Computer Fraud and Security
ISSN (online): 1873-7056

Safeguarding Personal Information Privacy in AI-Driven Data Engineering: Challenges and Protection Strategies

Anandan Dhanaraj

New Jersey Institute of Technology, USA

Abstract

The interation of data engineering processes with artificial intelligence has essentially transformed the handling of personal information by organizations, posing an enormous privacy risk in real-time data processing and automated workflow configurations. Ordinary security cannot be used to address the special issues that systems based on AI that continuously process streaming data and are distributed across architectures present. In the case of machine learning models operating on high-velocity data streams, they introduce points of exposure to the data lifecycle, specifically during ingestion and model inference. This implies that privacy protection strategies should be algorithmic and operate with strict time and computational constraints. Privacy-preserving technologies, such as differential privacy, federated learning, homomorphic encryption, and secure multi-party computation, among others, can be used to safeguard sensitive data and enable AI applications to work with the required data simultaneously. Privacy-by-design principles and privacy engineering methodology are organizational frameworks that provide systematic means of integrating privacy protections into every phase of system development. Legal requirements, including technical controls and governance processes, are outlined by the GDPR, CCPA, and new AI-specific rules to ensure data engineers comply with them. Privacy technologies do not only fulfill the legal requirements, but their impact on the economy and society is far more drastic. They radically alter the competitive landscape of the business, the level of trust that customers have in one another, and the operation of digital rights in societies that are becoming more data-driven. The decisions that organizations make about privacy protection must be balanced against their sustainability commitments, since privacy systems with high consumption of computing resources do have environmental impacts. Active privacy engineering strategies assist companies in achieving their innovation as well as privacy objectives simultaneously. This will provide them with a competitive advantage through greater trust of stakeholders and regulatory strength, as well as promote the ethical application of AI in accordance with democratic principles and individual autonomy.

Keywords: Artificial Intelligence Privacy, Real-Time Data Processing Security, Federated Learning, Differential Privacy Mechanisms, Privacy Engineering Frameworks

Introduction

1.1 The Convergence of AI and Data Engineering

The modern data environment has also experienced a basic shift following the integrative inclusion of artificial intelligence technology in organizational data engineering practices. Based on a thorough study of AI adoption patterns among organizations that are global organizations, it has been reported that organizations have increased their adoption of AI systems in various business functions faster by far, and adoption rates have soared at a large scale across various business sectors, including manufacturing and financial services [1]. This broad usage is accompanied by the rise of real-time data processing architectures, which allow a continuous ingestion, processing, and action on data as it is generated. The change is not only in technological enhancements but a paradigm shift in the way organizations have to conceive data value mining and decision-making processes in their operations.

The real-time data processing architectures have become fundamental to the contemporary operations of enterprises, as organizations are in a position to extract real-time insights and stay competitive in a world that is rapidly changing. Apache Kafka, Apache Flink, and cloud-native streaming platforms are used to handle large volumes of events and power an application that may be fraud detection or personalized recommendations, predictive maintenance, or autonomous decision-making systems. Nevertheless, this technological responsiveness creates severe privacy concerns that require systematic consideration. The introduction of AI into data engineering processes creates vulnerabilities that traditional security practices often do not manage appropriately. The speed of streaming data poses specific difficulties,

whereby the latency specifications in milliseconds limit the possibilities of full privacy validation and consent authorization prior to the functioning taking place [2]. With streaming data running across distributed systems, including edge devices to cloud providers, personal data is exposed to unauthorized access, inference attacks, and regulatory non-compliance, and requires urgent demands for systematic privacy protection measures that may run within the time and computational limitations of real-time processing systems.

Dimension	Organizational Context	Technical Implications
AI Deployment Scope	Enterprise-wide implementation across business functions	Integration complexity with legacy systems
Processing Architecture	Streaming platforms enable continuous data ingestion	Latency constraints limit validation opportunities
Data Velocity	Event processing measured in milliseconds	Privacy verification must execute within temporal windows
Infrastructure Distribution	Multi-zone deployments across geographic regions	Multiple exposure points throughout the processing pipeline
Consent Management	Real-time authorization enforcement requirements	Propagation delays across distributed nodes

Table 1: AI Adoption Patterns and Real-Time Processing Challenges [1,2]

2. Privacy Problems with AI-based Real-Time Data Processing.

This adoption of artificial intelligence in automated decision-making systems poses significant privacy issues beyond the conventional data protection issues. Automated decision-making refers to procedures such that AI systems arrive at conclusions regarding people with no significant involvement of a human, and they could potentially impact their rights in the law, their economic opportunities, or their rights to services in general. Regulatory guidance by data protection authorities dictates that automated decision-making systems that have a legal consequence or other impact of this magnitude on an individual necessitate certain protection measures, such as the right to seek human intervention, give opinions, and appeal decisions [3]. These needs present significant implementation ailments to data engineering experts developing real-time processing pipelines in which human supervision should be traded off against operational latency specifications.

