Towards privacy-preserving machine learning in sovereign data spaces: opportunities and challenges

Mehdi Akbari Gurabi[1,2], Felix Hermsen[1,2], Avikarsha Mandal[2], and Stefan Decker[1,2]

[1] RWTH Aachen University, Germany
[2] Fraunhofer FIT, Germany
mehdi.akbari.gurabi@fit.fraunhofer.de

Abstract. The world of big data has unlocked novel avenues for organizations to generate value via sharing data. Current data ecosystem initiatives such as Gaia-X and IDS are introducing data-driven business models that facilitate access to diverse data sources and automate data exchange processes among organizations. However, this also poses challenges for organizations and their customers in preserving control over their own data. This paper provides an overview of the extension requirements on current usage control concepts in data spaces through technical means to augment data privacy guarantees. Our analysis clarifies the deficiencies regarding privacy within the realms of data sovereignty and sovereign data spaces, as well as the risks and opportunities associated with the application of machine learning on sensitive data. This work identifies promising foundational elements and presents areas of research for the integration of privacy-enhancing technologies into usage control for remote data science.

Keywords: Privacy-Preserving Machine Learning · Privacy Enhancing Technologies · Data Sovereignty · Data Spaces

1 Introduction

User data is commonly governed by centralized organizations across various web platforms, presenting a solid contrast to the handling of physical documents where the necessity to maintain control over confidential records is broadly recognized, for instance, by enforcing secrecy, destroying documents, and restricting access. Regrettably, the digital domain often falls short in introducing equivalent privacy practices, resulting in users frequently being uninformed about violations of their privacy [33]. Service providers in the digital realm not only store significant amounts of user data but also maintain a firm grip on personal user

This is the authors' version of the manuscript. The final publication is available at Springer via https://doi.org/10.1007/978-3-031-57978-3_11. Please refer to the final publication for the proper citation of this work.

information through their ability to analyze it and share data with other service providers [10]. On the other hand, enforcing data protection regulations can reduce the utility of data for these service providers and incur financial and human resource costs.

Transitioning to the broader data ecosystem, its primary objective is to derive value from data sharing. However, a significant amount of this data is sensitive, leading data providers to refrain from sharing and sidestep potential benefits and services. The direct transfer of sensitive data to data consumers is not a viable option due to the need for protection. Moreover, for sensitive data, encompassing personal private data and business secrets, traditional security measures or policy-based protections often prove inadequate. [23] This presents a significant challenge in fields such as cybersecurity and medicine and in combating industrial espionage. The advent of privacy-preserving machine learning and remote data science facilitates leveraging data and generating value while minimizing the risk of exposing sensitive information. The risks and opportunities associated with this emerging field are not studied in the current usage control concepts of data ecosystems [15]. Our objective is to inspect the requirements and potential challenges of employing privacy-enhancing technologies (PETs) for machine learning in data spaces. This necessitates addressing the following questions:

- What are the privacy threats in sovereign data spaces?
- How can we specify appropriate PETs for specific use cases in data spaces?
- How can we determine suitable frameworks and nuances in the specifications and parameters of a selected PET for the realization of private computing in data spaces?

To achieve our goals, we first describe data sovereignty, sovereign data space initiatives, and privacy solutions. Subsequently, we identified use cases, opportunities, and challenges for privacy-preserving machine learning within data ecosystems. Additionally, we examined the requirements necessary to address the technological gap in the usage control concept. Lastly, we presented the building blocks for the integration of PETs into data ecosystems in the discussion section.

2 Background

2.1 Data Sovereignty

In a study of 341 publications, Hummel et al. [18] found that the terms "data sovereignty" and "cyber sovereignty" are most commonly associated with countries and indigenous populations. The concept of data sovereignty is discussed in relation to countries, indigenous populations, and various stakeholders, including governmental organizations, the private sector, non-governmental organizations, and private users. While stating theirs in an analysis of the subject, three core concepts of data sovereignty are categorized:

Data sovereignty in regards to indigenous populations: This concept primarily focuses on legislative issues, viewing data sovereignty as a right and

value that should be defined, realized, and upheld by state legislatures [18] [31] [38]. Topics of discussion include self-governance [31] and data ownership as a component of national sovereignty [38].

Data sovereignty in regards to nations: At the national and international levels, research on data sovereignty has primarily concentrated on legal concepts and their implications. This includes governments' rights concerning data stored within their borders [21] and potential conflicts with data privacy regulations [6] [22].

Data sovereignty on the consumer level: Within this domain, data sovereignty is broadly defined as users' control over their own data. This control can be exercised within a specific cloud instance or across a wider data ecosystem. Research in this area has primarily focused on developing technical solutions for challenges associated with big data implementations, such as cloud computing and storage [11], IoT applications [30] [39], and data distribution and sharing [2][17]. Architectural solutions like Gaia-X and International Data Spaces (IDS) are primarily designed to address this particular definition of data sovereignty.

