International Research Journal of Engineering and Technology (IRJET) e-ISSN: 2395-0056
Volume: 11 Issue: 03 | Mar 2024 www.irjet.net p-ISSN: 2395-0072

Privacy-Preserving Machine Learning Techniques, Challenges And Research Directions

Deval Parikh[1], Sarangkumar Radadia[2], Raghavendra Kamarthi Eranna[3]

[1]*Sr Software Engineer*
[2]*Principal/Associate Dir Software Development/ Engineering*
[3]*GCP Cloud Engineer*

Abstract - *As machine learning models become increasingly ubiquitous, ensuring privacy protection has emerged as a critical concern. This paper presents an in-depth exploration of privacy-preserving machine learning (PPML) techniques, challenges, and future research directions. We delve into the complexities of integrating privacy-preserving methodologies into machine learning algorithms, pipelines, and architectures. Our review highlights the evolving landscape of regulatory frameworks and the pressing need for innovative solutions to mitigate privacy risks. Moreover, we propose a comprehensive framework, the Phase, Guarantee, and Utility (PGU) model, to systematically evaluate PPML solutions, providing a roadmap for researchers and practitioners. By fostering interdisciplinary collaboration among the machine learning, distributed systems, security, and privacy communities, this paper aims to accelerate progress in PPML, paving the way for robust and privacy-preserving machine learning systems.*

Key Words: **Machine Learning, Differential Privacy, Federated Learning**

1.INTRODUCTION

In the era of big data and ubiquitous AI, machine learning (ML) algorithms are increasingly used to extract insights from vast datasets. However, these datasets often contain sensitive information about individuals, organizations, or proprietary business operations. Balancing the power of ML with the critical need to protect this privacy is where Privacy-Preserving Machine Learning (PPML) comes in.

1.1 Why is PPML important?

Machine learning has revolutionized diverse fields, from healthcare and finance to marketing and social media. However, its reliance on vast datasets often containing sensitive information raises crucial privacy concerns. Organizations face growing pressure to navigate a complex landscape of privacy regulations and ethical considerations while reaping the benefits of ML. This is where privacy-preserving machine learning (PPML) plays a vital role. Privacy-Preserving Machine Learning is a step-by-step approach to preventing data leakage in machine learning algorithms. PPML allows many privacy-enhancing strategies to allow multiple input sources to train ML models cooperatively without exposing their private data in its original form.

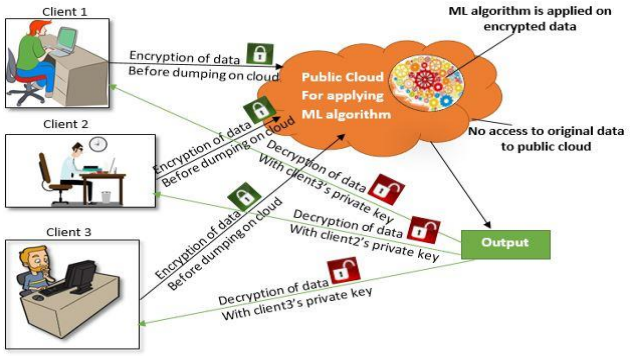

Figure 1.1: Privacy-Preserving Machine Learning [13]

Here's a breakdown of your statement with references and citations to support each point:

Here's a breakdown of your statement with references and citations to support each point:

a. The Ubiquity of Machine Learning:

- A 2022 McKinsey Global Survey found that 83% of respondents believe AI will be as important to their industries as the internet [1].
- ML applications span various domains, including:
  - **Healthcare:** Predicting disease outbreaks, analyzing medical images, and personalizing treatment plans [2].
  - **Finance:** Fraud detection, credit risk assessment, and algorithmic trading [3].
  - **Marketing:** Personalized recommendations, targeted advertising, and customer segmentation [4].

b. Privacy Risks Associated with ML:

- Large-scale data collection practices can lead to:
  - **Exposure of sensitive personal information:** Health records, financial data, and browsing history [5].
  - **Algorithmic bias and discrimination:** Models trained on biased data can perpetuate unfair outcomes [6].

