principle is addressed in (Wang et al., 2017; Pappachan et al., 2017; Fernandez et al., 2021b; Mhaidli et al., 2020), focusing on either minimizing data collection or processing data with fewer personal identifiers. Accuracy principle is addressed by Agarwal et al. (Agarwal et al., 2020) by means of ontologies and by Alkhariji et al. (Alkhariji et al., 2021) by synthesising privacy by design (PbD) knowledge from privacy patterns and privacy principles. Storage princi-

ple is addressed by (Lee and Kobsa, 2019; Sanchez et al., 2020; Toumia et al., 2020; Fernandez et al., 2019; Das et al., 2018) by using a variety of techniques such ontology, privacy architecture, privacy assistants. Confidentiality and integrity principle is addressed by the studies (Benhamida et al., 2021; Fernandez et al., 2021b; Koelle et al., 2018a; Ren et al., 2021; Arachchige et al., 2020; Bose et al., 2015; Fernández et al., 2020; Shi et al., 2021), by using a wide range of techniques such as anonymization, encryption, differential privacy. With 28 out of 36 studies addressing the data protection by design principle, it is clear that privacy is treated as a first class requirement from design to implementation of services and products. Which is a big leap for privacy from the conventional contractual paradigm to computational paradigm. For instance, (Lee and Kobsa, 2019; Fernandez et al., 2021b; Antignac and Le Métayer, 2014; Conzon et al., 2019; Fernandez et al., 2019; Foukia et al., 2016) can be categorized as privacy by architecture approaches which enable privacy at the architecture level. Furthermore, frameworks based on privacy by design principle are proposed in (Pappachan et al., 2017; Huang et al., 2021; Morel et al., 2019; Kammüller, 2018). The last principle, i.e., data subject access request, is to provide individuals with information and more control over their personal data. It is addressed in (Fernández et al., 2020; Toumia et al., 2020; Morel et al., 2019) by means of policy graph, ontology, framework, etc., and in (Alkhariji et al., 2021) by synthesising knowledge on data subject participation and access principle, and relevant privacy patterns. Confidentiality and integrity are prerequisites for privacy, data protection by design and accountability are key to designing and implementing the privacy controls, while data minimization and purpose limitation leans more towards policy decisions. Our literature review shows that research efforts are unevenly distributed. For instance, accountability is addressed in merely 5 studies, although accountability is explicitly included in the GDPR to evidence compliance as well as to strengthen the core privacy principles (P1-P6). This may be a result of the long history of contractual governance practices for privacy. Quite a number of studies are addressing confidentiality and integrity because there is a long history in developing access control mechanisms for specific security policies.

All privacy principles have their distinct roles in protecting personal data during its entire lifecycle, but does not necessarily ensure privacy at large. For instance, transparency is a prerequisite for accountability, but it does not ensure accountability. Similarly, lawfulness and purpose limitation does not en-

sure data protection by design, and data protection by design does not guarantee that data subject's access requests are fulfilled. We need to ascertain if implementations of these privacy principles are effective in providing users with awareness and control over personal data collection and processing. There are several research challenges that need to be addressed to achieve a reasonably privacy-aware IoT. In the next section, we identify specific challenges that needs to be addressed to empower users in making informed privacy choices.

5 CHALLENGES

Privacy of IoT device (used by the user) is prerequisite for privacy of user, and which is only possible when user can exercise appropriate control over their data. However, the current situation as revealed by, e.g., (CISCO, 2019) indicates lack of transparency, awareness, and control. Based on our analysis of primary studies addressing PETs in Section 4.1 and privacy principles in Section 4.2, we reflect in the following on challenges to a human-centered privacy. The analysis sheds lights on three research areas that needs to be addressed to empower users.

