



Review



Revealing the landscape of privacy-enhancing technologies in the context of data markets for the IoT: A systematic literature review

Gonzalo Munilla Garrido ^{a,*}, Johannes Sedlmeir ^b, Ömer Uludağ ^a, Ilias Soto Alaoui ^a, Andre Luckow ^c, Florian Matthes ^a

^a Technical University of Munich, Department of Informatics, Munich, Germany

^b Project Group Business & Information Systems Engineering of the Fraunhofer FIT, Bayreuth, Germany

^c Ludwig Maximilian University of Munich, Department of Computer Science, Munich, Germany

ARTICLE INFO

Keywords:

Anonymization
Big data
Copy problem
Data exchange
Marketplace
Platform
Secure computation

ABSTRACT

IoT data markets in public and private institutions have become increasingly relevant in recent years because of their potential to improve data availability and unlock new business models. However, exchanging data in markets bears considerable challenges related to disclosing sensitive information. Despite considerable research focused on different aspects of privacy-enhancing data markets for the IoT, none of the solutions proposed so far seems to find a practical adoption. Thus, this study aims to organize the state-of-the-art solutions, analyze and scope the technologies that have been suggested in this context, and structure the remaining challenges to determine areas where future research is required. To accomplish this goal, we conducted a systematic literature review on privacy enhancement in data markets for the IoT, covering 50 publications dated up to July 2020, and provided updates with 24 publications dated up to May 2022. Our results indicate that most research in this area has emerged only recently, and no IoT data market architecture has established itself as canonical. Existing solutions frequently lack the required combination of anonymization and secure computation technologies. Furthermore, there is no consensus on the appropriate use of blockchain technology for IoT data markets and a low degree of leveraging existing libraries or reusing generic data market architectures. We also identified significant challenges remaining, such as the copy problem and the recursive enforcement problem that – while solutions have been suggested to some extent – are often not sufficiently addressed in proposed designs. We conclude that privacy-enhancing technologies need further improvements to positively impact data markets so that, ultimately, the value of data is preserved through data scarcity and users' privacy and businesses-critical information are protected.

1. Introduction

IoT devices have been improved, mass-produced, and deployed in the past few decades through steady progress in information and communication technologies (ICTs) and motivated by a trend of data-driven decision-making, automation, and the opportunity for new business models. IoT devices' primary collective purpose is to interact with the physical world and enable the measurement and collection of events and interactions [S23]. These characteristics apply to IoT devices deployed in, for example, a factory or a powerline network and many devices employed by people, such as cell phones, laptops, or wearables. The *volume*, *velocity*, and *variety* of the information generated by the IoT is immense, which drove practitioners to coin the term *big data* and develop tools for their analysis [S4]. Public and private institutions use *big data* to promote the public good, innovations, and improve products

and services. Big data has become the foundation of the emerging data economy, which in Europe was worth nearly 2% of its GDP in 2016, close to 300 billion Euros [1]. However, the generation, collection, storage, processing, distribution, and analysis of *big data* to realize such economic potential also come with challenges for enterprises and responsibilities toward society.

Big data needs to be accessible to institutions that can harness their potential and develop innovations, lest society fails to materialize their advantages. Unfortunately, a significant share of the world's data is siloed and exploited solely by the institutions that host them [2], consequently locking the untapped potential of the data economy and hindering progress in science, business, and society. To surmount this obstacle, a paradigm shift toward openness emerged in the form of electronic data markets for the IoT, i.e., mediums for the trade

* Correspondence to: Boltzmannstrasse 3, 85748, Garching, Germany.

E-mail addresses: gonzalo.munilla-garrido@tum.de (G.M. Garrido), johannes.sedlmeir@fit.fraunhofer.de (J. Sedlmeir).

<https://doi.org/10.1016/j.jnca.2022.103465>

Received 26 July 2021; Received in revised form 23 May 2022; Accepted 7 July 2022

Available online 16 July 2022

1084-8045/© 2022 Elsevier Ltd. All rights reserved.

of information across the Internet based on electronic infrastructure [3]. This paradigm brings potential benefits, such as increasing the efficiency of business processes, facilitating growth by unlocking new business models [4], and profiting from trading. Decision-makers in governments and businesses have recognized the economic potential of data markets and hence recently supported significant projects that provide a shared digital infrastructure for data-sharing initiatives such as GAIA-X [5] or the automotive-related Catena-X, which promote the collaboration of large enterprises in data markets.

Despite data markets' promise to benefit society by fostering innovation and collaboration across enterprises, these markets hold the risk of exposing individuals' and businesses' sensitive information [6,7]. Moreover, confidence in privacy protection is an essential driver of users' willingness to share their data [8]. Similarly, businesses are unwilling to bear the risk of unintentionally leaking their customers' private or business-critical information. Consequently, the adoption of data markets is generally hampered. Additionally, while a corporation may have taken security measures to protect collected data from unauthorized access or unintended use, data buyers might not have the same standards. Hence, in this case, exchanging data entails an additional risk that the seller needs to mitigate *before* the data are shared. Furthermore, blockchains are expected to play an essential role in the ability of institutions to trade data in tokenized form [8], but their inherent transparency further increases the need to make data exchanged in markets less sensitive [S25]. Additional trends that aggravate the negative consequences of lacking data protection for institutions are recent privacy laws such as the GDPR in Europe or the CCPA in California with their increasingly expensive fines for data breaches [9].

As a first reaction to these risks, practitioners and corporations have increased their systems' security. However, if data are sold and, thus, replicated, confidentiality is not sufficient to protect privacy: Only managing or modifying the data in a way that enhances privacy while preserving as much utility as possible is effective [10]. Thus, institutions have started to allocate more resources to balance data utility and privacy, employing privacy-enhancing technologies (PETs) [S27]. The term PET was coined in 1995 in a report by the Dutch Data Protection Authority and the Ontario Information Commissioner [11] that explored a novel approach to privacy protection [186]. These technologies take the form of architectures built with privacy-by-design principles and policies [S39,S6], or data modifications based on heuristics or mathematical privacy guarantees. Prominent examples of PETs are differential privacy [12,13], syntactic anonymization definitions like k -anonymity [14], homomorphic encryption [15–17], trusted execution environments [18], secure multiparty computation [19], zero-knowledge proofs [20,21], and a set of conventional de-identification approaches such as masking, rounding, or hashing [22].

The relevance of PETs in data markets also increases with the growing adoption of IoT devices, such as in vehicles, wearables, smartphones, and the applications that stream data daily from millions of individuals' private lives to data marketplaces [23]. Despite their current relevance and growing attention [S27], researchers and institutions still find PETs challenging to understand, integrate, and deploy in IoT data markets because most PETs are technically complex and have a wide range of variations and combinations with different tradeoffs [24]. Regarding research addressing these challenges, primary studies are predominant, i.e., studies based on original designs developed or data collected by their authors, while secondary studies collecting and systematizing existing knowledge are less frequent. The applications proposed by primary studies range from funneling data from markets into machine learning (ML) algorithms [S2,S25,S43], crowdsourcing data into markets [S16,S22,S24,S47], adopting data markets for smart mobility ecosystems [S3], smart manufacturing [S19], smart homes [S11], and smart wearables in the health industry [S48].

On the other hand, 9 out of the 50 studies that we identified in our systematic literature review (SLR) are secondary, and out of

these, four studies [S19,S23,S35,S48] cover *some* of the PETs available for data markets for the IoT, yet without giving a detailed comparison of their functionalities, benefits, and limitations. The other five secondary studies [S4,S8,S14,S27,S38] perform high-level surveys revolving around challenges, non-technical privacy strategies, and user-centric perspectives on data markets for the IoT. However, none of these secondary studies provided a rigorous, *systematic* review that collected and mapped PETs and challenges comprehensively. Moreover, as we discuss in Section 9, we noted a low level of re-using existing components to build a more holistic architecture for data markets in related work, which may indicate the need for systematically analyzing the current seminal components, strengths, and weaknesses of solutions proposed for privacy-enhancing IoT data markets.

Consequently, we tackle the research gaps mentioned above with a comprehensive and detailed SLR that aims to guide decision-makers, privacy officers, policymakers, and researchers in the challenge of employing PETs to build or participate in privacy-enhancing IoT data markets. We guide these stakeholders by identifying, classifying, and describing how PETs are leveraged in the current body of scientific knowledge (see Sections 6 and 7) and presenting key findings from our SLR (see Section 9). Moreover, for the benefit of the reader, we distill terminology from the extant literature to differentiate and navigate the concepts of PETs in the scope of this SLR (see Section 5). We also organize related work into a reference model for the use of PETs in IoT data markets in distinct categories (see Fig. 10 and Fig. C.11) and identify narrow and broad challenges that PETs can tackle or circumvent (see Fig. 8). Through mapping PETs to the distilled terminology and the identified narrow challenges, we want to support practitioners in making informed decisions about the appropriate PETs to employ in the context of IoT data markets (see Table 3).

The remainder of the paper is structured as follows. Section 2 introduces the main concepts of privacy, data markets, and the IoT. Section 3 portrays how we conducted our SLR on publications dated before July 2020, followed by a discussion of related work in Section 4 and a distillation of terminology in Section 5. Sections 6 and 7 present the main results from analyzing the content of the studies in our SLR, followed by a structured review of challenges in Section 8. Based on these results and the studies' metadata, we extract a set of key findings and artifacts in Section 9, where we also provide future work and discuss the limitations of our research. Finally, Section 10 updates our study by including new selected publications from May 2022 to July 2020, and Section 11 concludes with a summary of the results.

2. Background

2.1. Privacy

Given the increased attention and relevance of privacy during the past decades, practitioners have provided many acknowledged definitions. For example, [25] stated that “[Privacy is] *the claim of individuals [...] to determine for themselves when, how and to what extent information about them is communicated [...]*”. Similar definitions have been given by other authors like [26] or [27]. Despite these efforts, D. J. Solove argued that any attempt to distill a unique, timeless definition is infeasible due to privacy's multifaceted concept [28]. However, in the field of computer science, a narrower definition may be possible by adopting an *attack model* perspective, as the concept of privacy would likely not have emerged if transgressors would not exist: attackers of one's secrets tacitly give meaning to privacy. Therefore, a helpful definition in the context of computer science may be F. T. Wu's [29]: “[Privacy] is defined not by what it is, but by what it is not – it is the absence of a privacy breach that defines a state of privacy”. F. T. Wu hence defined privacy as a product of a threat model, the one from [30], in which a practitioner needs to determine what information to hide, from whom, and what harms should be prevented before defining legal and technical tools.

Once IT architectures and tools enhance privacy appropriately, other advantages emerge. For example, from an economic perspective, privacy enables data utilization across organizations and applications to create new fair products and services and prevent price discrimination [31]. Furthermore, employing PETs may increase the number of sources and data harvested by institutions because PETs help to overcome regulatory barriers [32], in addition to mitigating the risk of fines and differentiating and appreciating a brand [4]. Moreover, political freedom and stability may only be achieved by unobtrusive forms of governments [33], privacy-enhancing journalism, and less pervasive forms of digital products such as social media that can enable malicious social engineering [S38]. Moreover, research indicates that compromising privacy can result in negative long-term economic effects [34].

Despite these benefits, and while consumers emphasize that privacy is important to them, they are typically not willing to make small additional efforts or pay for privacy [35], the so-called *privacy paradox*. Thus, in the past decades, governments have enacted rules and laws to protect consumers against violations of their privacy, specifically for data captured through advanced ICTs such as personal computers, the world wide web, or smartphones. Examples include the European Data Protection Directive in 1995, the HIPAA Privacy and Security Rule of 1996, the APEC cross-border privacy rules of 2011, the GDPR of 2016, and the Consumer Privacy Act of 2020 in the USA, which comprises Acts such as the CCPA of 2018 in California.

2.2. Data markets

According to [3], there is a misconception in everyday language between the terms “market” and “marketplace”: A marketplace is the implementation of a market in terms of *infrastructure*, time, and location (virtual or physical) where the participants transact. Markets are the *environments* where buyers and sellers set the price and quantity of a particular good. Marketplaces have evolved over millennia; however, the most drastic changes arguably have happened in the past few decades. ICT has driven the costs of instant and ubiquitous communication to an often negligible amount, which has led to the digitization of many existing transaction-based ecosystems, including marketplaces [3]. Moreover, ICT has enabled the creation of virtual marketplaces that did not exist before [36] – the most prominent example being e-commerce. In this context, data have become goods themselves [3]. Data markets incentivize institutions to collect more data and to profit from trading, and, in turn, the resulting improvements and innovations benefit the public good [37].

Beyond the above formal definition, from the selected studies, we can carve out several characteristics of data marketplaces: [S1] indicated that most of the data marketplaces in operation are centralized, where the platform is run by either a trusted third party (a broker) that coordinates buyers and sellers or by the data owner (e.g., a large institution) who is also selling the data. Another 15 selected studies also proposed decentralized architectures employing distributed ledger technology to counter the drawbacks of centralized systems (e.g., single point of failure or trust on a potentially malicious entity). On the other hand, [S46] took another perspective, characterizing data trading platforms depending on the number and type of data domains: general platforms include data from any source type, while specialized platforms focus on one domain, e.g., financial, healthcare, or social media. [S27] identified two categories for data markets based on the type of participant: companies or private individual customers, e.g., owners of a smart home. Altogether, we distilled three dimensions for characterizing data markets: (i) the degree of centralization, (ii) the types and number of data domains, and (iii) the types of sellers and consumers. These dimensions permeate most of the identified solutions in this study, and all exhibit individual privacy trade-offs of which practitioners need to be aware (see Section 9).

2.3. The internet of things

The IoT is considered a network of physical devices that leverages sensors to measure and collect information from the real world and support the access and exchange of data via the Internet instantly and ubiquitously [38]. IoT devices are considered essential for gathering big data [39], which in turn brings new opportunities such as targeted advertisement, predictive maintenance, and quality improvements. Consequently, many companies have introduced the IoT in their strategy for participating in the data economy [38] and make substantial investments in the technologies that make them possible: sensors, wireless networks, and cloud computing infrastructure [39]. In this SLR, the definition of the IoT includes any device with a CPU connected to the Internet, including sensors in factories, supply chains, or vehicles, and devices such as smartphones, wearables, and computers that people use daily. These devices act as data collectors and as the gateway to a plethora of applications that collect users’ actions and behavior, such as browsers, social media, e-commerce, or media entertainment, as well as sensor data generated in business processes like manufacturing and predictive maintenance, and use them for analyzes and predictions.

The design and implementation of data markets are dependent on the IoT. The ubiquity of IoT devices generates many constellations for different degrees of decentralization, with a myriad of possible sources and prosumer types. Furthermore, while such ubiquity will likely boost the data economy and its products and services, IoT devices also permeate many aspects of an individual’s life, e.g., dealing with highly sensitive healthcare data or capturing sensitive information from a business perspective. Hence, the sensitivity of the data gathered from IoT devices calls for the implementation of PETs.

3. Research process

3.1. Goal and research questions

We employed the Goal-Question-Metric paradigm [40] to formulate the focus of this study as follows: we systematically analyze peer-reviewed literature to provide an overview of the state-of-the-art concerning available research and trade-offs on privacy-enhancing data markets for the IoT as well as potential research gaps from the point of view of both scholars and practitioners. Based on this paradigm, the research questions (RQs) we pursued were:

RQ1. What relevant PETs enable IoT data markets?

By answering this RQ, we aim to reveal, describe, and classify PETs in the context of data markets for the IoT based on their fundamentals and applications to give researchers and practitioners an overview of the PETs researched and employed so far.

RQ2. What challenges and trade-offs hinder privacy-enhancing IoT data markets?

Through answering this RQ, we account for explicit and implicit challenges depicted and tackled in existing work so that researchers may quickly identify pain points in the field and focus their research.

3.2. SLR execution

We conducted a SLR based on the guidelines of [41]. SLRs aim to collect, structure, and summarize the existing evidence and gaps in a particular research field to pave the way for future research. Furthermore, SLRs need to provide a rigorous and auditable methodology that can be reviewed and replicated [42]. SLRs define research questions, and a set of predefined inclusion and exclusion criteria that assess potentially relevant primary studies to answer them [43,44]. Table F.10 of the Appendix F contains the criteria for this SLR related to focus, quality, and accessibility.

To conduct the *study search*, we identified the most relevant publications in the field of privacy-enhancing data markets for the IoT to

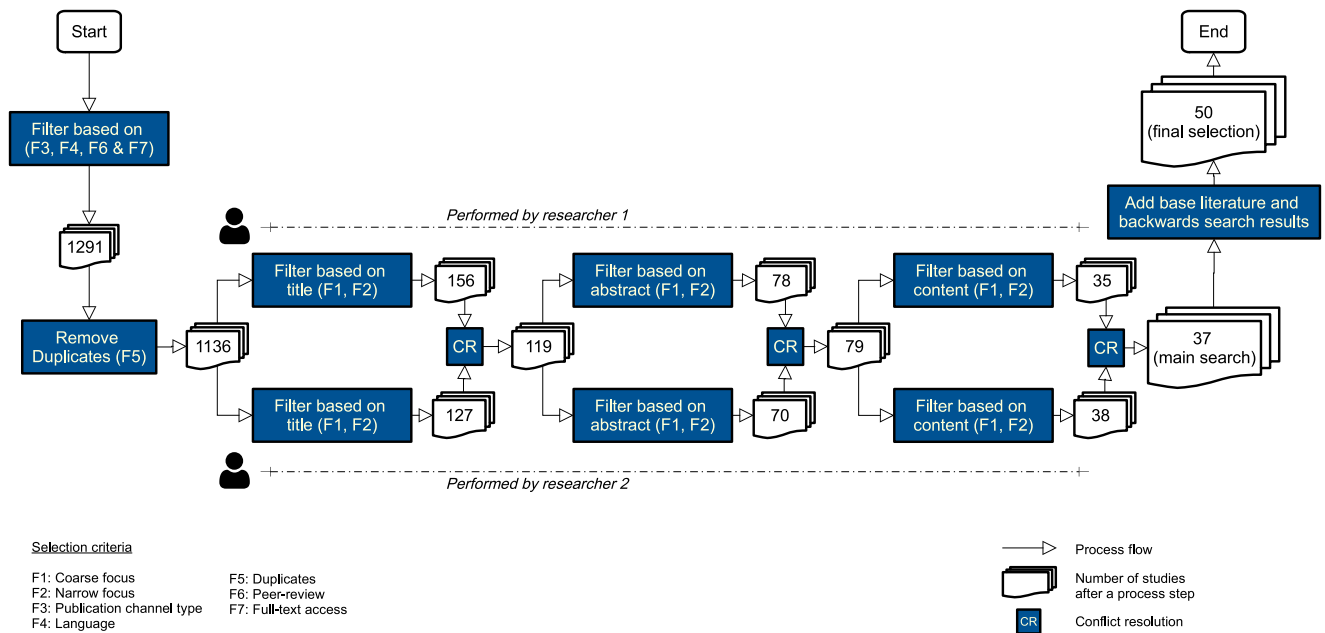


Fig. 1. Study selection process.

answer our research questions [45,46]. To obtain a corpus of high-quality publications, we defined a search strategy based on the work of [46]. Accordingly, our strategy consisted of three phases:

(i) A preliminary search of the base literature. The base literature includes representative papers (8) in the field of privacy-enhancing data markets for the IoT known to the researchers before the SLR, and some other publications found manually in the digital library of the researchers' university, which also complied with the criteria described in Table F.10. We created preliminary search strings based on identified keywords and synonyms that we found in the base literature and research questions. Afterward, we parsed our base literature with a tool to analyze frequent phrases and keywords. Using the results of this analysis, we refined our search terms [47]. Finally, we clustered the search terms into three strings based on the field of this SLR: Privacy, data markets, and IoT. Altogether, we composed the following search strings:

C1: *privacy OR private OR encryption OR encrypted OR encrypt OR data protection*

C2: *data market OR data marketplace OR data trading OR data broker OR data trader OR data auction*

C3: *Internet of things OR Internet of everything OR IoT OR sensor OR connected devices OR networked devices OR smart devices OR controller OR edge computing OR cloud infrastructure OR machine to machine OR M2M OR web-of-things OR WoT OR mobility OR automotive OR vehicle OR car OR automobile OR industry 4.0 OR smart grids OR V2V OR IIoT OR machine learning OR mobile OR cyber-physical OR microservice OR microcontroller OR micro-service OR micro-controller OR blockchain OR neural network OR smart learning OR automated driving OR autonomous driving OR smart city OR smart factory.*

Consequently, we defined our final search string as **C1 AND C2 AND C3**.

(ii) The main search. Since no single source may contain all the high-quality, relevant publications [48,49], we selected seven electronic databases (see Table F.9 that focus on computer science or software engineering and, according to L. Chen et al. cover the most relevant databases in these fields [50]. The time frame that we specified covers any publication included in the selected digital libraries before the 13th of July 2020. With the defined search string, time interval, databases, and following the process Fig. 1 depicts, we collected 1291 studies (1136 after duplicates removal), which two researchers filtered independently and redundantly by title (119 selected

out of 1136), abstract (79 selected out of 119), and body (37 selected out of 79) following the predefined inclusion and exclusion criteria of Table F.10 to reduce bias. After each of the three filtering phases, both researchers resolved conflicts in an informed discussion and attended to the criteria.

(iii) A backward search of the references of the 37 studies resulting from the main search. After filtering by title, abstract, and body, considering our inclusion and exclusion criteria, we included another 11 studies in our corpus. The process resulted in a total of 50 studies from which we subsequently extracted and synthesized data. Hence, the SLR yielded a considerable but not excessive number of results. Furthermore, thanks to the multiple synonyms in the search string, the 37 studies only missed two studies from the base literature. Moreover, the backward search only added a modest number of new works (11). Thus, the process suggests that the choice of search terms was suitable.

To answer the research questions, we performed a *data extraction* of key information from the 50 publications in a structured manner [51]. To reduce the degree of bias, two scientists defined and independently followed an extraction card, which contained the following twenty fields: Authors, cite count, year, country, publication channel, publication type, publication source, research type, research approach, contribution type, tags, topic, subtopic, sub-subtopic, research goal, research questions, study findings, privacy-enhancing architecture or technologies, challenges, and future work. After the two scientists completed the data extraction, they held an informed discussion to resolve any possible conflicts on the extracted information. For the *data synthesis* necessary to answer RQ1 and RQ2, we adapted the "narrative synthesis" method described by [41] and performed the following synthesis procedure: (i) we developed a preliminary synthesis of the findings, followed by (ii) exploring relationships in the data and (iii) refining the preliminary synthesis with the newly acquired knowledge. After the refinement, we returned to the second step until we deemed the RQs answered.

Finally, with the goal of including significant updates and reaffirm the findings extracted from our initial research process, we conducted the same systematic search process for publications dated between July 2020 and May 2022 (24 new publications) and included the findings in Section 10.

Table 1
Short description of the involved layers.

Layer	Description
<i>Storage</i>	Persists data for future use.
<i>Processing</i>	Uses data from storage and runs algorithms on it, typically to extract information.
<i>Communication</i>	Exchange of data with other devices.
<i>Verification</i>	Checking via processing whether the data received and the identities involved are authentic.
<i>Sovereignty</i>	Ability to govern which (sensitive) information the communication with others exposes.
<i>Consensus</i>	A special case of communication and verification in which data is compared and synchronized with other devices' data.

4. Related work

In our SLR, we found nine secondary studies, i.e., studies that systematize and organize existing knowledge, conducted in the context of privacy enhancements for IoT data markets [S4,S8,S14,S19,S23,S27,S35,S38,S48], which Table B.4 summarizes. Some of these studies focused on privacy-related challenges in data markets for the IoT [S14,S19,S27,S38], while others delved into PETs from a technical perspective and discussed their challenges and opportunities [S23,S48,S35]. Lastly, [S8] analyzed users' preferences in privacy-enhancing data markets, and [S4] listed technical design choices for data markets for the IoT.

These secondary studies provided valuable contributions and built the foundation of our work; however, they focused on different aspects of IoT data markets and, therefore, lacked depth in the concepts we present in this study. For example, [S23,S35] and [S48] briefly discussed, altogether, blockchain technology, secure and outsourced computation, k -anonymity, and differential privacy and sketched their implications without considering other PETs that we identified in our SLR. Furthermore, although [S19] provided an overview of the available technologies and their challenges, the authors did not discuss PETs in detail, e.g., the study mentioned anonymization but did not delve into k -anonymity or differential privacy. Additionally, the authors only discussed three out of the six challenges we found: the recursive enforcement problem, the utility and privacy trade-off, and attacks on privacy. Moreover, [S19] based their results on observations from exemplary use cases and, therefore, cannot provide the scientific rigor and comprehensiveness of a SLR. The remaining secondary studies focused on privacy strategies instead of technology, discussed digital rights, described challenges at a high level, or provided a user-centric view of data markets.

Regarding PETs classification, some selected papers included in our SLR provided a framework. Notably, [S48] considered two categories for classification (*outsourced computations* and *information sharing*), whereas [S19] provided a more involved classification than [S48] with five layers: *data security*, *data processing*, *proving support*, *platform capabilities*, and *external measures*. Furthermore, outside of our SLR, a notable framework developed by [10], which was heavily inspired by H. Nissenbaum's work on contextual integrity [52], dissected an information flow into input, computation, and output, and assessed privacy and verifiability in each step. They also wrapped their framework with flow governance, i.e., the information flow rules upon which participants agree. To provide an improved classification of PETs, we inspired some components of our classification from [10] and [S19] and distilled from the 50 selected papers the set of layers necessary to build a privacy-enhancing IoT data market: *verification*, *storage*, *communication*, *processing*, and *sovereignty*, which Table 1 describes succinctly. Moreover, we also considered important layers necessary for a functional data market that do not require PETs (see Fig. 10).