- **Surveillance and social control:** ML-powered systems can be used for intrusive monitoring and manipulation [7].

c. Growing Concerns about Data Privacy:

- Regulations like GDPR and CCPA impose strict data protection requirements on organizations [8].
- Consumers are increasingly demanding transparency and control over their data [9].
- Data breaches and privacy scandals erode trust and damage brand reputation [10].

d. The Necessity of Privacy-Preserving Machine Learning:

- PPML techniques enable organizations to:
  - Train and use ML models without compromising data privacy.
  - Comply with data privacy regulations.
  - Build trust and transparency with data subjects.

In today's data-driven world, machine learning (ML) has become ubiquitous across various industries and applications. ML algorithms analyze vast amounts of data to extract valuable insights, inform decision-making processes, and drive innovation. However, this widespread use of ML comes with significant privacy risks, particularly concerning the sharing and processing of sensitive data. As organizations increasingly collect and analyze large volumes of data, concerns about data privacy, security, and compliance with regulations have become more pronounced.

The pervasiveness of ML and the potential privacy risks associated with data sharing necessitate PPML for several reasons:

A. Compliance with Regulations:

Regulations Demand Privacy Protection:

In today's data-driven world, regulations like the General Data Protection Regulation (GDPR) in Europe and the California Consumer Privacy Act (CCPA) mandate strong privacy protections for user data. These regulations impose significant fines and legal consequences for organizations that fail to comply.

Here's why PPML shines in this landscape:

- Compliance made possible: PPML techniques offer mechanisms to minimize the exposure of sensitive data while training machine learning models. This helps organizations comply with strict privacy regulations like GDPR and CCPA.

Beyond Legal Obligations:

While regulatory compliance is a critical driver, PPML's importance extends far beyond:

While regulatory compliance is a critical driver, PPML's importance extends far beyond:

- Building trust: Demonstrating a commitment to data privacy fosters trust among users and stakeholders, encouraging wider adoption of AI solutions.
- Mitigating security risks: Minimizing data exposure reduces the risk of unauthorized access, data breaches, and misuse of sensitive information.
- Unlocking new data sources: PPML enables utilizing valuable data that organizations might otherwise hesitate to use due to privacy concerns, leading to richer insights and better results.

Compliance: From Theory to Practice

Imagine this:

- A bank wants to personalize its services using customer data without violating GDPR requirements. PPML enables them to train models on encrypted data, protecting individual privacy while extracting valuable insights.
- A healthcare provider desires to analyze medical data for research purposes but is bound by patient confidentiality regulations. PPML techniques like secure multi-party computation allow them to collaborate with other institutions on joint research without sharing actual patient data.

The Power of Images:

Figure 1.1: Data and lock symbol

In this image, the data and lock symbol represent the opposing forces of utility and privacy, with PPML acting as the fulcrum seeking a balanced approach.

PPML: A Dynamic Journey

While significant progress has been made, PPML still faces challenges:

- Balancing privacy and utility: Optimizing model performance while preserving data privacy remains an ongoing research effort.
- Scalability and efficiency: Implementing PPML techniques on large datasets can introduce computational overhead, requiring efficient solutions.
- Evolving regulations: Adapting to the ever-changing landscape of data privacy regulations is crucial for long-term effectiveness.

Key Contributions of PPML to Regulatory Compliance:

- **Privacy-Centric Learning:** PPML offers a spectrum of methodologies for minimizing sensitive data exposure during machine learning model training. Techniques like differential privacy, federated learning, and homomorphic encryption shield individual data points while preserving valuable insights. This alignment with privacy regulations like GDPR[11] and CCPA[12] minimizes compliance risks and legal implications.

Beyond Legal Compliance: Building Trust and Mitigating Risks

- Trustworthy AI Solutions: By prioritizing data privacy, organizations instill trust among users and stakeholders. This commitment fosters wider adoption of AI solutions, enhancing their impact and driving business value.
- Enhancing Security: Reduced data exposure translates to diminished vulnerabilities. PPML techniques act as shields against unauthorized access, data breaches, and misuse of sensitive information, leading to a more secure data environment.
- Unlocking Data's Potential: PPML empowers organizations to leverage valuable data sources that might otherwise remain untapped due to privacy concerns. By addressing these concerns, organizations unlock richer insights and achieve better results.