5.1 Lack of Technical Expertise in Privacy Notice Comprehension

Complicated privacy notices, low privacy-notice comprehension, and often hard-to-locate privacy notices lead to poor-quality consent. In 2008, McDonald and Craner estimated that reading privacy notices took about 244 hours/year for an average Internet user (McDonald and Cranor, 2008). A decade later, there is a manifold increase in the number of devices and the number of applications an average person uses in a day. The number of privacy notices that users need to consent has increased exponentially. Arguably, consent is often the most exploited legal ground for processing personal data. Research has shown that comprehending privacy notices imposes a high cognitive and time burden on users (Fabian et al., 2017). Despite a fairly large number of studies focused on lawfulness, fairness, and transparency, only a few studies focused on improving the comprehension of privacy notice text, which is key to obtaining "good quality" consent in IoT systems. For instance, Morel et al. (Morel et al., 2019) proposed a policy language with formal semantics which aims to enhance informed consent, which is more relevant for privacy-sensitive data subjects who are rare and willing to spend quite a significant time and go through

all the options to set privacy preferences. In our opinion, it does not significantly reduce the cognitive burden of notice comprehension. Feng et al. (Feng et al., 2021) proposed a design space for privacy choices based on a user-centered analysis to make privacy notices effective, and demonstrated its adoption through interfaces in the IoTA application. The design space is sketched for the practitioners to conceptualize and implement effective privacy notices. However, practitioners proactively using the design space is broadly not in line with *collect all - use all* business models. An alternative to expecting service providers to design and develop comprehensible notices, is to develop assistive technology that could assist users in comprehending privacy notices. However, it is challenging to extract information from privacy notices and present privacy notice text intuitively to users so that the privacy notice is more likely to be noticed, read, and understood.

5.2 Lack of Transparency and Control of Personal Data

It is problematic to provide users with the transparency and control of their personal data, while data exchange in the IoT continuum often results in data flows towards unknown IoT actors. Users usually have no way to know and control the collected personal data and associated data practices, such as the duration for which data will be retained, the third-parties the data is shared with, or the purpose of the data collection. Unsurprisingly, the privacy implications of such hidden data collection and processing are not conveyed to users, and the users do not feel uncomfortable with hidden monitoring. Worse still, privacy notices are hard to locate, and users do not have simple means to elicit consent in IoT settings, as most IoT devices do not have any user interface (Emami-Naeini et al., 2017; Morel et al., 2019). To remedy this situation, the GDPR mandates that data practices be disclosed to data subjects at the time of data collection or before data collection to improve transparency. Article 29 Working Party recommends introducing user-friendly tools through which the data subjects can be informed and engaged regarding the processing of their personal data. IoT systems need a privacy dashboard, i.e., a usable interface which provides users with privacy information and controls which are easy to understand and use in managing their privacy settings. Moreover, there are unknown IoT actors with whom data is often shared (by manufacturers or service providers) and it is difficult to discover such actors in long privacy notices. There are several studies that propose PETs that can increase

user's awareness of presence of nearby IoT devices, such as (Song et al., 2020; Koelle et al., 2018b). Besides, there are PETs which provide some level control in addition to awareness, such as IoTA (Wang et al., 2017) and PARA (Fernandez et al., 2021a). But the limitation is that it depends on device manufacturers or service providers to describe the privacy-related information in a standardized machine-readable format. And this brings us back to the reality of lack of interests and incentives for manufacturers and service providers. Given the number of IoT devices, such solutions relying on manually describing notices in a standardized format are not scalable. Automated notice processing is a research direction that has potential to address the issue of lack of transparency and control, as it can also bring to surface hidden and unfair data practices. Automated methods can be devised to process privacy notices using deep learning, natural language processing, and decentralized machine learning principles to automatically extract precise and nuanced information, such as data practices, opt-out choices. However, it is challenging to extract nuanced information with high accuracy from ambiguous notices. Equally challenging is to devise a single destination tool that provides users privacy information and controls (of their IoT devices), which is easy to utilize in managing their privacy settings.