2.2 Sovereign Data Spaces

Initially, dataspaces were primarily focused on addressing heterogeneity in data management. However, current approaches to dataspaces have a paradigm shift, emphasizing data exchange while preserving data sovereignty. [37] In recent years, numerous approaches have emerged that specifically address data protection, storage, and exchange, with a particular emphasis on data sovereignty. These approaches have been realized in varying degrees of adoption and implementation. One notable example is the International Data Spaces Reference Architecture Model (IDS-RAM) [28], which serves as an abstract model for a data ecosystem known as data space. The IDS-RAM facilitates data sharing among different users, their applications, and a range of supporting services. Similarly, the Gaia-X architecture shares similarities with the IDS-RAM; however, it aims to establish and interconnect multiple European Data Ecosystems within a federated framework rather than solely creating them [4]. Moreover, other initiatives such as Private Data System (PDS) and Fiware have their unique approach toward data management and utilization. PDS emphasizes the strict management and protection of data, providing mechanisms and frameworks to allow entities to share information while maintaining privacy and security controls [2]. On the other hand, Fiware, an open-source initiative, furnishes a set of APIs and tools designed to enable the development of smart solutions by the power of data sharing, specifically connected to the Internet of Things. This enables development in various domains such as Smart Cities, Smart Industry, and Smart Agriculture. [3]. All mentioned initiatives underscore the necessity of robust data management solutions and provide distinctive frameworks and tools to facilitate the secure utilization of data across various scenarios and applications, while diverse in their approaches.

Usage Control in Data Spaces Usage control can be seen as an extension of traditional access control mechanisms. Its purpose is to enforce usage restrictions on data even after access has been granted, effectively functioning as an audit mechanism that generates evidence of compliant data usage. This capability addresses security requirements that cannot be fully achieved through traditional access control measures, such as message forwarding. Diverse methods are available for data usage control, each emphasizing distinct advantages and limitations. Some of the most common methods are policy-based usage control, which employs specific policy languages like ODRL or XACML to govern data usage; data provenance, which chronologically tracks data access, providing insights into its usage patterns over time; and token-based usage control, wherein data owners issue tokens to regulate data access and monitor its use. [29] The selection of a suitable usage control method relies on the requirements of the given framework. Usage control is an essential building block in Data space initiatives. It establishes a connection between usage policies and data, enabling continuous monitoring and control over data processing, storage, aggregation, and forwarding activities.

In IDS, usage control sticks policy metadata to the transferred message. A simplified conceptual representation of data transfer with usage control in the IDS framework is depicted in Figure 1. In the broader context of data management and application integration, several technologies, notably LUCON [34], MYDATA, and Degree (D°) have been developed. These technologies vary in aspects [15], such as policy language, enforcement mechanisms, data provenance support, flexibility, and scalability.

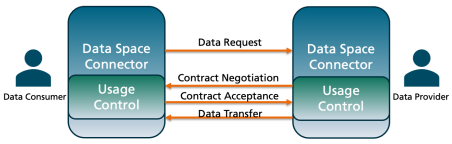

Fig. 1. A simple conceptual view of IDS regarding usage control

Despite all the advances, the existing usage policies provide inadequate measures for data protection. As soon as the data reaches the data consumer, the data provider must rely on the data consumer's promises to comply with the usage policies when processing the data. However, it is essential to ensure compliance with usage policies through technical guarantees for sensitive or regulated data [23].

2.3 Privacy Solutions

Issues surrounding the definition and resolution of privacy concerns are complicated, encompassing various perspectives that demand careful consideration.

First, the definition of privacy itself poses challenges, as it is a concept that evolves and adapts alongside societal, legal, and technological advancements. Protecting privacy can be categorized through these distinct approaches, each advanced in various domains [16]. First, the perspective of privacy as control is often referred to as "soft" privacy. It focuses on legal approaches such as those encapsulated in the GDPR [32], assuming that data controllers, once entrusted, will act responsibly and follow the prescribed guidelines. Second, privacy as confidentiality is often referred to as "hard" privacy. It underscores the technological approach under the assumption that data controllers cannot be inherently trusted, hence recommending the minimization of personal data collection. Third, the paradigm of privacy as practice, or "contextual" privacy. It is acknowledged in the social sciences, recognizes that privacy is contextually enacted. In this realm, technological designs need to be conscientiously aligned with the contexts within which privacy is practiced and perceived.

This trilogy of approaches collectively shapes the conceptualization and strategic protection of privacy across diverse platforms and contexts. In this paper, we primarily focus on the technological approach with the assumption that the data controller may violate legal guidelines and protocols (hard privacy). Despite this, choosing appropriate solutions to protect privacy requires a systematic approach. Decision-makers must evaluate and select technologies, policies, and practices that align with the specific context, requirements, and regulations. This implicates an accurate examination of different components such as PETs, access and usage controls, certification and trust solutions, and consent management systems. It must include assessing their suitability, scalability, transferability, and effectiveness in addressing privacy concerns. Additionally, comprehensive privacy impact assessments and ethical considerations should be included in the decision-making process to ensure the selection of solutions that prioritize privacy protection.