Furthermore, not all of the technologies included in [S19] enhance privacy, e.g., version control and most distributed ledger technologies. Therefore, unlike in [S19], we have introduced another branch




Fig. 2. Overview of the categories of our classification of the identified technologies among the selected studies. Note that some PETs also enhance authenticity.

for technologies focused on authenticity, which we call authenticity-enhancing technologies (AET). Note that some PETs accomplish data authenticity or integrity while enhancing privacy or confidentiality, e.g., zero-knowledge proofs, homomorphic encryption, or some digital signatures (see terminology in Section 5). Specifically, some PETs can also be AETs, but AETs are not always PETs. Lastly, we classified the identified AETs into the *verification* and *consensus* layers, which are strongly associated with distributed ledger technology, as they coordinate entities and provide verification guarantees. We display our classification framework in Fig. 2. Accordingly, we structure our key results into privacy-enhancing technologies (PETs, Section 6), and authenticity-enhancing technologies (AETs, Section 7). The authors of the selected 41 primary studies jointly employed the technologies included in our classification to create holistic or parts of data market architectures for the IoT; Tables D.5, D.6, and D.7 describe the most salient architectures.

Overall, none of the related work conducted a SLR to create a holistic view of the body of scientific knowledge in privacy-enhancing IoT data markets, which, therefore, indicates the lack of an academically rigorous secondary study in this field [41]. Furthermore, despite the efforts in [S48,S19] and [10], there is not yet a *comprehensive classification and fine-grained analysis of technologies and challenges* that researchers have studied in the context of privacy-enhancing IoT data markets (see Sections Section 6, 7, and 8). Lastly, unlike other secondary studies, we also provide a detailed mapping of technologies, IoT data market layers, and challenges in Table 3.

5. Terminology

To help the reader follow our SLR, we first provide some terminology. These definitions are the distillation of the concepts found in the 50 selected studies and other seminal studies regarding utility and integrity [53], and confidentiality and privacy [54]. When we use the word *assure*, a technology fully guarantees the quality of the data or computation. In contrast, when we use the term *enhance*, a technology improves the quality of the data or computation to some extent. These qualities concern with *authenticity*, *integrity*, *confidentiality*, *privacy*, and *utility*. In line with the definition of privacy of Section 2.1 in the context of computer science, we define these qualities by the absence of an attack against them, if applicable.

Data authenticity is preserved when a malicious entity has not tampered with the truthfulness of the original data; truthfulness covers both *provenance* and *integrity*. In the context of PETs, the degree of authenticity of data can be reduced to enhance privacy. Correspondingly, *identity authenticity* is preserved when a malicious entity has not impersonated another entity. If identity authenticity is assured, then the *provenance* of the data is also assured. In the context of PETs, the identity of an entity can be concealed to enhance privacy. *Data integrity* is preserved if data that have been copied and stored or are in motion are equal to the original [53]. In practical scenarios where data are exchanged, if *integrity* is not preserved, then *data authenticity* is also inherently violated. *Computational integrity* is preserved when, even in the presence of malicious entities, the output of an algorithm that runs on data is

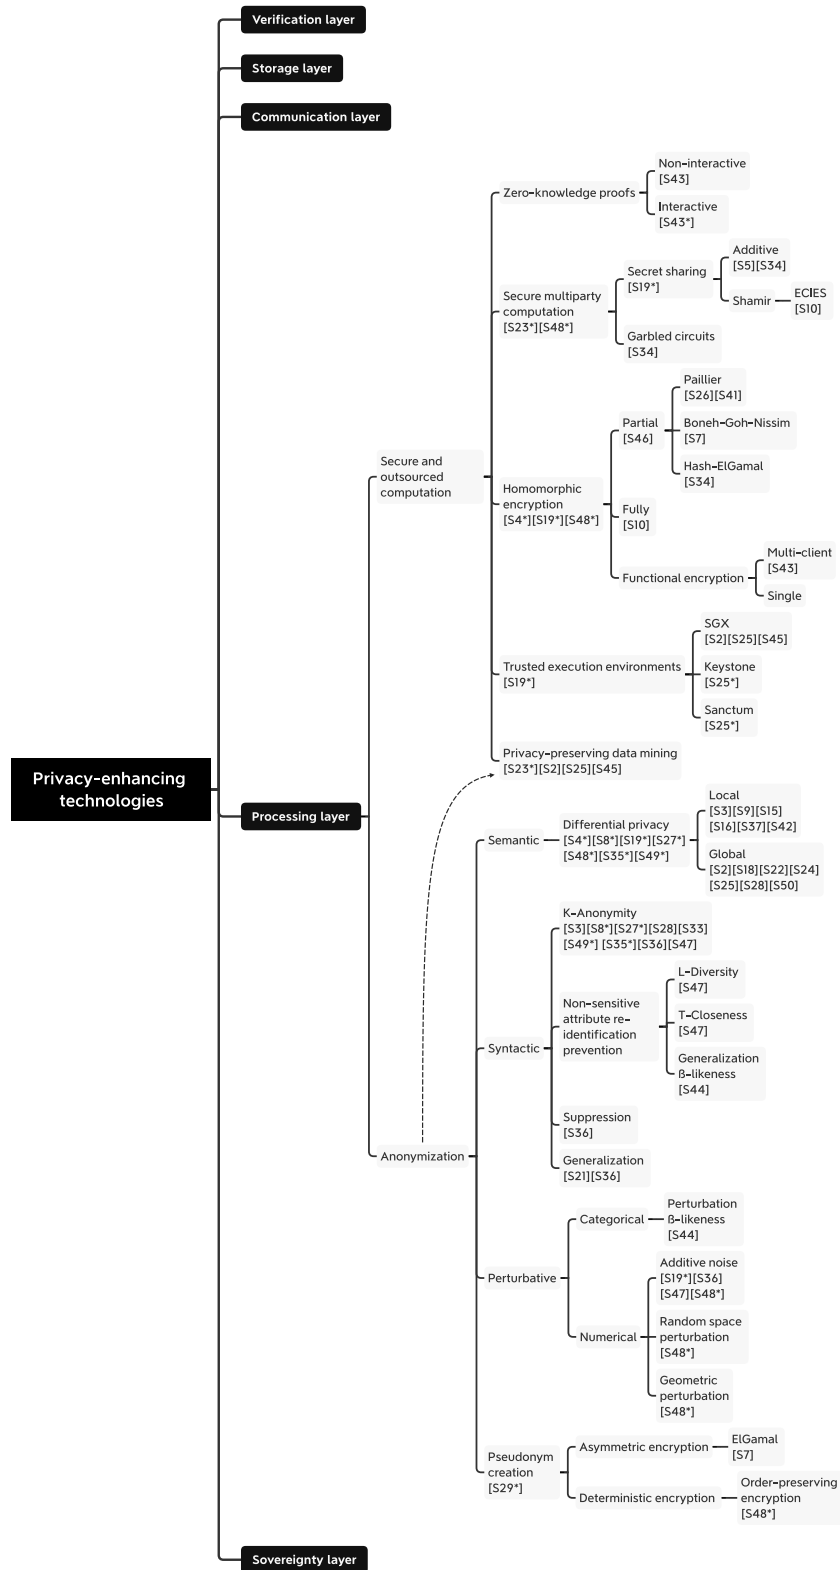


Fig. 3. Classification of PETs employed for data processing. Any other privacy approach encountered in the SLR without explicit inclusion of the underlying technology was either not included in a leaf node but in a parent node or completely dismissed if too vague. *The publication reviews or briefly comments on the technology without delving in-depth or using it as a building block of the architecture concept, e.g., included for future work.

computed correctly. In the context of PETs, the computation can be concealed to enhance privacy. Furthermore, some technologies *enhance confidentiality*, i.e., ensure that data or specific properties thereof are only shared only with the intended parties. Furthermore, we refer to *utility* as a measure of the usefulness of data for the successful

completion of a task; it is high when the data is *authentic*. While PETs reduce *authenticity*, they are helpful in the balancing act between *utility* and *privacy*, as PETs may help to facilitate the sharing of data, which otherwise would not have been revealed (zero utility).

We briefly give some illustrative examples of the interplay of concepts: A digital signature can *assure data integrity* provided the corresponding private key is not accessible to an adversary, and *identity authenticity* if the signature includes a digital certificate that a trusted third party issued; otherwise, digital signatures *cannot assure identity authenticity*. Distributed ledgers can *assure data and computational integrity* by replicated storage and computation [55], but these ledgers cannot enhance *data authenticity*; additionally, replication is often problematic regarding *confidentiality* and, hence, *privacy* [56]. Furthermore, zero-knowledge proofs can provide evidence for *data and identity authenticity and computational integrity* without violating *privacy*, and truth discovery can *enhance* these qualities independent of *privacy* considerations. Moreover, privacy-preserving data mining can *enhance* privacy; however, if the right PETs are not employed, qualities such as *computational integrity* may not be *enhanced* or *assured*. As last examples, technologies such as differential privacy or *k*-anonymization *enhance* privacy by reducing *data authenticity*, and onion routing or ring signatures *enhance* privacy by forgoing or reducing *identity authenticity*, respectively. These latter technologies consequently reduce *data utility* in exchange for *privacy*.

Lastly, we are mindful of the term *tackling*, which refers to a technology that *directly* and fully or partially solves a current challenge in the context of privacy enhancement, e.g., the copy problem or the recursive enforcement problem (REP) (see challenges in Section 8). We use the term *circumvent* when a technology bypasses a problem, i.e., the technology does not directly address the issue. However, the entities that leverage the circumventing technology are still not affected by the problem. For example, obscuring the data and computation in a third-party server with homomorphic encryption (HE) does *not tackle* the REP; instead, HE *circumvents* such problem because the third-party server cannot see the contents. On the other hand, distributed ledger technology *tackles* the REP with a redundant and hence tamper-evident storage and execution.

The following Sections 6 and 7 describe the technologies that we identified in our SLR. We provide a new categorization of these technologies based on the characteristics emphasized in the corresponding selected publications and the technical properties described in this Section.

6. Privacy-enhancing technologies

6.1. Processing layer

The PETs we included in data processing aim to enhance the privacy of either data inputs, outputs, the intermediate steps of a computation, or a combination thereof while maintaining a high degree of utility. This Section follows the structure of Fig. 3.

6.1.1. Secure and outsourced computation

Secure and outsourced computation comprises PETs that enhance privacy through confidentiality. Furthermore, if the PET also employs digital signatures and their cryptography primitives, then the PET can also assure the integrity of the data and computation and identity authenticity in the presence of a digital certificate.

Zero-knowledge proofs (ZKPs). With ZKPs, a technology firstly conceived in the 1980s by [20], a *verifier* can verify the authenticity of the data and the integrity of a computation conducted by a *prover* without the need to access the data or replicate the computation itself [21]. If the statement that is proven is about claims attested in a digital certificate signed by a trusted entity (e.g., age over 18), ZKPs can verify identity authenticity while keeping the information leaked about the identity minimal.

Specifically, ZKPs exhibit (i) zero-knowledgeness, i.e., the *verifier* learns nothing new from the *prover* beyond the correctness of their statement, (ii) completeness, i.e., the prover can convince the verifier of a correct statement with high probability, and (iii) soundness,

i.e., the *prover* cannot convince the *verifier* of a wrong statement with high probability [57,S43]. Furthermore, there are interactive and non-interactive ZKP protocols. With the latter, there is no need to engage in sequential message exchange, and the prover can convince multiple parties of a claim with a single, potentially short, message [57]. These characteristics make non-interactive ZKPs highly attractive for use in blockchains [56]. ZKPs are also the building block of many anonymous credentials, which are also known as privacy-preserving attributed-based credentials [58]. They allow the verification of information in a digital certificate without disclosing any unnecessary data, including the highly correlating value of the signature. Anonymous credentials were initially proposed in 1985 by [59], and developed further with ZKPs and blind signatures [60] chiefly by [61,62] and by [63]. Lately, anonymous credentials have seen renewed interest also in the context of digital wallets for end users' identity management [64,65].

Within our SLR in IoT data markets, [S43] employ non-interactive ZKPs to verify the correct computation of outputs, which, in turn, unlocks the payment from a smart contract in the Agora blockchain, eliminating a third-party verifier. While ZKPs have their limitations due to computational complexity, typically for the prover, and there is still a considerable gap between cryptographers and software engineers [66], we expect to see more publications such as [S43]. This projection is justified by the significant improvements in ZKPs' performance, and ease of use in recent years [67,68] and the availability of an increasing variety of domain-specific programming languages to implement ZKPs, such as *bellman* or *circom* in combination with *snarkjs*. Recently, first research has emerged that uses ZKPs to prove that a machine learning model was trained correctly on specific data [69], and there are many opportunities to leverage them in data markets, such as demonstrating that the input data of a computation was signed by a sensor that received a certificate from a trusted third party without revealing the sensor's identity or the data. In this case, the digital signature and certificate can be regarded as AETs, while their verification inside a ZKP enhances privacy and, hence, qualifies ZKPs as a PET.

Secure multiparty computation (MPC). In broad terms, MPC enables multiple parties to exchange information obliviously and jointly compute a function without revealing individual inputs to each other [19,70]. The MPC implementations that we observed in our SLR employ either secret sharing [S5,S10,S34] or garbled circuits [S34]. In secret-sharing-based MPC, each party first obfuscates the input by splitting it into shares. Secondly, this party distributes the shares among the other computing parties. Afterward, each party executes arithmetic operations independently on these shares, and finally, all parties share the outputs to reconstruct the result [19].

In Shamir's scheme [71], one can specify a minimum of shares that the recipient needs to reconstruct the output, and any combination of fewer shares does not reveal anything about the secret to the receiving entity [S19,187,188,S48]. On the other hand, in additive secret sharing, all the shares are needed. Outside MPC, Shamir's scheme has been commonly used for key management schemata for cryptographic systems so that if some shares that represent a private key are lost, one can still reconstruct the key with the remaining shares [71]. On the other hand, MPC can also be implemented by garbled circuits [19], for only two [72] or multiple [73] parties. Garbled circuits are protocols that enable secure computation by using functions translated into Boolean circuits, i.e., a sequence of basic logic gates such as AND, XOR, and OR that may be combined to construct any function [19,74,S48]. Garbled circuits make use of oblivious transfer [75], which in turn utilizes asymmetric encryption, and symmetric encryption for encrypting and decrypting each gate's truth table. Lastly, there are MPC hybrids that combine these approaches [76].

MPC allows computing functions without revealing the inputs to other participating parties. MPC protects inputs against brute force attacks and it is to date considered less computationally expensive than alternatives such as fully homomorphic encryption [77]. Drawbacks of

MPC include its high processing and communication costs [58] and sensitivity to network latency, which can considerably decrease the performance [S5,S9,S48]. Additionally, MPC protocols often need to be supplemented by mechanisms that prevent collusion [55]. Moreover, since the individually provided inputs are only locally available, one cannot stop malicious entities from jeopardizing the authenticity of the input with false inputs. MPC can only prevent curious entities from learning information. A countermeasure for this reduction in accountability is zero-knowledge proofs to enforce the authenticity of participants' local computations while maintaining them confidential [78].

Three of the papers in our SLR implement MPC in their architecture: [S5] uses additive secret sharing, [S10] employs Shamir's secret sharing, and [S34] leverages a combination of garbled circuits and additive secret sharing. Additionally, other publications acknowledge the importance of MPC schemata by including them in their review [S19,S23,S48]. We provide additional details in Table D.5. While several *frameworks for MPC* are available, the MPC solutions employed by these three publications were handcrafted. This may indicate that the integration of MPC into existing systems requires features that are not yet available with generic tools, such as performance aspects.

Homomorphic encryption (HE). HE allows performing operations on encrypted data (ciphertext) as if they were not encrypted. After the computation, the entities with the corresponding secret key can decrypt the output [189]. There are variations of HE depending on the diversity of operations it can perform [15,16]: Fully homomorphic encryption (FHE) schemata support addition and multiplication, while partially homomorphic encryption (PHE) schemata allow for only one of these alternatives — typically in exchange for drastically improved performance. Any other schema in-between is called somewhat homomorphic encryption [S48].

Five out of the six studies that use HE in our SLR use a PHE variation [S41,S26,S7,S34,S46]. Each of the former four specifies the name of the employed schema, namely the Paillier cryptosystem in the first two [17], Boneh–Goh–Nissim [79], and Hash–ElGamal [80]. The latter study only briefly mentions the additive homomorphic property of their handled data. On the other hand, [S10] uses FHE with a schema called fully homomorphic non-interactive verifiable secret sharing [81]. Several other articles in our SLR underline the importance of HE [S4,S19,S48]. On the other hand, [S43] suggested the use of multi-client functional encryption [82,83] instead of HE so that the scheme combines data from some individuals with others, and, in turn, malicious entities cannot trace back the output of the computation to a single user, as it may happen in HE. Furthermore, there is a related scheme called functional encryption [S43] that allows to retrieve a pre-specified function executed on a set of ciphertexts [82], e.g., decrypt only the mean of a set of encrypted numbers by deriving a function-specific decryption key from the secret key that was used for encrypting the data. A summary of papers from our SLR that mention or use HE is given in Table D.5.

The major limitation of FHE is its high computational complexity and the comparatively large storage needs of its ciphertext, which poses a significant challenge for its use and is aggravated in the context of IoT devices' limitations [S41,S7,S48,S24]. Therefore, the approach adopted by most authors is the use of PHE instead of FHE [84], which, while still not as efficient as other PETs, consumes significantly more computing resources than PHE [85].

As observed for the case of MPC, while there exist generic frameworks such as *SEAL*, *HElib*, or *TFHE*, the authors of the publications in our SLR utilized handcrafted solutions, which may indicate the lack of framework versatility or performance. Overall, practitioners and companies may use HE to perform lightweight functions on data privately on non-local resources, e.g., computing in the cloud, which otherwise would be too expensive to maintain in-house. MPC would usually be preferred over HE when the inputs to the function belong to multiple parties. Nonetheless, some selected publications also employ HE in these cases, e.g., when data brokers determine the winner of an auction [S26,S34,S41].

Trusted execution environments (TEE). TEEs were first defined in 2009 by the Open Mobile Terminal Platform as “hardware and software components providing facilities necessary to support applications” that are secure against attacks that aim to retrieve cryptographic key material or other sensitive information. These features include defense against more sophisticated hardware attacks such as probing external memory [18] or measuring execution times and energy consumption. Moreover, TEEs defend against adversaries who are legitimate owners of the hardware or remote access to the operating system that can run the code themselves. TEEs allow a user to define secure areas of memory (“enclaves”) that enhance confidentiality and assure data and computation integrity of the code and data loaded in the TEE [S2], i.e., any other program outside the enclave cannot act on the data. Specifically, TEEs associate unique encryption keys to computer hardware, making software tampering at least as hard as hardware tampering and certifying the computation results within the TEE. The reason is that the only way to hacking a TEE is physical access to the hardware and, consequently, performing manipulations so that the hardware provides false certifications to bypass remote attestation and sealed storage [S45]. Seal-stored data may not be accessed unless the user employs the correct hardware and software, and remote attestation is a process whereby a trusted third-party assures that the execution of a program in a specific piece of hardware is correct [S45].

Four of the selected papers in our SLR leverage TEEs [S2,S12,S25,S45], and a review mentions their importance [S19]. [S2,S25] and [S45] proposed TEEs to confidentially train and evaluate machine learning models on data available through a data market. While the role of TEEs in data markets overlaps with the use of HE and MPC, authors have preferred the latter technologies to enhance confidentiality in auctions and data processing, which may be due to the limited memory TEEs offered at the time. The reviewed four studies used Intel's Software Guard Extension (SGX) [86], where Intel is the trusted third party, and, therefore, the single point of failure. However, practitioners should be mindful of the numerous vulnerabilities present in TEEs [87–89], and Intel's SGX deprecation in 2022 [90], which affects many of the designs found in this SLR dated before July 2020. Therefore, we suggest practitioners to explore Sanctum [91], Keystone [92] and AWS Nitro [93]. Specifically, [S2,S25] and [S45] used SGX for confidential computing, and [S12], employed SGX for their blockchain architecture to perform “Proof of Useful Work”. In this type of consensus mechanism, nodes perform useful computations instead of computing hashes like in Bitcoin or Ethereum mining. Moreover, [S2] decided to use TEEs to enhance data and computation confidentiality for machine learning algorithms because of the low performance of MPC and HE on machine learning [94].

On the other hand, we noted that TEEs designed for resource-constrained devices – potentially at the cost of offering less functionality – were not prominently discussed in the selected papers. This includes, for instance, ARM TrustZone [95], which is relevant as many IoT devices run on ARM processors, and trusted execution modules [96].

Privacy-preserving data mining (PPDM). [S23] describe PPDM as a means to enhance privacy while extracting useful information from data mining. Data mining includes ML and conventional statistical analyses such as aggregations (e.g., mean or quantiles). PPDM is achieved by performing the computation where the data reside, protecting the computation with cryptographic or data perturbation means, or a combination thereof. As a comprehensive example, suppose the clients' local data and computation are cryptographically protected and the clients have the capability to perturb data. In that case, the computation can run anywhere, which is accomplished by deploying a ML model and input data in a trusted execution environment (TEE) or implementing a ML model using MPC or HE. With input or computation perturbation, the clients also enhance the privacy of the outputs.

A popular tool for PPDM is federated learning (FL) [97–99], as it avoids collecting users' data. Specifically, FL collaboratively trains

a seed ML model across multiple clients' local data, after which a server aggregates the resulting weights to form a unique model (process repeated across rounds). Researchers have increased the privacy of FL by sharing weights with secure aggregation protocols [100] (MPC, Shamir's secret sharing), and protected the privacy in client selection [101] and update parameter sharing [102] with additive HE (i.e., partial HE). Alternatively, split learning approaches [103,104] decompose neural networks' layers into elements and, thus, the input data and labels do not need to be within the same machine. Split learning presents advantages over FL when the local hardware for computations belongs to different network speeds or hardware configurations [105]. Additionally, gossip learning [106,107] proposes a framework whereby multiple models perform a random walk over clients, where they are trained and merged with other models they encounter.

Another way to achieve PPDM is by perturbing input data or weights of the ML model with anonymization techniques such as differential privacy (DP), resulting in privacy-enhancing optimization schemata like DP-stochastic-gradient-descent (DP-SGD) [108]. DP-SGD perturbs the weights' updates with noise and, therefore, one may not reconstruct the inputs based on the outputs, which may happen in ML or stand-alone FL [94]. Practitioners can plug in weight DP perturbation in central ML, FL, gossip learning, or split learning, in combination with MPC as well. We depict such leverage of anonymization technologies for PPDM with the dashed line connecting both elements in Fig. 3.

With PPDM, individuals may enjoy a higher degree of privacy than outsourcing the computation transparently to a trusted third party. Data markets can offer an infrastructure leveraged by PPDM, where data prosumers and consumers only need to provide the input data and ML models, much like the studies in our SLR propose [S2,S25,S45] using TEEs to train models with DP. Like data, trained ML models could also be exchanged in markets.

6.1.2. Anonymization

While the previously presented PETs hide sensitive data from un-solicited parties and, thus, provide confidentiality while enhancing or assuring data and computation integrity, the authorized receiver of the plaintext may reverse engineer the output and correlate data records with individuals (re-identification attack). Consequently, employing only secure and outsourced computation PETs is insufficient to provide the required degree of privacy in cases where the recipient may not be fully trusted. Anonymization technologies can help in these situations by protecting non-explicit identifiers and sensitive attributes [S47,S44]. The cost of this protection is forgoing data authenticity and thus decreasing utility. Given the frequency of re-identification attacks, anonymization should be a critical element of any survey or modern online application, and, in particular, IoT data markets [S8]. One may observe that anonymization technologies rely on statistics, probability theory, and heuristics, while secure and outsourced computation usually employs cryptography and trusted hardware.

This sub-section describes our findings for the most employed anonymization technologies identified in our SLR. We categorize most of them into two groups [190]. *Syntactic* technologies provide a numerical value to the degree of individuals' protection in a dataset, resulting in a perceptible perturbation of data, e.g., generalizing the values 42, 44, and 45 to the interval [40, 45] such that it is harder for an attacker to distinguish between the three individuals. *Semantic* technologies enforce a privacy definition to a learning mechanism executed over a dataset, namely differential privacy, whereby the output distribution of the mechanism should be insensitive to the removal or addition of an individual in the dataset. Typically, the property is fulfilled by adding calibrated noise to the output of a mechanism, yielding a result that is not syntactically different from the original value, e.g., 42 could become 45 after noise addition.

Semantic technologies have an advantage over syntactic technologies, as they provide a mathematical guarantee of privacy agnostic to

background information, i.e., an attacker cannot use related information to re-identify an individual in the dataset. Additionally, we discuss other anonymization technologies not covered in these two groups, namely noise perturbation and pseudonym creation. Perturbation, in this context, is not classified as semantic because its process does not necessarily provide a formal semantic privacy guarantee (such as in differential privacy) and, simultaneously, the outputs are not syntactically modified. In essence, anonymization techniques obfuscate information by perturbing the data during measurement or processing [S8]. From this perspective, anonymization can be understood as a kind of statistical disclosure control [109], and, thus, also encompasses semantic techniques such as differential privacy.

Syntactic technologies. We identify the implementation of the privacy definitions of k -anonymity and its variations l -diversity and t -closeness, a newly proposed model called β -likeness, and also their building-blocks: generalization, and suppression. The most frequently utilized model for syntactic anonymization in our SLR is k -anonymity [S3,S28,S33,S36,S47], which was also reviewed or highlighted by [S8,S27,S49] and [S35]. K -anonymity is a privacy model that guarantees any individual in a dataset to be indistinguishable from at least $k - 1$ others. K -anonymization, i.e., altering a dataset to fulfill k -anonymity, clusters a set of sensitive attribute values into equivalence classes of size k . However, finding an optimal value of k for minimum information loss is NP-hard. Thus, researchers have proposed alternative heuristics [191]. Nonetheless, some of the selected studies used the building blocks of k -anonymization (transformations): generalization [S21,S36] and suppression [S36]. Suppression deletes selected data points, while generalization substitutes data points for others that belong to a higher level in a manually pre-defined hierarchy, e.g., substituting a city by a country to make the location less detailed.

