AMCIS 2024 Proceedings

Americas Conference on Information Systems (AMCIS)

August 2024

Breaking Down Privacy by Design: A Threefold Perspective

Alexandra Klymenko
Technical University of Munich, alexandra.klymenko@tum.de

Stephen Meisenbacher
Technical University of Munich, stephen.meisenbacher@tum.de

Ali Asaf Polat
Technical University of Munich, ge95vuk@mytum.de

Florian Matthes
Technical University of Munich, matthes@tum.de

Breaking Down Privacy by Design: A Threefold Perspective

Completed Research Full Paper

Alexandra Klymenko
Technical University of Munich
alexandra.klymenko@tum.de

Stephen Meisenbacher
Technical University of Munich
stephen.meisenbacher@tum.de

Ali Asaf Polat
Technical University of Munich
aliasaf.polat@tum.de

Florian Matthes
Technical University of Munich
matthes@tum.de

Abstract

The age of big data has raised concerns over how organizations processing data can best safeguard the privacy of individuals. The framework of Privacy by Design (PbD) provides a foundation for the integration of privacy into systems proactively and not reactively, setting the standard for privacy as the default mindset. Particularly in the demonstration of compliance as mandated by data protection regulations, the seven principles proposed by PbD can serve as a useful starting point for responsible data processing. Nevertheless, the principles of PbD are intentionally open-ended and do not make a distinction between legal, technical, and organizational aspects. Based on existing literature, we address this gap by investigating PbD from these three perspectives, with a particular focus on mapping PbD principles to Privacy-Enhancing Technologies. We validate our findings in a series of iterative sessions with privacy professionals, who confirm the accuracy and practical relevance of our work.

Keywords

Data Privacy, Privacy by Design, Privacy-Enhancing Technologies

Introduction

Concerns over the safeguarding of individual privacy have grown steadily in recent years, particularly in the technological sphere with the recent rise of Artificial Intelligence. With this, the question becomes how to preserve privacy, and the answer has largely come in the form of both regulatory response and the development of Privacy-Enhancing Technologies (PETs).

A guiding framework for data privacy comes in the form of “Privacy by Design” (PbD) (Cavoukian 2009), a concept that advocates for privacy-conscious technology design, rather than post-hoc privacy handling. PbD can serve as a useful tool for practitioners in the demonstration of data privacy compliance, a challenging task involving the interdisciplinary effort between legal, practical, and business stakeholders (Klymenko et al. 2023), who must be informed of the value of data protection, as well as of the appropriate technical and organizational measures available to facilitate it (Klymenko et al. 2022).

While the PbD principles serve as a starting point for the road towards compliance, the framework itself is intentionally vague and thus leaves “many open questions about their application when engineering systems” (Gürses et al. 2011). Furthermore, PbD principles do not make the distinction between the legal, technical, or organizational aspects of data privacy and privacy compliance. This, therefore, can make it difficult for practitioners to discern which aspects of PbD are most relevant to their work, and accordingly, lead to confusion as to how they can practice PbD specifically in their role.

In this work, we aim to address this gap, namely, to investigate PbD in a threefold light: legal, technical, and organizational. We call these the *aspects* of PbD. Drawing upon insights from the body of existent literature on data protection regulation and PbD, we explore what PbD means for practitioners of differing expertise. Subsequently, we compile a set of technical-, legal-, and organizational-specific PbD aspects, which represent tangible characteristics and activities specifically related to each facet of PbD. These aspects are evaluated and validated by expert practitioners in the field of privacy.

Following this, we focus on the technical aspects of PbD, specifically on the integration of PbD principles and state-of-the-art PETs. The lack of understanding of how to move from elicitation of privacy requirements based on PbD principles to their implementation using PETs is one of the main challenges in the context of PbD (Diamantopoulou et al. 2017). Previous works have aimed to bridge the gap between PbD and their practical implementation via privacy engineering principles and privacy patterns (Rubenstein and Good 2013; Alkhariji et al. 2021), however, to the best of authors' knowledge, no works map PbD principles directly to PETs. Motivated by this, we map PbD principles to PETs, to demonstrate how implementing PETs can help organizations achieve PbD. This serves not only to connect the two concepts but also to increase the practical adoption of advanced PETs, which remains low, in particular, due to the lack of clear relevance and mapping to privacy regulations (Klymenko et al. 2023).

As such, the contributions of our work are as follows:

1. We analyze and clarify what PbD principles entail from the legal, technical, and organizational perspectives, extending the existing notion of PbD to be more specialized and practically oriented.
2. We place a particular focus on five advanced Privacy-Enhancing Technologies, showing how they aid in satisfying PbD principles and support the process of privacy compliance.