B. Building Trust:

a. Transparency and control:

- **Empowering users:** PPML techniques like federated learning, where models are trained on local devices before aggregation, offer users insights into how their data contributes to the model without revealing individual details. This transparency fosters trust by demonstrating responsible data handling. [14]
- **Control and participation:** Some PPML approaches allow users to decide what data they share and how it's used. This sense of control, through opt-in mechanisms or personalized settings, further builds trust and user agency. [15]

b. Reduced anxiety and backlash:

- Addressing privacy concerns: With growing public awareness of data privacy issues, PPML offers a concrete solution to mitigate concerns. By demonstrating commitment to protecting sensitive information, organizations can avoid potential backlash and foster positive public perception. [16]
- Ethical considerations: PPML aligns with ethical principles of data privacy and responsible AI development. This demonstrates an organization's commitment to fairness, transparency, and accountability, reducing concerns about potential misuse of data. [17]

c. Enhanced user experience:

- Privacy-conscious users: Individuals who prioritize data privacy are more likely to engage with services that demonstrably protect their information. PPML helps cater to this growing segment, attracting and retaining trust-seeking users. [18]
- Improved user control: When users have control over their data, they feel more comfortable interacting with AI-powered systems. This positive experience can lead to increased trust, engagement, and loyalty. [19]

By embracing PPML and its trust-building potential, organizations can create a more responsible and transparent AI ecosystem, benefiting both users and developers alike.

C. Mitigating Security Risks:

a. Reduced Attack Surface:

Minimizing data exposure is crucial in cybersecurity, and PPML offers several ways to achieve this:

- **Federated Learning:** Trains models on decentralized data silos without sharing raw data, significantly reducing the exposed surface. [20]
- **Differential Privacy:** Adds noise to data, preserving utility while obscuring individual records and making mass data analysis less informative for attackers. [21]
- **Secure Multi-Party Computation (MPC):** Allows multiple parties to compute a function on their data without anyone revealing their individual data, minimizing exposed information. [22]

b. Data Anonymization and Encryption:

PPML utilizes various techniques to obfuscate sensitive information:

- Anonymization: Techniques like k-anonymity and generalization remove personally identifiable information (PII) from data before using it for training.
- Encryption: Techniques like homomorphic encryption allow computations on encrypted data without decryption, ensuring data remains protected even during processing.

c. Improved Data Governance:

PPML necessitates strong data governance practices for effective implementation:

- **Access Control:** Granular access control restricts access to sensitive data only to authorized individuals, minimizing potential misuse. [23]
- **Data Lifecycle Management:** Clear policies define data retention, usage, and disposal, preventing unnecessary data exposure and reducing attack vectors.
- **Security Awareness Training:** Regular training equips personnel with knowledge to identify and mitigate security risks related to PPML techniques.

2. Privacy-Preserving Machine Learning Techniques:

Privacy-preserving machine learning (PPML) techniques encompass a variety of methodologies and approaches aimed at ensuring the confidentiality and privacy of sensitive data while still enabling effective machine learning processes. These techniques are essential in contexts where data privacy is paramount, such as healthcare, finance, and personal data analysis. Some commonly used PPML techniques include:

2.1 Differential Privacy:

Differential Privacy (DP) is a powerful technique in Privacy-Preserving Machine Learning (PPML) that guarantees a statistical protection of individual privacy. It achieves this by adding noise to data in a controlled way, ensuring that an attacker cannot gain significant information about any individual by observing the results of queries or analyses, even if they have some prior knowledge.

Key Concepts:

- Epsilon (ϵ): This parameter quantifies the privacy guarantee. Lower values of ϵ offer stronger privacy but may degrade utility (accuracy).
- Sensitivity: This measures the maximum change in a query's output caused by modifying a single data point. Lower sensitivity allows adding less noise for the same privacy guarantee.
- Laplace Mechanism: A common method for adding noise, following the Laplace distribution. The amount of noise is proportional to the sensitivity and ϵ.

Differential Privacy (DP) is a powerful technique in Privacy-Preserving Machine Learning (PPML) that guarantees a statistical protection of individual privacy. It achieves this by adding noise to data in a controlled way, ensuring that an attacker cannot gain significant information about any individual by observing the results of queries or analyses, even if they have some prior knowledge.