The selected studies [S28,S33] applied k -anonymization to aggregate data from a set of entities. [S47] innovated upon [S28] and [S33] by also employing the l -diversity model to ensure at least l different values in sensitive attributes, and t -closeness so that the distribution of the sensitive attributes within each equivalence class was at most at a distance t from the overall dataset distribution of that attribute. These two models have their own limitations, they prevent homogeneity and external knowledge attacks (l -diversity) and skewness and similarity attacks (t -closeness) [190], to which k -anonymity is vulnerable. Furthermore, [S36] tailored the use of k -anonymity based on record history, privacy policies, and disclosure context. Their new approach prevented a significant decrease in the data utility compared to homogeneously applying k -anonymity to all individuals' records equally.

Nonetheless, there are detractors of syntactic technologies in data markets because of the need for a centralized intermediary that sees and aggregates the data in a, e.g., k -anonymous fashion [S18]. Moreover, [S44] stated that these conventional syntactic approaches are not sufficient because they lack an attacker perspective in the model. For this reason, they designed a novel model called β -likeness that explicitly bounds the additional knowledge that an adversary gains from seeing the released data.

In the context of this SLR, k -anonymization is mainly employed before sharing data in an IoT data market. However, researchers also employ k -anonymity for privacy-enhancing location-based services that exchange location data in IoT data markets, whereby similar fake locations hide the real ones. This type of approach fits well with IoT devices embedded in phones, vehicles, and laptops, among other mobile *things*. Some of the approaches named by [S3] were *cloaking*, which consists of sending a more extensive region that encompasses the real one, and *geomasking*, whereby the real location is randomly displaced outside of an inner circle but within an outer one. [S3] adopted *geomasking* for situations where low accuracy is sufficient, and a high degree of privacy is required. A modern alternative to releasing anonymized data is synthetic data generation, which creates data by randomly sampling

from a distribution representative of the real data. Practitioners can employ generative adversarial networks (GANs) [110], or GANs with differential privacy for a higher protection [111] to synthesize data. Synthetic data could be helpful in some contexts as they “look” similar to the real data (unlike fulfilling k -anonymity), e.g., developing applications before testing them with the real data.

Altogether, anonymization technologies and PPDM compose the building blocks of *statistical disclosure control* [112], which organizations may leverage *internally* in their privacy-preserving data management and analysis solutions, or *externally*, by using privacy-enhancing publishing solutions [113] in, e.g., data markets. Among past surveys focused on the latter solutions (namely k -anonymity and related models, in addition to a few cryptographic primitives), the reader may refer to [113–116] and [117] for further specialized reading. Most notably, [115] provide a comprehensive survey of syntactic models and differential privacy and their *attacks*, and [117] compile a helpful *mapping* of data types to the appropriate syntactic and semantic techniques for anonymization.

Semantic technologies. Introduced by [12] in 2006, differential privacy (DP) proposes a formal guarantee of privacy that has become the golden standard for researchers [S37]. DP appears in its pure form or one of its flavors in 13 of the 35 studies that propose a solution in our SLR. Furthermore, another seven studies refer to DP to underline its importance or drawbacks. The potential reasons behind the high number of references and use of DP are multifaceted. While HE or MPC may protect the inputs’ and computations’ confidentiality, they do not prevent reverse-engineering the outputs (re-identification attacks). Moreover, syntactic technologies or other conventional anonymization technologies, e.g., additive noise, lack a mathematical guarantee of privacy and are subject to background knowledge attacks. DP, however, tackles these issues.

In broad terms, DP guarantees that the output distribution of an analysis (a statistical query or a ML model) over a dataset is “essentially” identical, irrespective of the presence or absence of an individual in the dataset. Additionally, DP is agnostic to auxiliary information available in the present or the future. DP is *typically* achieved by adding random noise sampled from a probability distribution such as the Laplacian or the Gaussian. Specifically, the noise limits the output distribution *difference* of an analysis executed over two datasets (one *with* and one *without* an individual) to be no greater than an upper bound, making the outputs “differentially” indistinguishable. Overall, DP bounds the amount of new information gained by an attacker after observing the output of an analysis.

The set of selected studies of our SLR that used DP in their proposed solutions are [S2,S3,S9,S15,S16,S18,S22,S24,S25,S28,S37,S42,S50]. Tables D.6 and D.7 summarize their proposed architectures. Six of these studies employ DP locally [S3,S9,S15,S16,S37,S42], i.e., the noise is added to the data of an individual on the client-side. In contrast, the rest of the studies apply DP centrally, i.e., on aggregated data on the server-side. Furthermore, we can cluster the studies into those that focus on a data trading design for data markets [S22,S37,S15,S50], crowdsensing data markets [S42,S9,S16,S24], and architectures that host a data market in an attempt to achieve end-to-end privacy [S2,S3,S18,S25,S28].

While DP offers a mathematical privacy guarantee, DP is not a panacea. DP still holds flaws in its real-world implementations [118] that the research community and practitioners should address. Moreover, DP’s combination with ML needs further improvements regarding balancing privacy and accuracy [109]. In our SLR, [S48] and [S36] identify two specific problems with DP: firstly, DP cannot be used when a high level of accuracy is required [S48], e.g., analyzing data from the brakes of vehicles to improve safety. Secondly, releasing an entire dataset with current DP approaches is troublesome. Despite these challenges, the authors of [S2] and [S50] argue that the benefits of DP predominate, as DP can adapt to many use-cases and allows a practitioner to fine-tune the added noise to enhance privacy.

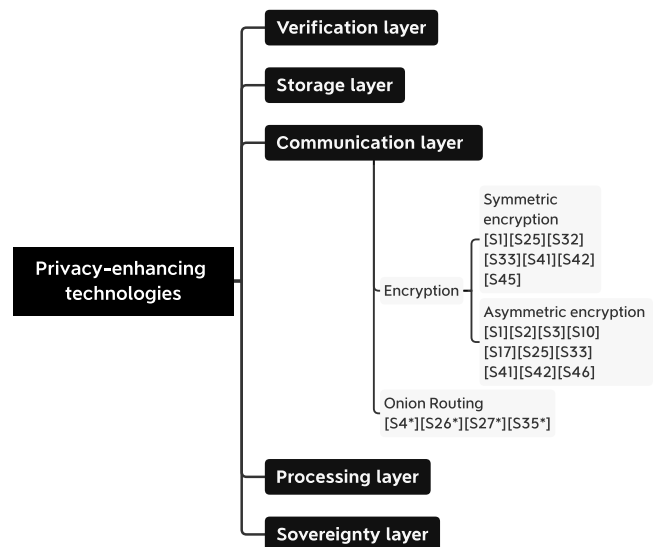


Fig. 4. Classification of PETs employed for communication.

As we already noted with ZKP, MPC and HE, the authors of the selected publications that used DP did not employ open-source DP libraries such as *OpenDP*, *Google-DP*, *diffprivlib*, *TensorFlowprivacy*, or *Chorus*. Instead, they used handcrafted implementations of DP. Aside from syntactic and semantic technologies, other anonymization technologies are simpler to implement, e.g., sampled data release, character masking, truncation, rounding, top and bottom coding, data swapping, randomization, creating pseudonyms, character scrambling, microaggregation, or noise perturbation [22,119]. Moreover, [120] designed an algorithm that combines the syntactic definition of k -anonymity with DP. The two other anonymization technologies employed by another three studies were noise perturbation [S36,S47] and pseudonym creation [S7].

Perturbative anonymization. Perturbation relies on the use of noise to obfuscate sensitive information. One of the simplest forms of perturbation is additive noise, employed in [S36] and [S47]. Additive noise consists of adding to a deterministic value a random value sampled from a uniform distribution whose bounds are set by a specific percentage of the deterministic value. Furthermore, [S48] reviews two novel perturbative technologies: First, random space perturbation [121], which strives to protect the privacy of cloud-stored data by utilizing a confluence of order-preserving encryption, dimensionality expansion, random noise injection, and projection. Second, geometric perturbation [122], which is motivated by the idea of protecting the geometric transformations that a machine learning model may perform on a dataset rather than the data itself. While perturbative technologies aim to tackle the same problems, unlike DP, they do not provide mathematical guarantees of privacy, even though some are also based on noise addition.

Pseudonym creation. Pseudonym creation is applied to direct identifiers, e.g., names or social security numbers, to enhance privacy while uniquely identifying each record. Practitioners create pseudonyms by hashing or deterministically encrypting an identifier, e.g., using order-preserving encryption [S48], or by applying asymmetric key encryption like ElGamal [S7]. However, researchers have demonstrated that pseudo-anonymization falters against some attacks like profiling, task tracing, or re-identification [S35].

6.2. Communication layer

The PETs included in this Section enhance the confidentiality of data in transit or of the sender’s identity (see Fig. 4). These PETs rely on cryptography.

Encryption. Encryption is one of the most fundamental technologies to enhance confidentiality [S19] because after encrypting a piece of data (cipher), only the anointed holders of a decryption key can decipher such data. We underline that encryption cannot guarantee privacy because nothing stops an intended receiver from publicly sharing the decrypted message; this also emphasizes employing anonymization PETs. Encryption may be symmetric (one key to both encrypt and decrypt data) or asymmetric, known as public-key cryptography (two keys, a public key to encrypt and a private key to decrypt, or vice versa). Encryption is the building block of virtually every secure communication established through a network and takes a key role in digital signatures.

While some publications from our SLR employed asymmetric encryption for the confidential communication of data [S1,S2,S3,S10,S17,S25,S33,S41,S42,S46] (most of these publications also employed asymmetric encryption for digital signatures, hence the high frequency of digital signatures in Fig. G.14), other publications such as [S1,S25,S32,S33,S41] and [S45] employed symmetric encryption to also confidentially store data. Naturally, encryption is also a building block for digital signatures (verification layer) and for the storage layer to maintain data at rest confidential. We depict this relationship with the dashed lines connecting these elements in Fig. 5.

Onion routing. Onion routing, the backbone of the P2P network resulting from the Tor project [123], consists of a series of re-transmission steps through the network's nodes. A sender's message is encrypted once for each step. The intermediaries decrypt only their appointed encryption layer. Thus, the node only knows the immediate sender and receiver but not the origin of the chain of messages. As the messages are encrypted, the nodes cannot see the contents either. Overall, onion routing renders one's messages unreadable and untraceable. The paper that suggests employing onion routing in a data market context is [S26], which some of the identified reviews equally appreciate [S4,S27,S35].

However, some drawbacks exist. Implementations backed by Tor have high-latency and redundant communication that challenges bandwidth, which can be hard to align with high transactional environments, such as IoT data markets. Moreover, if an architecture decides to use Tor, the network is often blocked by IT departments within organizations or even subject to state-level censorship by some governments [123]. Therefore, practitioners can use alternative technologies such as a VPN to enhance entities' privacy in a network in these contexts. However, these typically centralized alternatives generally offer lower anonymity guarantees, e.g., a VPN provider can identify a user [58].

Because there is no central authority to set privacy policies unilaterally, one must remember that onion routing enhances privacy by preventing malicious entities from collecting IP addresses to identify users. Onion routing will not help if the data that users submit to the network is intrinsically sensitive or correlating. To tackle these limitations, practitioners may use onion routing as a building block of a more extensive privacy-enhancing system that leverages other PETs [S26].

Notably, the publications surveyed in our SLR did not include many other untraceability protocols, which would include mixnet-based alternatives to onion routing (e.g., anonymous remailers, Chaum's mixes [124]), DC-nets [125], or peer-to-peer anonymous communication systems. An extensive overview of these systems is given by [126].

6.3. Storage layer

The authors of the selected papers that propose confidential storage functionality in their architecture leverage symmetric encryption, mostly AES [S1,S25,S32,S33,S41,S45] (encryption is described in the *communication layer*). Furthermore, researchers could leverage InterPlanetary File Systems (IPFS) [127] to compensate for the lack of storage capacity in blockchains to some extent. Specifically, IPFS is a

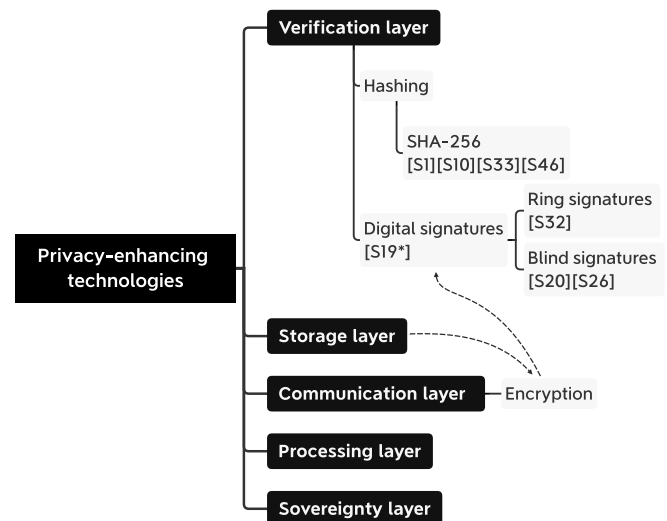


Fig. 5. Classification of PETs employed for verification.

peer-to-peer protocol for data storage and access in a distributed file system. Among the PETs in the processing layer, practitioners could employ homomorphic encryption [128] to encrypt and store certain types of data, so that data are readily available to compute confidential operations. Furthermore, unless strictly necessary, practitioners should store encrypted data that are, in turn, *anonymized* with syntactic or semantic technologies. In case of a breach that leaks the decryption key, anonymized data would reduce the likelihood of attackers re-identifying individuals.

6.4. Verification layer

Some of the PETs that support data processing cannot verify the authenticity of data, identities, or the integrity of data [S19] by themselves. The PETs we include in this Section accomplish these verifications with different levels of privacy enhancement. The data processing PETs that can assure identity authenticity and data integrity use the digital signatures of the *verification layer* and the encryption technologies of the *communication layer* as building blocks. Furthermore, the credibility associated with verifying the information exchanged, analysis outputs, and identities can increase the willingness of users to share data [S8]. To navigate this Section, we refer to Fig. 5.

Privacy-enhancing digital signatures (DSs). DS schemata assure data integrity and identity authenticity if accompanied by a digital certificate. As a consequence, DSs also provide non-repudiation [S1], i.e., actions that an entity cannot deny later. The steps that usually constitute a DS scheme are private and public key generation, encrypting a digest of data with a private key, and a signature verifier that employs the public key to check whether the sender signed the data with the private key.

DSs and the encryption primitives of the communication layer are so fundamental that one of the selected studies solely relies on HTTPS for their data market architecture [S17]. However, this architecture does not consider privacy beyond data in transit. Hence, most selected studies rely on multiple PETs. Moreover, although not all of the selected studies explicitly mention DSs, we can safely assume that since DSs are already a living part of virtually any enterprise IT system, most selected studies employ them in their architectures (hence the high frequency of DS utilization in Fig. G.14). Nonetheless, while DSs allow verifying the integrity of data or the authentic identity of the sender, users still need to trust the sender with the authenticity of the data.

So far, we have only described DS as an authenticity-enhancing technology. However, some of the studies selected in this SLR employed

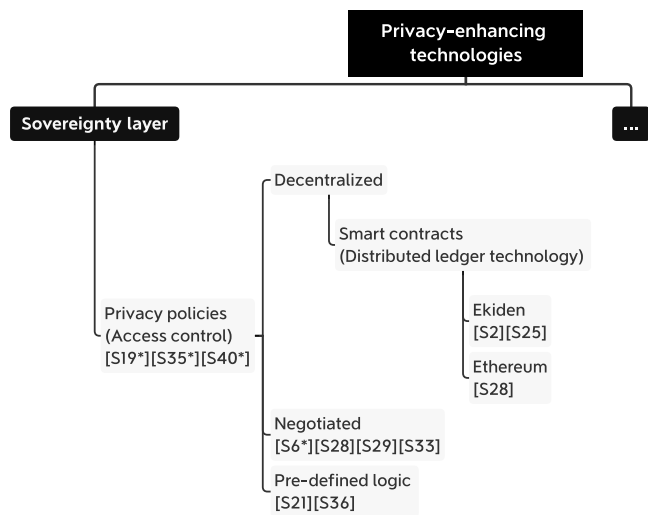


Fig. 6. Classification of PETs used for sovereignty purposes.

two DS schemata based on asymmetric encryption primitives that make DSs privacy-enhancing:

- **Ring signatures** [S32], whereby any party within a pre-defined set of parties could have been the signer of a message. Thus, the identity of the authentic signer is kept hidden [192,193].
- **Blind signatures** [S20,S26], whereby the signer does not have access to the content being signed [60]. It is possible to use blind signatures in combination with zero-knowledge proofs to convince the signer that the content to be signed has the expected properties. Also, one can make an entity sign multiple contents and allow for spot checks to detect fraud. The latter procedure has been employed in the first approaches toward privacy-enhancing payments [59].

Hashing. Hashing is a tool to deterministically map data of an arbitrary length to a fixed output length. In the context of privacy and verification, and aligned explicitly with some of the selected studies [S1,S10,S33,S46], hashing is used to verify the integrity of transferred data by hashing the data and making the hash public before transferring the data. In this manner, the recipient can verify the integrity of the confidentially transferred data by comparing the hash of the received data with the previously published hash. Provided the entropy of the data is sufficiently high, nobody except the intended recipient can determine the data from the published hashed value. Hence, hashing can be considered a form of version control with a privacy component.

The hash function employed by the publications mentioned above was SHA-256. Their authors commonly use the published hashed data on distributed ledger technologies to ensure immutability and availability. In this setting, hashing enhances the confidentiality of the sender's data while the parties (network nodes) ensuring the ledger's integrity (and inherently the persisted hash) cannot unveil the original data. The original data is only viewed by the intended receiver, which validates the integrity of the data received through another channel with the hash persisted in the ledger.

6.5. Sovereignty layer

The *sovereignty layer* deals with the concept of information control, the perceived ability to govern what is exposed from one's data [54]. Specifically, based on an entity's requirements, this layer defines the entity's rules and guidelines regarding ownership and management that can indirectly govern data processing, verification, and other IoT data market layers. Furthermore, to prevent entities' violation of privacy,

practitioners should map these rules to the PETs capable of fulfilling them. For example, GAIA-X's high-level architecture contemplates privacy policies in their data *sovereignty layer* [5]. Privacy policies (closely related to access controls), have been predominant across publications, and, thus, dominate the sovereignty layer depicted in Fig. 6, which illustrates the three identified types.

Privacy policies and privacy by design. Privacy policies embody the requirements and guidelines of a data governance model and are meant to be part of any privacy-enhancing application. To define them, given the regulatory and human aspects of privacy policies, it is also helpful to adopt perspectives from definitions beyond computer science, such as the one underlined in Section 2. [129] indicates that the privacy requirements depend on the recipient of the information, e.g., an individual can have different reservations when disclosing information to a family member than to the government. Privacy policies should reflect this definition, which means that individuals should express the privacy policies they expect. Several studies in our SLR explicitly proposed policies as part of their solution [S6,S21,S25,S28,S29,S33,S36], while many others reviewed privacy policies [S19,S35,S40] or mentioned similar ideas. For example, [S8] did not provide a concrete implementation or explicitly named privacy policies. However, they mentioned that, in data sharing scenarios, the data owners should be able to control some fundamental aspects: data types to share, with whom to share, the required degree of trust in another party, the purpose of sharing, and for which benefit.

Among the publications discussing privacy policies, there is a discernible classification. Four of these publications [S6,S28,S29,S33] considered privacy policies as a negotiation between the user and a third party, such as the data consumer or data broker. [S6] provided a set of legal requirements and high-level technical solutions that facilitated the introduction of policies in international data markets, e.g., writing policies in a standard language. On the other hand, [S29] presented an approach where the data owner could choose among a set of four privacy policies, which included how data was aggregated, and [S33] relied on a *privacy policy manager* that acted as a gatekeeper and managed the privacy settings from a set of users. Furthermore, another set employs a pre-defined logic to execute PETs based on the desires and track record of the data shared by the individual [S21,S36], while the last set relied on smart contracts for decentralized pre-defined [S2,S25] or negotiated [S28] policies.

Nonetheless, the implementation of policies faces challenges. Firstly, there may be multiple colliding policies, i.e., applications must prioritize policies depending on the context [S21]. Secondly, there is no uniformly accepted global standard for electronic privacy policies [S6]. [S40] investigate how practitioners model privacy policies in different domains, focusing on IoT applications. They point to the lack of a uniform standard and propose to utilize ontology-based privacy-knowledge modeling. Thirdly, policy enforcement also causes overhead and an increase in latency due to the need for compliance checks and a lack of automation [S21]. Furthermore, conventional users should include their privacy preferences with minimal manual effort, as they could be overwhelmed otherwise. [S40] suggested using recommender systems based on similar users' data to address this issue. However, this solution may incur a biased recommendation. Moreover, data acquisition expenditure for privacy policies should not incur costly computational resources as they scale to a growing number of transactions [S40].

Privacy policies are crucial to protect users' privacy; however, they are not enough. Organizations must consider privacy issues at each stage of the data pipeline (i.e., processing data end-to-end with the extract-transform-load framework [130]), contemplating aspects that escape user-defined or mutually-agreed policies, and taking into account that typically, neither users nor data brokers will be privacy experts. If a user does not know the potential harms of sharing sensitive information such as DNA data, a data consumer may take advantage

of the user. Therefore, while privacy policies are a stepping stone toward end-to-end privacy, practitioners must develop systems with a privacy-by-design philosophy [S36,S4].

Privacy by design is a term coined in the '90s by the former information and privacy commissioner for the Canadian province of Ontario, A. Cavoukian, who created seven principles [131]. Privacy by design claims that privacy goes beyond current regulations and must be an ever-present concern in the minds of organizations [131]. Following privacy-by-design principles entails, for example, preventing sensitive information extraction by default [S36], minimizing the amount of shared data at each exchange (proportionality) [194], and increasing the price of large data packages [194], among others. However, adopting these design principles comes with effort, forcing developers to adapt their design patterns. For example, current homomorphic encryption techniques force data scientists to express their analysis in terms of additions and multiplications, and differential privacy requires new software engineering design patterns that track the privacy budget of individuals or data scientists.

Smart contracts (SCs). A SC alone is mainly equivalent to conventional scripts. Nevertheless, because SCs are executed in distributed-ledger-technology-based architectures (DLT) (see Section 7.1), SCs inherit from DLT their enhanced availability and integrity guarantees [55]. DLTs execute SCs synchronously on every node of a P2P network if an arbitrary transaction demands a function's execution. Once deployed, no one can change the script, not even the creators (unless there is an intended call of the script that enables modification), and the script will remain in the network as long as the network exists unless specified differently (e.g., through a self-destruct call). This inherited integrity property of SCs makes them a unique tool to specify and enforce policies between parties or any other process where no trusted third party is available.

Within this review, all the studies that used the Ethereum, Quorum, Hyperledger Iroha, or Ekiden blockchains relied on SCs to declare privacy policies [S2,S25] (Ekiden) [S28] (Ethereum), fair auctions [S3] (Hyperledger Iroha), or payments or incentives [S32] (Ethereum) [S43] (Agora) [S3] (Hyperledger Iroha). However, while SCs ease verification and enable democratic proposals of privacy policies, SCs also inherit the privacy flaws of DLT, i.e., SCs by default imply the disclosure of data and computations to all DLT network nodes [56,132]. For example, the architecture from [S25] employed SCs to set user-defined policies, yet it relied on trusted execution environments to enforce them. *SCs alone cannot enforce privacy policies without relying on other PETs.* The only privacy-related feature that a SC can offer to an IoT data market is declaring privacy policies.

Data access control. Data access control refers to allowing an organization or an individual to choose *who* has access to *which* data. Access control represents a subset of privacy policies in data markets and may utilize different PETs to enforce access rights. While access control is a long-established approach, [S25] propose a novel method, using a key-rotation system [195] in combination with a key manager. Thus, the potential impact of a leaked key is only temporal, with the downside of shorter access permissions.

7. Authenticity-enhancing technologies

The included authenticity-enhancing technologies (AET) focus on enhancing the authenticity of data and identities and also cover data integrity as described in Section 5. Some of the AETs that we describe incorporate privacy-enhancing features, while others do not address or even aggravate privacy protection issues and, thus, need to be combined with PETs.

7.1. Consensus layer

Distributed ledger technology (DLT). While DLT may take different forms, most architectures follow the blockchain design pattern, except for IOTA, which uses the so-called Tangle [196]. A blockchain is a tamper-proof distributed database whose state is stored, synchronized, and replicated by nodes in a P2P network following a consensus algorithm [55]. By its distributed nature, the shared ledger becomes a medium to verify claims, data, payments, or contracts, as once an entity writes something on the ledger, it is practically impossible to modify or erase this record in the future. This property makes blockchain a decentralized and highly reliable alternative to conventional auditing methods like version control [S19].

Benefits of DLT in IoT data markets are the ability to represent the governance, distribution, and roles of authorities on a technical basis [S3], and the enforcement or transparent storage of pre-defined rules by the architects of the respective platform [S25]. Other benefits include eliminating the need for a trusted third party, which removes a single point of failure, improves censorship resistance, and provides more robust data and computational integrity guarantees. DLTs also enable payments through their often built-in cryptocurrencies or other payment systems implemented via smart contracts [S32,S46].

However, some of the studies in our SLR also point at the challenges of current DLT designs: IOTA fails to deliver regarding throughput [S28], is still centralized [S20], and provably has security flaws [133]. Furthermore, blockchains exhibit low transaction throughput [S25], high latency [S11], limited storage [S1] and scalability [S11,S25], computational overhead [S25], high energy consumption [S12], and, most importantly, excessive information exposure that can entail a privacy violation [134]. However, some of these aspects can be mitigated. For example, the energy consumption issue only concerns proof-of-work blockchains [135], and performance can be improved to some extent by private permissioned blockchains that restrict participation in consensus and read access to a small number of nodes in a consortium [136].