Background and Related Work

Privacy by Design

Privacy by Design represents a proactive approach to embedding privacy considerations into the design of IT systems and incorporates seven foundational principles (Cavoukian 2009):

P1. **Proactive not Reactive; Preventative not Remedial**. The first principle entails a preventative approach to privacy, i.e., proactively preventing privacy risks from materializing in the first place, rather than remediating damage from the already occurred privacy breaches.

P2. **Privacy as the Default Setting**. The second principle ensures maximum privacy by default, i.e., automatically setting privacy to the highest level of protection without the need for users to adjust their privacy settings. This includes, in particular, limiting the collection, use, retention, and disclosure of personal data to what is strictly necessary for the specified lawful purposes.

P3. **Privacy Embedded into Design**. PbD must be integrated into software systems and business processes, ensuring privacy is treated as an essential aspect from the outset rather than added on as an afterthought. This must also be done without impairing any of its other functions.

P4. **Full Functionality – Positive-Sum, not Zero-Sum**. Implementing PbD early in the lifecycle of a system can help to avoid unnecessary compromises within the system (i.e., zero-sum), for example with security vs. privacy. Instead, PbD aims to address all objectives in a mutually beneficial "win-win" manner, achieving a positive-sum outcome.

P5. **End-to-End Security – Full Lifecycle Protection**. PbD requires strong security measures to be embedded into the system throughout the data processing lifecycle, including its collection, use, and eventual destruction. The security principle is important to PbD, as according to Cavoukian (2006), "without strong security, there can be no privacy."

P6. **Visibility and Transparency – Keep it Open**. PbD requires that organizations handle data according to the stated promises and that their privacy policies are openly communicated to all stakeholders, thus establishing accountability and trust, as well as ensuring compliance.

P7. **Respect for User Privacy – Keep it User-Centric**. PbD should be tailored to the needs and interests of individuals, empowering users to have control over how their data is managed.

While PbD provides a fundamental framework for approaching privacy in IT systems, it has often been criticized for the lack of a clear definition and the abstract nature of some of the principles, which make it "unclear what a request for PbD practically means" (van Rest et al. 2014). Guided by this challenge, we aim to deconstruct the vague notion of PbD by breaking it down into actionable items, thus providing practitioners with clearer guidance on the steps required for adhering to PbD.

Privacy-Enhancing Technologies

The concepts of PbD and PETs are closely linked, with both their origins tracing back to a collaboration between Canadian and Dutch researchers and authorities in the 1990s (Hustinx 2010; Hes and Borking 1995). Formally defined, PETs represent “a coherent system of ICT measures that protects privacy by eliminating or reducing personal data or by preventing unnecessary and/or undesired processing of personal data, all without losing the functionality of the information system” (van Blarkom et al. 2003).

PETs play an important role in embedding the principles of PbD in IT systems. As outlined by the UK Information Commissioner’s Office (ICO 2023), PETs can help to follow PbD by (1) adhering to the principle of data minimization, (2) ensuring an appropriate level of security, (3) implementing strong anonymization and pseudonymization techniques, and (4) mitigating the risk of potential data breaches by making personal data incomprehensible to unauthorized individuals.

Despite the high potential of these technologies in protecting privacy, as well as supporting adherence to PbD and complying with regulations, many advanced PETs such as Differential Privacy or Homomorphic Encryption remain predominantly in academia and are not widely adopted in practice (Hansen et al. 2015). Among the reasons for this is the limited awareness and understanding of complex PETs, as well as the unclear applicability of advanced technologies to regulations, which results in the lack of incentive to surpass the bare minimum required for being compliant “on paper” and invest in more advanced solutions (Klymenko et al. 2023). As such, the second contribution of our work aims to specify the ways in which some of the most researched state-of-the-art PETs can be practically utilized in line with PbD.

Methodology

To guide this work, the following two research questions were defined:

RQ1. What do the principles of PbD entail from the legal, technical, and organizational perspectives?

RQ2. How can the principles of PbD be supported by individual PETs?

Our research started with a Systematic Literature Review (SLR), following the methodology of Kitchenham et al. (2015). The following search strings were employed:

S0. (“Privacy by Design”) AND author:“Cavoukian”

S1. (“Privacy by Design”) AND (“Legal” OR “Technical” OR “Organizational” OR “Organization” OR "Organisational" OR "Organisation")

S0 and S1 were designed to answer RQ1, where the goal of S0 was to discover basis literature with the author Ann Cavoukian, who is regarded as the creator of PbD. To answer RQ2, a combination of search queries *(“Privacy by Design”) AND (“Privacy-Enhancing Technologies” OR “PETs”),* as well as *(“Privacy by Design”) AND “[PET Name]”* was initially employed; however, this yielded only one result (Zhang and Kamel Boulos 2022). As such, to provide a technological basis of study for RQ2, the detailed overview of the selected PETs as given by Fantaye (2023) was used.