5.3 Lack of Personalized Privacy Recommendations

Given the pervasiveness of IoT devices/services and therefore the pervasiveness of privacy notices connected with them, there remains the challenge to provide consent on a near constant basis. *Privacy recommendations* can assist IoT users in making meaningful privacy decisions, but the challenge is two-fold: first, it is hard to generate personalized privacy recommendations, and second it is harder to generate the recommendations in a privacy preserving manner. Privacy-preference recommender systems generate privacy recommendations based on users' behaviors with a similar propensity for privacy. To do so, users' diverse privacy preferences (e.g., level of privacy awareness, sensitivity towards privacy, understanding of utility-privacy trade-offs) needs to be captured and modeled. An emerging paradigm in privacy research is the use of machine-learning to make privacy preference recommendations (Das et al., 2018). For instance, Lee and Kobsa demonstrates the feasibility to computationally model and predict privacy preferences in IoT (Lee and Kobsa, 2016). In another study, they also present a preference model that predicts users' binary privacy decisions on whether

to accept or reject a specific IoT scenario, with an accuracy of 77% (Lee and Kobsa, 2017). Zhang et al. (Zhang et al., 2020) developed a predictive model (with *purpose* as a feature) using clustering techniques to assist users in consenting to allow or deny personal data processing in video analytics deployments. Their model was able to predict 93.9% of the binary decisions (allow/deny) with an accuracy of 88.9%. However, a weak point of the conventional machine learning approaches employed for privacy recommendations is the need to process data at the cloud, potentially exposing privacy preference data, e.g., privacy awareness level, privacy sensitivity, privacy choices, sensitivity towards privacy, understanding of utility-privacy trade-offs. Therefore, a decentralized approach involving learning privacy-preference models locally on devices would be a better privacy-preserving approach. Federated learning (FL) (Lim et al., 2020) has the potential to design a privacy preserving privacy-preference model using edge computing. For example, Wang et al. (Wang et al., 2020) presents a support vector machine (SVM) based FL, which enables each mobile user to learn its SVM model using its own local dataset and then collaboratively train a global SVM model, while keeping the training data locally.

6 RELATED WORK

Aleisa et al. (Aleisa and Renaud, 2017) reports on a systematic literature review of privacy-preserving solutions in IoT. The solutions are analyzed in terms of techniques they deployed and extent to which they support privacy principles. The authors consider OECD's eleven privacy principles, while we consider GDPR's privacy principles. Furthermore, the authors also assessed the literature for privacy threats arising from IoT data collection, e.g., profiling, linkage, location tracking.

We and Aleisa et al. (Aleisa and Renaud, 2017) share two critical findings: first, the need to involve data subjects from the outset, and second, the need to design for privacy awareness. Akil et al. (Akil et al., 2020) performed a systematic literature review to analyze types of privacy-preserving identifiers for the IoT environment, particularly implementing pseudonymity. Their classification scheme includes: application domains, types of pseudonyms, and system architecture. Akil et al. (2020) and our study differs in terms of research focus, as they focus only on the data minimization principle while our classification and analysis includes all the seven (GDPR) privacy principles. Li et al. (Li and Palanisamy,

2018) study the privacy protection problem in IoT through a comprehensive review of the state-of-the-art by jointly considering the architecture of the IoT system, the principles of privacy laws, and the representative PETs. In terms of PETs, Li et al. (2018) does an in-depth analysis and evaluation of the various PETs at different layers (Perception, Networking, Middleware, Application) of an IoT architecture. We and Li et al. (2018) share the same motivation, i.e., to provide a broader understanding of privacy principles and state-of-the-art PETs, and understand how the privacy principles may drive the design of PETs. Li et al. (2018) considers eight principles extended from the five basic Fair Information Practices was the foundation of most EU privacy laws. They are to some extent similar to privacy principles we also focus on.