Despite the possible operational improvements, employing a DLT for a privacy-enhancing IoT data market needs in-depth consideration. Firstly, through highly replicated storage, a DLT is not suitable for storing large amounts of data produced by IoT devices, not even in a privacy-compliant manner. Consequently, most architectures of the selected studies transfer data through interplanetary file systems [S28], employ a hashing verification approach as described in the *communication layer* [S1,S10,S33] or use Merkle trees [S46]. Secondly, while DLT allows for disintermediation and verification in a trust-less manner, it exposes to the network whatever information someone writes on the ledger for as long as the network exists, which may, among others, violate GDPR's Article 17 "*Right to be forgotten*" for personally identifiable information [137]. Lastly, even if an organization uses a DLT only for the matching and clearing steps of an auction, potentially sensitive business information such as turnover can become available to other network participants, which can conflict with antitrust regulation.

Despite the privacy and performance issues of DLT, 31% of the included papers implemented a DLT as the backbone of IoT data market architectures, employing the Ethereum blockchain [S1,S13,S20,S28,S31,S32,S45,S46], Quorum [S18], the Agora blockchain [S43], Hyperledger Iroha [S3], Hyperledger Fabric [S10,S33], IOTA [S20,S28,S30], Intel's TEE-based consensus Rem [S12], and Ekiden [S2,S25]. Other publications only considered them agnostically [S11,S49] or in a review [S19,S23]. The most salient architectures are described in Table D.7.

Some of the selected studies included privacy-enhancing features in their stack. For example, Quorum supports private transactions and private contracts through a public-private state separation and P2P encrypted message exchange for the direct transfer of private data [S18]. However, the interaction between the private and public ledgers is thus naturally limited and cannot be directly applied, for example, to an on-chain payment system. Another example is Ekiden, which offers a

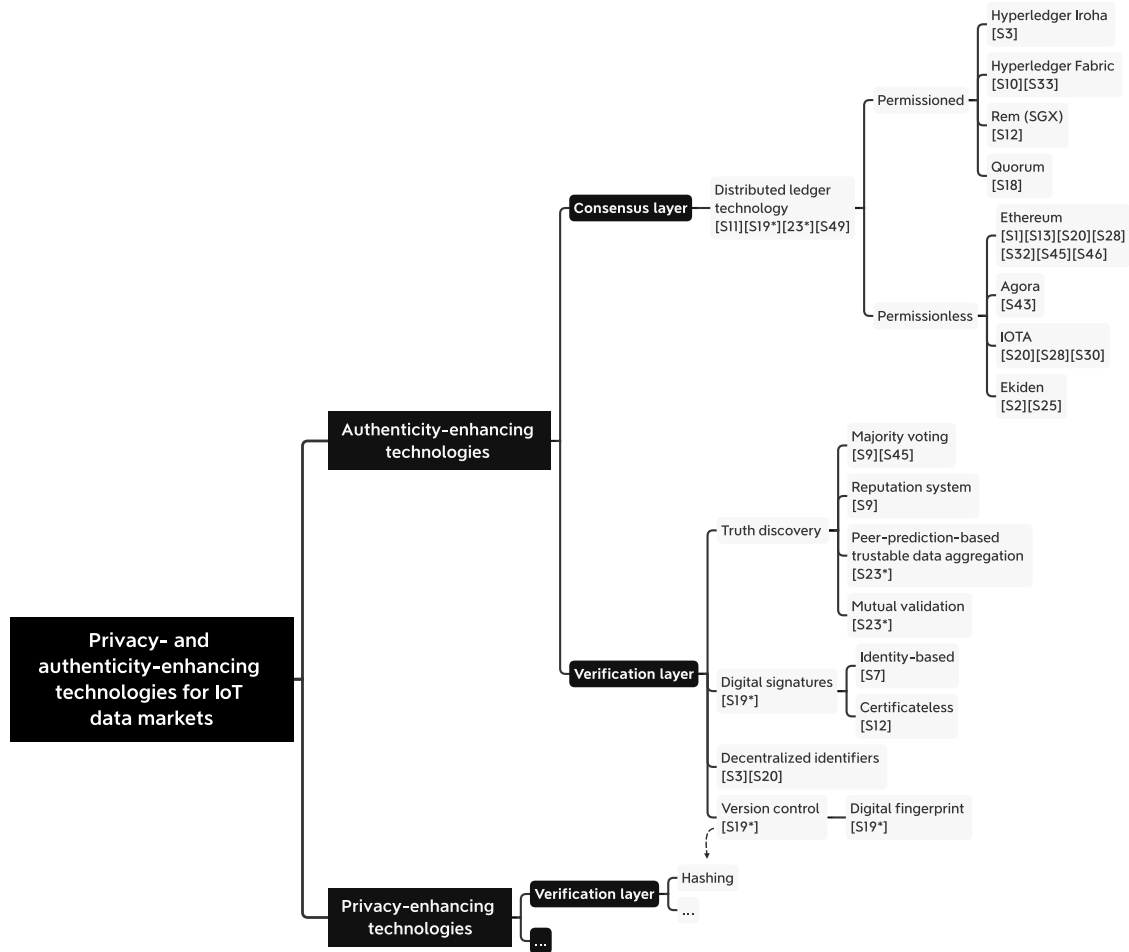


Fig. 7. Classification of the identified authenticity-enhancing technologies in the selected studies of this SLR.

horizontally scalable blockchain potentially capable of hosting end-to-end privacy-enhancing applications through key management protocols and Intel's TEEs [S25] (Note that [S12] uses these TEEs only for consensus, not privacy). Like the solution that [S45] presented, Ekiden allows for smart contracts to execute data analysis in TEEs. However, it is essential to note that these DLTs accomplish the described privacy and integrity functionalities not because of the DLT characteristics but by leveraging the PETs described throughout Section 6.

7.2. Verification layer

This verification layer corresponds to AETs that can be employed for the verification of data and identities. We structure this Section according to Fig. 7.

Truth discovery (TD). TD encompasses algorithms aiming to find the authentic value when different data sources provide conflicting information. As a consequence, TD enhances data and computation integrity and can also enhance identity authenticity, e.g., through reputation systems [S9]. In our SLR, we found that TD takes different forms. For example, the survey by [S23] mentioned a mechanism called peer-prediction-based trustable data aggregation [197], in which the system administrator rewards participants for predicting outcomes of arbitrary events based on other participants' data. This design created incentives for honest reports and therefore enhanced data correctness, resulting in almost all participants choosing to report their bids truthfully [S23]. Moreover, [S23] also proposed mutual validation in which an IoT device compares its data with that of other nearby IoT devices. However, this only applies to specific measurements that are positively correlated

for neighboring devices, e.g., temperature, speed of a vehicle, or location in particular settings. It also seems challenging to establish generic handling of differences. Other TD approaches are majority voting, implemented by [S9] in their crowdsourcing architecture and by [S45] in their data processing-as-a-service model. Specifically, given the use of differential privacy in the former approach, they systematically discovered high-quality data with an *estimated* measure of utility that compares individual data points with an aggregate (the "majority"). The latter publication created a reputation system based on the quality of previously sold data.

While most TD approaches leverage transparency to enhance data and identity authenticity and data and computation integrity, TDs are flexible to include PETs such as ZKPs, MPC, HE, TEEs, and DP such as in [S9]. Furthermore, TD can also tackle the oracle problem of DLT, i.e., nodes within the network cannot assure the authenticity of data from outside the network, e.g., the price of a physical asset or the result of an election. For example, ChainLink [138] is an initiative that utilizes incentives to create a trusted oracle network and incorporates many of the principles of TD.

Digital signatures (DS). While DS¹ schemata are commonplace for authentication purposes in today's IT architectures, we have found in selected studies the use of two notable public key cryptography (PKC) schemata that offer some convenience-related advantages over conventional PKC systems:

¹ We introduced the fundamentals of DSs in the *verification layer* within the PETs branch.

- **Identity-based** [S7], where a key generation center (KGC) creates a secret key in a way that the entity's public key can be a publicly available unique string, e.g., the entity's email address. The KGC must be trusted because it holds the master secret key from which all parties' secret keys can be derived. This digital signature assures identity authenticity as the signature is digitally certified by the KGC.
- **Certificateless** [S12] DS schemata are a special form of identity-based PKC whereby an entity's private key is generated by both the entity and a KGC so that the KGC is not aware of the private key of the entity. However, the entity can prove that the KGC was involved in the key generation [139]. This approach assures the authenticity of an entity while tackling the single point of failure of the KGC.

Decentralized identifiers (DIDs). Identifiers can link an entity electronically across multiple IT systems, such as mobile phone numbers, ID cards, user names, or emails. These links are sometimes but not always unique and are facilitated by identity providers that centrally host registries of these identifiers' [140]. In contrast, DIDs are globally unique (with certainty through publishing them on a DLT or probabilistically through randomized generation) identifiers decoupled from centralized registries. DIDs essentially correspond to URLs linked to a file containing one or several public keys and associated metadata that specifies the policies of controlling or interacting with the associated identity. There are two studies in our SLR that employed DIDs in combination with DLT in their conceptual frameworks [S3,S20], described in Table D.7.

Digital fingerprints (DF). DFs are unique physical identifiers that can be attached to or are inherent of items, and thus, one can be sure to interact with, e.g., the right IoT device [S19]. DFs can be seen as a form of version control at a high level. However, attaching an identifier securely to a physical object is difficult unless it has a unique property, e.g., unique metal patterns in the soldering of a chip. However, even if the attachment is relatively tamper-proof, e.g., with a crypto-chip, the same problem also pervades the items that interact with the digital fingerprinted item, e.g., tracking scanners. Therefore, despite the authenticity assurance of DFs, their authentication can only be as truthful as the honesty of the devices that scan the DF.

8. Privacy challenges in IoT data markets

This Section aims to answer RQ2 by distilling the implicit and explicit challenges unveiled in our SLR and other seminal studies [10,24] concerning privacy in the context of IoT data markets. We further classify them into narrow and broad challenges depending on the scope of their definition. Fig. 8 summarizes and outlines the structure of this Section.

8.1. Narrow challenges

Aside from the inherent complexity and low maturity of some PETs and the compatibility issues with legacy systems [24], we identified another specific set of challenges tackled or circumvented by the selected studies.

8.1.1. The trade-off between utility and privacy

Practitioners working with personal data face the challenge of balancing the enhancement of individuals' privacy with the preservation of data's utility [S14]. This challenge is explicitly mentioned by some of the selected studies [S6,S9,S13] and implicitly tackled by others [S2,S24,S25,S35,S47]. This dichotomy is the underlying reason behind the tension between data owners and consumers: the former aim to maximize privacy while the latter intends to maximize utility, which, in turn, is frequently determined by data authenticity (see terminology in Section 5). Furthermore, privacy officers should consider balancing this

trade-off at each stage of an information flow [10]: input, computation, output, in transit, and at rest, which increases the complexity of the task. On the other hand, decision makers' or data scientists' quality of judgment depends on computational integrity and the authenticity of data and identities, which is affected by the privacy-utility trade-off.

While PETs from the secure and outsourced computation category seem to circumvent the utility-privacy trade-off by concealing inputs, computation, and outputs, the anointed recipients of these outputs can still perform a re-identification attack. Thus, anonymization PETs, such as differential privacy, should also be included in the stack as they lower the probability of successful re-identification attacks [10].

Data and identity authenticity and accountability bring another problem in the utility-privacy trade-off. Some PETs, namely anonymization technologies, increase plausible deniability at the expense of reducing authenticity and, therefore, accountability [S14,S43]. If the data is fuzzy, the data owner may claim that such a result is not resembling the truth, which is favorable for individual users. However, such protection is not beneficial for society in some contexts, e.g., in criminal contexts. Regarding authentic data from fuzzy identities, if an authority cannot trace data back to the origin, an individual could try to claim plausible deniability, which would hinder processes such as tracking COVID-19 patients to improve pandemic countermeasures. For practitioners to find a balance in these contexts, the requirements of any application or platform should first define the accountability of the involved entities to strike an optimal balance between utility and privacy.

Regarding accountability in data markets specifically, mechanisms to punish misbehavior, such as banning an entity for re-selling or not selling authentic data [S7], can be beneficial to enhance the utility of the market. While AETs such as truth discovery, e.g., majority voting [S9,S45] or reputation systems [S9], incentivize market participants to report honestly about the exchanged data, their identity, and computation integrity, among others, the privacy of the entities is not necessarily enhanced. Additional related issues that may arise are verifying the purchased data's authenticity without violating the individuals' privacy [S7]. For example, if a data broker sells analysis outputs (insights) and not the (privacy-enhanced) original data, the original data owner's digital signature is not valid to authenticate the insights [S7]. Nevertheless, systems could use zero-knowledge-proof-based authentication of data and computation, coupled with value deposits locked in smart contracts to hold participants accountable through enforcing reimbursement across intermediaries. Moreover, other nascent solutions exploit the ubiquity and proximity of IoT devices in specific contexts because the data gathered is likely to be correlated, which allows for mutual data authenticity verification among IoT devices [S23]. However, exploiting correlation can only be used for specific measurements, e.g., weather conditions, vehicle speed, and location, among others.

8.1.2. The recursive enforcement problem

Trust is an essential component in distributed systems that involve different stakeholders and, thus, trust is specifically relevant in the context of IoT data markets. Definitions of trust generally refer to a "[...] directional relationship between two entities" [S49] where one entity (the trustor) has subjective expectations on the behavior of the other entity (the trustee) based on previously observed behavior (reputation-based trust) [141] or the belief in competencies and corresponding actions — often within a specific context [142,143] and incentivized by joint interests [144]. Following this definition, we can consider users as trustors of application owners protecting their data when they engage in digital activity. However, the number of data breaches [9] and privacy scandals such as Cambridge Analytica indicate that this trust is not always deserved. The recursive enforcement problem (REP) encompasses the underlying problem of third-party trust with more nuance: Given a third-party authority (A), there ought to be another authority (B) to supervise A , so that A can be trusted. In turn, there should be yet another authority C to supervise B [10], and so forth.

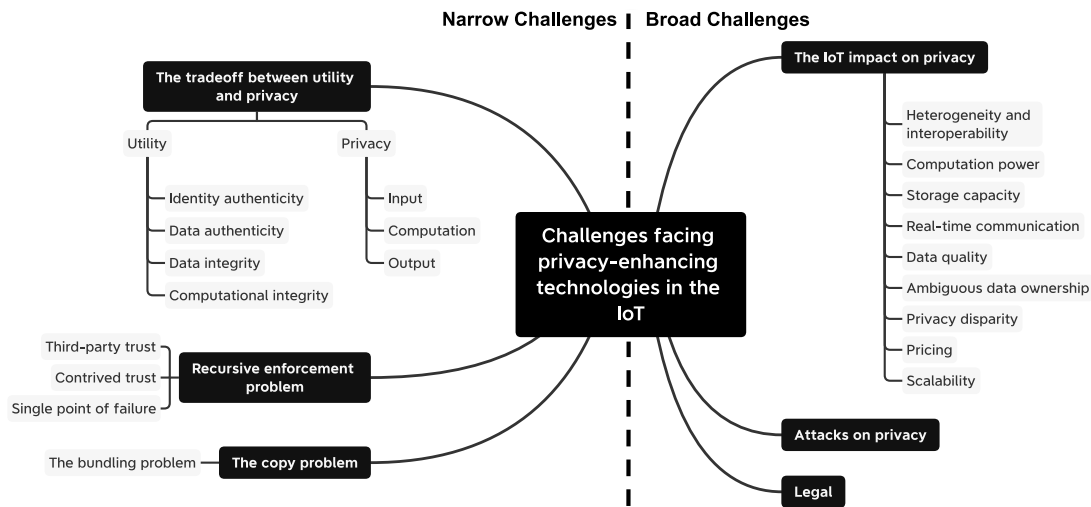


Fig. 8. Overview of the narrow and broad challenges facing privacy-enhancing IoT data markets.

The REP is a significant challenge that has been covered and tackled implicitly by some of the studies in our review [S1,S3,S4,S7,S10,S16,S25,S41]. Additionally, others tackle a sub-set of the REP, which is the single point of failure of trusting a unique third party [S6,S12]. According to [S8], the hesitation in trusting third parties is one of the main reasons for the slow adoption of IoT data markets. Indeed, it is hard to technically ensure and prove that the third party will not use one's data for purposes other than those agreed [S8]. Additionally, adoption is further slowed down because the use of third parties to supervise other parties incurs costs [S46]. Furthermore, users' daily interactions with "trusted" third parties can be regarded as a product of *contrived trust*, another form of the REP. For instance, applications from large service providers with negligible competitors push users to accept the sometimes poor privacy conditions, e.g., GPS apps. Note that *contrived trust* is different from the trust users have on cryptography, open-source code, or consensus mechanisms that the broad scientific community has audited over the years.

Tackling the REP requires reducing the power and the responsibility of the third party in a particular aspect of a specific service by, e.g., distributing such responsibility among other parties or distributing the power among multiple parties that enforce rules on each other. These measures can ease the hesitation to trust a single third party, tackle *contrived trust*, and reduce the single point of failure because the third party would be supervised and held accountable by other third parties in a flat hierarchy. Fortunately, the PETs included in this study can also circumvent – only onion routing can tackle – the REP, which, in turn, reduces the need for third-party trust and, therefore, reduce *contrived trust* and a single point of failure. Additionally, 5 of the 7 AETs included in this study can tackle the REP, primarily distributed ledger technology, whose architecture was purposefully built to tackle the byzantine generals' problem [145,146], a manifestation of the REP.

8.1.3. The copy problem

Once an entity releases data freely or for profit-seeking, the data is no longer under the original owner's control. Consequently, the recipients of such data can *copy* and, e.g., re-sell or use the entity's data for a non-agreed purpose without informing or acknowledging the original owner [10]. Beyond the privacy threats the copy problem (CP) entails for users of service providers under poor privacy conditions, the CP is a major obstacle for organizations to engage in data markets, which some of the selected studies implicitly tackle [S2,S25,S45]. The CP leads companies to either hoard or sell data as fast as organizations obtain the data, lest its value drops [10]. Nonetheless, secure and outsourced computation PETs such as trusted execution environments or homomorphic encryption can tackle the CP by allowing other entities

to extract value without losing control over the input data beyond the specific information sold, such as an algorithm's evaluated output on this data. This paradigm is profound because tackling the CP makes data scarce (to some extent), as the original data is not shared, and the data owner would not allow a non-agreed computation. Thus, selling the *access* to data can be more attractive to companies, as data would preserve their value longer than releasing the data.

A subset of the CP is the *bundling problem* (BP) [10], which is an attack vector different from re-identification that occurs when an entity requests actively or passively more data than strictly needed to (i) prove a claim or (ii) perform an analysis. Harvesting more information than needed worsens data breaches' consequences for individuals and companies and indicates questionable business ethics. For instance, (i) to prove one's age with an ID card, the prover usually shares all the information in the ID instead of only the age and proof of the card's authenticity. The BP is a subset of the CP because if one tackles the CP, neither necessary nor additional information is released beyond the required computation. For example, tackling the CP by restricting verification and processing to a trusted execution environment also tackles the BP. In this setting, the data consumers cannot copy the necessary data or metadata for other unsolicited analyses, despite being able to process metadata to verify the authenticity of the data and the integrity of the computation and obtaining the desired outputs of the analysis. Additionally, (ii) anonymization-based PETs such as differential privacy or *k*-anonymity reduce data authenticity to tackle the BP. For instance, in a demographic analysis that only requires the first digits of the ZIP code to perform clustering, data curators can generalize the ZIP codes with *k*-anonymity, so only the strictly necessary information is revealed to data scientists. Nonetheless, anonymization can suffer from background-knowledge-based attacks [S48,S156] and does not solve the CP because the data consumers can replicate the privacy-enhanced data.

8.2. Broad challenges

8.2.1. The IoT impact on privacy

The paradigm brought by the IoT brings significant amounts of data to markets. However, this paradigm also bears some of the shortcomings of IoT devices [198]. Table 2 contains an overview of these challenges and briefly discusses their impact on privacy. In summary, privacy is always affected by the context and employed technologies, which underlines the importance of adhering to privacy-by-design principles [131] and the need for practitioners in other fields such as software engineering, economics, law, and politics to tackle together the diverse issues that IoT entails for privacy.

Table 2

Overview of challenges brought by the IoT paradigm into data markets explicitly covered by some of the studies included in this SLR.

Challenge	Studies	Description	Impact on privacy
<i>Heterogeneity and Interoperability</i>	[S36] and [S39]	The IoT consists of billions of IoT devices from different manufacturers, running different software on different local networks and geographic regions, with different computation power and storage capacity [147]. Furthermore, different communications standards, connectivity and availability aggravate the interplay of IoT devices.	An IoT data market should be agnostic to these differences and minimize any additional requirements; however, it is unclear how global data markets should harmonize data coming from different jurisdictions with different privacy regulations and how an IoT device can interact with another whose, e.g., verification schemata are considered inadequate. In addition to these obstacles, a lack of interoperability may restrain PETs that involve the communication between many devices, e.g., MPC.
<i>Computation power</i>	[S5], [S11] and [S48]	Manufacturers produce many IoT devices designed to consume low energy and require minimal volume, limiting these IoT devices to the core functionalities of monitoring and communication [S11].	Any additional computation requires a higher investment in resources and manufacturing, and running some PETs becomes infeasible without this extra investment. Consequently, a set of PETs is excluded without more computation power, e.g., cryptography-based PETs such as HE, MPC, ZKP, some digital signatures, or consensus algorithms. This limitation, however, may only apply to contexts where it might not be possible to connect IoT devices acting as clients with proprietary or trusted third-party nodes where these PETs are executed.
<i>Storage capacity and real-time communication</i>	[S1], [S5], [S11], [S17], [S29] and [S48]	Minimizing the physical volume of an IoT device reduces their price but limits their storage capacity, forcing IoT devices to transmit the data to a data warehouse or a data market as quickly as possible. This tendency intensifies in some IoT applications where the time delay tolerance is low to enhance the utility of real-time information [S17].	Processing time constrains the number of usable PETs, excluding those that require long execution times, such as fully HE or creating a ZKP.
<i>Data quality</i>	[S14]	An unreliable IoT design may afflict thousands of IoT devices mass-produced by a manufacturer, which at deployment may lead to millions of unreliable data points. Furthermore, networks may also be unreliable, further worsening the quality [S14].	The impact may seem beneficial in terms of privacy; however, unreliable data leads to verification and secure computation schemata to fail and anonymization technologies to over-perturb the data as the underlying data is not entirely truthful.
<i>Ambiguous data ownership</i>	[S4], [S10] and [S14]	When purchasing a device or a cluster of IoT devices, e.g., a phone, consumers also expect to own the data they are generating. However, the phone manufacturer and service providers expect to receive parts of this data nowadays with meager consent. In addition to this clash of interests, there are scholars that ponder whether data belongs to anyone in the first place, like [148].	Having unclear data ownership leads to a misguided deployment of PETs, which may cause detrimental consequences if the privacy measures fall short. On the other hand, if the practitioner knows who has the right to the data and what the owner is reticent to share with a third party, then selected PETs and their privacy tuning can be optimized accordingly.
<i>Privacy disparity</i>	[S4] and [S40]	Depending on the IoT devices' deployment location, the degree of privacy measures should be higher or lower, e.g., sensors in vehicles, smart homes, phones, and wearables. Furthermore, IoT deployments should adapt the monitoring time to an adequate amount depending on the context [S40].	Some PETs, such as semantic and syntactic technologies, allow adjusting the degree of privacy; however, others are more rigid. Selecting and adapting a PET to the IoT devices' deployment context requires expertise.
<i>Pricing</i>	[S4], [S14], [S17] and [S50]	There are multiple variables imposing the price of data aside from supply and demand: the truthfulness, the source, either purchasing the data or the access [S17], and the privacy level. These factors add additional complexity to pricing, e.g., the sources have become disparate with the IoT, which drives pricing to a more granular task than before, when aggregated data could be sold as a unit [S17].	Aside from payment enforcement mechanisms, pricing involves negotiations, which frequently must ensure privacy. This adds an extra layer of complexity to the deployment of PETs. Furthermore, as data markets trade with more granular data points, PETs that need aggregation might be excluded in some contexts, e.g., syntactic technologies such as k-anonymity.
<i>Scalability</i>	[S25] and [S36]	The number of IoT devices and streamed data grow exponentially across industries [147,149], which extends data collection and improves analytics across different domains, e.g., health, insurance, or finance. To gain these benefits, there is a need to increase networks' communication and overall storage capacity as well as interoperability and security efforts.	As the IoT scales, analysts will access more datasets from different domains to create new products and services, e.g., linking driving behavior with insurance in pay-how-you-drive schemata [150]. Such innovations stem from the "mosaic effect" [151], where disparate datasets with limited information value can obtain significance when combined with other datasets. However, malicious entities can leverage such an effect to extract sensitive personal information not explicitly contained in a dataset [152].

8.2.2. Attacks on privacy

Adversaries can be malicious, actively trying to breach users' privacy through hacking, or honest but curious, passively gathering data from users to reveal hidden insights [S35]. Both of these entities can carry re-identification attacks with the collected information. Within the context of the IoT, the list of security and privacy attacks is extensive (sniffing, cache poisoning, DoS/DDoS, sinkhole attacks, replay attacks, among others) [153]. Furthermore, within our SLR, [S45] discuss some of the additional attack vectors these malicious or curious entities may execute in the context of IoT data markets to learn sensitive information from users. Notable ones include: *Data forwarding*, which is one way the *copy problem* materializes; *roles collision*, where data brokers and buyers may be the same or collaborating entities, and, therefore, the broker could rig the auction for its benefit and access the sold data; and *side channel attacks*, where attackers exploit the physical properties of the hardware or its power consumption to extract knowledge from the hidden computations (trusted execution environments suffer mainly from this attack).

Such attacks make the possession of data intrinsically risky because if attackers are successful, data re-identification is possible [S38], even if data have undergone some form of privacy enhancement [154]. There are common attacks used to re-identify data, e.g., reconstruction, tracing, or linkage attacks [118,155]. Some of the most famous re-identification *white-hat* attacks involve [156] who deanonymized the Netflix Prize dataset with IMDB's public dataset in 2008, [157] who performed the same feat with Amazon's public review data, and [6] (the inventor of *k*-anonymity) who re-identified participants within a genome sequence dataset in 2013. Furthermore, in 2014, [7] tracked drivers with home address and vehicle speed as inputs, and in 2020, [154] matched users with large-scale mobility datasets from a mobile network operator and transportation smart card usage.