To conduct the SLR, we utilized Google Scholar with each of the search strings, with the constraint that the strings must appear in the title. Google Scholar was used due to its comprehensive indexing of relevant knowledge bases, including ACM Digital Library, AIS eLibrary, and many others. After searching, results were tabulated and deduplicated. They were then screened according to two primary inclusion/exclusion criteria: (1) must be openly or institutionally accessible, and (2) must be written in English. Finally, the remaining sources were filtered for relevance to research questions.

Once the abovementioned sources were found, an analysis of these sources was conducted to identify relevant information regarding the three aspects of PbD and their relation to PETs. Sources from S0 and S1 were read by a team of researchers and coded according to PbD principle and aspect (legal, technical, or organizational). Subsequently, each PET as highlighted by Fantaye (2023) was analyzed in light of the findings from RQ1. These codes were systematized, and the structured results are found in Tables 3 and 4.

	Hits	After Dedup.	After Exclusion	After Filtering	Final Sources
S0	168	73	36	16	(Cavoukian 2009; Cavoukian 2011; Cavoukian et al. 2014; Cavoukian et al. 2020)*
S1	18	13	9	3	(Cavoukian et al. 2010; Pencarrick et al. 2013; Rachovitsa 2016)

*In this case, *representative sources*

Table 1. The SLR Results

These results were then validated and evaluated in a series of iterative sessions with experts serving as privacy professionals. The goal of these sessions was to verify that our structured results were accurate, as well as to obtain further insights on the topic. The evaluation sessions were held either live via online meetings, or asynchronously via written communication. Depending on the expertise and proficiency of the expert in question, either the PbD aspects (RQ1) or PET mappings (RQ2) were evaluated.

The experts were acquired through a combination of personal contacts, referrals, and LinkedIn. The candidates were selected based on their expertise and the relevance of their background and experience to the research topic. In addition, preference was given to professionals certified by the International Association of Privacy Professionals (IAPP). Contact with the candidates was initiated through personalized emails or direct messages on LinkedIn. The candidates were provided in advance with relevant information regarding the study and their role in it, ensuring informed consent. All personally identifiable information about the experts was anonymized to protect their privacy.

ID	Role	Industry Domain	Org. size	Country	Experience	Feedback*
I1	Senior Privacy Consultant	IT Consulting	Large	Germany	10-20 years	Aspects (v/w)
I2	LL.M. Candidate	Academy	–	Germany	0-5 years	Aspects (v)
I3	Privacy/Compliance Executive	Legal Compliance	Large	Germany	10-20 years	Aspects (w)
I4	Privacy Consultant	Legal Compliance	Medium	Canada	5-10 years	Aspects (w)
I5	Data Privacy Officer	Energy	Large	Germany	10-20 years	Aspects (w)
I6	Privacy Manager	Legal Compliance	Large	Turkey	10-20 years	Aspects (v/w)
I7	Founder, Consultant	Legal Compliance	Medium	Germany	25+years	Aspects (w)
I8	Lawyer, Computer Engineer	IT	Small	Turkey	10-20 years	PETs Mapping (w)
I9	Data Protection Manager	Legal Compliance	Medium	USA	10-20 years	PETs Mapping (w)
I10	Privacy Engineer	Legal Compliance	Large	Germany	10-20 years	PETs Mapping (v)
I11	Privacy Engineer	Legal Compliance	Small	Switzerland	25+ years	PETs Mapping (w)

***Feedback Type**: v: verbal communication; w: written communication

Table 2. Experts Involved in the Evaluation of PbD Aspects and PbD-PETs Mapping

Legal, Technical, and Organizational Aspects of PbD

In this section, we discuss the legal, technical, and organizational aspects of PbD, summarizing our findings in Table 3. The aspects were derived from the publications listed in Table 1. It is important to note that as these aspects are inherently connected, they often overlap with and contribute to multiple PbD principles. With this in mind, based on the feedback from I1, we placed focus on finding the most related principle for the given aspect and, accordingly, removed duplicate bullet points in Table 3. Following a brief introduction to each aspect, some of the received expert feedback is summarized.

Legal Aspects

Legal practitioners are often at the forefront of reading and interpreting relevant regulations and laws, which represents a proactive approach to understanding the necessary privacy requirements an organization must follow. According to Cavoukian (2011), PbD serves “as the foundation for a second generation approach to privacy regulation”, incorporating flexible privacy principles into law, which can be observed in regulations such as the General Data Protection Regulation (GDPR).

In practice, this often involves the conduction of Data Protection Impact Assessments (DPIAs) or the creation of legally binding contracts with third-party vendors to ensure compliance with the regulatory mandates. The concept put forth by P2 is itself a legal requirement, for example under Article 25 of the EU GDPR. Therefore, P2 has close ties to the legal mandates of purpose limitation, data minimization, and storage limitation. In a similar way, legal activities such as DPIAs and legal requirements such as data minimization directly influence the design and implementation of systems (Oetzel and Spiekermann 2013). By following such practices, the protection of privacy can become ingrained into technology design.