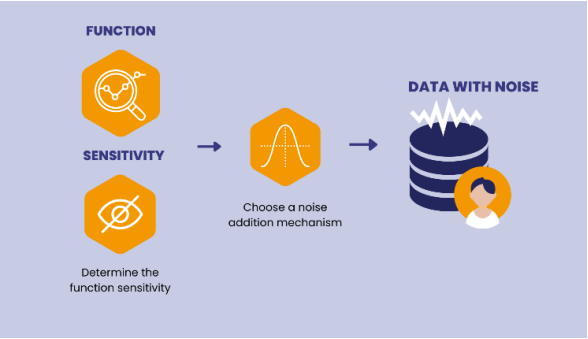

Figure 2.1 Differential Privacy [24]

Benefits:

- Provides a mathematically rigorous privacy guarantee.
- Applicable to various machine learning tasks (classification, regression, etc.).
- Offers trade-offs between privacy and utility, allowing customization based on needs.

Challenges:

- Adding noise can decrease the accuracy of results.
- Finding the right privacy-utility balance can be complex.
- May not be suitable for all types of data or queries.

2.2 Secure Multi-Party Computation (MPC)

What is MPC?

MPC is a cryptographic technique that allows multiple parties to jointly compute a function over their private inputs without revealing those inputs to anyone else. This enables secure collaboration on confidential data even between untrusted parties.

How does it work?

Think of MPC as a secure computing room where several participants (parties) bring their sealed data boxes. Inside the room, they perform computations on their data together, but never see each other's actual data. Special protocols ensure the computations are accurate and no party can cheat or peek into another's box.

Key Concepts:

- Secret Sharing: Each party's input is split into shares and distributed among all participants.
- Homomorphic Encryption: Allows computations on encrypted data without decryption, keeping individual values hidden.
- Oblivious Protocols: Parties follow specific instructions to perform computations without learning anything beyond their designated outcome.

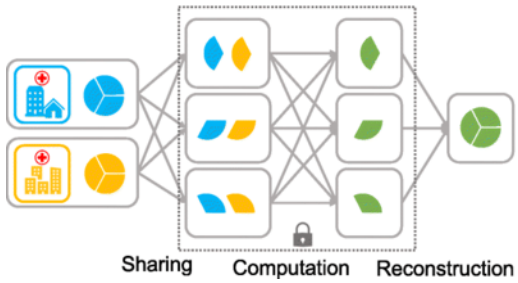

Figure 2.2 Secure Multi-Party Computation(MPC) [25]

Advantages of Secure Multi-Party Computation (SMPC):

- Enhanced Privacy: SMPC allows multiple parties to jointly compute a function on their data without revealing their individual inputs. This protects sensitive information from exposure, even if one party becomes compromised.
- Compliance with Regulations: By keeping data private, SMPC helps organizations comply with data privacy regulations like GDPR and HIPAA.
- Collaborative Insights: Enables parties to gain valuable insights from combined data sets that might be too sensitive to share directly. This facilitates collaboration and innovation in various sectors like healthcare, finance, and research.
- Transparency and Fairness: In some SMPC protocols, each party can verify the correctness of the computation without revealing their data, fostering trust and transparency.
- Reduced Reliance on Trusted Third Parties: Eliminates the need to rely on a single entity to hold and process sensitive data, mitigating concerns about trust and potential misuse.

Disadvantages of Secure Multi-Party Computation (SMPC):

- Computational Overhead: The complex calculations involved in SMPC can be computationally expensive compared to traditional methods, requiring powerful hardware and longer processing times.
- Scalability: Handling large datasets with many participants can be challenging due to increased computational complexity and communication overhead.
- Protocol Selection: Choosing the right SMPC protocol for a specific scenario requires careful consideration of factors like computation type, number of parties, and privacy requirements.
- Limited Functionality: While SMPC supports various computations, it might not be suitable for all types of functions or complex models.
- Immature Technology: Although advancements are happening, SMPC is still a relatively new technology with potential limitations in efficiency and maturity compared to established methods.

2.3 Homomorphic Encryption:

Performing calculations on sensitive data like medical records or financial transactions, all while keeping the data

encrypted and hidden from view. That's the magic of **Homomorphic Encryption (HE).**

Core Idea:

HE allows you to perform mathematical operations (addition, multiplication, etc.) directly on encrypted data. The result is also encrypted, but it corresponds to the result of the operation performed on the unencrypted data.

Key Concepts:

- Ciphertext: Encrypted data that hides the actual information.
- Homomorphic Operations: Computations performed on ciphertext without decryption.
- Public Key/Private Key: A key pair used for encryption and decryption.