Cha et al. (Cha et al., 2018) conducted a comprehensive literature review on PETs in IoT. The review is based on 120 primary studies published between 2014-2017. The authors present an effectiveness comparison of PETs in terms of three types of indicators: functional requirements, privacy properties, and user-centered properties. Indicators of functional requirements are informed consent, anonymity and pseudonymity mechanisms, and policy enforcement. Indicators of privacy properties are unlinkability and unobservability. Indicators of user-centered properties are transparency and usability. In terms of privacy principles, Cha et al. (2018) have a broader research focus, as their review cover privacy principles from the GDPR as well as ISO/IEC 29100:2011. A review of prominent PETs supporting the fifteen principles (of the GDPR and ISO/IEC 29100:2011) is also presented. Loukil et al. (Loukil et al., 2017) reports on a systematic literature review of privacy-preserving solutions in IoT. Their classification scheme consists of six facets: application domain, IoT architectures, security properties, data life cycle, and privacy-preserving techniques. Loukil et al. (2017) also investigated eleven ISO privacy principles, focusing on architecture and data life cycle. Loukil et al. (2017) and Cha et al. (2018) share a finding that information security related principle (confidentiality and integrity) is the most supported principle, while Loukil et al. (2017) and our study share a finding that accountability is the least supported principle.

7 CONCLUSION

In this article, we investigated the extent to which the GDPR's privacy principles are supported by privacy solutions in IoT. Our review is based on 36 primary studies, which were analyzed along two di-

mensions: PETs and privacy principles. Through meticulous analysis of the primary studies, we find that despite a growing stream of research in privacy-preserving IoT there is an apparent chasm between achieved privacy and desired privacy. Studies and surveys have also shown that consumers intrinsically care about privacy and want to take actions to protect their privacy. This is also reflected by the increasing consumer-oriented PETs. For instance, several studies have started to explore innovative PETs (e.g., IoTA, PARA) for consumers to govern data and its use. In addition to the analysis of PETs and principles, as well as through a reflection of current state of notice & consent paradigm, we derive three main challenges to privacy-aware IoT. Namely, the lack of technical expertise in privacy notice comprehension, the lack of transparency and control of personal data, and the lack of personalized privacy recommendations. Without addressing these challenges, there will be unfair (personal) data collection and processing, unfair commercial practices, privacy violations and reduced consumer trust in organizations.

To fully realise the ambition of privacy-aware IoT, additional approaches and solutions are required to present privacy notices to users in ways that enhance their awareness, understandability, and control without overburdening them with notice comprehension and privacy configurations. A potential enabler for such solutions is automated analysis of notices to extract precise information, e.g., data practices, using natural language processing and deep learning. Privacy recommendation is another research direction where the recommender learns from the user's privacy preferences and generates personalized privacy recommendations. In order to protect metadata about users' privacy decision making information, privacy recommendations needs to be generated in privacy-preserving manner, e.g., using federated learning.

ACKNOWLEDGEMENTS

This work has been carried out as part of the EU projects CyberSec4Europe (830929) and CyberKit4SME (883188) funded by the EU Horizon 2020 Research and Innovation Programme.

REFERENCES

- Acquisti, A., Brandimarte, L., and Loewenstein, G. (2020). Secrets and likes: the drive for privacy and the difficulty of achieving it in the digital age. *Journal of Consumer Psychology*, 30(4):736–758.