The pursuit of legal requirements such as fairness and transparency should not come at the cost of functionality. Thus, from a legal perspective, the goal becomes to promote techniques that enhance privacy while still achieving the main functional requirements of a system. P5 emphasizes the protection of personal information along the whole lifecycle of data processing. This relates to the confidentiality, integrity, and availability of data subjects' personal information by requiring protection against unauthorized or unlawful processing and accidental loss, destruction, or damage. Informing individuals about their rights, the purpose of data processing, and the existence of any third parties are some of the legal obligations falling under visibility and transparency principle P6. This encapsulates a notion of accountability, which becomes important to understanding the legal perspective of PbD. In line with P7, respecting individual privacy is inherently user-centric, and the requirements set forth by many regulations put this person at the forefront. By adhering to informed consent, ensuring data accuracy, allowing for opting out and deletion of data, and providing redress in the event of incidents, the user is made the focal point of legal measures for PbD.

Technical Aspects

In the modern technical landscape, many organizations collect personal information and therefore have compliance requirements. Cavoukian, in her technical documentation for PbD for software engineers (Cavoukian et al. 2014), provides a strong basis for incorporating the seven principles into the technical aspects of privacy, claiming that it "has become increasingly apparent that software systems need to be complemented by a set of governance norms that reflect privacy dimensions".

P1 implies that technical measures for privacy protection should be implemented proactively. This includes adopting PETs, such as encryption, identity and access management, network protection, and others. In addition, establishing checks to recognize poor privacy designs within a system becomes crucial to ensuring technical measures are soundly implemented. By incorporating default privacy settings within systems, products, and services (P2), data subjects are protected by principle rather than in a post hoc manner. From a technical view, P3 entails implementing privacy measures and PETs within the design and architecture of a system, as well as considering data minimization, secure data transmission, and consent management systems while in the design process.

P4 encourages keeping the full functionality of the system, and from a technical perspective, this introduces the task of designing and developing systems, products, or services in a way that enables both robust privacy protection and full functionality. As a potential solution, PETs offer a balance of functionality and robust privacy protection, thus presenting a promising option to achieve P4. In a similar vein, P5 aims for full life cycle protection; therefore, the implementation of technical security measures and other technologies to protect personal information at all stages of the software development life cycle should be followed. In line with P6, organizations could incorporate features that provide transparency, design their systems to allow users to view and control their privacy settings, and provide tools for individuals to exercise their rights. Lastly, the implementation of PETs facilitates respecting users and their privacy, particularly in the design of systems that prioritize user privacy above all else.

Organizational Aspects

Ensuring compliance is essential when collecting and processing personal information, becoming "very useful in building business and competitive advantages" (Cavoukian 2011), especially in regulated environments. PbD aids in demonstrating compliance, where the organizational aspect plays a key role.

A clear commitment at the highest levels of an organization to safeguard privacy will help organizations to address privacy risks proactively. In addition, it is important to monitor third-party vendors' processing activities and their security measures regularly. Likewise, external validation of an organization's practices via audits, seals, and certification programs can help to identify gaps and improve transparency. In doing so, data violations might be avoided before they happen. Based on P2, privacy should be prioritized within the organization, and privacy considerations should be incorporated into the organizational culture. Furthermore, organizations should update their privacy policies regularly. Organizational measures for P3 could include assigning privacy responsibilities, such as a Data Protection Officer (DPO), while developing systems and services, as well as the building up of Privacy Risk Management (PRM) processes. PRM allows organizations to identify, analyze, and evaluate privacy risks.