The fact that profiling activities are used to examine personal information to assess, anticipate, or classify personal attributes, behaviors, or preferences presents further privacy risks in AI-driven systems. Many machine learning algorithms have a high degree of opaqueness that makes it difficult to offer meaningful explanations of automated decisions, which is a requirement in several data protection frameworks. The tension between the complexity of algorithms and the transparency requirements in organizations deploying AI systems requires that organizations solve the problem of letting data subjects comprehend the rationale on which decisions affecting them are made. Also, dynamic profiling risks come about in the streaming data environments where models continuously train on incoming data since individual profiles can be updated in real-time on the basis of recent behavior patterns without sufficient prospects of data subjects to access or challenge such characterizations.

The privacy issue with automated decision-making involves the issue of algorithm biases and discrimination. The existence of historical data, which is biased, can result in discriminatory outcomes, and subsequently, the AI models that are trained on these biases can be used to further increase them or reinforce the already existing biases in society. In data engineering systems, processes to identify and address bias need to be added at all stages of the model lifecycle, including training data selections and continuous monitoring of the production predictions. Livestream processing systems make bias detection processes more difficult because the properties of streaming data can change over time, and these changes bring in new sources of bias that cannot be automatically detected by traditional validation methods. The

organizations are forced to adopt ongoing monitoring systems that evaluate model fairness indicators among demographic subgroups and ensure processing throughput demands required to operate real-time applications.

The problem of offering enough transparency in automated systems of decision-making is aggravated by the technical limitations of real-time data processing. The process of generating explanations of why AI decisions should be made must operate simultaneously with predictions, so it can serve to provide immediate user feedback or response to appeals, though producing high-quality explanations puts an extra computational load on the predictions that can exceed latency requirements. Data engineers should be cautious of coming up with the description mechanisms that would be meaningful enough to give insights into the decision logic without impacting the performance of the system itself. This usually necessitates intermediate solutions that use pre-computed explanation templates along with on-the-fly customization to particular input characteristics so as to guarantee transparency requirements whilst allowing reasonable response times to production systems serving end-users or downstream applications.

2.2 Latency Constrainment of Privacy and Validation.

The real-time data streaming systems are run based on strict latency requirements, which essentially define the privacy protection policies. The explanation of time properties of distributed systems becomes critical in achieving good privacy control in streaming systems. A study of latency of data engineering in data streaming engineering has shown that delays that are network round-trip time, disk access patterns, and memory operations can be measured and accumulated over multiple stages of processing pipelines [4]. In situations where privacy validation mechanisms are required to operate within these temporal limits (e.g., consent verification, data quality checks, anomaly detection), then engineers have challenging tradeoffs between complete privacy protection and acceptable system performance.

The decentralized streaming architecture increases privacy concerns beyond latency concerns. Data moving through distributed systems passes through many processing points in various compute nodes, availability zones, and possibly different geographic locations, with each point implying a possible point of exposure to unauthorized access or data leakage. The control of privacy needs to be applied at every level without creating cascaded delays that impact the end-to-end latency specifications. In systems with sensitive personal data like financial transactions or medical records, the combined latency of privacy measures in all pipeline stages should be limited to acceptable levels, which can be in the hundreds of milliseconds range with user-facing applications or in the seconds range with analytics pipelines at the back end.

Smart patterns of in-memory processing used by streaming systems to provide low-latency performance also pose further privacy implications. The speed of processing data stored in memory and not in persistent storage is better, but it complicates encryption-at-rest solutions, and can put sensitive data at risk of memory-based attacks or accidental logging. Organizations need to carefully analyze whether data in memory needs to be encrypted or not, and they need to understand that cryptographic operations add computational overhead, which can affect throughput capacity. Other methods, such as isolation of memory containing and hardware-enhanced trusted execution environments, are privacy-protecting and have less effect on performance, but these technologies require the addition of more complexity in architecture and operational management overhead.