- Agarwal, R., Elsaleh, T., and Tragos, E. (2020). Gdpr-inspired iot ontology enabling semantic interoperability, federation of deployments and privacy-preserving applications. *arXiv preprint arXiv:2012.10314*.
- Akil, M., Islami, L., Fischer-Hübner, S., Martucci, L. A., and Zuccato, A. (2020). Privacy-preserving identifiers for iot: a systematic literature review. *IEEE Access*, 8:168470–168485.
- Aleisa, N. and Renaud, K. (2017). Privacy of the internet of things: a systematic literature review. In *Hawaii International Conference on System Sciences 2017*, pages 5947–5956.
- Alkharaji, L., Alhirabi, N., Alraja, M. N., Barhamgi, M., Rana, O., and Perera, C. (2021). Synthesising privacy by design knowledge toward explainable internet of things application designing in healthcare. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 17(2s):1–29.
- Altman, I. (1977). Privacy regulation: Culturally universal or culturally specific? *Journal of social issues*, 33(3):66–84.
- Antignac, T. and Le Métayer, D. (2014). Privacy architectures: Reasoning about data minimisation and integrity. In *International Workshop on Security and Trust Management*, pages 17–32. Springer.
- Arachchige, P. C. M., Bertok, P., Khalil, I., Liu, D., Camtepe, S., and Atiquzzaman, M. (2020). A trustworthy privacy preserving framework for machine learning in industrial iot systems. *IEEE Transactions on Industrial Informatics*, 16(9):6092–6102.
- Benhamida, F. Z., Navarro, J., Gómez-Carmona, O., Casado-Mansilla, D., López-de Ipiña, D., and Zaballos, A. (2021). Pyff: A fog-based flexible architecture for enabling privacy-by-design iot-based communal smart environments. *Sensors*, 21(11):3640.
- Bose, T., Bandyopadhyay, S., Ukil, A., Bhattacharyya, A., and Pal, A. (2015). Why not keep your personal data secure yet private in iot?: Our lightweight approach. In *IEEE 10th International Conference on Intelligent Sensors, Sensor Networks and Information Processing*, pages 1–6.
- Cha, S.-C., Hsu, T.-Y., Xiang, Y., and Yeh, K.-H. (2018). Privacy enhancing technologies in the internet of things: perspectives and challenges. *IEEE Internet of Things Journal*, 6(2):2159–2187.
- CISCO (2019). Consumer privacy survey: The growing imperative of getting data privacy right. https://www.cisco.com/c/dam/global/en_uk/products/collateral/security/cybersecurity-series-2019-cps.pdf. Accessed: 2022-01-06.
- Conzon, D., Rashid, M. R. A., Tao, X., Soriano, A., Nicholson, R., and Ferrera, E. (2019). Brain-iot: Model-based framework for dependable sensing and actuation in intelligent decentralized iot systems. In *2019 4th International Conference on Computing, Communications and Security (ICCCS)*, pages 1–8. IEEE.
- Das, A., Degeling, M., Smullen, D., and Sadeh, N. (2018). Personalized privacy assistants for the internet of things: Providing users with notice and choice. *IEEE Pervasive Computing*, 17(3):35–46.