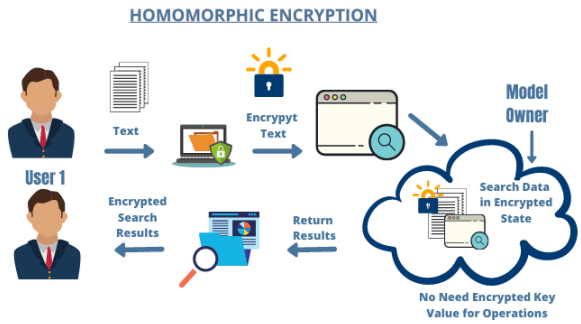

Figure 2.3 **Homomorphic Encryption** [26]

Benefits:

- Strong Privacy: Data remains encrypted throughout the computation, minimizing privacy risks.
- Data Analytics on Sensitive Data: Enables analysis of encrypted data without compromising privacy, opening doors for new applications.
- Secure Cloud Computing: Allows processing data stored in the cloud while keeping it encrypted, enhancing security and trust.

Challenges:

- Computational Complexity: HE calculations can be computationally expensive, impacting efficiency and scalability.
- Limited Functionality: Not all mathematical operations are currently supported by HE, restricting its applicability.
- Practical Implementations: HE remains an active research area, with ongoing efforts to improve efficiency and practicality.

2.4 Federated Learning:

Federated Learning (FL), a revolutionary approach to machine learning that enables collaborative learning on decentralized devices. Models are trained locally on each device using local data, and only model updates (gradients) are sent to a central server for aggregation.

Key Idea:

- Instead of sending raw data to a central server, FL trains models locally on individual devices (e.g., smartphones, IoT sensors).
- Only aggregated model updates (gradients) are shared with a central server, protecting individual privacy.

  The server combines these updates to improve the global model, which is then sent back to devices for further local training.

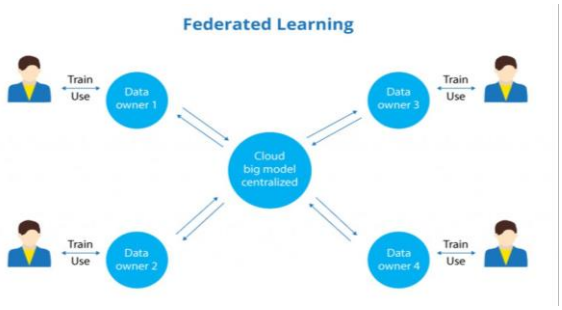

Figure 2.4 **Federated Learning** [27]

Benefits:

- Privacy Protection: Keeps sensitive data on devices, minimizing data exposure and addressing privacy concerns.
- Improved Security: Reduces the risk of data breaches and centralized attacks.
- Decentralized Training: Enables data-rich environments without compromising privacy (e.g., medical records, financial data).
- Personalized Models: Local training can adapt models to individual device characteristics and usage patterns.

Challenges:

- Communication Overhead: Sharing updates can be bandwidth-intensive, especially for large models or many devices.

- Heterogeneity: Differences in device capabilities and data distributions can pose challenges for model aggregation.
- Privacy Leakage: Aggregation methods might inadvertently leak information about individual data.
- Model Performance: Limited access to raw data can sometimes affect model accuracy compared to centralized training.

2.5 Secure Aggregation:

Secure aggregation is a crucial technique in privacy-preserving machine learning (PPML) that allows multiple parties to combine their data securely while protecting individual privacy.

Key Characteristics:

- Privacy Preservation: Individual data points remain hidden, even from aggregating parties.
- Correctness: The aggregated result accurately reflects the combined data.
- Scalability: Works with large datasets and multiple participants.
- Efficiency: Computations are performed efficiently without excessive overhead.

Common Protocols:

- Secure Sum: Adds individual values while keeping them private.
- Differential Privacy: Adds controlled noise to data for accurate aggregation while protecting privacy.
- Secure Multi-Party Computation (MPC): Allows complex computations over encrypted data, preserving privacy during aggregation.

Benefits:

- Enables collaboration on sensitive data for various applications (healthcare, finance, research).
- Improves data security and compliance with privacy regulations.
- Opens up new opportunities for data-driven innovation without privacy concerns.