Overall, IoT data markets will facilitate access to large quantities of data from different domains, including biometrics, which will increase the impact of these attacks and the potential harms to individuals, e.g., insurance, employment, or price discrimination. Therefore, IoT data markets require a more robust adoption of PETs and security standards.

8.2.3. Legal challenges

Progressively along the past decades, governmental institutions have released laws to protect the privacy of their citizens (see Section 2.1). These laws also refer to an individual's and businesses' right to exploit their data commercially, which provides leeway for data markets [S28] and aims to uncover the untapped potential of data for innovations.

Nonetheless, research points out the sometimes unrealistic expectation to monitor the entirety of the Internet for privacy violations [S45], and the dexterity of hackers to find novel deception methods [S3], and that laws are more reactive than preventative. Well-known networks of illegal proprietary digital asset exchanges, e.g., scientific works and how users of digital services give away data, tacitly provide testimony of the failure of data-related legal measures today, and the problems will likely increase with the accruing number of IoT devices [S38]. Moreover, privacy regulations can strangle free markets and innovations if they are too stringent [S45].

Aligned with these deficiencies, [S38] introduced privacy regulation pitfalls that the IoT unfolds in data markets in 2013. They note that (i) definitions of personally identifiable information will be deprecated as unprecedented amounts of data can be aggregated, easing re-identification, (ii) the development and audit of PETs is costly, which may limit business models and potentially make disregarding privacy regulation profitable [199], (iii) privacy violations result on small fines or remain unpunished, (iv) technology tends to outpace regulation, and (v) the ubiquity of IoT devices will yield more illegal secondary personal data markets. After almost a decade of further research, (i) seems valid, at least in some scenarios. The ambiguity of privacy

regulation is a barrier in some cases, as practitioners may default to weaker forms of privacy if their architecture appears to comply. This leads to re-identification – an attack that is also more practical with the increasing number of IoT devices [156] – being more likely to succeed. However, in defense of these practitioners, while PETs have improved since 2013, some PETs that offer better privacy enhancements are still complex and not yet performant in 2021.

Based on the prior arguments, pitfall (ii) seems to hold; however, (iii) is no longer a strong pitfall. Since the enforcement of GDPR [137] in 2018, GDPR has punished multiple corporations with considerable fines ranging between €20 million and up to 4% of a corporation's annual worldwide turnover of the preceding financial year. As of the writing of this publication, GDPR has harvested considerable fines assigned to Google in France on two occasions [200], Amazon [158], H&M [159] or the telecommunications operator TIM [160]. These fines alone accumulate to €282 million. These statistics are a sign that PETs are not appropriately introduced in production applications even by big technology companies and that not complying with privacy regulations in an IoT data market has dire economic consequences. While these fines could indicate how profitable it still is to violate privacy regulation (iii), one can no longer vigorously defend (iii). Pitfall (iv) seems to materialize as long as the nature of law-making does not change. Lastly, pitfall (v) is concerning, given the existence of legal personal data markets that store up to 750 million user profiles and trade 75 million online auctions daily like BlueKai [23], whose data could leak to the increasing number of illegal *shadow markets* [S6].

9. Discussion

This Section presents a set of key findings (KF) distilled from the two research questions answered in Sections 6 and 7 as well as 8, the content and metadata of the 50 publications included in our SLR, and other seminal studies that we encountered throughout our SLR but which do not necessarily address IoT data markets directly. Lastly, we cover the limitations of this study and future work.

9.1. Key findings

(KF1) *The attention of scientists toward privacy-enhancing technologies in the field of data markets for iot devices has increased notably in recent years.* The selected publications are modern, as 49 of the 50 studies were published between 2012 and 2020, and 34 of them (68%) were published either in 2018, 2019, or during the first half of 2020. While the absolute number of publications in 2020 is lower than in 2019 because we captured only the first seven months of 2020, Fig. 9 illustrates the arguably accelerating trend of the cumulative curve of publications in the field of privacy-enhancing IoT data markets.

(KF2) *The most frequent research type (design and creation) and least common research contribution (lessons learned) suggest that privacy-oriented iot data markets are still maturing and have not faced many production-grade implementations yet.* According to Fig. G.13, around 76% of the publications use a *design and creation* research approach, while only 4% perform a case study. A further indication of field novelty is that only one out of the 50 publications had the contribution type *lessons learned* [S40]. Furthermore, while 35 studies (70%) were of research type *solution proposal*, to the best of the author's knowledge, only one solution appears to have an implemented *system that is applied in production* [S25].

(KF3) *The selected studies rarely leverage existing libraries that provide PETs and often only build upon architectures developed in previous work to a small degree. Therefore, to gain more practical relevance, it may be beneficial for researchers to improve and extend existing work instead of reinventing the wheel.* The research community and industry have developed many open-source libraries to employ zero-knowledge proofs,

Table 3

A mapping of privacy- and authenticity-enhancing technologies (PET and AET) to the narrow challenges using the terminology defined for this review. The extent of enhancement of privacy, utility, and characteristics of the different PETs and AETs varies from significantly increasing ++, over +, +-, - to significantly decreasing --. na denotes *not applicable*. w/ denotes *with*. *Considering a digital certificate when using digital signatures, if applicable. The privacy column assumes data and identity are authentic.

Layer	Technology category	Technology	Narrow challenges							
			Privacy-utility tradeoff					Recursive enforcement problem	Copy problem [Bundling problem (BP)]	
			Privacy [Confidentiality (Conf)]			Utility				
Input [Data (D)] [Identity (I)]	Computation	Output	Authenticity [Data (D)] [Identity (I)]	Integrity [Data (D)] [Computation (C)]						
PET										
Processing	Secure and outsourced computation (SOC)	Zero knowledge proofs for computational integrity	++ D	++	++	++ D	++ D, C	Circumvents	Tackles (BP)	
		ZKPs of anonymous credentials	++ I	++	++	++ I	++ D, C	Circumvents	Tackles (BP)	
		Trusted execution environments	++ D	++	+ Conf	na	++ D, C	Circumvents	Tackles	
		Partially homomorphic encryption	++ D	++	+ Conf	na	++ D, C	Circumvents	Tackles	
		Fully homomorphic encryption	++ D	++	+ Conf	na	++ D, C	Circumvents	Tackles	
		Secure multiparty computation	++ D	++	+ Conf	na	++ D, C	Circumvents	Tackles	
	Anonymization (AN)	Privacy-preserving data mining	W/ SOC (w/ AN)	++ D (+ D)	++ (na)	+ Conf (+)	na (+- D)	++ D, C (na)	Circumvents	Tackles
			Federated (w/ AN)				na (+- D)	na (na)		
		Differential privacy (DP)	+ D	na	+	+ D	na	na	Tackles (BP)	
		K-anonymity	+ D	na	+ D	- D	na	na	Tackles (BP)	
Perturbative	+ D	na	+ D	- D	na	na	Tackles (BP)			
Pseudonym creation	+ I	na	+ D	- I	na	na	Tackles (BP)			
Storage Communication	Encryption	Storage layer: ++ D, Conf Communication layer: ++ D, Conf			na	++ D	Circumvents	na		
	Onion routing	++ I D, Conf	na	na	na	++ D	Tackles	na		
Verification	Hashing	++ D, Conf	na	na	na	++ D	Circumvents	na		
	Ring digital signatures	+ I	na	na	- I	++ D	na	Tackles (BP)		
	Blind digital signatures	-- I D	na	na	++ I D	++ D	na	Tackles (BP)		
Sovereignty	Smart contracts (for privacy policies)	Characteristics and challenges are pegged to distributed ledger technology.								
	Privacy policies (Access control)	Characteristics and challenges are pegged to the selected PETs and AETs employed to fulfil the privacy requirements.								
AET										
Consensus	Distributed ledger technology (permissioned)	- I D, Conf	+ Conf	+ Conf	+ I	na D	+ D, C	Tackles	na	
	Distributed ledger technology (permissionless)	++ I D	-	-	+ I	na D	++ D, C	Tackles	na	
Verification	Truth discovery	- I D	-	-	+ I	+ D	++ D, C	Tackles	na	
	Decentralized identifiers	- I	na	na	++ I	na	Tackles	na		
	Digital fingerprint (Version control)	- I D	na	na	++ I	++ D	++ D	na	na	
	Identity-based digital signatures	- I	na	na	++ I	++ D	na	na		
	Certificateless digital signatures	- I	na	na	++ I	++ D	Tackles	na		

homomorphic encryption, secure multiparty computation, or differential privacy (see Section 6). However, none of the studies have indicated their use. Furthermore, studies often do not build upon each other, leading to overlapping further. For example, [S41] and [S34] both showcase an auction that obscures the bids by employing partially homomorphic encryption. However, W. Gao et al. only refers to the work from Z. Chen et al. in one line, noting that “[...] *there is only few literatures on designing privacy-preserving schemes in data market auctions*”. Moreover, [S42] builds upon [S27] and [S2] upon [S25], but each of these two sets belongs to the same group of researchers. In conclusion, it may be beneficial for researchers to incorporate building blocks

from previous data market architectures to advance privacy-oriented research.

Moreover, many studies included in Table D.7 aim to create an IoT data marketplace employing distributed ledger technology (DLT). However, there seems not to be a consensus about which DLT to use for IoT data markets, as the authors build upon Ethereum, IOTA, Hyperledger Iroha, Fabric, Agora, or Quorum, among others. Specifically, as an example, [S32] uses Ethereum smart contracts for payments while [S28] only uses these contracts for safelisting and employs IOTA for payments instead.




Fig. 9. Publications in the field of privacy-enhancing data markets for the IoT from January 2002 to July 2020.

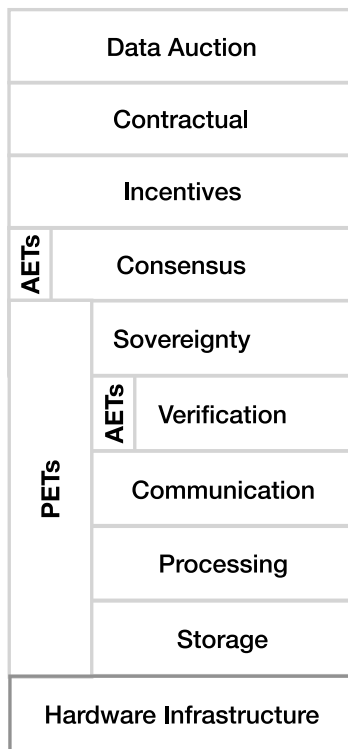


Fig. 10. Reference model for the layers of a privacy-enhancing IoT data market.

(KF4) *The content of the selected studies can be categorized into two main orthogonal research streams within the context of privacy-enhancing IoT data markets: architectures and data trading schemata.* The first research stream is dedicated to the design of privacy-enhancing architectures for the exchange of data in IoT data markets (25 studies, 50%), and the second one focuses on the design of privacy-enhancing data trading such as auctions (12 studies, 24%). The remaining studies can be associated with domains like legal [S6], user preferences [S8], or IoT data market challenges [S14,S19]. The selected studies, and also international initiatives such as the European GAIA-X [5], hence envision data markets beyond matchmaking and auction capabilities. Specifically, the studies that we analyzed structure the software, hardware, abstract entities, and their coordination, data processing, storage, communication, and the offered services to build a holistic or part of a privacy-enhancing IoT data market that includes PETs to tackle some of the challenges described in Section 8.

(KF5) *Despite the acknowledged need for combining anonymization and secure and outsourced computation techniques, none of the researchers behind*

the 12 studies proposing data trading schemata, and only two publications out of the 25 designing data market architectures employ both PET categories in combination. Although PETs such as homomorphic encryption (HE) or secure multiparty computation conceal inputs and computation, the outputs can leak information about the underlying data and hence may be exposed to re-identification attacks [10]. Combining secure and outsourced computation techniques with anonymization-based PETs like differential privacy (DP) can help to make the outputs less sensitive. Moreover, leveraging only anonymization PETs does not sufficiently address the copy problem.

Within the 12 selected studies focused on data trading, DP is the most frequently used PET in auctions to enhance the privacy of the exchanged data. [S16,S22,S24] and [S37] employ DP in various forms to set the privacy levels and, subsequently, the price of the traded IoT data. Researchers might choose DP over other anonymization technologies because DP is the only PET with a mathematical guarantee of privacy [13]. At the same time, partially HE (PHE) is the PET of choice to enhance the privacy of the bidding process. A group of authors [S26,S34,S41] chose PHE primarily for hiding the bids, confidentially computing the winner, and only revealing the output to the auction's winner. Researchers might decide to use PHE over other forms of HE, secure multiparty computation, or trusted execution environments (TEEs) despite PHE's significantly less general scope because PHE has relatively high performance and is conceptually simple.

Together, DP and PHE can holistically enhance the privacy of auctions, which is a contribution we have not found in this review. [S48] emphasize that some HE schemata, such as Paillier's, must complement other methods to guarantee more protection. Moreover, while HE protects the input and the computation itself, if the intended recipients of the decrypted output are malicious, they may reverse engineer the output to learn properties about the input. An additional modification employing, e.g., differential privacy, of inputs or decrypted outputs before sharing may help prevent this attack in exchange for accuracy and thus utility. The same argument applies to other secure and outsourced computation methods when used in isolation. We consequently point to a lack of combination in the research stream of data market architectures, except for two publications from the same group of researchers [S2,S25], which use TEEs to train machine learning models with DP.

(KF6) *The selected studies employ three dimensions to characterize data markets that entail privacy concerns: the degree of decentralization, the types and number of data domains, and the types of sellers and consumers.* Each of these dimensions, for example, characterized by [S1,S27] and [S46] respectively, brings privacy concerns. Data may be stored by the seller, the platform provider, or a decentralized platform using, e.g., a combination of commercial cloud storage, interplanetary file systems, or blockchains. Depending on the degree of decentralization and replication, practitioners need to consider different leakage risks. In particular, if the architecture relies on a blockchain, PETs are particularly important [134].

An increase in the number and types of data domains opens additional attack vectors and more possibilities for malicious entities to link an individual's data across databases. This hyper-connectivity between datasets can render the definitions of de-identified data, such as HIPAA's, obsolete and suggests that privacy enhancements in the data economy should be defined globally and not locally.

The degree of privacy enhancement should depend on the type of seller and consumer, e.g., consumers may expect higher privacy guarantees when a health insurance company gathers their data than when the collector is a renowned health research institution.

(KF7) *Based on our classifications in Sections 6 and 7 and inspired by a set of seminal selected studies, we have created a reference model for the design of IoT data markets in Fig. 10, and detailed in Table 3.* Most of the studies included in this SLR proposed solutions without following a reference model, except for [S3] and [S7], who developed their own without a systematic research. [S7] condense their architecture into two layers: data acquisition and trading. On the other hand, [S3] present a more holistic view of privacy-enhancing IoT data markets with six layers (*identification, privacy, contractual, communication, consensus, and incentive*) inspired by the Open System Interconnection model and heavily conditioned by the use of blockchain technology. This model, however, lacks essential steps of an IoT data market that several publications in our SLR focused on, namely storage [S1,S12,S25,S28] and processing [S7,S19,S20,S45]. Furthermore, the *identification layer* [S3] can be regarded as a subset of verification, which also includes data verification. Other studies base their market design on the type of participants [S15,S24,S27,S33,S43], e.g., sellers, aggregators, brokers, among others, and the type of data domain [S46], e.g., health, financial, or a combination. However, these categories cannot be transferred to other contexts as easily as a reference model agnostic to entity and data domain types.

Our reference model hence combines and generalizes some of the layers from [S3,S7] and complements them with additional layers such as the data auction, storage, verification, processing, and sovereignty layer (see Fig. 10). Most of these layers need multiple PETs, as there is no "one-size-fits-all" technology to enhance privacy. To navigate these layers in detail, refer to Table 3 and Fig. C.11. Furthermore, we distinguish between a contractual and sovereignty-related design to separate formal agreements from privacy and ownership policies. Furthermore, given the distinct purpose and implementation that auction schemata play in a data market, they should be respected by a unique IoT data market layer (auction dedicated studies: S16,S22,S24,S37, among others). Lastly, incentives are necessary to encourage behavior that preserves the pre-defined qualities of the IoT data market, e.g., optimized prices [S3,S16], data authenticity [S16,S23], or maintaining the infrastructure like a permissionless DLT.

(KF8) *Aside from the ubiquity of digital signatures in IT systems, in this slr, distributed ledger technology (DLT) is most frequently employed as the backbone of IoT data market design (see fig. g.14), despite the lack of consensus on its use and DLT-based applications in production.* Although centralized systems seem more efficient and easier to deploy, and despite the seemingly few industrial applications running on blockchain today, many researchers in this SLR still advocate for distributed systems using DLT. Within the 35 solution proposals, around 31% chose permissionless DLT, 14% consortium DLTs, and the authors of the remaining 55% either reviewed DLT, implemented a centralized solution, or focused on designing narrow features. However, we noted that within the 45% of DLT-based designs, many authors still relied on single entities for data processing or storage. Specifically, only one of the 50 studies [S25] has a public blockchain-based ecosystem in production, yet without a real-world use case running. These statistics indicate a lack of adoption despite substantial research efforts.

Furthermore, while blockchains enhance authenticity, assure integrity, and enable payments without the need for a trusted third party, blockchains are limited in storage capacity [S1], computation

power [S25], and can exacerbate privacy issues because of their tamper proof-quality and inherent data and computation replication [56,132,136,161]. Consequently, almost all studies that include blockchain technology to support an IoT data market require PETs to protect users' data and identities. These studies go as far as creating innovative privacy-enhancing blockchain architectures with other PETs as building blocks, e.g., trusted execution environments [S25,S45], or adding a privacy layer to their market design based on differential privacy [S3]. However, within the literature, there are also questionable statements such as "[...] researchers and technologists have found that blockchain can be a potential solution to the privacy problem by decentralizing information [...] Blockchain can be used to securely share private information [...]" [S3], "Blockchain-based approaches provide decentralized security and privacy [...]" [S11], or "Blockchain has been proven to possess security, immutability, and privacy properties, which has caused a lot of researchers to introduce it into the privacy and security concerned IoT" [S23]. These statements, coupled with the current excitement around blockchain, can lead practitioners in the industry to wrongfully push blockchain for "privacy". Therefore, the community would benefit from clear explanations of why authors employ blockchain and clearly state the need for other technologies to enhance privacy.

9.2. Limitations

Even though we have adopted a rigorous research design and paid particular attention to the selection and analysis of published studies, SLRs have limitations that may have undermined our effectiveness. These threats include (i) incompleteness of study search, (ii) bias in study selection and (iii) inaccuracy of data extraction.

(i) Some relevant publications might be absent. To mitigate this limitation, we searched in several highly reputed digital libraries, performed a preliminary search to determine suitable search strings, conducted a backward search to identify additional related work, and included studies in advance that met the standards and filters of this SLR. These measures reduce the probability of missing relevant publications. (ii) The experience and knowledge of the researchers may drive the study selection with an inherent bias. Nonetheless, following Kitchenham [43], we aimed to create a set of explicit inclusion and exclusion criteria to maximize the degree of objectivity. To mitigate different appreciations of these criteria, we conducted a preliminary search to ensure researchers have a consistent understanding of the requirements. Furthermore, two researchers conducted the selection process independently and resolved the conflicts between their decisions interactively. (iii) There might be a bias in selecting the extracted data, which may affect the classification results of the selected studies. To mitigate this potential limitation, the two researchers specified a set of data extraction cards (see Section 3.2) to eliminate any misalignment in the data extraction process results.

9.3. Future work

The opportunities and need for future work in the context of privacy and data markets for the IoT highlighted by the selected studies resonate with the challenges covered in Section 8. Most notably, there is a need to solve the copy problem [10,S4,S17] and to lessen IoT devices' limitations regarding computation [S5,S11], storage and capacity [S29,S48] to tackle or circumvent the constraints PETs may induce. Moreover, to decrease the probability of re-identification attacks, further work is needed to advance the maturity of PETs and combine them, e.g., bringing together differential privacy and secure and outsourced computation efficiently. Additional research is also necessary to create standards for data markets, such as a language to describe privacy requirements, universal APIs to interact between different IoT devices with various degrees and techniques for privacy protection, and machine-readable definitions of privacy, e.g., using ontologies [S40]. In this context, a more detailed description and

classification of the layers that we found relevant for classifying privacy and authenticity enhancing technologies (see Fig. 10) constitutes a promising and relevant avenue for future research. Nonetheless, we want to emphasize that privacy is not the only challenge that needs to be addressed, as future research must also consider, for instance, scalability.

If society considers privacy a necessity, it should be enhanced by default and optimally in any system without attaching price tags to one's privacy, as some of the selected studies pursued suggest [S16,S15, S22]. We find this posture a worthy research endeavor and encourage researchers to ponder whether monetizing privacy in a competitive market ultimately benefits society. Furthermore, legal practitioners have ample ground to develop legislation specifically around privacy in IoT data markets and for economists to delve into data pricing and decentralized market interactions. Legal researchers could investigate how stringent privacy regulations should be, as heavy regulation may strangle free markets and innovations [S45]. Additionally, the legal, pricing and privacy aspects hinge around data sovereignty. As long as ownership is ambiguous, researchers' efforts will struggle to maximize impact. Furthermore, the relevance of our results may reach beyond IoT data markets, as the analysis of PETs and derived insights, e.g., how IoT impacts privacy, can permeate other research areas such as privacy-by-design software engineering, policy-making, and data governance, politics, and economics. Moreover, most PETs have specific performance-, complexity- or utility-related shortcomings (which we describe in Section 6) that researchers can address.

Lastly, we recommend that researchers derive decision trees based on Table 3 to enhance the decision-making of privacy officers beyond our work. Moreover, we could not find any formulation of an information-theoretic quantification of the data leaked from a data market. We also encourage social scientists to focus on questions related to data sovereignty. To realize a vision of data markets that benefit society, we suggest researchers concentrate on roadblocks such as the copy problem. Finally, institutions should consider updating their privacy-enhancing processes to effectively participate in IoT data markets.

10. Reassessment of the results

This section provides and discusses new key publications since the research process ended. Accordingly, we conducted a research process as per Section 3 for studies dated between July 2020 and May 2022 and, among them, picked for discussion the ones providing the most significant updates to our systematic literature review or, on the contrary, underlining our previous findings. Note that the references included in this section correspond *only* to the newly found publications.

In our new search, we again selected primary and secondary studies. Overall, our new search resulted in 24 publications: 3 more from 2020, 14 from 2021, and – as of May – 7 from 2022. These statistics indicate that the trend depicted in Fig. 9 (consolidating KF1) has not reversed. Notably, we could still not find publications discussing production-ready deployments of privacy-enhancing architectures or auction schemata and no reference to open-source tooling despite PETs being more mature since July 2020 (underlining KF2 and KF3).

Among the *secondary studies*, [162] explored the concept of privacy in the digital economy more broadly and pointed out the need for interdisciplinary research to supplement the purely technical PET constructions with the economic (tradeoff between accuracy and privacy) and governance perspectives (privacy policies) that we elaborate on in our paper. [163] systematically investigated privacy-oriented identity management in the context of the IoT, such as anonymous credentials and other techniques that our review covers. Moreover, [164] conducted a less systematic survey of standards and future challenges, including discussions regarding authentication and access control, and highlighted a subset of the privacy challenges of IoT that we present in Table 2. Additionally, [165] presented a recent systematic literature

review on designing data markets. However, their work did not focus on privacy.

More secondary studies, such as from [166] considered blockchain and smart contracts beneficial for privacy. Furthermore, [167] reviewed PETs in the context of crowdsensing and emphasized the privacy issues with smart contracts, along with practical challenges in security and feed-in of reliable data. They suggested a subset of the anonymization techniques that we present in Section 6, such as privacy-oriented digital signatures, anonymous networking, k -anonymity, t -diversity, t -closeness, and differential privacy (DP), and some more specific ones in the context of location. While they also mentioned ZKPs, there is no detailed discussion of the secure computing techniques we survey. [168] considered privacy mechanisms in data sharing for collaborative forecasting and discussed the tradeoff between privacy and accuracy. They distinguished between perturbative techniques (“data transformation”), MPC-based protocols, and distributed or federated approaches (“decomposition”) combined with DP. Lastly, [169] provided a survey that examined the privacy risks that machine learning poses on IoT data markets supported by blockchains. Hence, our SLR, with its comprehensive focus on privacy, still fills the gaps that we discussed in Section 4.

The *primary studies* followed a similar pattern to the previously collected studies. Above all, many still employed blockchains and often did not provide clear explanations the corresponding benefits and acknowledgments of the corresponding challenges, specifically regarding privacy (reaffirming KF8). We again encountered questionable claims such as “[...] a decentralized approach based on distributed ledger technologies (DLT) enables data trading while ensuring trust, security, and privacy” [170], without discussing why DLT enhances privacy in the rest of the publication about benchmarking IoT data trading protocols in blockchains. Others followed suit on the use of blockchain to support electric vehicle trading marketplaces with IPFS and a scheme to hide payment sources [171] and cloaking location with k -anonymity [172] or proposing a new blockchain architecture with permissioned domains to enhance privacy for data market places [173]. Another presented several building blocks (blockchain, trusted execution environments, gossip learning) without an evaluation of the proposal [174].

A notable exception is the comprehensive details provided by [175] in their blockchain architecture. Their architecture stores encrypted sensor data in cloud storage, and smart contracts support sensor registration, data auctioning, and payments. While the smart contract emits notifications and displays the endpoint for retrieving proxy re-encrypted data, the data are exchanged off-chain confidentially via proxy re-encryption. This construction addresses transparency and scalability issues regarding sensor data. Nevertheless, bidding and payment processes may still reveal sensitive information and require future research by combining this approach with some of the PETs we surveyed. [176] also acknowledged the aggravation of privacy issues on blockchains and combined federated learning with DP to obfuscate clients' weights and use ZKPs to prove the integrity of the training and evaluation process, which they required for providing fair incentives managed by a smart contract. Another related publication by [177] presented a blockchain-based solution for tracking IoT sensor data across marketplaces and, thus, only detecting but not preventing illegitimate replication and resale.