Principles	Legal Aspects	Technical Aspects	Organizational Aspects
(P1) Proactive not Reactive; Preventative not Remedial	• Legal requirement to implement security measures • Conducting DPIA in case of high-risk activities • Establishing a legally binding contract with third-party vendors such as service providers and data processors	• Implementing preventative measures within systems from the beginning • Adopting PETs, encryption, identity and access management, and other technical mechanisms • Establishing methods to recognize poor practices and anticipate undesired outcomes	• Establishing privacy policies, procedures, and guidelines that help employees proactively address privacy risks • Commitment from management to uphold high standards of privacy • Regularly monitoring third-party vendors' processing activities and security measures
(P2) Privacy as the Default Setting	• Legal principle itself; Privacy by Default • Compliance with purpose limitation, data minimization, and storage limitation as privacy defaults • Lawfulness, fairness, and transparency requirements	• Implementing default privacy settings within systems, products, and services • Putting into place technical controls for each phase of the software development cycle	• Prioritizing privacy as the default option within the organization • Incorporating privacy considerations into the organizational culture, decision-making processes, and overall privacy governance framework • Regularly updating privacy policies
(P3) Privacy Embedded into Design	• Meeting all legal obligations related to individuals' privacy protection • Data minimization, purpose limitation, consent, lawfulness, fairness, transparency, DPIA for high-risk activities, etc.	• Implementing privacy measures and PETs within the design and architecture of systems, products, and services • Including data minimization, secure data transmission, consent management, user data access, and deletion, etc. during design and implementation	• Assigning privacy responsibilities, such as a DPO when developing systems, products, and services • Establishing a risk management process, which includes identifying, analyzing, and evaluating the privacy risks • Conducting external audits, obtaining certifications and seals
(P4) Full Functionality Positive-Sum not Zero-Sum	• Ensuring lawfulness, fairness, transparency, accuracy, and other privacy-related obligations while striving for functionality that enhances privacy	• Designing and developing systems, products, or services in a way that enables both robust privacy protection and full functionality • Employing PETs, such as encryption, access controls, anonymization techniques, or other technical mechanisms	• Ensuring privacy protections without hindering the overall functionality of the system • Establishing cross-functional teams that bring together privacy experts, developers, and managers • Establishing monitoring, analysis, and evaluation mechanisms to check functionality
(P5) End-to-End Security – Full Life Cycle Protection	• Safeguarding individuals' personal information during the whole processing lifecycle • Ensuring confidentiality, integrity, and availability of data subjects' personal information	• Technical implementation of security measures and technologies to protect personal data at all stages of its lifecycle • Employing encryption, access controls, secure data storage practices, using PETs and data minimization techniques	• Establishing policies and practices that prioritize end-to-end privacy • Assigning responsibilities for security, conducting regular security assessments • Conducting internal audits at planned intervals to check the security management system
(P6) Visibility and Transparency – Keep it Open	• Legal requirement to be transparent • Informing individuals about their rights, the purpose of data processing, and the existence of any third parties • Accountability of organizations • Data breach notifications	• Incorporating features that provide transparency • Designing systems that offer privacy dashboards or controls which allow users to view and control their privacy settings • Provide tools for individuals to exercise their privacy rights	• Establishing policies and procedures that prioritize transparency • Providing clear and accessible privacy notices • Informing individuals about the purpose and scope of data processing
(P7) Respect for User Privacy – Keep it User-Centric	• Legal requirements for explicit user consent, exercising individual rights, and withdrawing the given consent • Keeping individuals' data up-to-date • Providing individuals with the right to redress	• Implementing PETs and design systems, products, or services that prioritize user privacy • Encouraging direct data subject access to their data	• Fostering a user-centric mindset, where user feedback and preferences regarding privacy are actively considered and incorporated into privacy related decision-making processes • Creating organizational culture, policies, and practices that involve a user-centric approach

Table 3. Legal, Organizational and Technical Aspects of PbD

Establishing cross-functional teams within an organization brings together privacy experts, designers, developers, etc. Fostering such interdisciplinary interaction can be pivotal in ensuring that functionality is not lost with the protection of privacy (P4). In line with P6, organizational measures could include establishing policies and procedures within the organization to prioritize transparency and providing clear and accessible privacy notices to individuals. Lastly, user feedback and preferences are considered and incorporated into privacy-related decision-making within the organization. Fostering organizational culture, policies, and practices to involve a user-centric approach to privacy is much aligned with P7.

Expert Feedback

All results presented in Table 3 were supplemented and evaluated by privacy professionals, as discussed in Methodology. The initial breakdown of each PbD principle was assessed and corresponding changes were made. For example, for P1, I7 suggested mentioning other stakeholders such as third-party vendors, service providers, and data processors in the legal aspect. For the technical aspect of P1, I6 recommended mentioning privacy requirements from the beginning of the software development. For the organizational aspect of P1, I7 advised adding monitoring requirements for the third parties.

Privacy by Design and Privacy-Enhancing Technologies

While PbD principles require the adoption of relevant technical safeguards, their descriptions do not mention any specific technologies that must be implemented. As such, it becomes important for practitioners to understand the different ways of applying these principles within their systems. In this section, we show how technical aspects of PbD can be implemented by leveraging PETs. It is important to emphasize that the presented results provide a general overview of how the selected PETs can contribute to supporting PbD principles and should not be viewed as a direct mapping of PbD principles to specific PETs. In order to guarantee alignment with the technical aspects of PbD principles, a multitude of factors within the system, such as its architecture and implementation specifics, must be considered.

Selected Privacy-Enhancing Technologies

In a recent study on the most predominant PETs in the academic literature, Fantaye (2023) describes the top five most researched PETs for data processing and analysis. As such, we utilize this subset of five PETs as a representative sample for our mapping of PbD principles to PETs, presented in the following.