- Elkhodr, M., Shahrestani, S., and Cheung, H. (2013). A contextual-adaptive location disclosure agent for general devices in the internet of things. In *38th Annual IEEE Conference on Local Computer Networks-Workshops*, pages 848–855. IEEE.
- Emami-Naeini, P., Bhagavatula, S., Habib, H., Degeling, M., Bauer, L., Cranor, L. F., and Sadeh, N. (2017). Privacy expectations and preferences in an iot world. In *Thirteenth Symposium on Usable Privacy and Security ({SOUPS} 2017)*, pages 399–412.
- Emami-Naeini, P., Dheenadhayalan, J., Agarwal, Y., and Cranor, L. F. (2021). Which privacy and security attributes most impact consumers' risk perception and willingness to purchase iot devices? In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 1937–1954.
- EP and CEU (2016). The General Data Protection Regulation (GDPR). <https://eur-lex.europa.eu/eli/reg/2016/679/oj>. Accessed: 2019-11-24.
- Escher, S., Weller, B., Köpsell, S., and Strufe, T. (2020). Towards transparency in the internet of things. In *Annual Privacy Forum*, pages 186–200. Springer.
- Fabian, B., Ermakova, T., and Lentz, T. (2017). Large-scale readability analysis of privacy policies. In *Proc. of the international conference on web intelligence*, pages 18–25.
- Feng, Y., Yao, Y., and Sadeh, N. (2021). A design space for privacy choices: Towards meaningful privacy control in the internet of things. In *Proc. of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–16.
- Fernandez, C. B., Lee, L. H., Nurmi, P., and Hui, P. (2021a). Para: Privacy management and control in emerging iot ecosystems using augmented reality. In *ACM International Conference on Multimodal Interaction*. Association for Computing Machinery (ACM).
- Fernandez, C. B., Nurmi, P., and Hui, P. (2021b). Seeing is believing? effects of visualization on smart device privacy perceptions. In *Proc. of the 29th ACM International Conference on Multimedia (MM'21)*.
- Fernández, M., Franch Tapia, A., Jaimunk, J., Martínez Chamorro, M., and Thuraisingham, B. (2020). A data access model for privacy-preserving cloud-iot architectures. In *Proc. of the 25th ACM Symposium on Access Control Models and Technologies*, pages 191–202.
- Fernandez, M., Jaimunk, J., and Thuraisingham, B. (2019). Privacy-preserving architecture for cloud-iot platforms. In *2019 IEEE International Conference on Web Services (ICWS)*, pages 11–19. IEEE.
- Foukia, N., Billard, D., and Solana, E. (2016). Pisces: A framework for privacy by design in iot. In *2016 14th Annual Conference on Privacy, Security and Trust (PST)*, pages 706–713. IEEE.
- Godinho de Matos, M. and Adjerid, I. (2021). Consumer consent and firm targeting after gdpr: The case of a large telecom provider. *Management Science*.
- Huang, Y., Cai, Z., Pang, J., Xie, Z., and Bourgeois, A. G. (2021). Game theory based privacy protection for context-aware services with long-term time series data. In *ICC 2021-IEEE International Conference on Communications*, pages 1–6. IEEE.