Privacy by Design Privacy by Design (PbD) is an essential approach in privacy engineering, aimed at integrating privacy considerations throughout the product development lifecycle to ensure user privacy protection [5]. Core principles that guide the PbD methodology are [35]: proactive prevention of breaches, privacy as the default setting, privacy integration into the design, full functionality without trade-offs, end-to-end security, visibility and transparency, and respect for user privacy. These principles facilitate the integration of privacy into system design, promoting proactive protection, comprehensive measures, and user-centered privacy considerations. With the increasing importance of data sharing in data ecosystems, adopting the PbD framework has become crucial for regulatory compliance and following privacy standards.

Privacy Engineering Understanding stakeholder requirements is crucial for implementing effective privacy approaches and, consequently, ensuring organizational strategies proficiently address privacy concerns and adopt security measures [24]. The encompassing comprehension of these needs enables organizations

to pinpoint risks and choose suitable privacy engineering tactics, considering that system activities such as data transfer, storage, and processing have substantial implications for user privacy. Privacy impact varies based on factors such as data type, transfer mode, user access, and the three layers of privacy: user sphere, joint sphere, and recipient sphere [36]. The Requirement Engineering Process, as described by [20], mandates the integration of various inputs including existing system information, stakeholder needs, organizational standards, external regulations, and domain information. By incorporating these inputs into the requirement engineering process, organizations can formulate comprehensive requirements and design privacy mechanisms that align with stakeholder needs, comply with organizational standards and external regulations, and are tailored to the unique aspects of the operational domain.

Privacy Enhancing Technologies Even in situations where there is a lack of trust in the identity and certification of the other party, PETs offer potential solutions. These PETs facilitate the exchange of modified, encrypted, or masked data instead of raw data. Examples of PETs include anonymization techniques and specialized data processing procedures. Additionally, cryptographic protocols such as secure multiparty computation (MPC) or homomorphic encryption (HE) can be utilized. These cryptographic techniques enable AI and ML to be conducted without granting the data user access to the raw data or requiring them to disclose their algorithms or models to the data provider. By employing PETs, higher levels of security can be achieved compared to traditional usage control mechanisms that operate on raw data.

The list of traditional PETs includes the basic encryption techniques for data confidentiality, cookie blockers, and communication anonymizers [14]. Nonetheless, to apply PETs effectively in data spaces, PETS must enable data science functionalities ranging from data publishing to data processing, data mining, and machine learning. This paper focuses specifically on PETs for distributed computing and privacy-preserving machine learning (PPML), for instance, Secure Multiparty Computation (MPC) protocols, Federated Learning, or Differential Privacy approaches. MPC involves multiple parties collaboratively evaluating a public function using private inputs without revealing those inputs to each other [12]. In contrast, Federated Learning transfers the machine learning model to the data provider side instead of transmitting raw data to the data consumer. Differential privacy aims to protect privacy by enabling a theoretical privacy guarantee that adds noise to the function output [9].

3 Opportunities and Challenges

The utilization of PETs for ML presents both opportunities and challenges in the subject of data-driven services. In this section, we first explore the use cases for PPML in data spaces, and continue with the potential benefits and drawbacks associated with PETs, categorizing them as opportunities and challenges, respectively. By recognizing these opportunities and challenges, organizations

can make informed decisions and develop strategies to leverage the benefits of PETs while mitigating potential drawbacks and risks.

3.1 PPML in Data Ecosystems

In use cases where highly sensitive data is involved, there are situations where any data exchange through a data space is categorically excluded. This restriction consequently limits the application of AI and ML. This applies to domains such as medicine, energy, banking cybersecurity, and defense against industrial espionage, for instance, in a supply chain use case [19] [13]. There are also use cases in which sharing knowledge is beneficial but challenging due to the confidential and organization-specific nature of the information involved [27]. These scenarios require tailored data protection measures and modeling, for instance, in cybersecurity incident handling and response use cases [25] [1].

Within different domains, there exist numerous examples where the implementation of AI and ML technologies can greatly enhance services. However, data protection concerns pose significant obstacles to leveraging these technologies effectively. General use cases illustrate applications and methods for secure data handling and utilization: creating data-driven services in a secure data space, utilizing multiple databases from both public and private sectors to provide e-services to individuals through techniques such as MPC or HE; linking and analyzing data from varied sources to forge new models and insights, applying SMPC or Federated Learning; and publishing data for ML research via synthetic data generation, differential privacy, and data anonymization.

3.2 Opportunities

Here is a list of main opportunities for enabling PPML in data spaces:

OP1: Supporting of Privacy by Design: PETs fundamentally support PbD by embedding privacy protections into processes. The integrative approach of PPML with PETs ensures that privacy is not an afterthought, but is initially built into the design and architecture of systems and practices. PETs provide mechanisms and techniques that protect user data throughout its lifecycle. Therefore, the adoption of PETs to support PbD helps comply with privacy regulations and emphasizes the commitment to protect user privacy.