The streaming data processing is continuous, posing special challenges to privacy impact assessment and compliance validation. Streaming systems, in contrast to batch processing systems, where discrete processing jobs may be comprehensively analyzed in terms of their privacy consequences, run continuously with potentially changing processing logic in response to runtime state or the evolving structure of adaptive models. Privacy tests should consider dynamic system behavior, which may not be completely defined during the design phase and requires continuous monitoring and validation methods that track privacy breaches in production contexts. Checking of privacy compliance. Automated privacy compliance checking at streaming pipelines allows spotting of possible problems like unauthorized data access patterns, unexpected data quality degradation, potentially exposing sensitive data, or consent violations with the data being processed to uses other than those allowed by user authorizations. These computerized systems need to have a low latency overhead and have a high accuracy rate to prevent false positives that will lead to unneeded investigations or system errors.

Regulatory Aspect	Individual Rights	Implementation Considerations
Human Intervention	Right to obtain human review of automated decisions	Workflow integration balancing oversight with latency
Decision Contestation	Ability to express views and challenge outcomes	Appeal mechanisms with audit trail preservation
Transparency Obligations	Meaningful explanations of decision logic	Computational overhead for synchronous explanation generation
Profiling Safeguards	Protection against discriminatory characterizations	Continuous monitoring across demographic subgroups
Temporal Constraints	Real-time explanation generation requirements	Hybrid approaches combining templates with runtime customization

Table 2: Automated Decision-Making Privacy Requirements [3,4]

3. Privacy-Protecting Systems and Design Methodologies

3.1 Intelligent Privacy-Protecting AI Systems.

Creation and application of privacy-sustaining methods that are specifically crafted to be applied to artificial intelligence systems has grown at an ever quicker pace since organizations have become aware of the need to safeguard personal data throughout machine learning lifecycles. A close examination of privacy-saving solutions to AI has shown that several technical features can be used to implement a defense-in-depth solution to a range of threat models and attack patterns [5]. Differential privacy is one of the most mathematically precise methods, in which formal assurances are given that individual data input cannot be differentiated by aggregate outcomes in a methodical addition of noise. Federated learning represents another foundational privacy-preserving technique enabling collaborative machine learning without centralizing raw data across organizational boundaries. By training local models on decentralized data sources and aggregating only model parameters rather than raw records, federated learning substantially reduces privacy risks associated with data centralization. Research examining federated learning implementations demonstrates significant architectural considerations for deployment at scale, including communication efficiency challenges when coordinating model updates across distributed participants, strategies for handling non-independently and identically distributed data across federated nodes, and security mechanisms preventing malicious participants from poisoning the global model through adversarial local updates [6]. The integration of federated learning into data engineering pipelines requires careful consideration of network topology, bandwidth constraints, and synchronization protocols, ensuring model convergence while maintaining acceptable training iteration times. Homomorphic encryption techniques enable computations on encrypted data without requiring decryption, providing strong confidentiality guarantees even when processing occurs in untrusted environments such as public cloud infrastructure. While computationally intensive, partially homomorphic encryption schemes supporting specific operations such as addition or multiplication have become increasingly practical for production deployment in privacy-sensitive applications. Organizations implementing homomorphic encryption must carefully evaluate which portions of their processing pipelines benefit most from encrypted computation, recognizing that the substantial performance overhead makes universal application impractical for high-throughput streaming systems. Hybrid architectures combining homomorphic encryption for the most sensitive operations with conventional processing for less critical computations provide a balanced approach, achieving acceptable privacy protection without completely sacrificing system performance. Secure multi-party computation protocols enable multiple parties to jointly compute functions over their private inputs while preventing any participant from learning others' data beyond what can be inferred from the final computation result. These cryptographic protocols prove particularly valuable for cross-organizational collaborations where competitive concerns or regulatory constraints prevent direct data sharing.