Differential Privacy

Differential Privacy (DP) roots itself in a mathematical guarantee for the protection of individuals within a dataset. Traditionally, this is quantified by the epsilon (ϵ) parameter, also known as the *privacy budget*. By proactively adding calibrated noise to the output of a query on the database, differentially private mechanisms ensure that individuals remain indistinguishable within some bound, such that participation in a dataset is anonymized, while still allowing for useful analysis of data.

Homomorphic Encryption

Homomorphic Encryption (HE) enables performing computations on encrypted data without the need to decrypt it first. This is advantageous in comparison to symmetric or asymmetric encryption, which requires decryption before any meaningful computations can be performed. In this way, HE is a promising technology that allows for privacy-preserving data sharing, protecting the identity and information of users while still preserving data utility for data processors.

Federated Learning

Federated Learning (FL) allows to train Machine Learning models collaboratively without exchanging or centralizing raw data. In an FL setting, each client conducts training locally, thus keeping sensitive information private, and shares only its model updates with a central server. The server then aggregates the results, updates the global model, and sends the improved global model back to the clients. The scheme enables privacy preservation because sensitive data is never shared with a central aggregator, which minimizes the potential for privacy breaches.

Zero-Knowledge Proofs

In the case where validation is required, for example, to prove age or identity, Zero-Knowledge Proofs (ZKPs) provide a mechanism whereby a *prover* can validate requested information without ever sharing the raw information or data with the *verifier*. Thus, ZKPs are an important application of cryptography that facilitates privacy-preserving authentication and validation.

Secure Multiparty Computation

Secure Multiparty Computation (SMPC) allows for two or more parties to perform computation in a collaborative way, without the need for any of the individuals to share their own information. In essence, SMPC employs a cryptographic technique called "secret sharing", which requires each party to fragment and distribute their personal data in such a way that the desired computation can be jointly calculated, but the individual data points cannot be reconstructed with certainty. In this way, SMPC protects against information disclosure and exposure of potentially sensitive information.

Mapping of PbD Principles to PETs

In order to create a bridge between PbD and PETs, we align PbD principles with PETs, showing how these technologies may be helpful in achieving the goals set out by PbD. Table 4 presents the summary of the results, which were evaluated by privacy professionals who possess both technical and legal expertise.

Note that P2 (Privacy as the Default Setting) is not directly mapped to any PET, as the requirement of P2 relies on the organizational decisions made before the implementation of PETs; therefore, P2 can be seen as a precursor to implementing any PET. In the case of P6, no mapping is made to PETs, as none of the studied PETs operate in a way that fulfills the principle. Rather, P6 can be satisfied through various mechanisms that provide users the option to exercise their privacy rights, such as cookie consent.

Supported Principle	Supported By
Differential Privacy	
(P1) Proactive not reactive, preventative not remedial	Adding calibrated noise to dataset queries through Differential Privacy techniques, organizations can ensure privacy and confidentiality from the very beginning in a proactive manner.
(P3) Privacy Embedded into Design	By lessening the likelihood of reidentification of individuals, organizations can ensure that privacy is protected from the moment data is collected.
(P4) Full Functionality – Positive-sum, not zero sum	By controlling the amount of noise added to the data, organizations can achieve an appropriate balance that protects privacy while still enabling meaningful analysis and decision-making.
(P5) End-to-end Security – Full Life Cycle Protection	Ensuring privacy protection by adding carefully calibrated noise to the query results or aggregated information during the collection[*] and processing of data, and protecting data against reidentification or information leakage, it contributes to full life cycle protection.
(P7) Respect for User Privacy – Keep it User-Centric	Providing an assurance that participation in a database, as well as the data itself, is less likely to be traced back to individuals.
Homomorphic Encryption	
(P1) Proactive not reactive, preventative not remedial	Enabling computations on encrypted data, it allows for data analysis and processing while maintaining the privacy and confidentiality of the sensitive information, thus enabling organizations to protect sensitive data proactively.
(P3) Privacy Embedded into Design	Utilizing Homomorphic Encryption, organizations can ensure that sensitive data remains encrypted throughout the processing pipeline. This ensures that individuals' privacy is protected.
(P4) Full Functionality – Positive-sum, not zero sum	Allowing organizations to perform complex computations on encrypted data without sacrificing privacy or the accuracy of the analysis, it contributes to full functionality.
(P5) End-to-end Security – Full Life Cycle Protection	Encrypting data after collection, ensuring the confidentiality of sensitive information before and during data transfer and storage, it contributes to full life cycle protection.
(P7) Respect for User Privacy – Keep it User-Centric	Ensuring that individuals' data remains private and confidential throughout computations, safeguarded through encryption.
Federated Learning	
(P1) Proactive not reactive, preventative not remedial	Allowing model training without transferring the raw data to the central server, it proactively protects data with minimizing the risk of data breaches or unauthorized accesses.
(P3) Privacy Embedded into Design	Incorporating FL into the design of systems, with model distribution and local model training, organizations can ensure that the training of a model respects individuals' privacy.