OP2: protecting business data in addition to personal data: Enabling PETs protects not only personal data but also critical business secrets, intertwining comprehensive data protection within organizational operations. With an eye on confronting data-centric threats, categorizing them based on the nature of the disclosed data is a valid strategy. This categorization enables constructing mitigation measures that align with the unique characteristics of each threat, managing privacy risks, and thereby ensuring the protection of different categories of sensitive data. Within this schema, the sensitive data disclosure can be categorized into Personal Data Disclosure, which can further delve into Customer and Employee Personal Data Disclosures, and Business Data Disclosure, breaking into Proprietary and Operational Data Disclosures. Through this

structured framework, PETs can be strategically implemented to guard diverse data types, from personally identifiable information (PII) such as customer financial and health data to business-oriented data like proprietary innovations and confidential operational information.

OP3: Increasing data availability and collaboration: This is a vital utility of data ecosystems. This can be achieved by strategically leveraging PETs to facilitate the availability of more data, enabling the sharing of sensitive information that would otherwise remain inaccessible due to privacy constraints. Furthermore, the collaboration is motivated by the reduction of the extensive communication often mandated by complex data-sharing contracts and negotiations among stakeholders. It can attract potential partners who value the ease of accessing a broader range of data without encountering bureaucratic obstacles. Thus, utilizing PETs in a data-sharing ecosystem promises to bring a synergy of amplified data availability and robust collaborations, while protecting sensitive data.

OP4: Enabling the creation of new services: Enabling PPML in data spaces unlocks the potential for the development of innovative services by ensuring that data can be utilized for machine learning without compromising privacy controls. Specifically, through leveraging PPML technologies such as federated learning, HE, and SMPC, data remains protected during processing, mitigating risks associated with data exposure and misuse. Consequently, organizations can explore and create novel, data-driven services. For instance, in sectors such as healthcare or finance, where data sensitivity is vital, PPML facilitates the extraction of valuable insights from data, without revealing individual data points, thereby easing the development of tailored services while following regulatory and ethical obligations.

OP5: Enhancing public acceptance and trust: In the era of digital society, fundamental rights protection, particularly privacy, emerges as an essential principle, shaping the evolution towards a human-centered digital state and opening avenues for a transformative digital developmental leap. Therefore, the deployment of PETs enhances public trust and acceptance, ensuring distributed data protection without the exclusive reliance on major technology corporations. This not only boosts public confidence but also reflects a commitment to privacy, underpinning measurements for the protection of sensitive information. It can be conducive to public support and enhanced user trust in the digital landscape.

3.3 Challenges

Despite all the potential, there is a list of challenges to enabling PETs for ML in data spaces:

CH1: Balancing data utility and privacy protection: The first challenge involves maintaining an equilibrium between data utility and privacy protection. The integration of PETs has the potential to introduce noise or impose limitations on data usage, thereby impacting the overall utility of the data. Ensuring a balance wherein privacy measures do not restrict scenarios of data analysis and utilization is necessary for both the selection and implementation

of PETs, ensuring that the essence and utility of data are not unreasonably compromised.

CH2: Managing computational and communication overhead: Computational and communication overhead is an important challenge in the adoption of PETs, more specifically for cryptographic-based approaches such as MPC and HE. The inherent increase in computational and communication requirements upon employing PETs results in augmented resource demands. This overhead has the potential to impact system performance, response time, and scalability, and make computations infeasible for real-world scenarios. In this context, it is necessary to focus on the optimization and management of resources to uphold system efficiency and efficacy.

CH3: Navigating complexity in usage control architecture integration: Complexity in integrating PETs into existing usage control architectures forms another challenge. The need to ensure that PETs are not only correctly integrated but also compliant with privacy regulations introduces additional layers of complexity. The implementation of PETs goes beyond a single-step solution, demanding a comprehensive, multi-step approach that necessitates various considerations at each stage. This requires additional design and technical efforts to adapt and optimize the architectural frameworks to incorporate PETs.

CH4: Ensuring reputation and robustness amidst privacy measures: The other challenge is regarding upholding reputation and robustness in the application of PETs. Ensuring the absolute robustness and efficacy of PETs is vital to fostering a trusted environment and protecting sensitive data. A comprehensive pre-implementation analysis of system requirements, privacy risks, and desired outcomes is vital, as is the establishment of robust monitoring and auditing processes to ensure sustained compliance and to identify potential vulnerabilities. Furthermore, the protection of data and privacy measures against successful attacks is necessary, not only to protect the data but also to prevent reputational damage.

4 Conceptual Approach to Include PPML into a Data Space

4.1 Mapping PETs and Protection Objectives

To provide an overview of various privacy protection capabilities and facilitate the implementation of technical privacy guarantees, it becomes imperative to categorize PETs according to their respective protection goals and conditions. It is important to note that various PETs can be applied to achieve different outcomes in the same use case, depending on the specific conditions and objectives. As demonstrated in recent studies [26] and [8], different privacy-preserving approaches were employed for a similar classification problem, resulting in varying performance and accuracy levels. This highlights the flexibility and adaptability of PETs, allowing organizations to select the most suitable techniques based on their unique requirements. Table 1 offers a comprehensive categorization of PETs in this context.