Mechanism	Privacy Guarantee Type	Deployment Considerations
Differential Privacy	Mathematical indistinguishability of individual contributions	Noise calibration balancing privacy and utility
Federated Learning	Decentralized training without raw data centralization	Communication efficiency across distributed participants
Homomorphic Encryption	Computation on encrypted data	Performance overhead limiting universal application
Secure Multi-Party Computation	Joint computation preserving input privacy	Protocol selection based on threat models
Hardware Enclaves	Isolated execution environments	Memory-based attack mitigation strategies

Table 3: Privacy-Preserving Technique Characteristics [5,6]

4. Compliance and Implementation Problems with Regulations.

4.1 California Privacy Protection Agency Enforcement Landscape.

The regulatory landscape of AI and data privacy has become particularly heated, with enforcement agencies increasing their operations and setting precedents on what is expected in terms of compliance. The latest imposition measures of the California Privacy Protection Agency show more advanced methods to detect and punish privacy breaches, especially those concerned with automated decision-making systems and AI-based data processing [8]. The enforcement landscape shows several new trends, such as increased scrutiny over the data deletion practices of organizations, scrutiny over whether privacy notice reflects the reality of actual data collection and use practices, and scrutiny of whether organizations have sufficient mechanisms that allow consumers to exercise their privacy rights, such as access, deletion, and opt-out preferences. The priority of enforcement is increasingly in the compliance of the organizations with the privacy decisions of the consumers, ie, real-time systems. The Agency has insisted that automated processes of personal information should have real-time or near-real-time enforcement of consumer opt-out preferences, instead of batch processing privacy requests that may permit further processing hours or days after consumers exercise their rights. The Agency's enforcement actions reveal particular concerns regarding third-party data sharing and the sale of personal information. Organizations employing AI systems that share data with vendors, partners, or service providers face scrutiny regarding whether these transfers constitute "sales" under California law, requiring explicit consumer opt-out rights. The technical implementation of opt-out mechanisms must extend beyond the organization's direct systems to encompass downstream recipients, requiring contractual provisions and technical controls ensuring third parties also honor consumer privacy preferences. Data engineering architectures must incorporate mechanisms tracking data lineage across organizational boundaries, enabling organizations to demonstrate compliance with consumer deletion requests by identifying and remediating all copies of personal information, including those held by external parties.

4.2 Economic and Social Impact of Privacy Technology.

The wider meaning of privacy-enhancing technologies is far more than the regulatory compliance and covers the inherent aspects of economic relations and the nature of society in a more data-intensive environment. In-depth review of effects of privacy-enhancing technologies on business, individuals, and society presents complex results in terms of economic, social, and technological outcomes [9]. Economically, those organizations that have invested in privacy technologies gain competitive advantages in terms of increased customer confidence, less exposure to regulatory risks, and more data partnership possibilities hitherto restricted because of privacy barriers. Privacy-saving methods allow new business models in which organizations have the opportunity to share machine learning projects or data analytics initiatives without disclosing organizational proprietary or sensitive information, generating value through the joint knowledge and retaining data sovereignty.

The positive social impacts of privacy-enhancing technologies are expressed in a greater sense of personal autonomy and resistance against the dynamics of surveillance capitalism, in which individual data are commodified without providing

any benefit to the user or sharing services with others. These technologies promote democratic values and digital rights that are critical to building healthy information societies by offering technical systems that allow people to play their part in collective intelligence with AI systems without jeopardizing personal privacy. Healthcare applications are especially shown to have strong benefits to society, in which privacy-enhancing collaborative learning provides a means of medical research and development of a diagnostic model based on a wide range of patients, and keeps sensitive health data confidential, ensuring that patients have faith in the medical institutions.

Nevertheless, implementing privacy-protecting technologies is also associated with significant equity concerns that need to be tackled by organizations and policymakers. The technical complexity and computational costs of sophisticated privacy controls can generate digital divides in which organizations with financial means and those with higher incomes have better privacy safeguards, and society's underserved groups are subject to heavy surveillance and exploitation of data. Companies need to think about how privacy protection measures unwillingly form a system of tiers where higher-quality services secure high privacy guarantees and the free or low-cost offerings save little or no privacy, which may further widen disparities and ensure the right to privacy only to the parties able to afford it.

The long-run economic effects of privacy technologies are data market and value distribution patterns in digital ecosystems. With privacy-saving methods becoming more advanced and broadly implemented, the business model of the past, which relies on the unlimited collection and monetization of data, is under threat, and organizations need to find other ways to create value and take into account the privacy of users and ensure their economic sustainability. The change provides a challenge to existing organizations that have established their business models on data aggregation and an opportunity to new entrants that can come with privacy-centred alternatives appealing to privacy-conscious consumers. Privacy technologies, as such, have not only emerged as a result of technical or regulatory change but as a fundamental reorganisation of digital economy dynamics with far-reaching effects on the manner in which value is created, distributed, and captured among economic actors.