These publications fall into the category of *architectures* identified in KF4. We also found papers in the *data trading schemata* category (or related): two new data auction schemata enhanced with DP [178, 179], a task assignment scheme in crowdsensing that hides the tasks' content with homomorphic encryption (HE) for crowdsensing [180], and another where they employ DP on billing data [181]. The latter publication, however, does not discuss *fairness*, which is critical in monetary use cases as a noisy bill can make data prosumers profit less from their data on some occasions. Furthermore, [182] and [183] focused on determining fair prices for end users' datasets that are

anonymized with DP according to their accuracy and, correspondingly, risks of revealing sensitive information.

One interesting development that explored the paradigm in ML markets comes from [184]. They developed a privacy-enhanced framework to evaluate the quality of ML models and data for sale with functional encryption, achieving improvements over similar schemata implemented with HE. Another novel concept by [185] strives to empower users with tools to help them determine the risk of sharing their information and, accordingly, make an informed decision about their data framework. It is thus closely related to enforcing privacy policies in the sovereignty layer. Except for [176], who focused narrowly on federated learning, the rest of the new (notable) primary studies did not leverage the combination of anonymization and outsourced computation technologies, underlining KF5. Notably, the new publications have not altered the data market characterization of KF6 or the reference model of KF7.

11. Conclusion

With this review, we reveal the landscape of PETs in data markets for the IoT. We have conducted a systematic literature review (SLR) to identify and filter the studies aiming to solve this landscape's challenges. Consecutively, we formulated terminology to dissect the selected studies' architectures and findings and identified the PETs that related work employed and which specific challenges they addressed.

The authors of the selected studies in this SLR have devised proposals for privacy-enhancing IoT data marketplaces to comply with privacy requirements while maintaining utility, profitability, and fair and seamless data exchange. Since this is a relatively new, multidisciplinary research field, the optimal combination of technologies and theoretical foundations employed in these proposals is still in the development phase. Therefore, no proposal has established itself as canonical yet. Moreover, we observed that the research community needs to further explore the balancing act of utility and privacy before data markets flourish. We conclude that the practicality of PETs needs to advance further to positively impact data markets for the IoT. Additionally, we suggest researchers solve the copy problem and improve privacy-enhancing verification as their absence discourages data markets from forming. We also discovered that research on privacy-oriented data markets could benefit from increased reuse of components from previous articles and existing open-source libraries and a more explicit description of critical objectives. For example, the benefits of utilizing distributed ledger technology (DLT) in data markets for IoT architectures often remain unclear, and authors do not sufficiently consider DLT's lack of maturity and inherent privacy challenges.

The IoT's particular characteristics bring new challenges for privacy enhancement, most notably, the consequences of a lack of interoperability, computation and storage constraints, and the privacy disparity across jurisdictions. We have also observed the importance of *first* determining the sovereignty layer in data market design, as the participants' ownership and management rules impact the PETs in the rest of the layers. We also must underline that there is no "one-size-fits-all" PET. Only a combination may tackle the various privacy challenges facing data markets for the IoT. Lastly, we recommend that institutions invest resources in the research and adoption of PETs to remain competitive in the advent of a more privacy-enhancing IoT.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We would like to thank the Bayerisches Forschungsinstitut für Digitale Transformation for supporting our research on differential privacy, and the Bavarian Ministry of Economic Affairs, Regional Development and Energy for their funding of the project "Fraunhofer Blockchain Center (20-3066-2-6-14)" that made this paper possible.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jnca.2022.103465>.

Selected studies

- [S1] Y.N. Li, X. Feng, J. Xie, H. Feng, Z. Guan, Q. Wu, A decentralized and secure blockchain platform for open fair data trading, *Concurr. Comput.* 32 (7) (2019a) 1–11, <http://dx.doi.org/10.1002/cpe.5578>.
- [S2] N. Hynes, D. Dao, D. Yan, R. Cheng, D. Song, A demonstration of sterling: A privacy-preserving data marketplace, *Proc. VLDB Endow.* 11 (12) (2018) 2086–2089, <http://dx.doi.org/10.14778/3229863.3236266>.
- [S3] D. López, B. Farooq, A multi-layered blockchain framework for smart mobility data-markets, *Transp. Res. C* 111 (June 2019) (2020) 588–615, <http://dx.doi.org/10.1016/j.trc.2020.01.002>.
- [S4] F. Liang, W. Yu, D. An, Q. Yang, X. Fu, W. Zhao, A survey on big data market: Pricing, trading and protection, *IEEE Access* 6 (May) (2018) 15132–15154, <http://dx.doi.org/10.1109/ACCESS.2018.2806881>.
- [S5] D. Bogdanov, R. Jagomägis, S. Laur, A universal toolkit for cryptographically secure privacy-preserving data mining, in: *LNCs 7299 - Intelligence and Security Informatics*, Vol. 7299, 2012, URL <https://link.springer.com/content/pdf/10.1007/978-3-642-30428-6.pdf>.
- [S6] S. Spiekermann, A. Novotny, A vision for global privacy bridges: Technical and legal measures for international data markets, *Comput. Law Secur. Rev.* 31 (2) (2015) 181–200, <http://dx.doi.org/10.1016/j.clsr.2015.01.009>.
- [S7] C. Niu, Z. Zheng, F. Wu, X. Gao, G. Chen, Achieving data truthfulness and privacy preservation in data markets, *IEEE Trans. Knowl. Data Eng.* 31 (1) (2019) 105–119, <http://dx.doi.org/10.1109/TKDE.2018.2822727>, arXiv:1812.03280.
- [S8] E.M. Schomakers, C. Lidynia, M. Ziefle, All of me? Users' preferences for privacy-preserving data markets and the importance of anonymity, *Electron. Mark.* (2020) <http://dx.doi.org/10.1007/s12525-020-00404-9>.
- [S9] Y. Li, C. Miao, L. Su, J. Gao, Q. Li, B. Ding, Z. Qin, K. Ren, An efficient two-layer mechanism for privacy-preserving truth discovery, in: *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2018, pp. 1705–1714, <http://dx.doi.org/10.1145/3219819.3219998>.
- [S10] L. Zhou, L. Wang, T. Ai, Y. Sun, BeeKeeper 2.0: Confidential blockchain-enabled IoT system with fully homomorphic computation, *Sensors (Switzerland)* 18 (11) (2018a) <http://dx.doi.org/10.3390/s18113785>.
- [S11] A. Dorri, S.S. Kanhere, R. Jurdak, P. Gauravaram, Blockchain for IoT security and privacy: The case study of a smart home, in: *2017 IEEE International Conference on Pervasive Computing and Communications Workshops, PerCom Workshops 2017*, IEEE, 2017, pp. 618–623, <http://dx.doi.org/10.1109/PERCOMW.2017.7917634>.
- [S12] R. Li, T. Song, B. Mei, H. Li, X. Cheng, L. Sun, Blockchain for large-scale internet of things data storage and protection, *IEEE Trans. Serv. Comput.* 12 (5) (2019b) 762–771, <http://dx.doi.org/10.1109/TSC.2018.2853167>.
- [S13] J. Wei, M. Sabonuchi, R. Roche, Blockchain-enabled peer-to-peer data trading mechanism, in: *2018 IEEE International Conference on Internet of Things (IThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData)*, 2018, pp. 1349–1354, <http://dx.doi.org/10.1109/Cybermatics>.
- [S14] Z. Zheng, W. Mao, F. Wu, G. Chen, Challenges and opportunities in IoT data markets, in: *SocialSense 2019 - Proceedings of the 2019 4th International Workshop on Social Sensing*, 2019, pp. 1–2, <http://dx.doi.org/10.1145/3313294.3313378>.
- [S15] M.M. Khalili, X. Zhang, M. Liu, Contract design for purchasing private data using a biased differentially private algorithm, in: *Proceedings of NetEcon 2019: 14th Workshop on the Economics of Networks, Systems and Computation - in Conjunction with ACM EC 2019 and ACM SIGMETRICS 2019*, 2019, <http://dx.doi.org/10.1145/3338506.3340273>.
- [S16] L. Yang, M. Zhang, S. He, M. Li, J. Zhang, Crowd-empowered privacy-preserving data aggregation for mobile crowdsensing, in: *Proceedings of the International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, 2018, pp. 151–160, <http://dx.doi.org/10.1145/3209582.3209598>.

- [S17] K. Mišura, M. Žagar, Data marketplace for internet of things, in: Proceedings of 2016 International Conference on Smart Systems and Technologies, SST 2016, IEEE, 2016, pp. 255–260, <http://dx.doi.org/10.1109/SST.2016.7765669>.
- [S18] X. Zheng, Data trading with differential privacy in data market, ACM Int. Conf. Proc. Ser. (8) (2020) 112–115, <http://dx.doi.org/10.1145/3379247.3379271>.
- [S19] J. Pennekamp, M. Henze, S. Schmidt, P. Niemietz, M. Fey, D. Trauth, T. Bergs, C. Brecher, K. Wehrle, Dataflow challenges in an internet of production, in: ACM Workshop on Cyber-Physical Systems Security & Privacy (CPS-SPC'19), November 11, 2019, ACM, London, United Kingdom, 2019, pp. 27–38, <http://dx.doi.org/10.1145/3338499.3357357>.
- [S20] Z.J. Wang, C.H.V. Lin, Y.H. Yuan, C.C.J. Huang, Decentralized data marketplace to enable trusted machine economy, in: 2019 IEEE Eurasia Conference on IOT, Communication and Engineering, ECICE 2019, IEEE, 2019a, pp. 246–250, <http://dx.doi.org/10.1109/ECICE47484.2019.8942729>.
- [S21] M. Guerriero, D.A. Tamburri, E. Di Nitto, Defining, enforcing and checking privacy policies in data-intensive applications, in: Proceedings - International Conference on Software Engineering, 2018, pp. 172–182, <http://dx.doi.org/10.1145/3194133.3194140>.
- [S22] M. Shi, Y. Qiao, X. Wang, Differentially private auctions for private data crowdsourcing, in: Proceedings - 2019 IEEE Intl Conf on Parallel and Distributed Processing with Applications, Big Data and Cloud Computing, Sustainable Computing and Communications, Social Computing and Networking, ISPA/BDCloud/SustainCom/SocialCom 2019, IEEE, 2019, pp. 1–8, <http://dx.doi.org/10.1109/ISPA-BDCloud-SustainCom-SocialCom-48970.2019.00013>.
- [S23] J. Du, C. Jiang, E. Gelenbe, L. Xu, J. Li, Y. Ren, Distributed data privacy preservation in IoT applications, IEEE Wirel. Commun. 25 (December) (2018) 68–76, <http://dx.doi.org/10.1109/MWC.2017.1800094>.
- [S24] G. Gao, M. Xiao, J. Wu, S. Zhang, L. Huang, G. Xiao, DPDT: A differentially private crowd-sensed data trading mechanism, IEEE Internet Things J. 7 (1) (2020a) 751–762, <http://dx.doi.org/10.1109/JIOT.2019.2944107>.
- [S25] R. Cheng, F. Zhang, J. Kos, W. He, N. Hynes, N. Johnson, A. Juels, A. Miller, D. Song, Eکیدen: A platform for confidentiality-preserving, trustworthy, and performant smart contracts, in: Proceedings - 4th IEEE European Symposium on Security and Privacy, EURO S and P 2019, 2019, pp. 185–200, <http://dx.doi.org/10.1109/EuroSP.2019.00023>.
- [S26] T. Jung, X.Y. Li, Enabling privacy-preserving auctions in big data, in: Proceedings - IEEE INFOCOM, Vol. 2015-August (BigSecurity), IEEE, 2015, pp. 173–178, <http://dx.doi.org/10.1109/INFOCOMW.2015.7179380>, arXiv:1308.6202.
- [S27] C. Perera, R. Ranjan, L. Wang, End-to-end privacy for open big data markets, IEEE Cloud Comput. 2 (4) (2015b) 44–53, <http://dx.doi.org/10.1109/MCC.2015.78>.
- [S28] M. Zichichi, M. Contu, S. Ferretti, V. Rodríguez-Doncel, Ensuring personal data anonymity in data marketplaces through sensing-as-a-service and distributed ledger technologies, in: CEUR Workshop Proceedings, Vol. 2580, 2020, URL https://www.researchgate.net/publication/340183476_Ensuring_Personal_Data_Anonymity_in_Data_Marketplaces_through_Sensing-as-a-Service_and_Distributed_Ledger.
- [S29] S. Duri, M. Gruteser, X. Liu, P. Moskowitz, R. Perez, M. Singh, J.M. Tang, Framework for security and privacy in automotive telematics, in: Proceedings of the ACM International Workshop on Mobile Commerce, 2002, pp. 25–32, <http://dx.doi.org/10.1145/570709.570711>.
- [S30] P. Tzianos, G. Pipelidis, N. Tsiamitros, Hermes: An open and transparent marketplace for iot sensor data over distributed ledgers, in: ICBC 2019 - IEEE International Conference on Blockchain and Cryptocurrency, IEEE, 2019, pp. 167–170, <http://dx.doi.org/10.1109/BLOC.2019.8751331>.
- [S31] K. Li, L. Tian, W. Li, G. Luo, Z. Cai, Incorporating social interaction into three-party game towards privacy protection in iot, Computer Networks 150 (2019) 90–101, <http://dx.doi.org/10.1016/j.comnet.2018.11.036>.
- [S32] Y. Zhao, Y. Yu, Y. Li, G. Han, X. Du, Machine learning based privacy-preserving fair data trading in big data market, Inform. Sci. 478 (2019) 449–460, <http://dx.doi.org/10.1016/j.ins.2018.11.028>.
- [S33] S. Kiyomoto, M.S. Rahman, A. Basu, On blockchain-based anonymized dataset distribution platform, in: 2017 IEEE 15th International Conference on Software Engineering Research, Management and Applications (SERA), IEEE, 2017, pp. 85–92, URL <https://ieeexplore.ieee.org/document/7965711>.
- [S34] Z. Chen, L. Chen, L. Huang, H. Zhong, On privacy-preserving cloud auction, in: Proceedings of the IEEE Symposium on Reliable Distributed Systems, IEEE, 2016, pp. 279–288, <http://dx.doi.org/10.1109/SRDS.2016.045>.
- [S35] L. Pournajaf, D.A. Garcia-Ulloa, L. Xiong, V. Sunderam, Participant privacy in mobile crowd sensing task management, ACM SIGMOD Rec. 44 (4) (2016) 23–34, <http://dx.doi.org/10.1145/2935694.2935700>.
- [S36] D. Sánchez, A. Viejo, Personalized privacy in open data sharing scenarios, Online Inf. Rev. 41 (3) (2017) 298–310, <http://dx.doi.org/10.1108/OIR-01-2016-0011>.
- [S37] K. Jung, S. Park, Privacy bargaining with fairness: Privacy-price negotiation system for applying differential privacy in data market environments, in: 2019 IEEE International Conference on Big Data, IEEE, 2019, pp. 1389–1394, <http://dx.doi.org/10.1109/BigData47090.2019.9006101>.
- [S38] J.H. Ziegeldorf, O.G. Morchon, K. Wehrle, Privacy in the internet of things: Threats and challenges, Secur. Commun. Netw. 7 (12) (2014) 2728–2742, <http://dx.doi.org/10.1002/sec.795>.
- [S39] C. Perera, C. McCormick, A.K. Bandara, B.A. Price, B. Nuseibeh, Privacy-by-design framework for assessing internet of things applications and platforms, ACM Int. Conf. Proc. Ser. 07-09-Nove (2016b) 83–92, <http://dx.doi.org/10.1145/2991561.2991566>.
- [S40] C. Perera, C. Liu, R. Ranjan, L. Wang, A. Zomaya, Privacy-knowledge modeling for the internet of things: A look back, Computer 49 (12) (2016a) 60–68, <http://dx.doi.org/10.1109/MC.2016.366>.
- [S41] W. Gao, W. Yu, F. Liang, W.G. Hatcher, C. Lu, Privacy-preserving auction for big data trading using homomorphic encryption, IEEE Trans. Netw. Sci. Eng. 7 (2) (2020b) 776–791, <http://dx.doi.org/10.1109/TNSE.2018.2846736>.
- [S42] S. Park, K. Park, J. Lee, K. Jung, PRIVATA: Differentially private data market framework using negotiation-based pricing mechanism, in: Proceedings of ACM CIKM Conference (CIKM'19), November 3–7, 2019, Beijing, China, 2019, pp. 156–157, http://dx.doi.org/10.1007/978-3-663-10915-0_47.
- [S43] V. Koutsos, D. Papadopoulos, D. Chatzopoulos, S. Tarkoma, P. Hui, Agora: A privacy-aware data marketplace, 2020, p. 13, URL <https://eprint.iacr.org/2020/865.pdf>.
- [S44] J. Cao, P. Karras, Publishing microdata with a robust privacy guarantee, Proc. VLDB Endow. 5 (11) (2012) 1388–1399, <http://dx.doi.org/10.14778/2350229.2350255>, arXiv:1208.0220.
- [S45] W. Dai, C. Dai, K.K.R. Choo, C. Cui, D. Zou, H. Jin, SDTE: A secure blockchain-based data trading ecosystem, IEEE Trans. Inf. Forensics Secur. 15 (2020) 725–737, <http://dx.doi.org/10.1109/TIFS.2019.2928256>.
- [S46] Z. Guan, X. Shao, Z. Wan, Secure, fair and efficient data trading without third party using blockchain, in: 2018 IEEE International Conference on Internet of Things (Things) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData), IEEE, 2018, pp. 1349–1354, doi:10.1109/Cybermatics.
- [S47] M.A. Alsheikh, Y. Jiao, D. Niyato, P. Wang, D. Leong, Z. Han, The accuracy-privacy trade-off of mobile crowdsensing, IEEE Commun. Mag. 55 (6) (2017) 132–139, <http://dx.doi.org/10.1109/MCOM.2017.1600737>, arXiv:1702.04565.
- [S48] S. Sharma, K. Chen, A. Sheth, Toward practical privacy-preserving analytics for IoT and cloud-based healthcare systems, IEEE Internet Comput. 22 (2) (2018) 42–51, <http://dx.doi.org/10.1109/MIC.2018.112102519>.
- [S49] A. Colman, M.J.M. Chowdhury, M. Baruwat Chhetri, Toward a trusted marketplace for wearable data, in: Proceedings - 2019 IEEE 5th International Conference on Collaboration and Internet Computing, CIC 2019 (Cic), 2019, pp. 314–321, <http://dx.doi.org/10.1109/CIC48465.2019.00044>.
- [S50] Z. Cai, Z. He, Trading private range counting over big IoT data, in: Proceedings - International Conference on Distributed Computing Systems, Vol. 2019-July, IEEE, 2019, pp. 144–153, <http://dx.doi.org/10.1109/ICDCS.2019.00023>.

References

- [1] IDC, Open Evidence, European data market SMART, 2017, URL https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=44400.
- [2] A.R. Miller, C. Tucker, Health information exchange, system size and information silos, 2013, p. 29, URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1457719.
- [3] F. Stahl, F. Schomm, G. Vossen, L. Vomfell, A classification framework for data marketplaces, Vietnam J. Comput. Sci. 3 (3) (2016) 137–143, <http://dx.doi.org/10.1007/s40595-016-0064-2>.
- [4] McKinsey & Company, Four ways to accelerate the creation of data ecosystems, 2020, <https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/four-ways-to-accelerate-the-creation-of-data-ecosystems>.
- [5] G. Eggers, B. Fondermann, B. Maier, K. Ottradovetz, J. Pformmer, R. Reinhardt, H. Rollin, A. Schmieß, S. Steinbuß, P. Trinius, A. Weis, C. Weiss, S. Wilfling, GAIA-X: Technical architecture, 2020, URL https://www.data-infrastructure.eu/GAIA-X/Redaktion/EN/Publications/gaia-x-technical-architecture.pdf?_blob=publicationFile&v=5.
- [6] L. Sweeney, A. Abu, J. Winn, Identifying participants in the personal genome project by name, SSRN Electron. J. (2013) <http://dx.doi.org/10.2139/ssrn.2257732>.
- [7] X. Gao, B. Firner, S. Sugrim, V. Kaiser-Pendergrast, Y. Yang, J. Lindqvist, Elastic pathing: your speed is enough to track you, in: Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing - UbiComp '14 Adjunct, ACM Press, Seattle, Washington, 2014, pp. 975–986, <http://dx.doi.org/10.1145/2632048.2632077>.
- [8] A. Sunyaev, N. Kannengießer, R. Beck, H. Treiblmaier, M. Lacity, J. Kranz, G. Fridgen, U. Spankowski, A. Luckow, Token economy, Bus. Inf. Syst. Eng. (2021) URL <https://link.springer.com/article/10.1007/s12599-021-00684-1>.
- [9] IBM Security and Ponemon Institute LLC, 2018 Cost of a data breach study: Global overview, 2018, p. 47, URL https://www.intlxolutions.com/hubfs/2018_Global_Cost_of_a_Data_Breach_Report.pdf.
- [10] A. Trask, E. Bluemke, B. Garfinkel, C.G. Cuervas-Mons, A. Dafoe, Beyond privacy trade-offs with structured transparency, 2020, arXiv:2012.08347. URL https://www.researchgate.net/publication/347300876_Beyond_Privacy_Trade-offs_with_Structured_Transparency.