Privacy Goal	**PETs**	**Conditions**
Preserve anonymity while publishing data	Data anonymization	Data needs to be shared with an external party, and the adversary model is known. The use case always requires real data (closed-world assumption).
	Synthetic data generation	Data needs to be shared with an external party. Real-world data is not necessary for the analysis, but a quantifiable level of utility is required. [7] The use case only requires stochastic similarities to the original data, and the data can contain a bound amount of hallucinated information (open-world assumption).
Protect privacy in algorithms by adding noise to the output to gain plausible deniability	Differential privacy	A function must be evaluated on privacy-sensitive input. Private data is not shared; only the result of the function is shared with an external party. The use case requires strong provable privacy protection.
Enable computations on encrypted data	Secure multi-party computation	Raw/encoded/anonymized data should not leave the data provider due to high sensitivity. The use case requires aggregated results or computations on encrypted data. Communication resources are not a concern, and response time is not critical.
	Homomorphic encryption	Raw/encoded/anonymized data should not leave the data provider due to high sensitivity. The use case requires only aggregated results or computations on encrypted data. Computational resources are not a concern, and response time is not critical for the use case.
	Functional encryption	Raw/encoded/anonymized data should not leave the data provider due to sensitivity. The use case requires computations on a specific function using sensitive encrypted data. Computational resources are not a concern, and response time is not critical for the use case. It offers medium privacy guarantees due to partial data leakage in the protocols. Involvement of a trusted third party is necessary for key generation.
Train machine learning models in a decentralized manner without sharing data	Distributed learning (E.g., federated learning)	Data does not leave the data provider. Further protections are required for privacy, such as secure aggregation of the results. Usually, a trusted third party is required for model aggregation.
Establish a secure execution environment	Trusted execution environments	Operations on data must be secured with an additional hardware layer when classical cloud services are not secure enough.

Table 1. Mapping PETs to privacy goals and conditions

4.2 PETs Integration into Data Protection Procedure

In the context of PETs integration, several key considerations arise. First, it is imperative that raw sensitive data remains within the Data Provider's control and does not leave their connector, ensuring data privacy. Additionally, the establishment of rules governing the protection of sensitive data is crucial, with agreement between the Data Consumer and Data Provider on protection methods, algorithm hyperparameters, and privacy metrics. Monitoring and logging mechanisms should be implemented to track compliance with the rules and ensure the proper protection of sensitive data. Effective communication between the data provider and data consumer during computations is also necessary for certain PETs, enabling collaboration and facilitating secure computation protocols. Finally, the application of the agreed protection methods should be executed efficiently, whether in relation to the data itself or the computational process. These considerations collectively contribute to the successful implementation and integration of PETs.

Establishing the connection between PETs and data usage processes is crucial because PETs cannot be directly implemented across an entire system. Instead, PETs are applied to specific data usage processes. For example, when a privacy risk is identified and traced back to the data publishing phase, anonymization can be employed to address it. While differential privacy may only be applied during the data processing phase. Therefore, it is vital to identify the relevant data usage processes where PETs can be effectively applied to mitigate privacy risks. Figure 2 illustrates the main idea of the conceptual approach to integrate privacy-preserving machine learning into IDS as the remote data science component. The development of an interface between usage control and PETs is crucial in facilitating the exchange of protected data and enabling secure computations without crossing the data provider's boundaries. Its design should focus on the integration of PETs into the data processing workflow, aiming to enhance privacy protection while preserving the data utility for ML tasks.

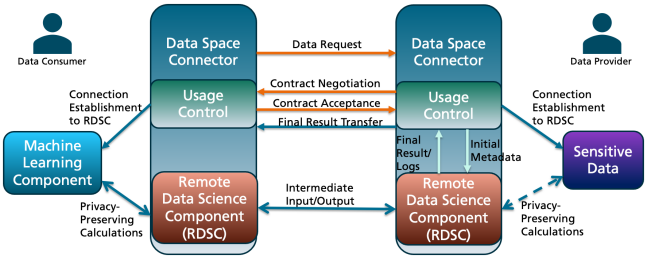

Fig. 2. Conceptual approach to include remote data science component in a data space: this figure depicts the core components to enable PPML as remote data science in addition to a simplified illustration of existing usage control in data space connectors.

When utilizing PETs for a specific scenario, it is crucial to consider multiple factors. Figure 3 shows the conceptual workflow for utilizing PPML in a dataspace. The steps outline the process of facilitating PPML for any scenario in data spaces. First, identifying and classifying sensitive data sets the stage for precise, purpose-driven data processes while recognizing all engaged stakeholders ensures well-rounded data-sharing practices. Also, defining the objectives of data-sharing provides a benchmark for evaluating its success. Afterward, a thorough risk evaluation step seeks to spotlight potential privacy risks and further, it should be used as the input to assess the conceivable consequences of data leakage, protecting against possible exposure of critical information. In the next step, the selection of data exchange type can be done. In this stage, evaluating and selecting optimal PETs (PET Suitability assessment and PET selection) ensures that chosen solutions address identified privacy risks and align with the specific demands of data sharing. Meanwhile, the selection and execution of proper ML approaches should be done. The data exchange type selection can be revisited based on the feedback from PET selection and ML requirements steps. Further, adjusting PETs involves negotiating and fine-tuning hyperparameters to meet desired privacy preservation levels, followed by establishing a secure communication channel to enable streamlined utilization of specific PETs and facilitate efficient data exchanges and computations in PPML (E.g., for SMPC). The application of optimization and pre-computations seeks to augment PPML performance, reducing computational loads and enhancing scalability for specific PETs. Finally, implementing PETs adheres to predetermined protocols, ensuring compliance, and following protective measures, while subsequent monitoring and auditing of execution phases ensure continual regulatory alignment and provide valuable insights for future PPML scenarios.