Consideration	System Impact	Organizational Response
Communication Overhead	Bandwidth requirements are proportional to model complexity	Network topology optimization and compression
Computational Intensity	Orders of magnitude increases for cryptographic operations	Hardware acceleration through specialized processors
Memory Utilization	Increased footprints for encrypted intermediate results	Infrastructure provisioning and capacity planning
Enforcement Priorities	Real-time consumer privacy choice implementation	Consent management platform architecture
Third-Party Controls	Data lineage tracking across organizational boundaries	Contractual provisions and technical safeguards

Table 4: Performance and Enforcement Trade-offs [7,8]

5. Future Prospects and Environmental Consciousness.

5.1 Environmental Implications of Privacy-Preserving Computation.

The environmental facets of privacy-preserving technologies are also becoming of great concern, with organisations weighing the privacy preservation goals and their sustainability and environmental accountability. The proposed environmental footprint of privacy-preserving computation, research indicates that there are serious energy consumption considerations in the cryptographic methods and distributed learning algorithms [10]. Encryption methods with homomorphic encryption, such that computations can be done on encrypted data, introduce significant computational load over plaintext methods, and energy usage is proportional to the complexity of the encryption scheme and the complexity of computations done on encrypted data. The type of organizations implementing such methods needs to take a close look at the environmental expenses and privacy advantages, which may restrict homomorphic encryption to the least privacy-intensive tasks in which a high level of confidentiality is warranted, with a higher energy usage.

The tradeoffs in federated learning methods, computing distributed to edge devices, include complicated environmental tradeoffs relative to centralized data center training. Although distributed computation can save the cost of data transmission energy, and could also possibly harness the idle edge devices' capacity, it raises the total energy consumption of millions of distributed computers, which may be less energy efficient than optimized data center facilities. The overall environmental impact is based on issues such as the energy efficiency of the devices involved, the carbon intensity of the electricity of the distributed infrastructure compared to the central data centers, and how often the model needs to update itself with distributed computing. Federated learning organizations should also carry out full lifecycle analyses that help them analyze the overall effects they have on the environment, as opposed to looking at data center power usage in isolation.

The privacy requirements and data retention policies interact to form more environmental considerations. The unnecessary storage can be minimized by privacy laws that require personal information to be deleted as soon as it is no longer used for its initial collection, which will lower the energy requirements of storage facilities and the cooling systems installed to support them. Automated data lifecycle with privacy-based retention policies has two purposes of compliance and environmental sustainability. Privacy methods like differential privacy, however, might demand that historical records of privacy budget spending be retained to avoid privacy leakage due to repeated queries, which creates a tension between the principles of data minimization and privacy accounting that organizations must tread carefully.

Privacy-preserving computation methods are increasingly approaching environmental sustainability as a design (as well as privacy protection and performance) requirement. Future research directions involve in-house design of energy-efficient cryptographic protocols based on modern processor architectures, hardware acceleration schemes that minimize the amount of compute required on privacy-intensive operations, and hybrid architectures that strategically assign privacy-intensive workloads to energy-efficient hardware. The emergence of environmentally-friendly privacy technologies is a pressing research and engineering problem with far-reaching long-term implications on the sustainability of privacy-respecting data processing on a large scale, as organizations are increasingly pressured to act with greater environmental responsibility and simultaneously ensure privacy.