- [11] R. Hes, J.J. Borking, Netherlands, I.a.P. Commissioner/Ontario (Eds.), Privacy-enhancing Technologies: the Path to Anonymity, rev. ed., in: *Achtergrondstudies en Verkenningen*, (11) Registratiekamer, The Hague, 1998, URL https://www.researchgate.net/publication/243777645_Privacy-Enhancing_Technologies_The_Path_to_Anonymity.
- [12] C. Dwork, F. McSherry, K. Nissim, A. Smith, Calibrating noise to sensitivity in private data analysis, in: S. Halevi, T. Rabin (Eds.), *Theory of Cryptography*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2006, pp. 265–284, URL https://link.springer.com/chapter/10.1007/11681878_14. Online; accessed 30 December 2021.
- [13] C. Dwork, A. Roth, The algorithmic foundations of differential privacy, *Found. Trends® Theor. Comput. Sci.* 9 (3–4) (2013) 211–407, <http://dx.doi.org/10.1561/04000000042>.
- [14] P. Samarati, L. Sweeney, 1998. Protecting privacy when disclosing information: k-anonymity and its enforcement through generalization and suppression. 19. URL https://epic.org/privacy/reidentification/Samarati_Sweeney_paper.pdf.
- [15] M.A. Will, R.K. Ko, *A Guide to Homomorphic Encryption*, Elsevier Inc., 2015, p. 101, <http://dx.doi.org/10.1016/B978-0-12-801595-7.00005-7>.
- [16] P. Chaudhary, R. Gupta, A. Singh, P. Majumder, Analysis and comparison of various fully homomorphic encryption techniques, in: 2019 International Conference on Computing, Power and Communication Technologies, GUCON 2019, Galgotias University, 2019, pp. 58–62, URL <https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=8940577>.
- [17] P. Paillier, Public-key cryptosystems based on composite degree residuosity classes, *Eurocrypt* (1999) http://dx.doi.org/10.1007/3-540-48910-X_9.
- [18] OMTTP, Advanced trusted environment: OMTTP TR1, 2009, p. 204, URL <http://www.gsma.com/newsroom/wp-content/uploads/2012/03/omtpadvancedtrustedenvironmentomtptr1v11.pdf>.
- [19] A.C. Yao, Protocols for secure computations, in: 23rd Annual Symposium on Foundations of Computer Science (Sfcs 1982), IEEE, Chicago, IL, USA, 1982, pp. 160–164, <http://dx.doi.org/10.1109/SFCS.1982.38>.
- [20] S. Goldwasser, S. Micali, C. Rackoff, The knowledge complexity of interactive proof systems, *SIAM J. Comput.* 18 (1) (1989) 186–208, <http://dx.doi.org/10.1137/0218012>.
- [21] O. Goldreich, Y. Oren, Definitions and properties of zero-knowledge proof systems, *J. Cryptol.* 7 (1) (1994) 1–32, <http://dx.doi.org/10.1007/BF00195207>.
- [22] G. Bondel, G.M. Garrido, G. Baumer, F. Matthes, 2020. Towards a privacy-enhancing tool based on de-identification methods. 8. URL <https://aisel.aisnet.org/pacis2020/157/>.
- [23] S. Spiekermann, R. Böhme, A. Acquisti, K.-L. Hui, Personal data markets, *Electron. Mark.* 25 (2) (2015) 91–93, <http://dx.doi.org/10.1007/s12525-015-0190-1>.
- [24] P.B. Anne Zöll, C.M. Olt, Privacy-sensitive business models: Barriers of organizational adoption of privacy-enhancing technologies, 2021, p. 22, URL https://aisel.aisnet.org/ecis2021_rp/34/.
- [25] A.F. Westin, *Privacy and Freedom*, IG Publishing, New York, 1967, URL <https://scholarlycommons.law.wlu.edu/wlulr/vol25/iss1/20/>.
- [26] G.A. Fink, H. Song, S. Jeschke (Eds.), *Security and Privacy in Cyber-Physical Systems: Foundations, Principles, and Applications*, first ed., Wiley IEEE Press, Hoboken, NJ, 2018, URL <https://ieeexplore.ieee.org/servlet/opac?bknumber=8068866>.
- [27] K. Renaud, D. Galvez-Cruz, Privacy: Aspects, definitions and a multi-faceted privacy preservation approach, in: *Proceedings of the 2010 Information Security for South Africa Conference*, ISSA 2010, 2010, pp. 1–8, <http://dx.doi.org/10.1109/ISSA.2010.5588297>.
- [28] D.J. Solove, The meaning and value of privacy, in: B. Roessler, D. Mokrosinska (Eds.), *Social Dimensions of Privacy*, Cambridge University Press, Cambridge, 2015, pp. 71–82, <http://dx.doi.org/10.1017/CBO9781107280557.005>.
- [29] F.T. Wu, Defining privacy and utility in data sets, in: 84 University of Colorado Law Review 1117 (2013); 2012 TRPC, 2012, pp. 1117–1177, <http://dx.doi.org/10.2139/ssrn.2031808>.
- [30] M. Deng, K. Wuyts, R. Scandariato, B. Preneel, W. Joosen, A privacy threat analysis framework: supporting the elicitation and fulfillment of privacy requirements, *Requir. Eng.* 16 (1) (2011) 3–32, <http://dx.doi.org/10.1007/s00766-010-0115-7>.
- [31] R. Garratt, M.R.v. Oordt, Privacy as a public good: A case for electronic cash, *J. Polit. Econ.* (2018) <http://dx.doi.org/10.1086/714133>.
- [32] N. Kaaniche, M. Laurent, Attribute-based signatures for supporting anonymous certification, in: I. Askoxylakis, S. Ioannidis, S. Katsikas, C. Meadows (Eds.), *Computer Security – ESORICS 2016*, Springer International Publishing, Cham, 2016, pp. 279–300, URL <https://www.semanticscholar.org/paper/Attribute-Based-Signatures-for-Supporting-Anonymous-Kaaniche-Laurent-Maknavicu/3b0624ff32b9258ca2351c894d320d83a546fcd6>.
- [33] J.E. Campbell, M. Carlson, Panopticon.com: Online surveillance and the commodification of privacy, *J. Broadcast. Electron. Media* 46 (4) (2002) 586–606, http://dx.doi.org/10.1207/s15506878jobem4604_6.
- [34] A. Lichter, M. Löffler, S. Siegloch, The long-term costs of government surveillance: Insights from stasi spying in east Germany, *J. Eur. Econom. Assoc.* 19 (2) (2020) 741–789, <http://dx.doi.org/10.1093/jeea/jvaa009>, arXiv:https://academic.oup.com/jeea/article-pdf/19/2/741/37108669/jvaa009.pdf.
- [35] S. Kokolakis, Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon, *Comput. Secur.* 64 (2017) 122–134, <http://dx.doi.org/10.1016/j.cose.2015.07.002>.
- [36] J. Coppel, E-Commerce: Impacts and Policy Challenges, OECD Economics Department Working Papers 252, 2000, <http://dx.doi.org/10.1787/801315684632>, Series: OECD Economics Department Working Papers Volume: 252.
- [37] J. Kennedy, Big data's economic impact, 2021, [Online]. Available: <https://www.ced.org/blog/entry/big-datas-economic-impact>, [Accessed on 04 Jul. 2021].
- [38] A.M. Oberländer, M. Röglinger, M. Rosemann, A. Kees, Conceptualizing business-to-things interactions – a sociomaterial perspective on the internet of things, *Eur. J. Inf. Syst.* 27 (4) (2018) 486–502, <http://dx.doi.org/10.1080/0960085X.2017.1387714>.
- [39] I. Lee, K. Lee, The internet of things (IoT): Applications, investments, and challenges for enterprises, *Bus. Horiz.* 58 (4) (2015) 431–440, <http://dx.doi.org/10.1016/j.bushor.2015.03.008>.
- [40] V. Basili, G. Caldiera, D. Rombach, The goal question metric approach, *Encycl. Softw. Eng.* (1994) 528–532, URL <http://www.cs.toronto.edu/~sme/CSC444F/handouts/GQM-paper.pdf>.
- [41] B.A. Kitchenham, D. Budgen, Evidence-Based Software Engineering and Systematic Reviews, Chapman and Hall/CRC, 2015, URL <https://dl.acm.org/doi/book/10.5555/2994449>.
- [42] B. Kitchenham, Procedures for Performing Systematic Reviews, Joint Technical Report, 2004, <http://dx.doi.org/10.5144/0256-4947.2017.79>.
- [43] D.C.B. Mariano, C. Leite, L.H.S. Santos, R.E.O. Rocha, R.C. de Melo-Minardi, A guide to performing systematic literature reviews in bioinformatics, 2017, arXiv:1707.05813.
- [44] T. Dybå, T. Dingsoyr, G. Hanssen, Applying systematic reviews to diverse study types: An experience report, in: *Proceedings - 1st International Symposium on Empirical Software Engineering and Measurement, ESEM 2007* (7465), 2007, pp. 126–135, <http://dx.doi.org/10.1109/ESEM.2007.59>.
- [45] O. Dieste, A. Grimán, N. Juristo, Developing search strategies for detecting relevant experiments, *Empir. Softw. Eng.* 14 (5) (2009) 513–539, <http://dx.doi.org/10.1007/s10664-008-9091-7>.
- [46] H. Zhang, M.A. Babar, P. Tell, Identifying relevant studies in software engineering, *Inf. Softw. Technol.* 53 (6) (2011) 625–637, <http://dx.doi.org/10.1016/j.infsof.2010.12.010>.
- [47] A. Kilgarriř, V. Baisa, J. Buřta, M. Jakubiřek, V. Kovář, J. Michelfeit, P. Rychlý, V. Suchomel, The sketch engine: ten years on, *Lexicography* (2014) URL https://www.researchgate.net/publication/271848017_The_Sketch_Engine_Ten_Years_On.
- [48] O.P. Brereton, B.A. Kitchenham, D. Budgen, M. Turner, M. Khalil, Lessons from applying the systematic literature review process within the software engineering domain, *J. Syst. Softw.* 80 (4) (2007) 571–583, <http://dx.doi.org/10.1016/j.jss.2006.07.009>.
- [49] B.A. Kitchenham, O.P. Brereton, A systematic review of systematic review process research in software engineering, *Inf. Softw. Technol.* 55 (12) (2013) 2049–2075, <http://dx.doi.org/10.1016/j.infsof.2013.07.010>.
- [50] L. Chen, M.A. Babar, H. Zhang, Towards an evidence-based understanding of electronic data sources (january 2015), 2010, <http://dx.doi.org/10.14236/ewic/ease2010.17>.
- [51] C. Wohlin, P. Runeson, M. Höst, M.C. Ohlsson, B. Regnell, A. Wesslén, *Experimentation in Software Engineering*, Springer Science & Business Media, 2012, <http://dx.doi.org/10.1007/978-3-642-29044-2>.
- [52] H. Nissenbaum, *Privacy in Context: Technology, Policy, and the Integrity of Social Life*, Stanford University Press, 2009, URL <https://www.sup.org/books/title/?id=8862>.
- [53] R.Y. Wang, D.M. Strong, Beyond accuracy: What data quality means to data consumers, *J. Manage. Inf. Syst.* 12 (4) (1996) 5–33, <http://dx.doi.org/10.1080/07421222.1996.11518099>.
- [54] T. Dinev, H. Xu, J.H. Smith, P. Hart, Information privacy and correlates: an empirical attempt to bridge and distinguish privacy-related concepts, *Eur. J. Inf. Syst.* 22 (3) (2013) 295–316, <http://dx.doi.org/10.1057/ejis.2012.23>.
- [55] B.-J. Butijn, D.A. Tamburri, W.-J.v.d. Heuvel, Blockchains: a systematic multivocal literature review, *ACM Comput. Surv.* 53 (3) (2020) 1–37, URL <https://dl.acm.org/doi/abs/10.1145/3369052>.
- [56] R. Zhang, R. Xue, L. Liu, Security and privacy on blockchain, *ACM Comput. Surv.* 52 (3) (2019) 1–34, URL <https://dl.acm.org/doi/10.1145/3316481>.
- [57] G.I. Simari, A primer on zero knowledge protocols, 2002, p. 12, URL <http://cs.uns.edu.ar/~gjs/publications/zkp-simari2002.pdf>.
- [58] N. Kaaniche, M. Laurent, S. Belguith, Privacy enhancing technologies for solving the privacy-personalization paradox: Taxonomy and survey, *J. Netw. Comput. Appl.* (2020).
- [59] D. Chaum, Security without identification: transaction systems to make big brother obsolete, *Commun. ACM* 28 (10) (1985) 1030–1044, <http://dx.doi.org/10.1145/4372.4373>.
- [60] J.L. Camenisch, J.-M. Piveteau, M.A. Stadler, Blind signatures based on the discrete logarithm problem, in: A. De Santis (Ed.), *Advances in Cryptology — EUROCRYPT'94*, Springer Berlin Heidelberg, Berlin, Heidelberg, 1995, pp. 428–432, URL <https://link.springer.com/chapter/10.1007/BFb0053458>.

- [61] J. Camenisch, A. Lysyanskaya, Dynamic accumulators and application to efficient revocation of anonymous credentials, in: G. Goos, J. Hartmanis, J. van Leeuwen, M. Yung (Eds.), *Advances in Cryptology — CRYPTO 2002*, Vol. 2442, in: *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2002, pp. 61–76, http://dx.doi.org/10.1007/3-540-45708-9_5.
- [62] J. Camenisch, T. Groß, Efficient attributes for anonymous credentials, 2010, p. 29, URL <https://eprint.iacr.org/2010/496.pdf>.
- [63] S.A. Brands, *Rethinking Public Key Infrastructures and Digital Certificates: Building in Privacy*, MIT Press, Cambridge, MA, USA, 2000, URL <https://direct.mit.edu/books/book/1912/Rethinking-Public-Key-Infrastructures-and-Digital>.
- [64] J. Sedlmeir, R. Smethurst, A. Rieger, G. Fridgen, Digital identities and verifiable credentials, *Bus. Inf. Syst. Eng.* 63 (5) (2021) 603–613.
- [65] V. Schlatt, J. Sedlmeir, S. Feulner, N. Urbach, Designing a framework for digital KYC processes built on blockchain-based self-sovereign identity, *Inf. Manage.* (2021) 103553.
- [66] E. Bangert, S. Barzan, S. Krenn, A.-R. Sadeghi, T. Schneider, J.-K. Tsay, Bringing zero-knowledge proofs of knowledge to practice, 2009, p. 12, URL <https://eprint.iacr.org/2009/211.pdf>.
- [67] M. Hoffmann, M. Kloof, A. Rupp, Efficient zero-knowledge arguments in the discrete log setting, revisited, in: *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, ACM, London United Kingdom, 2019, pp. 2093–2110, <http://dx.doi.org/10.1145/3319535.3354251>.
- [68] T. Nakanishi, H. Yoshino, T. Murakami, G.-V. Policharla, Efficient zero-knowledge proofs of graph signature for connectivity and isolation using bilinear-map accumulator, in: *Proceedings of the 7th ACM Workshop on Asia Public-Key Cryptography*, ACM, Taipei Taiwan, 2020, pp. 9–18, <http://dx.doi.org/10.1145/3384940.3388959>.
- [69] Y. Zhang, Zero-knowledge proofs for machine learning, in: *Proceedings of the 2020 Workshop on Privacy-Preserving Machine Learning in Practice*, Association for Computing Machinery, 2020, p. 7, URL <https://doi.org/10.1145/3411501.3418608>.
- [70] M. Stadler, Publicly verifiable secret sharing, in: *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Vol. 1070, 1996, pp. 190–199, http://dx.doi.org/10.1007/3-540-68339-9_17.
- [71] A. Shamir, How to share a secret, *Commun. ACM* 22 (11) (1979a) 612–613, <http://dx.doi.org/10.1145/359168.359176>.
- [72] Y. Lindell, B. Pinkas, A proof of security of Yao's protocol for two-party computation, *J. Cryptol.* 22 (2) (2009) 161–188, <http://dx.doi.org/10.1007/s00145-008-9036-8>.
- [73] A. Ben-David, N. Nisan, B. Pinkas, Fairplaymp: a system for secure multi-party computation, in: *Proceedings of the 15th ACM Conference on Computer and Communications Security - CCS '08*, ACM Press, Alexandria, Virginia, USA, 2008, p. 257, <http://dx.doi.org/10.1145/1455770.1455804>.
- [74] S. Yakoubov, A gentle introduction to Yao's garbled circuits, 2017, URL <https://web.mit.edu/sonka89/www/papers/2017ygc.pdf>.
- [75] Z.A. Genç, V. Iovino, A. Rial, The simplest protocol for oblivious transfer, *Inform. Process. Lett.* 161 (2020) 1–12, <http://dx.doi.org/10.1016/j.ipl.2020.105975>.
- [76] P. Pullonen, S. Siim, Combining secret sharing and garbled circuits for efficient private IEEE 754 floating-point computations, in: M. Brenner, N. Christin, B. Johnson, K. Rohloff (Eds.), *Financial Cryptography and Data Security*, Vol. vol. 8976, in: *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2015, pp. 172–183, http://dx.doi.org/10.1007/978-3-662-48051-9_13.
- [77] Y. Yang, X. Huang, X. Liu, H. Cheng, J. Weng, X. Luo, V. Chang, A comprehensive survey on secure outsourced computation and its applications, *IEEE Access* 7 (2019a) 159426–159465, URL <https://ieeexplore.ieee.org/document/8884162/>.
- [78] E. Boyle, N. Gilboa, Y. Ishai, A. Nof, Practical fully secure three-party computation via sublinear distributed zero-knowledge proofs, in: *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, ACM, London United Kingdom, 2019, pp. 869–886, URL <https://dl.acm.org/doi/10.1145/3319535.3363227>.
- [79] D. Boneh, E.-J. Goh, K. Nissim, Evaluating 2-DNF formulas on ciphertexts, 3378, 2005, pp. 325–341, URL https://www.researchgate.net/publication/221354138_Evaluating_2-DNF_Formulas_on_Ciphertexts,
- [80] V. Nikolaenko, S. Ioannidis, U. Weinsberg, M. Joye, N. Taft, D. Boneh, Privacy-preserving matrix factorization, in: *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security*, Association for Computing Machinery, 2013, pp. 801–812, URL <https://doi.org/10.1145/2508859.2516751>.
- [81] L. Zhou, L. Wang, Y. Sun, T. Ai, AntNest: Fully non-interactive secure multi-party computation, *IEEE Access* 6 (2018b) 75639–75649, URL <https://ieeexplore.ieee.org/document/8550709/>.
- [82] D. Boneh, A. Sahai, B. Waters, Functional encryption: Definitions and challenges, in: Y. Ishai (Ed.), *Theory of Cryptography*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2011, pp. 253–273, URL <https://eprint.iacr.org/2010/543.pdf>.
- [83] J. Chotard, E. Dufour Sans, R. Gay, D.H. Phan, D. Pointcheval, Decentralized multi-client functional encryption for inner product, in: T. Peyrin, S. Galbraith (Eds.), *Advances in Cryptology – ASIACRYPT 2018*, Springer International Publishing, Cham, 2018, pp. 703–732, URL <https://eprint.iacr.org/2017/989.pdf>.
- [84] Z. Brakerski, C. Gentry, V. Vaikuntanathan, (Leveled) fully homomorphic encryption without bootstrapping, *ACM Trans. Comput. Theory* 6 (3) (2014) <http://dx.doi.org/10.1145/2633600>.
- [85] W. Wang, Y. Hu, L. Chen, X. Huang, B. Sunar, Exploring the feasibility of fully homomorphic encryption, *IEEE Trans. Comput.* 64 (3) (2015) 698–706, <http://dx.doi.org/10.1109/TC.2013.154>.
- [86] I. Anati, S. Gueron, S.P. Johnson, V.R. Scarlata, Innovative technology for CPU based attestation and sealing, 2013, p. 7, URL <https://software.intel.com/content/dam/develop/external/us/en/documents/hasp-2013-innovative-technology-for-attestation-and-sealing-413939.pdf>.
- [87] F. Khalid, A. Masood, Vulnerability analysis of qualcomm secure execution environment, *Comput. Secur.* 116 (2022) 102628, <http://dx.doi.org/10.1016/j.cose.2022.102628>, URL <https://www.sciencedirect.com/science/article/pii/S016740482200027X>.
- [88] F. Alder, J. Van Bulck, J. Spielman, D. Oswald, F. Piessens, Faulty point unit: ABI poisoning attacks on trusted execution environments, *Digit. Threats Res. Pract.* 3 (2) (2022) <http://dx.doi.org/10.1145/3491264>, URL <https://doi-access.uib.tum.de/10.1145/3491264>.
- [89] D. Skarlatos, M. Yan, B. Gopireddy, R. Sprabery, J. Torrellas, C.W. Fletcher, MicroScope: Enabling microarchitectural replay attacks, in: *Proceedings of the 46th International Symposium on Computer Architecture*, ACM, 2019, pp. 318–331, <http://dx.doi.org/10.1145/3307650.3322228>.
- [90] Intel, 12th generation intel core processors, 2022, URL <https://cdrdv2.intel.com/v1/dl/getContent/655258>.
- [91] V. Costan, I. Lebedev, S. Devadas, Sanctum: Minimal hardware extensions for strong software isolation, 2016, p. 19, URL <https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/costan>.
- [92] D. Lee, D. Kohlbrenner, S. Shinde, K. Asanović, D. Song, Keystone: an open framework for architecting trusted execution environments, in: *Proceedings of the Fifteenth European Conference on Computer Systems*, ACM, Heraklion Greece, 2020, pp. 1–16, <http://dx.doi.org/10.1145/3342195.3387532>.
- [93] AWS, Nitro system, 2022, Online; accessed 4 May 2022. URL <https://aws.amazon.com/ec2/nitro/>.
- [94] W. Wei, L. Liu, M. Loper, K.-H. Chow, M.E. Gursoy, S. Truex, Y. Wu, A framework for evaluating gradient leakage attacks in federated learning, 2020, [arXiv:2004.10397](https://arxiv.org/abs/2004.10397). URL <https://www.semanticscholar.org/paper/A-Framework-for-Evaluating-Gradient-Leakage-Attacks-Wei-Liu/9853a348f61aec83b410f307ab905a4ae001fcd4>.
- [95] S. Pinto, N. Santos, Demystifying ARM TrustZone: A comprehensive survey, *ACM Comput. Surv.* 51 (6) (2019).
- [96] V. Costan, L.F. Sarmiento, M.v. Dijk, S. Devadas, The trusted execution module: Commodity general-purpose trusted computing, in: *International Conference on Smart Card Research and Advanced Applications*, Springer, 2008, pp. 133–148.
- [97] J. Konečný, B. McMahan, D. Ramage, Federated optimization: Distributed optimization beyond the datacenter, 2015, [arXiv:1511.03575](https://arxiv.org/abs/1511.03575). URL <https://docplayer.net/15450695-Federated-optimization-distributed-optimization-beyond-the-datacenter.html>.
- [98] T. Li, A.K. Sahu, A. Talwalkar, V. Smith, Federated learning: Challenges, methods, and future directions, *IEEE Signal Process. Mag.* 37 (3) (2020) 50–60, <http://dx.doi.org/10.1109/MSP.2020.2975749>.
- [99] Q. Yang, Y. Liu, Y. Cheng, Y. Kang, T. Chen, H. Yu, Federated learning, *Synth. Lect. Artif. Intell. Mach. Learn.* 13 (3) (2019b) 1–207, <http://dx.doi.org/10.2200/S00960ED2V01Y201910AIM043>.
- [100] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H.B. McMahan, S. Patel, D. Ramage, A. Segal, K. Seth, Practical secure aggregation for privacy-preserving machine learning, in: *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, in: *CCS '17*, Association for Computing Machinery, New York, NY, USA, 2017, pp. 1175–1191, <http://dx.doi.org/10.1145/3133956.3133982>.
- [101] S. Zhang, Z. Li, Q. Chen, W. Zheng, J. Leng, M. Guo, Dubhe: Towards data unbiasedness with homomorphic encryption in federated learning client selection, in: *50th International Conference on Parallel Processing*, in: *ICPP 2021*, Association for Computing Machinery, New York, NY, USA, 2021a, <http://dx.doi.org/10.1145/3472456.3473513>, URL <https://doi-org.eaccess.uib.tum.de/10.1145/3472456.3473513>.
- [102] W. Yang, B. Liu, C. Lu, N. Yu, Privacy Preserving on Updated Parameters in Federated Learning, *ACM TURC'20*, Association for Computing Machinery, New York, NY, USA, 2020, pp. 27–31, <http://dx.doi.org/10.1145/3393527.3393533>.
- [103] P. Vepakomma, O. Gupta, A. Gupta, T. Swedish, R. Raskar, Split learning for health: Distributed deep learning without sharing raw patient data, 2018, URL https://aiforsocialgood.github.io/iclr2019/accepted/track1/pdfs/31_aig_iclr2019.pdf.
- [104] O. Gupta, R. Raskar, Distributed learning of deep neural network over multiple agents, *J. Netw. Comput. Appl.* 116 (2018) 1–8, <http://dx.doi.org/10.1016/j.jnca.2018.05.003>.