[*] In a legal context, data protection laws typically focus on data that has been collected, transmitted, and stored by the service provider rather than data that remains solely on the user's device and has not been shared with the service provider or any third parties.

(P4) Full Functionality – Positive-sum, not zero sum	Minimizing the risk of exposing sensitive information by sharing model updates instead of raw data while still enabling the model training, it contributes to full functionality.
(P5) End-to-end Security – Full Life Cycle Protection	Allowing data to remain on user devices or edge servers, ensuring that sensitive data is not centralized during the collection and processing phases, it contributes to full life cycle protection.
(P7) Respect for User Privacy – Keep it User-Centric	Keeping individuals' data decentralized and on device, reducing the need to share sensitive information with a central server.
Zero-Knowledge Proofs	
(P1) Proactive not reactive, preventative not remedial	Enabling individuals or entities to prove assertions or statements without revealing unnecessary details, it proactively protects privacy.
(P3) Privacy Embedded into Design	Embedding Zero-Knowledge Proofs into the design of authentication systems, organizations can ensure that privacy is protected during the authentication process.
(P4) Full Functionality – Positive-sum, not zero sum	Demonstrating the validity of information without revealing sensitive details, it preserves privacy while enabling necessary verification processes and contributes to full functionality.
(P5) End-to-end Security – Full Life Cycle Protection	Enabling secure authentication and data validation without revealing underlying data, it safeguards against privacy breaches and contributes to full life cycle protection.
(P7) Respect for User Privacy – Keep it User-Centric	Proving the validity of individuals' information or claims without revealing the actual data and minimizing the amount of sensitive information shared while achieving the intended outcome.
Secure Multiparty Computation	
(P1) Proactive not reactive, preventative not remedial	Allowing multiple parties to collaborate and jointly perform computations on their combined datasets without disclosing the individual data parts, it proactively protects individuals' privacy.
(P3) Privacy Embedded into Design	Embedding Secure Multiparty Computation into the design of systems, it ensures that collaboration among multiple parties respects privacy with joint computation techniques.
(P4) Full Functionality – Positive-sum, not zero sum	Employing cryptographic techniques, it ensures that privacy is preserved while achieving the desired computation objectives and contributes to full functionality.
(P5) End-to-end Security – Full Life Cycle Protection	Enabling organizations to mitigate risks associated with centralized data storage, strengthen data confidentiality, and ensure the protection of sensitive information throughout the entire data life cycle, from collection to result sharing, it contributes to full life cycle protection.
(P7) Respect for User Privacy – Keep it User-Centric	Enabling the derivation of insights from combined data without revealing individuals' data points and fostering collaboration while preserving privacy.

Table 4. Supporting Privacy by Design Principles with PETs

Conclusion

We portray Privacy by Design in a threefold manner, showing that PbD principles contain important legal, technical, and organizational directives for professionals of differing backgrounds and expertise. In addition, we place a particular focus on PETs and demonstrate how they can support PbD principles. The threefold PbD aspects and the PbD-PETs mapping were verified with 11 privacy professionals, which validates its practical relevance. In this way, we make an important connection between the principles of PbD, the promise of PETs, and the transition of these theoretical concepts to actionable steps in practice.

The limitations of our work follow from the methodology by which the aspects were extracted and the mapping was created, possibly being threatened by researcher bias. Although we attempted to mitigate this bias by working in a team of three researchers and conducting the evaluation sessions with external professionals, our findings can certainly be strengthened by follow-up studies and further validation.

As future work, we propose the creation of a framework for integrating PbD principles in the design and implementation of PETs. This would serve to solidify the approach we followed in this paper, and ultimately, to facilitate a better understanding of novel PETs and their transition into practice.

The findings of this work have direct practical implications, aiding in making the concept and principles of PbD become more practically relevant, especially in shedding light on how PETs can cross the boundaries of the academic sphere into practical usability. By employing the accepted standard of PbD and making it more accessible to role-specific responsibilities in practice, we hope to accelerate the speed at which PbD becomes a default, and furthermore, to lay the groundwork for practically usable PETs.

REFERENCES

Alkhariji, L., Alhirabi, N., Alraja, M. N., Barhamgi, M., Rana, O., and Perera, C. 2021. "Synthesising Privacy by Design Knowledge toward Explainable Internet of Things Application Designing in Healthcare," *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* (17:2s), ACM New York, NY, pp. 1–29.

van Blarkom, G. W., Borking, J. J., and Olk, J. G. E. 2003. “Handbook of Privacy and Privacy-Enhancing Technologies,” *Privacy Incorporated Software Agent (PISA) Consortium, The Hague* (198), Citeseer, p. 14.