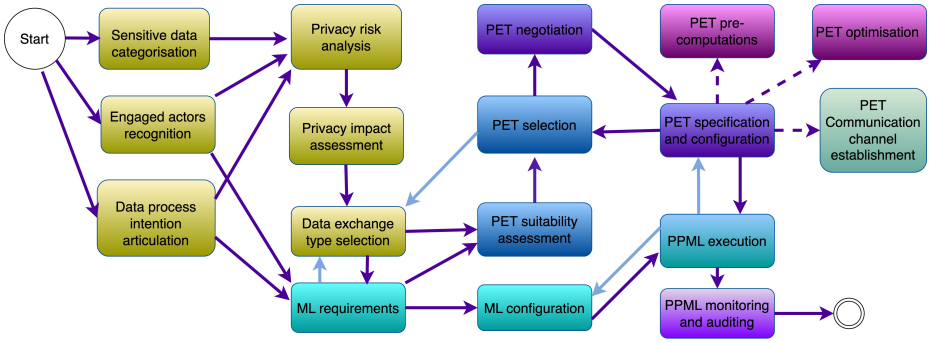

Fig. 3. Conceptual workflow for utilizing PPML in a data space: different categories of steps are shown in different colors, also optional steps are connected via dashed arrows.

Organizations can systematically select and deploy the appropriate PETs to enhance privacy and achieve the goals of their data-sharing initiatives for PPML by following the offered steps.

5 Discussion

In this section, we present the building blocks for leveraging PETs for ML within data spaces and discuss the forthcoming research directions for each building block.

PETs Negotiation To ensure the successful integration of PETs into the usage control framework of data ecosystems, the establishment of a configurable interface is imperative. This interface serves as a platform for negotiation and implementation between the data provider and data consumer, facilitating the selection of trust-building mechanisms that align with the specific characteristics of the data being exchanged, including its type, flexibility, and intended use.

In this context, the presence of sticky policies regarding the choice of PETs facilitates the decision-making process. Sticky policies are policies that persist and remain associated with the data throughout its lifecycle, influencing how the data is handled and processed. However, the absence of a standardized approach for describing PET decisions poses a challenge in this regard. Currently, there is no universally recognized format or framework for documenting and communicating PET decisions consistently. Addressing this issue requires the development of standardized protocols or guidelines that can provide a structure for PET decisions. This would enhance transparency, interoperability, and understanding among stakeholders involved in the usage control framework.

PETs Risk Assessment The objective of this building block is to create a traceable privacy risk assessment framework for sovereign data spaces, allowing for the identification and tracing of data usage processes associated with predicted privacy risks. This can be achieved by developing a privacy risk tree that leverages predictive techniques to identify potential risks. By establishing connections between privacy risks and data usage processes, a comprehensive understanding can be gained regarding the emergence of privacy risks within sovereign data spaces and the specific processes that contribute to these risks. Ultimately, this framework will enable the formulation of more targeted and efficient strategies for mitigating privacy risks in the context of usage control.

PETs Selection Once we have gained a comprehensive understanding of the privacy risks associated with a data-centric information system, including their type, severity, impact, and likelihood, as determined through the implementation of a PETs Risk Assessment building block, our focus shifts toward developing a methodology for the selection of suitable PETs and their respective hyperparameters. This process can be initiated by multi-criteria decision-making. This selection process takes into account the practical use case and the specific privacy risks identified, to effectively address the targeted privacy risks.

PETs Monitoring and Evaluation It is necessary to establish a monitoring and evaluation system that enables the measurement of the impact resulting from the application of PETs to the information system. This evaluation system plays a critical role in determining whether the implementation of PETs successfully eliminates or mitigates the identified privacy risks by reducing their likelihood or severity. By evaluating the effectiveness of the applied PETs, organizations can assess the level of privacy protection achieved and make informed decisions regarding risk mitigation strategies. All in all, by adopting a systematic approach that encompasses pre-implementation analysis, thorough measurement, and continuous monitoring and auditing, organizations can effectively enable PETs and enhance their privacy protection efforts. Evaluating the impact of PETs on privacy risks and data sovereignty via conducting thorough audits enables organizations to learn from the evaluation process and make improvements to enhance the effectiveness and efficiency of the protocol in future implementations.