Conclusion

The convergence of artificial intelligence and data engineering is a paradigm shift in the management of information in organizations, in which technological possibilities of concluding personal information have increased many times at the same time, while the risk of privacy invasion has also increased significantly, necessitating systematic reduction. The security dilemmas of defending personal data in AI-based real-time data processing systems are based on inherent properties of streaming systems, such as strict latency policies that restrict privacy verification chances, distributed processing that introduces multiple points of exposure, continuous training of models that facilitate a probability of an inference attack, and the use of automated decision-making systems that impact people without sufficient disclosure or human supervision. Privacy-preserving algorithms have now evolved to the level of providing mathematically rigorous privacy guarantees via differential privacy, supporting collaborative machine learning via federated learning networks, supporting encrypted computations via homomorphic encryption systems, and supporting secure multi-organizational computing via secure multi-party computation protocols. To achieve successful implementation of these mechanisms, performance implications should be carefully considered, along with computational overhead, communication bandwidth needs, memory utilization patterns, and latency properties that could limit their applicability to the real-time case. The enforcement of regulations has become very aggressive, such that the agencies show advanced skills in detecting privacy breaches, and they impose huge sums of money, which reflect the extent of the breaches and the harm they have demonstrated against the individuals harmed. The enforcement priorities of the California Privacy Protection Agency focus on real-time execution of consumer privacy preferences, overall third-party data distribution authorities, and correct representation of data procedures in privacy statements, posing technical implementation difficulties for businesses running complicated distributed systems. Economic sides of the privacy technologies include competitive benefits due to increased customer confidence, minimized exposure to regulatory risk, and new opportunities to collaborate to create value through shared understanding and data sovereignty. The societal benefits are reflected in individual autonomy and the ability to counteract the dynamics of surveillance capitalism, as well as supporting democratic engagement in cyberspace, but equity aspects need to be taken into consideration to make sure that the benefits of privacy are not restricted to the wealthy groups, instead of being the prerogative of those populations. The additional complexity is posed by environmental dimensions because computationally intensive privacy mechanisms consume significantly more energy

than plaintext processing does, and organizations must trade off privacy protection goals with sustainability promises in terms of careful mechanisms selection and optimization of infrastructure. The direction of privacy engineering in data engineering settings will be determined by new technology such as quantum-resistant cryptography, zero-knowledge proofs, and privacy-preserving synthetic data generation, as well as regulatory development toward coherent systems and AI-specific regulation needs. Privacy engineering must become a core capability of the organization, and privacy-by-design concepts need to be part of system design. Privacy impact evaluation needs to be performed regularly, an automated compliance infrastructure should be deployed, and stakeholders should be informed about the data practices of the organization.

References

[1] OECD, "The Adoption of Artificial Intelligence in Firms," 2025. [Online]. Available: https://www.oecd.org/content/dam/oecd/en/publications/reports/2025/05/the-adoption-of-artificial-intelligence-in-firms_8fab986b/f9ef33c3-en.pdf

[2] David Kjerrumgaard, "Latency numbers every data streaming engineer should know," DEV Community, 2025. [Online]. Available: https://dev.to/david_kjerrumgaard_d31d7e/latency-numbers-every-data-streaming-engineer-should-know-h91

[3] Office of the Data Protection Ombudsman, Finland, "Automated decision-making and profiling," Tietosuoja.fi, 2024. [Online]. Available: https://tietosuoja.fi/en/automated-decision-making-and-profiling

[4] Transcend, "AI and privacy: Navigating the intersection of artificial intelligence and data protection," Transcend, 2024. [Online]. Available: https://transcend.io/blog/ai-and-privacy

[5] Wencheng Yang et al., "Deep learning model inversion attacks and defenses: a comprehensive survey," SpringerNature Link, 2025. [Online]. Available: https://link.springer.com/article/10.1007/s10462-025-11248-0

[6] Shuya Feng et al., "DPI: Ensuring Strict Differential Privacy for Infinite Data Streaming," IEEE, 2024. [Online]. Available: https://ieeexplore.ieee.org/document/10646789

[7] Kunal Chandiramani et al., "Performance analysis of distributed and federated learning models on private data," ResearchGate, 2020. [Online]. Available: https://www.researchgate.net/publication/339541559_Performance_Analysis_of_Distributed_and_Federated_Learning_Models_on_Private_Data

[8] Arsen Kourinian et al., "California Privacy Protection Agency intensifies enforcement: Recent enforcement actions and trends," Mayer Brown Publications, 2025. [Online]. Available: https://www.mayerbrown.com/en/insights/publications/2025/05/california-privacy-protection-agency-intensifies-enforcement-recent-enforcement-actions-and-trends

[9] Jon Jacobson, et al "The impact of privacy-enhancing technologies (PET) on business, individuals and society," World Economic Forum, 2023. [Online]. Available: https://www.weforum.org/stories/2023/10/the-impact-of-privacy-enhancing-technologies-pet-on-business-individuals-and-society/

[10] Aitor Gómez-Goiri et al., "Privacy-preserving computation meets quantum computing: a scoping review," ScienceDirect, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2352864825000744