Cavoukian, A. 2009. “Privacy by Design: The 7 Foundational Principles,” *Information and Privacy Commissioner of Ontario, Canada* (5), p. 12.

Cavoukian, A. 2011. *Privacy by Design in Law, Policy and Practice*, desLibris.

Cavoukian, A. 2020. “Understanding How to Implement Privacy by Design, One Step at a Time,” *IEEE Consumer Electronics Magazine* (9:2), IEEE, pp. 78–82.

Cavoukian, A., Carter, F., Jutla, D., Sabo, J., Dawson, F., Fieten, S., Fox, J., Brown, P., Janssen, G., Jutla, D. N., and others. 2014. *Privacy by Design Documentation for Software Engineers Version 1.0*.

Cavoukian, A., Taylor, S., and Abrams, M. E. 2010. “Privacy by Design: Essential for Organizational Accountability and Strong Business Practices,” *Identity in the Information Society* (3), Springer, pp. 405–413.

Diamantopoulou, V., Kalloniatis, C., Gritzalis, S., and Mouratidis, H. 2017. “Supporting Privacy by Design Using Privacy Process Patterns,” in *ICT Systems Security and Privacy Protection: 32nd IFIP TC 11 International Conference, SEC 2017, Rome, Italy, May 29-31, 2017, Proceedings 32*, Springer, pp. 491–505.

Fantaye, J. 2023. *An Introduction and Overview of Privacy-Enhancing Technologies for Data Processing and Analysis*.

Gürses, S., Troncoso, C., and Diaz, C. 2011. “Engineering Privacy by Design,” *Computers, Privacy & Data Protection* (14:3), p. 25.

Hansen, M., Hoepman, J.-H., Jensen, M., and Schiffner, S. 2015. “Readiness Analysis for the Adoption and Evolution of Privacy Enhancing Technologies: Methodology, Pilot Assessment, and Continuity Plan,” *Technical Report: ENISA*.

Hes, R., and Borking, J. 1995. *Privacy-Enhancing Technologies: The Path to Anonymity*.

Hustinx, P. 2010. “Privacy by Design: Delivering the Promises,” *Identity in the Information Society* (3:2), Springer, pp. 253–255.

ICO. 2023. “Privacy-Enhancing Technologies (PETs),” Information Commissioner’s Office.

Kitchenham, B. A., Budgen, D., and Brereton, P. 2015. *Evidence-Based Software Engineering and Systematic Reviews*, (Vol. 4), CRC press.

Klymenko, O., Kosenkov, O., Meisenbacher, S., Elahidoost, P., Mendez, D., and Matthes, F. 2022. “Understanding the Implementation of Technical Measures in the Process of Data Privacy Compliance: A Qualitative Study,” in *Proceedings of the 16th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*, pp. 261–271.

Klymenko, O., Meisenbacher, S., and Matthes, F. 2023. “Identifying Practical Challenges in the Implementation of Technical Measures for Data Privacy Compliance,” *AMCIS 2023 Proceedings. 2.*

Oetzel, M. C., and Spiekermann, S. 2014. “A Systematic Methodology for Privacy Impact Assessments: A Design Science Approach,” *European Journal of Information Systems* (23:2), Taylor & Francis, pp. 126–150.

Pencarrick Hertzman, C., Meagher, N., and McGrail, K. M. 2013. “Privacy by Design at Population Data BC: A Case Study Describing the Technical, Administrative, and Physical Controls for Privacy-Sensitive Secondary Use of Personal Information for Research in the Public Interest,” *Journal of the American Medical Informatics Association* (20:1), BMJ Group BMA House, Tavistock Square, London, WC1H 9JR, pp. 25–28.

Rachovitsa, A. 2016. “Engineering and Lawyering Privacy by Design: Understanding Online Privacy Both as a Technical and an International Human Rights Issue,” *International Journal of Law and Information Technology* (24:4), Oxford University Press, pp. 374–399.

van Rest, J., Boonstra, D., Everts, M., van Rijn, M., and van Paassen, R. 2014. “Designing Privacy-by-Design,” in *Privacy Technologies and Policy: First Annual Privacy Forum, APF 2012, Limassol, Cyprus, October 10-11, 2012, Revised Selected Papers 1*, Springer, pp. 55–72.

Rubinstein, I. S., and Good, N. 2013. “Privacy by Design: A Counterfactual Analysis of Google and Facebook Privacy Incidents,” *Berkeley Tech. LJ* (28), HeinOnline, p. 1333.

Zhang, P., and Kamel Boulos, M. N. 2022. “Privacy-by-Design Environments for Large-Scale Health Research and Federated Learning from Data,” *International Journal of Environmental Research and Public Health* (19:19), MDPI, p. 11876.