PETs Execution and Communication The PETs execution building block is responsible for implementing the PETs and conducting client-side computations. Throughout this process, it should continuously send status reports, which are closely monitored to detect any anomalies or irregularities. Once the execution is complete, the final results are transmitted. These results can take various forms, such as protected data through anonymization techniques or the outcome of a secure computation. The communication building block plays a critical role in facilitating distributed and remote computing, such as MPC and federated learning, by enabling intermediate communications between parties to support collaborative learning.

PETs Pre-computation and Optimization Another important building block focuses on enhancing the efficiency of the system. PETs often face challenges related to their computationally intensive processes or the communication overhead involved. However, the efficiency can be improved through precomputations, especially in crypto-based approaches like MPC. Additionally, various optimization techniques are available, for instance, with regard to privacy-preserving data publishing via anonymization or synthetic data generation, which can further improve system performance. These optimization approaches play a crucial role in mitigating the computational challenges associated with PETs, making them more efficient, scalable, and practical for real-world applications.

6 Conclution

The main objective of data spaces is to facilitate secure data exchange and utilization, which can be enhanced by incorporating technical data protection guarantees through PETs. This paper investigates the concepts of data sovereignty, sovereign data space initiatives, and privacy solutions. We also present use cases

for privacy-preserving machine learning in data ecosystems and address the existing technological gaps in the usage control concept. In this regard, we analyze the opportunities and challenges of utilizing PETs and propose an initial concept that enables the integration of PETs into the usage control, allowing for the enforcement of privacy requirements on the data. We discuss the conceptual approach and outline the main building blocks necessary for integrating PETs into data spaces for privacy-preserving machine learning. The list of building blocks consists of negotiation, risk assessment, selection, monitoring, execution, and optimization components, each of which requires further scientific research and technological development to bridge the existing gaps. In future work, we aim to refine the initial concept to meet all requirements of PETs. We will also focus on developing an interface between the usage control and remote data science component while conducting research on the mentioned building blocks to realize the proposed approach.

7 Acknowledgements

This is the authors' version of the manuscript. The final publication is available at Springer via https://doi.org/10.1007/978-3-031-57978-3_11. Please refer to the final publication for the proper citation of this work. This work was funded by the TANGO project and partly supported by the BMBF-ANR-funded project Crypto4Graph-AI (funding number 01IS21100A). TANGO project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 101070052.

References

1. Akbari Gurabi, M., Mandal, A., Popanda, J., Rapp, R., Decker, S.: Sasp: a semantic web-based approach for management of sharable cybersecurity playbooks. In: Proceedings of the 17th International Conference on Availability, Reliability and Security. pp. 1–8 (2022)
2. Alboaie, S., Cosovan, D.: Private data system enabling self-sovereign storage managed by executable choreographies. In: Distributed Applications and Interoperable Systems: 17th IFIP WG 6.1 International Conference, DAIS 2017, Held as Part of the 12th International Federated Conference on Distributed Computing Techniques, DisCoTec 2017, Neuchâtel, Switzerland, June 19–22, 2017, Proceedings 17. pp. 83–98. Springer (2017)
3. Araujo, V., Mitra, K., Saguna, S., Åhlund, C.: Performance evaluation of fiware: A cloud-based iot platform for smart cities. Journal of Parallel and Distributed Computing **132**, 250–261 (2019)
4. Autolitano, S., Pawlowska, A.: Europe's quest for digital sovereignty: Gaia-x as a case study. IAI papers **21**(14), 1–22 (2021)
5. Caiza, J.C., Martín, Y.S., Guamán, D.S., Del Alamo, J.M., Yelmo, J.C.: Reusable elements for the systematic design of privacy-friendly information systems: A mapping study. IEEE Access **7**, 66512–66535 (2019)
6. Courtney, M.: Regulating the cloud crowd. Engineering & Technology **8**(4), 60–63 (2013)