- [105] M.G. Poirot, P. Vepakomma, K. Chang, J. Kalpathy-Cramer, R. Gupta, R. Raskar, 2019. Split learning for collaborative deep learning in healthcare. 9. URL <https://arxiv.org/abs/1912.12115>.
- [106] L. Giarretta, S. Girdzijauskas, Gossip learning: Off the beaten path, in: 2019 IEEE International Conference on Big Data (Big Data), 2019, pp. 1117–1124, <http://dx.doi.org/10.1109/BigData47090.2019.9006216>.
- [107] R. Ormándi, I. Hegedűs, M. Jelasity, 2012. Gossip learning with linear models on fully distributed data: Efficient p2p ensemble learning with linear models on fully distributed data. 25 (4), 556–571. <http://dx.doi.org/10.1002/cpe.2858>. URL <https://onlinelibrary.wiley.com/doi/10.1002/cpe.2858>.
- [108] M. Abadi, A. Chu, I. Goodfellow, H.B. McMahan, I. Mironov, K. Talwar, L. Zhang, Deep learning with differential privacy, in: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, ACM, 2016, <http://dx.doi.org/10.1145/2976749.2978318>.
- [109] J. Domingo-Ferrer, D. Sánchez, A. Blanco-Justicia, The limits of differential privacy (and its misuse in data release and machine learning), *Commun. ACM* 64 (7) (2021) 33–35, <http://dx.doi.org/10.1145/3433638>, URL <https://doi.org/eaccess.um.tum.de/10.1145/3433638>.
- [110] E. Dikici, L.M. Prevedello, M. Bigelow, R.D. White, B.S. Erdal, Constrained generative adversarial network ensembles for sharable synthetic data generation, 2020, arXiv:2003.00086. URL https://www.researchgate.net/publication/339642358_Constrained_Generative_Adversarial_Network_Ensembles_for_Sharable_Synthetic_Data_Generation.
- [111] A. Torfi, E.A. Fox, C.K. Reddy, Differentially private synthetic medical data generation using convolutional GANs, 2020, URL https://www.researchgate.net/publication/347624671_Differentially_Private_Synthetic_Medical_Data_Generation_using_Convolutional_GANs.
- [112] 2010. Nin, J., Herranz, J. (Eds.), Privacy and Anonymity in Information Management Systems. In: Advanced Information and Knowledge Processing, Springer, London. <http://dx.doi.org/10.1007/978-1-84996-238-4>. URL <http://link.springer.com/10.1007/978-1-84996-238-4>.
- [113] V. Puri, S. Sachdeva, P. Kaur, Privacy preserving publication of relational and transaction data: Survey on the anonymization of patient data, *Comput. Sci. Rev.* 32 (C) (2019) 45–61, <http://dx.doi.org/10.1016/j.cosrev.2019.02.001>.
- [114] B.C.M. Fung, K. Wang, R. Chen, P.S. Yu, Privacy-preserving data publishing: A survey of recent developments, *ACM Comput. Surv.* 42 (4) (2010) <http://dx.doi.org/10.1145/1749603.1749605>.
- [115] C.C. Aggarwal, P.S. Yu (Eds.), Privacy-Preserving Data Mining - Models and Algorithms, in: Advances in Database Systems, vol. 34, Springer, 2008, <http://dx.doi.org/10.1007/978-0-387-70992-5>.
- [116] P. Ram Mohan Rao, S. Murali Krishna, A.P. Siva Kumar, 2018. Privacy preservation techniques in big data analytics: a survey. 5 (1), 33. <http://dx.doi.org/10.1186/s40537-018-0141-8>. URL <https://journalofbigdata.springeropen.com/articles/10.1186/s40537-018-0141-8>.
- [117] M. Cunha, R. Mendes, J. ao P. Vilela, A survey of privacy-preserving mechanisms for heterogeneous data types, *Comp. Sci. Rev.* 41 (2021) 100403, <http://dx.doi.org/10.1016/j.cosrev.2021.100403>, URL <https://www.sciencedirect.com/science/article/pii/S1574013721000435>.
- [118] C. Dwork, A. Smith, T. Steinke, J. Ullman, Exposed! A survey of attacks on private data, *Annu. Rev. Stat. Appl.* 4 (1) (2017) 61–84, <http://dx.doi.org/10.1146/annurev-statistics-060116-054123>.
- [119] ISO, Privacy enhancing data de-identification terminology and classification of techniques, 2018, URL <https://www.iso.org/standard/69373.html>.
- [120] N. Li, W. Qardaji, D. Su, On sampling, anonymization, and differential privacy or, k-anonymization meets differential privacy, in: Proceedings of the 7th ACM Symposium on Information, Computer and Communications Security, in: ASIACCS '12, Association for Computing Machinery, New York, NY, USA, 2012, pp. 32–33, <http://dx.doi.org/10.1145/2414456.2414474>.
- [121] H. Xu, S. Guo, K. Chen, Building confidential and efficient query services in the cloud with RASP data perturbation, 2013, <http://dx.doi.org/10.1109/TKDE.2012.251>.
- [122] K. Chen, L. Liu, Geometric data perturbation for privacy preserving outsourced data mining, *Knowl. Inf. Syst.* 29 (3) (2011) 657–695, URL <http://link.springer.com/10.1007/s10115-010-0362-4>.
- [123] R. Henry, A. Herzberg, A. Kate, Blockchain access privacy: Challenges and directions, *IEEE Secur. Priv.* 16 (4) (2018) 38–45, <http://dx.doi.org/10.1109/MSP.2018.3111245>.
- [124] D. Chaum, Untraceable electronic mail, return addresses, and digital pseudonyms, *Commun. ACM* 24 (2) (1981) 84–90.
- [125] D. Chaum, The dining cryptographers problem: Unconditional sender and recipient untraceability, *J. Cryptol.* 1 (1) (1988) 65–75.
- [126] J. Ren, J. Wu, Survey on anonymous communications in computer networks, *Comput. Commun.* 33 (4) (2010) 420–431.
- [127] M.S. Ali, K. Dolui, F. Antonelli, IoT data privacy via blockchains and IPFS, in: Proceedings of the Seventh International Conference on the Internet of Things, in: IoT '17, Association for Computing Machinery, New York, NY, USA, 2017, <http://dx.doi.org/10.1145/3131542.3131563>.
- [128] M. Kesarwani, A. Kaul, S. Braghin, N. Holohan, S. Antonatos, Secure k-anonymization over encrypted databases, in: 2021 IEEE 14th International Conference on Cloud Computing (CLOUD), 2021, pp. 20–30, <http://dx.doi.org/10.1109/CLOUD53861.2021.00015>.
- [129] A.F. Westin, Privacy and freedom, 1970, URL <https://www.worldcat.org/title/privacy-and-freedom/oclc/792862>.
- [130] A. Raj, J. Bosch, H.H. Olsson, T.J. Wang, Modelling data pipelines, in: 2020 46th EuroMicro Conference on Software Engineering and Advanced Applications (SEAA), 2020, pp. 13–20, <http://dx.doi.org/10.1109/SEAA51224.2020.00014>.
- [131] A. Cavoukian, The 7 foundational principles, 2011, p. 2, URL <https://sites.psu.edu/digitalshred/2020/11/13/privacy-by-design-pbd-the-7-foundational-principles-cavoukian/>.
- [132] J. Sedlmeir, J. Lautenschlager, G. Fridgen, N. Urbach, The transparency challenge of blockchain in organizations, *Electron. Mark.* (2022) <http://dx.doi.org/10.1007/s12525-022-00536-0>.
- [133] E. Heilman, N. Narula, G. Tanzer, J. Lovejoy, M. Colavita, M. Virza, T. Dryja, Cryptanalysis of curl-p and other attacks on the IOTA cryptocurrency, *IACR Trans. Symm. Crypt.* (2020) 367–391, <http://dx.doi.org/10.46586/tosc.v2020.i3.367-391>.
- [134] D. Wang, J. Zhao, Y. Wang, A survey on privacy protection of blockchain: The technology and application, *IEEE Access* 8 (2020) 108766–108781, <http://dx.doi.org/10.1109/ACCESS.2020.2994294>, URL <https://ieeexplore.ieee.org/document/9093015/>.
- [135] J. Sedlmeir, H.U. Buhl, G. Fridgen, R. Keller, The energy consumption of blockchain technology: beyond myth, *Bus. Inf. Syst. Eng.* 62 (6) (2020) 599–608, URL https://www.researchgate.net/publication/342313238_The_Energy_Consumption_of_Blockchain_Technology_Beyond_Myth.
- [136] N. Kannengießer, S. Lins, T. Dehling, A. Sunyaev, Trade-offs between distributed ledger technology characteristics, *ACM Comput. Surv.* 53 (2) (2020) 1–37, URL <https://dl.acm.org/doi/10.1145/3379463>.
- [137] European Parliament, European Parliament and Council of the European Union, Regulation (EU) 2016/679 directive 95/46/EC (general data protection regulation): General data protection regulation, *Off. J. Eur. Union* (L119) (2016) 1–88, URL <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32016R0679>.
- [138] S. Ellis, A. Juels, S. Nazarov, Chainlink a decentralized oracle network, 2021, URL <https://link.smartcontract.com/whitepaper>.
- [139] S.S. Al-Riyami, K.G. Paterson, Certificateless public key cryptography, in: Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), Vol. 2894, 2003, pp. 452–473, URL <https://eprint.iacr.org/2003/126.pdf>.
- [140] D. Reed, M. Sporny, D. Longley, C. Allen, A. Grant, M. Sabadello, Decentralized identifiers (dids) v1.0, 2021, URL <https://w3c.github.io/did-core/>.
- [141] L. Mui, M. Mohtashemi, A. Halberstadt, A computational model of trust and reputation, in: Proceedings of the 35th Hawaii International Conference on System Sciences, IEEE, 2002, pp. 2431–2439.
- [142] T. Grandison, M. Sloman, A survey of trust in internet applications, *IEEE Commun. Surv. Tutor.* 3 (4) (2000) 2–16.
- [143] D. Artz, Y. Gil, A survey of trust in computer science and the semantic web, *J. Web Semant.* 5 (2) (2007) 58–71.
- [144] K.S. Cook, R. Hardin, M. Levi, Cooperation Without Trust? Russell Sage Foundation, 2005.
- [145] L. Ismail, H. Hameed, M. AlShamsi, M. AlHammadi, N. AlDhanhani, Towards a blockchain deployment at UAE university: Performance evaluation and blockchain taxonomy, in: Proceedings of the 2019 International Conference on Blockchain Technology, ACM, Honolulu HI USA, 2019, pp. 30–38, <http://dx.doi.org/10.1145/3320154.3320156>.
- [146] L. Lamport, R. Shostak, M. Pease, The Byzantine generals problem, in: Concurrency: The Works of Leslie Lamport, Association for Computing Machinery, New York, NY, USA, 2019, pp. 203–226, <http://dx.doi.org/10.1145/3335772.3335936>.
- [147] C. Perera, C.H. Liu, S. Jayawardena, The emerging internet of things marketplace from an industrial perspective: A survey, *IEEE Trans. Emerg. Top. Comput.* 3 (4) (2015a) 585–598, URL <https://ieeexplore.ieee.org/document/7004800>.
- [148] L. Determann, No one owns data, *UC Hast. Law Rev.* 70 (2018) 44, <http://dx.doi.org/10.2139/ssrn.3123957>.
- [149] J.H. Nord, A. Koohang, J. Paliszkievicz, The internet of things: Review and theoretical framework, *Expert Syst. Appl.* 133 (2019) 97–108, <http://dx.doi.org/10.1016/j.eswa.2019.05.014>, URL <https://www.sciencedirect.com/science/article/pii/S0957417419303331>.
- [150] S. Arumugam, R. Bhargavi, 2019. A survey on driving behavior analysis in usage based insurance using big data. 6 (1), 86. <http://dx.doi.org/10.1186/s40537-019-0249-5>. URL <https://journalofbigdata.springeropen.com/articles/10.1186/s40537-019-0249-5>.
- [151] D.E. Pozen, 2005. The mosaic theory, national security, and the freedom of information act. 52.
- [152] M. Archie, S. Gershon, A. Katcoff, A. Zeng, Who's watching? De-anonymization of netflix reviews using amazon reviews, 2018b, Online; accessed 30 December 2021. URL <https://www.oasislabs.com/how-it-works>.
- [153] L.M. Zagi, B. Aziz, Privacy attack on IoT: A systematic literature review, in: 7th International Conference on ICT for Smart Society: AIoT for Smart Society, ICIS 2020 - Proceeding, Institute of Electrical and Electronics Engineers Inc., 2020, URL <https://ieeexplore.ieee.org/document/9307568>.

- [154] D. Kondor, B. Hashemian, Y.-A. de Montjoye, C. Ratti, Towards matching user mobility traces in large-scale datasets, *IEEE Trans. Big Data* 6 (4) (2020) 714–726, <http://dx.doi.org/10.1109/TBDA.2018.2871693>.
- [155] A. Wood, M. Altman, A. Bembek, M. Bun, M. Gaboardi, J. Honaker, K. Nissim, D. O'Brien, T. Steinke, S. Vadhan, Differential privacy: A primer for a non-technical audience, *SSRN Electron. J.* (2018) <http://dx.doi.org/10.2139/ssrn.3338027>.
- [156] A. Narayanan, V. Shmatikov, Robust de-anonymization of large sparse datasets, in: 2008 IEEE Symposium on Security and Privacy (Sp 2008), IEEE, Oakland, CA, USA, (ISSN: 1081-6011) 2008, pp. 111–125, <http://dx.doi.org/10.1109/SP.2008.33>.
- [157] M. Archie, S. Gershon, A. Katcoff, A. Zeng, De-anonymization of netflix reviews using amazon reviews, 2018a, p. 5, URL <https://www.readkong.com/page/de-anonymization-of-netflix-reviews-using-amazon-reviews-1439089>.
- [158] N. Lomas, France fines google \$120M and amazon \$42M for dropping tracking cookies without consent, 2020, URL <https://dataprotection.news/france-fines-google-120m-and-amazon-42m-for-dropping-tracking-cookies-without-consent/>.
- [159] BBC, H&M fined for breaking GDPR over employee surveillance - BBC news, in: BBC, 2020, URL <https://www.bbc.com/news/technology-54418936>.
- [160] Marketing : the Italian SA fines TIM EUR27.8 million, 2020, URL https://edpb.europa.eu/news/national-news/2020/marketing-italian-sa-fines-tim-eur-278-million_en.
- [161] Q. Feng, D. He, S. Zeadally, M.K. Khan, N. Kumar, A survey on privacy protection in blockchain system, *J. Netw. Comput. Appl.* 126 (2019) 45–58, <http://dx.doi.org/10.1016/j.jnca.2018.10.020>.
- [162] C. Wang, N. Zhang, C. Wang, Managing privacy in the digital economy, *Fund. Res.* 1 (5) (2021) 543–551, <http://dx.doi.org/10.1016/j.fmre.2021.08.009>.
- [163] M. Akil, L. Islami, S. Fischer-Hübner, L.A. Martucci, A. Zuccato, Privacy-preserving identifiers for IoT: A systematic literature review, *IEEE Access* 8 (2020) 168470–168485, <http://dx.doi.org/10.1109/ACCESS.2020.3023659>.
- [164] T. Gebremichael, L.P.I. Ledwaba, M.H. Eldefrawy, G.P. Hancke, N. Pereira, M. Gidlund, J. Akerberg, Security and privacy in the industrial internet of things: Current standards and future challenges, *IEEE Access* 8 (2020) 152351–152366, <http://dx.doi.org/10.1109/ACCESS.2020.3016937>.
- [165] S.W. Driessen, G. Monsieur, W.-J. Van Den Heuvel, Data market design: A systematic literature review, *IEEE Access* 10 (2022) 33123–33153, <http://dx.doi.org/10.1109/ACCESS.2022.3161478>.
- [166] N. Deepa, Q.-V. Pham, D.C. Nguyen, S. Bhattacharya, B. Prabadevi, T.R. Gadekallu, P.K.R. Maddikunta, F. Fang, P.N. Pathirana, A survey on blockchain for big data: Approaches, opportunities, and future directions, *Future Gener. Comput. Syst.* 131 (2022) 209–226.
- [167] A.J. Perez, S. Zeadally, Secure and privacy-preserving crowdsensing using smart contracts: Issues and solutions, *Comp. Sci. Rev.* 43 (2022) 100450, <http://dx.doi.org/10.1016/j.cosrev.2021.100450>.
- [168] C. Gonçalves, R.J. Bessa, P. Pinson, A critical overview of privacy-preserving approaches for collaborative forecasting, *Int. J. Forecast.* 37 (1) (2021) 322–342, <http://dx.doi.org/10.1016/j.ijforecast.2020.06.003>.
- [169] Y. Wu, Z. Wang, Y. Ma, V.C.M. Leung, Deep reinforcement learning for blockchain in industrial IoT: A survey, *Comput. Netw.* 191 (2021) 108004, <http://dx.doi.org/10.1016/j.comnet.2021.108004>, URL <https://www.sciencedirect.com/science/article/pii/S1389128621001213>.
- [170] L.D. Nguyen, I. Leyva-Mayorga, A.N. Lewis, P. Popovski, Modeling and analysis of data trading on blockchain-based market in IoT networks, *IEEE Internet Things J.* 8 (8) (2021) 6487–6497, <http://dx.doi.org/10.1109/JIOT.2021.3051923>.
- [171] A. Sadiq, M.U. Javed, R. Khalid, A. Almogren, M. Shafiq, N. Javaid, Blockchain based data and energy trading in internet of electric vehicles, *IEEE Access* 9 (2021) 7000–7020, <http://dx.doi.org/10.1109/ACCESS.2020.3048169>.
- [172] Y. Long, Y. Chen, W. Ren, H. Dou, N.N. Xiong, DePET: A decentralized privacy-preserving energy trading scheme for vehicular energy network via blockchain and K-anonymity, *IEEE Access* 8 (2020) 192587–192596, <http://dx.doi.org/10.1109/ACCESS.2020.3030241>.
- [173] R. Xu, Y. Chen, Fed-DDM: A federated ledgers based framework for hierarchical decentralized data marketplaces, in: 2021 International Conference on Computer Communications and Networks, 2021.
- [174] L. Giarretta, I. Savvidis, T. Marchioro, S. Girdzijauskas, G. Pallis, M.D. Dikaiakos, E. Markatos, PDS2: A user-centered decentralized marketplace for privacy preserving data processing, in: 2021 IEEE 37th International Conference on Data Engineering Workshops (ICDEW), 2021, pp. 92–99, <http://dx.doi.org/10.1109/ICDEW53142.2021.00024>.
- [175] A. Manzoor, A. Braeken, S.S. Kanhere, M. Ylianttila, M. Liyanage, Proxy re-encryption enabled secure and anonymous IoT data sharing platform based on blockchain, *J. Netw. Comput. Appl.* 176 (2021) 102917, <http://dx.doi.org/10.1016/j.jnca.2020.102917>.
- [176] T. Rückel, J. Sedlmeir, P. Hofmann, Fairness, integrity, and privacy in a scalable blockchain-based federated learning system, *Comput. Netw.* 202 (2022) 108621, <http://dx.doi.org/10.1016/j.comnet.2021.108621>.
- [177] P. Gupta, V. Dedeoglu, S.S. Kanhere, R. Jurdak, TrailChain: Traceability of data ownership across blockchain-enabled multiple marketplaces, *J. Netw. Comput. Appl.* (2022) 103389, <http://dx.doi.org/10.1016/j.jnca.2022.103389>.
- [178] Y. Tian, B. Song, T. Ma, A. Al-Dhelaan, M. Al-Dhelaan, Bi-tier differential privacy for precise auction-based people-centric IoT service, *IEEE Access* 9 (2021) 55036–55044, <http://dx.doi.org/10.1109/ACCESS.2021.3067138>.
- [179] M. Zhang, L. Yang, S. He, M. Li, J. Zhang, Privacy-preserving data aggregation for mobile crowdsensing with externality: An auction approach, *IEEE/ACM Trans. Netw.* 29 (3) (2021b) 1046–1059, <http://dx.doi.org/10.1109/TNET.2021.3056490>.
- [180] X. Xu, Z. Yang, Y. Xian, ATM: Attribute-based privacy-preserving task assignment and incentive mechanism for crowdsensing, *IEEE Access* 9 (2021) 60923–60933, <http://dx.doi.org/10.1109/ACCESS.2021.3074142>.
- [181] F. Kserawi, S. Al-Marri, Q. Malluhi, Privacy-preserving fog aggregation of smart grid data using dynamic differentially-private data perturbation, *IEEE Access* 10 (2022) 43159–43174, <http://dx.doi.org/10.1109/ACCESS.2022.3167015>.
- [182] Y. Shen, B. Guo, Y. Shen, X. Duan, X. Dong, H. Zhang, C. Zhang, Y. Jiang, Personal big data pricing method based on differential privacy, *Comput. Secur.* 113 (2022) 102529, <http://dx.doi.org/10.1016/j.cose.2021.102529>.
- [183] Y. Hu, C. Li, A. Hu, A. Hu, J. Zhao, Trading off data resource availability and privacy preservation in multi-layer network transaction, *Phys. Commun.* 46 (2021) 101317, <http://dx.doi.org/10.1016/j.phycom.2021.101317>.
- [184] Q. Song, J. Cao, K. Sun, Q. Li, K. Xu, Try before you buy: Privacy-preserving data evaluation on cloud-based machine learning data marketplace, in: Annual Computer Security Applications Conference, ACM, 2021, pp. 260–272, <http://dx.doi.org/10.1145/3485832.3485921>.
- [185] M.N. Alraja, H. Barhamgi, A. Rattrout, M. Barhamgi, 2021. An integrated framework for privacy protection in IoT — Applied to smart healthcare. 91, 107060. <http://dx.doi.org/10.1016/j.compeleceng.2021.107060>. URL <https://www.sciencedirect.com/science/article/pii/S0045790621000744>.
- [186] R. Oppliger, Privacy-enhancing technologies for the world wide web, *Comput. Commun.* 28 (16) (2005) 1791–1797, <http://dx.doi.org/10.1016/j.comcom.2005.02.003>.
- [187] T.P. Pedersen, Non-interactive and information-theoretic secure verifiable secret sharing, in: Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), in: LNCS, vol. 576, 1992, pp. 129–140, http://dx.doi.org/10.1007/3-540-46766-1_9.
- [188] A. Shamir, How to share a secret, *Publ. ACM* (1979b) http://dx.doi.org/10.1007/978-3-642-15328-0_17.
- [189] X. Chen, Introduction to Secure Outsourcing Computation, Morgan & Claypool publishers, 2016, p. 94, URL https://www.researchgate.net/publication/295681472_Introduction_to_Secure_Outsourcing_Computation.
- [190] S. De Capitani Di Vimercati, S. Foresti, G. Livraga, P. Samarati, Data privacy: Definitions and techniques, *Int. J. Uncertain. Fuzz. Knowl.-Based Syst.* 20 (6) (2012) 793–817, <http://dx.doi.org/10.1142/S0218488512400247>, URL <https://www.worldscientific.com/doi/abs/10.1142/S0218488512400247>.
- [191] A. Meyerson, R. Williams, On the complexity of optimal k-anonymity, in: Proceedings of the ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, 23, 2004, pp. 223–228, <http://dx.doi.org/10.1145/1055558.1055591>, URL <https://dl.acm.org/doi/10.1145/1055558.1055591>.
- [192] R. Rivest, A. Shamir, Y. Tauman, How to leak a secret, in: Lecture Notes in Computer Science, Vol 2248, Springer, Berlin, Heidelberg, 2001, URL <https://ieeexplore.ieee.org/document/6032224> (<https://cryptolab.com/ethereum-network-congestion-doubles-gas-fees-as-game>).
- [193] M. Bellare, D. Micciancio, B. Warinschi, Foundations of group signatures: Formal definitions, simplified requirements, and a construction based on general assumptions, *Eurocrypt* 2656 (2003) 1–27, URL <https://cseweb.ucsd.edu/~mihir/papers/g.pdf>.
- [194] S. Gürses, C. Troncoso, C. Diaz, Engineering: Privacy by design, *Science* 317 (5842) (2011) 1178–1179, URL <https://www.esat.kuleuven.be/cosic/publications/article-1542.pdf>.
- [195] M. Yung, S. Jarecki, H. Krawczyk, A. Herzberg, Proactive secret sharing or : How to cope with perpetual leakage, *Communication* (1995) URL https://www.researchgate.net/publication/221355399_Proactive_Secret_Sharing_Or_How_to_Cope_With_Perpetual_Leakage.
- [196] IOTA-Foundation, About the tangle, 2020, URL <https://legacy.docs.iota.org/docs/getting-started/1.1/the-tangle/overview>.
- [197] A. Ghosh, K. Ligett, A. Roth, G. Schoenebeck, Buying private data without verification, in: EC 2014 - Proceedings of the 15th ACM Conference on Economics and Computation, 2014, pp. 931–948, <http://dx.doi.org/10.1145/2600057.2602902>, arXiv:1404.6003.
- [198] C. Perera, R. Ranjan, L. Wang, S.U. Khan, A.Y. Zomaya, Big data privacy in the internet of things era, *IT Prof.* 17 (3) (2015c) 32–39, <http://dx.doi.org/10.1109/MITP.2015.34>.
- [199] The Wall Street Journal, Google to pay \$22.5 million in FTC settlement, 2012, URL <https://www.wsj.com/articles/SB1000087239639044304004577579232818727246>.
- [200] J. Porter, Google fined €50 million for GDPR violation in France, 2019, URL <https://www.theverge.com/2019/1/21/18191591/google-gdpr-fine-50-million-euros-data-consent-cnll>.

Further reading

- [1] G. Goos, J. Hartmanis, J. van Leeuwen, D. Hutchison, T. Kanade, J. Kittler, J.M. Kleinberg, F. Mattern, J.C. Mitchell, M. Naor, O. Nierstrasz, C.P. Rangan, B. Steffen, 1973. Lecture notes in computer science. 556. URL https://doi.org/10.1007/978-3-319-70139-4_56.
- [2] P. Cramton, Y. Shoham, R. Steinberg, An overview of combinatorial auctions, ACM SIGCOM Exch. 7 (1) (2007) 3–14, URL <https://dl.acm.org/doi/10.1145/1345037.1345039>.
- [3] R. Dingledine, N. Mathewson, P. Syverson, Tor: The Second-Generation Onion Router, Tech. Rep., Defense Technical Information Center, Fort Belvoir, VA, 2004, URL <http://www.dtic.mil/docs/citations/ADA465464>.
- [4] D. Bogdanov, S. Laur, J. Willemson, 2008. Sharemind: a framework for fast privacy-preserving computations. 15. URL https://link.springer.com/chapter/10.1007/978-3-540-88313-5_13.
- [5] N. Wang, X. Xiao, Y. Yang, J. Zhao, S.C. Hui, H. Shin, J. Shin, G. Yu, Collecting and analyzing multidimensional data with local differential privacy, 2019b, p. 13, URL <https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=8731512>.
- [6] R. Bassily, A. Smith, Local, private, efficient protocols for succinct histograms, in: Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing, ACM, 2015, <http://dx.doi.org/10.1145/2746539.2746632>.
- [7] A. Herzberg, S. Jarecki, H. Krawczyk, M. Yung, Proactive secret sharing or: How to cope with perpetual leakage, in: D. Coppersmith (Ed.), Advances in Cryptology — CRYPTO' 95, Springer Berlin Heidelberg, Berlin, Heidelberg, 1995, pp. 339–352, URL https://link.springer.com/chapter/10.1007/3-540-44750-4_27.
- [8] A. Bellet, A. Habrard, M. Sebban, A survey on metric learning for feature vectors and structured data, 2014, URL <https://arxiv.org/abs/1306.6709>.
- [9] B. Poettering, D. Stebila, Double-authentication-preventing signatures, Int. J. Inf. Secur. 16 (1) (2017) URL <http://link.springer.com/10.1007/s10207-015-0307-8>.
- [10] S. Yu, C. Wang, K. Ren, W. Lou, Achieving secure, scalable, and fine-grained data access control in cloud computing, in: 2010 Proceedings IEEE INFOCOM, 2010, pp. 1–9, <http://dx.doi.org/10.1109/INFCOM.2010.5462174>.
- [11] R. Wieringa, N. Maiden, N. Mead, C. Rolland, Requirements engineering paper classification and evaluation criteria: A proposal and a discussion, Requir. Eng. 11 (1) (2006) 102–107, <http://dx.doi.org/10.1007/s00766-005-0021-6>.

Gonzalo Munilla Garrido joined the chair for Software Engineering of Business Information Systems (Sebis) at the Technical University Munich in October 2019. In cooperation with The BMW Group, he currently investigates how privacy-preserving techniques enable new iterations of data marketplaces. He received a M.Sc. mult. in Mechanical Engineering from the TUM and the Polytechnique University of Madrid.

Johannes Sedlmeir is a researcher at the Project Group Business & Information Systems Engineering of the Fraunhofer FIT and currently pursuing his Ph.D. in Information

Systems at the FIM Research Center, University of Bayreuth. In his research, Mr. Sedlmeir works on the energy consumption and performance benchmarking of different blockchains, decentralized digital identities, and the application of cryptographic methods such as Zero-Knowledge Proofs for scalability and privacy on blockchains. He received his M.Sc. in Theoretical and Mathematical Physics.

Ömer Uludağ joined the chair for Software Engineering of Business Information Systems at the Technical University Munich in April 2016. He holds a M.Sc. in Information Systems and currently works at Allianz SE.

Ilias Soto is currently a M.Sc. student in Information Systems at the Technical University of Munich and holds a B.Sc. in the same field.

Andre Luckow work currently focuses on applications of machine learning, quantum computing and blockchain. Mr. Luckow drives the application of computing technologies to problems in business and science bridging cross-functional gaps to create value via process improvements or the enablement of new products. Previously, Mr. Luckow held several functions at BMW Group IT in Munich, Germany. Furthermore, he teaches several courses at different universities, e.g. the Ludwig-Maximilians-University Munich, Germany, and Clemson University, SC, USA. Andre Luckow holds a Ph.D. in the field of distributed computing from the University of Potsdam.

Florian Matthes Since 2002 Florian Matthes holds the chair for Software Engineering for Business Information Systems at Technische Universität München.

The current focus of his research is on technologies driving the digital transformation of enterprises and societies: Enterprise architecture management, service platforms and their ecosystems, semantic analysis of legal texts and executable contracts on blockchains.

He is co-founder of CoreMedia, infoAsset and Tr8cy, co-founder and chair of Blockchain Bayern e.V. scientific advisor of Noumena Digital, member of the advisory board of the Ernst Denert-Stiftung für Software Engineering, and initiator and organizer of international conferences and workshops in software and enterprise engineering

Earlier stations of his academic career are the Goethe-University Frankfurt (Diploma 1988) the University of Hamburg (Ph.D. 1992), the Digital Systems Research Center (now HP SRC Classic) in Palo Alto, USA (Researcher 1992–1993), and the Technical University Hamburg-Harburg (Associate Professor 1997–2002).

Until 2010 he served as dean of studies at the Faculty for Informatics and member of the teaching board of TU München. Since 2020 he serves as "venture ambassador" of TUM venture labs.