- Kammüller, F. (2018). Formal modeling and analysis of data protection for gdpr compliance of iot healthcare systems. In *2018 IEEE International Conference on Systems, Man, and Cybernetics*, pages 3319–3324.
- Koelle, M., Ananthanarayan, S., Czupalla, S., Heuten, W., and Boll, S. (2018a). Your smart glasses' camera bothers me! exploring opt-in and opt-out gestures for privacy mediation. In *Proc. of the 10th Nordic Conference on HCI*, pages 473–481.
- Koelle, M., Wolf, K., and Boll, S. (2018b). Beyond led status lights-design requirements of privacy notices for body-worn cameras. In *Proc. of the 12th International Conference on Tangible, Embedded, and Embodied Interaction*, pages 177–187.
- Lee, H. and Kobsa, A. (2016). Understanding user privacy in internet of things environments. In *2016 IEEE 3rd World Forum on Internet of Things*, pages 407–412.
- Lee, H. and Kobsa, A. (2017). Privacy preference modeling and prediction in a simulated campuswide iot environment. In *2017 IEEE International Conference on Pervasive Computing and Communications (PerCom)*, pages 276–285. IEEE.
- Lee, H. and Kobsa, A. (2019). Confident privacy decision-making in IoT environments. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 27(1):1–39.
- Li, C. and Palanisamy, B. (2018). Privacy in internet of things: From principles to technologies. *IEEE Internet of Things Journal*, 6(1):488–505.
- Lim, W. Y. B., Luong, N. C., Hoang, D. T., Jiao, Y., Liang, Y.-C., Yang, Q., Niyato, D., and Miao, C. (2020). Federated learning in mobile edge networks: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, 22(3):2031–2063.
- Loukil, F., Ghedira-Guegan, C., Benharkat, A. N., Boukadi, K., and Maamar, Z. (2017). Privacy-aware in the iot applications: a systematic literature review. In *OTM Confederated International Conferences" On the Move to Meaningful Internet Systems"*, pages 552–569. Springer.
- McDonald, A. M. and Cranor, L. F. (2008). The cost of reading privacy policies. *Journal of law and policy for the information society*, 4:543.
- Mehrotra, S., Kobsa, A., Venkatasubramanian, N., and Rajagopalan, S. R. (2016). Tippers: A privacy cognizant iot environment. In *2016 IEEE International Conference on Pervasive Computing and Communication Workshops (PerCom Workshops)*, pages 1–6. IEEE.
- Mhaidli, A., Venkatesh, M. K., Zou, Y., and Schaub, F. (2020). Listen only when spoken to: Interpersonal communication cues as smart speaker privacy controls. *Proc. on Privacy Enhancing Technologies*, 2020(2):251–270.
- Morel, V., Cunche, M., and Le Métayer, D. (2019). A generic information and consent framework for the iot. In *2019 18th IEEE International Conference On Trust, Security And Privacy In Computing And Communications*, pages 366–373. IEEE.
- Norwegian Consumer Council (2021). Surveillance-based advertising: Consumer attitudes to surveillance-based advertising. <https://fil.forbrukerradet.no/wp-content/uploads/2021/06/consumer-attitudes-to-surveillance-based-advertising.pdf>.
- O'Connor, Y., Rowan, W., Lynch, L., and Heavin, C. (2017). Privacy by design: informed consent and internet of things for smart health. *Procedia computer science*, 113:653–658.
- Pappachan, P., Degeling, M., Yus, R., Das, A., Bhagavatula, S., Melicher, W., Emami-Naeini, P., Zhang, S., Bauer, L., Kobsa, A., 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)*, pages 193–198. IEEE.
- Ren, W., Tong, X., Du, J., Wang, N., Li, S., Min, G., and Zhao, Z. (2021). Privacy enhancing techniques in the internet of things using data anonymisation. *Information Systems Frontiers*, pages 1–12.
- Sadique, K. M., Rahmani, R., and Johannesson, P. (2020). Enhancing data privacy in the internet of things (iot) using edge computing. In *International Conference on Computational Intelligence, Security and Internet of Things*, pages 231–243. Springer.
- Sanchez, O. R., Torre, I., and Knijnenburg, B. P. (2020). Semantic-based privacy settings negotiation and management. *Future Generation Computer Systems*, 111:879–898.
- Shi, N., Tang, B., Sandhu, R., and Li, Q. (2021). Duce: Distributed usage control enforcement for private data sharing in internet of things. In *IFIP Annual Conference on Data and Applications Security and Privacy*, pages 278–290. Springer.
- Song, Y., Huang, Y., Cai, Z., and Hong, J. I. (2020). I'm all eyes and ears: Exploring effective locators for privacy awareness in iot scenarios. In *Proc. of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–13.
- The Royal Society (2019). Protecting privacy in practice: The current use, development and limits of privacy enhancing technologies in data analysis. <https://royalsociety.org/topics-policy/projects/privacy-enhancing-technologies/>.
- Toumia, A., Szonieczky, S., and Saleh, I. (2020). Colpri: Towards a collaborative privacy knowledge management ontology for the internet of things. In *2020 Fifth International Conference on Fog and Mobile Edge Computing (FMEC)*, pages 150–157. IEEE.
- Wang, J., Amos, B., Das, A., Pillai, P., Sadeh, N., and Satyanarayanan, M. (2017). A scalable and privacy-aware iot service for live video analytics. In *Proc. of the 8th ACM on Multimedia Systems Conference*, pages 38–49.
- Wang, S., Chen, M., Saad, W., and Yin, C. (2020). Federated learning for energy-efficient task computing in wireless networks. In *IEEE International Conference on Communications*, pages 1–6. IEEE.
- Zhang, S., Feng, Y., Das, A., Bauer, L., Cranor, L. F., and Sadeh, N. (2020). Understanding people's privacy attitudes towards video analytics technologies. *Proc. of FTC PrivacyCon*, pages 1–18.