7. Dankar, F.K., Ibrahim, M.: Fake it till you make it: Guidelines for effective synthetic data generation. Applied Sciences **11**(5), 2158 (2021)
8. Drichel, A., Akbari Gurabi, M., Amelung, T., Meyer, U.: Towards privacy-preserving classification-as-a-service for dga detection. In: 2021 18th International Conference on Privacy, Security and Trust (PST). pp. 1–10. IEEE (2021)
9. Dwork, C., Roth, A., et al.: The algorithmic foundations of differential privacy. Foundations and Trends® in Theoretical Computer Science **9**(3–4), 211–407 (2014)
10. Ernstberger, J., Lauinger, J., Elsheimy, F., Zhou, L., Steinhorst, S., Canetti, R., Miller, A., Gervais, A., Song, D.: Sok: Data sovereignty. Cryptology ePrint Archive (2023)
11. Esposito, C., Castiglione, A., Choo, K.K.R.: Encryption-based solution for data sovereignty in federated clouds. IEEE Cloud Computing **3**(1), 12–17 (2016)
12. Evans, D., Kolesnikov, V., Rosulek, M., et al.: A pragmatic introduction to secure multi-party computation. Foundations and Trends® in Privacy and Security **2**(2-3), 70–246 (2018)
13. Gaia-X: Gaia-x usecases, https://gaia-x.eu/use-cases/ (last accessed 30/11/2023)
14. Giaconi, G., Gunduz, D., Poor, H.V.: Privacy-aware smart metering: Progress and challenges. IEEE Signal Processing Magazine **35**(6), 59–78 (2018)
15. Gil, G., Arnaiz, A., Higuero, M., Diez, F.J.: Assessment framework for the identification and evaluation of main features for distributed usage control solutions. ACM Transactions on Privacy and Security **26**(1), 1–28 (2022)
16. Gürses, S.: Can you engineer privacy? Communications of the ACM **57**(8), 20–23 (2014)
17. Hoffmann, A., Wagner, A., Huyeng, T., Shi, M., Wengzinek, J., Sprenger, W., Maurer, C., Rüppel, U.: Distributed manufacturer services to provide product data on the web. In: EG-ICE (2018)
18. Hummel, P., Braun, M., Tretter, M., Dabrock, P.: Data sovereignty: A review. Big Data & Society **8**(1), 2053951720982012 (2021)
19. IDS: International data spaces usecases overview, https://internationaldataspaces.org/make/use-cases-overview/ (last accessed 30/11/2023)
20. Inflectra.com: Principles of requirements engineering or requirements management 101. (March 2018), https://www.inflectra.com/Ideas/Whitepaper/Principles-of-Requirements-Engineering.aspx. (last accessed 14/07/2023)
21. Irion, K.: Government cloud computing and national data sovereignty. Policy & Internet **4**(3-4), 40–71 (2012)
22. König, P.D.: The place of conditionality and individual responsibility in a "data-driven economy". Big Data & Society **4**(2), 2053951717742419 (2017)
23. Lohmöller, J., Pennekamp, J., Matzutt, R., Wehrle, K.: On the need for strong sovereignty in data ecosystems. Universitätsbibliothek der RWTH Aachen (2022)
24. Mead, N.R., Miyazaki, S., Zhan, J.: Integrating privacy requirements considerations into a security requirements engineering method and tool. International Journal of Information Privacy, Security and Integrity **1**(1), 106–126 (2011)
25. Nitz, L., Gurabi, M.A., Mandal, A., Heitmann, B.: Towards privacy-preserving sharing of cyber threat intelligence for effective response and recovery. ERCIM NEWS **126**, 33 (2021)
26. Nitz, L., Mandal, A.: Dga detection using similarity-preserving bloom encodings. In: European Interdisciplinary Cybersecurity Conference. pp. 116–120 (2023)

27. Nitz, L., Zadnik, M., Gurabi, M.A., Obrecht, M., Mandal, A.: From collaboration to automation: A proof of concept for improved incident response. ERCIM NEWS **129** (2022)
28. Otto, B., Steinbuss, S., Teuscher, A., Lohmann, S., et al.: Ids reference architecture model(version 3.0). international data spaces association (2019)
29. Pretschner, A., Hilty, M., Schütz, F., Schaefer, C., Walter, T.: Usage control enforcement: Present and future. IEEE Security & Privacy **6**(4), 44–53 (2008)
30. Qarawlus, H., Hellmeier, M., Pieperbeck, J., Quensel, R., Biehs, S., Peschke, M.: Sovereign data exchange in cloud-connected iot using international data spaces. In: 2021 IEEE Cloud Summit (Cloud Summit). pp. 13–18. IEEE (2021)
31. Rainie, S.C., Schultz, J.L., Briggs, E., Riggs, P., Palmanteer-Holder, N.L.: Data as a strategic resource: Self-determination, governance, and the data challenge for indigenous nations in the united states (2017)
32. Regulation, P.: Regulation (eu) 2016/679 of the european parliament and of the council. Regulation (eu) **679**, 2016 (2016)
33. Saleem, H., Naveed, M.: Sok: Anatomy of data breaches. Proc. Priv. Enhancing Technol. **2020**(4), 153–174 (2020)
34. Schütte, J., Brost, G.S.: Lucon: Data flow control for message-based iot systems. In: 2018 17th IEEE international conference on trust, security and privacy in computing and communications/12th IEEE international conference on big data science and engineering (TrustCom/BigDataSE). pp. 289–299. IEEE (2018)
35. Semantha, F.H., Azam, S., Yeo, K.C., Shanmugam, B.: A systematic literature review on privacy by design in the healthcare sector. Electronics **9**(3), 452 (2020)
36. Spiekermann, S., Cranor, L.F.: Engineering privacy. IEEE Transactions on software engineering **35**(1), 67–82 (2008)
37. Theissen-Lipp, J., Kocher, M., Lange, C., Decker, S., Paulus, A., Pomp, A., Curry, E.: Semantics in dataspaces: Origin and future directions. In: Companion Proceedings of the ACM Web Conference 2023. pp. 1504–1507 (2023)
38. Walter, M., Suina, M.: Indigenous data, indigenous methodologies and indigenous data sovereignty. International Journal of Social Research Methodology **22**(3), 233–243 (2019)
39. Yin, H., Guo, D., Wang, K., Jiang, Z., Lyu, Y., Xing, J.: Hyperconnected network: A decentralized trusted computing and networking paradigm. IEEE Network **32**(1), 112–117 (2018)