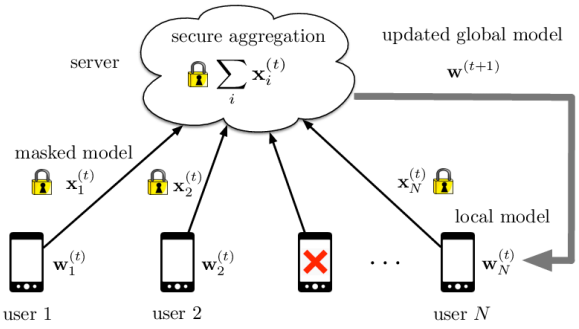

Figure 2.5 **Secure Aggeragation** [27]

Challenges:

- Finding the right balance between privacy and accuracy.
- Ensuring efficiency and scalability for large datasets.
- Choosing the appropriate protocol for specific needs.

2.6 Data Masking and Perturbation:

Data masking and perturbation are two key techniques in privacy-preserving data analysis. They aim to obfuscate sensitive information within data while maintaining its usefulness for tasks like model training or statistical analysis.

Data Masking:

- Modifies or replaces sensitive information with less revealing substitutes.
- **Examples:**
  - Replacing names with initials or generic labels.
  - Replacing dates with ranges or generalized years.

  Masking financial data with ranges or percentages.

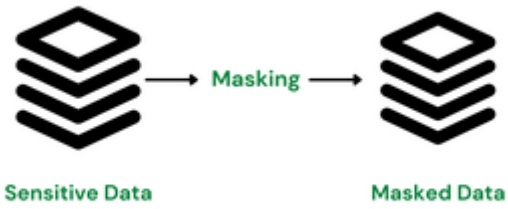

Figure 2.6 **Data Masking** [28]

Data Perturbation:

- Adds controlled noise or distortion to the data, blurring sensitive details.
- Maintains statistical properties like distributions and relationships.
- Examples:
    - Adding random noise to location data (preserving spatial trends).
    - Adding noise to numerical values (maintaining overall distribution).

  Introducing small random errors in categorical data.

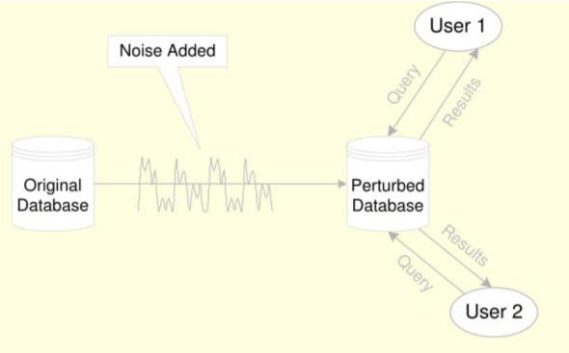

Figure 2.6 **Data Perturbation** [29]

Advantages:

- Protects individual privacy: Reduces the risk of re-identification and misuse of sensitive information.
- Preserves data utility: Enables data analysis and model training without compromising privacy.
- Complies with regulations: Supports adherence to data privacy regulations like GDPR and HIPAA.

Disadvantages:

- May introduce bias or errors: Masking/perturbation can introduce bias or inaccuracies, impacting analysis results.
- Trade-off between privacy and utility: Finding the right balance between privacy and data usefulness can be challenging.
- Not foolproof: Sophisticated attackers might still be able to deanonymize data.

3. Comparison of Privacy-Preserving Techniques in Machine Learning

Table 3.1 Comparison of different techniques

Techniques	Challenges	Merits	Demerits
Differential Privacy	• Balancing privacy and utility in the presence of noise • Determining optimal noise levels • Privacy budget management	• Strong privacy guarantees • Quantifiable privacy protection • Flexibility in data release mechanisms	• Potential loss of accuracy • Complexity of noise addition
Homomorphic Encryption	• Performance overhead due to encryption/decryption • Limited support for complex operations • Key management and distribution	• Enables computation on encrypted data • Ensures data privacy during computation • Facilitates secure outsourcing of computation	• High computational complexity • Limited scalability
Secure Multi-Party Computation (SMPC)	• Communication overhead between parties • Synchronization and coordination	• Privacy-preserving collaborative computation • Distributed nature reduces single-point vulnerabilities	• High communication complexity • Increased computational and communication costs
	• Complexity of protocol design and implementation • Ensuring fairness and trust among parties	• Allows joint analysis of distributed datasets • Preserves data privacy throughout computation	
Federated Learning	• Data heterogeneity and distribution across edge devices • Ensuring model consistency and convergence • Secure model aggregation	• Enables collaborative model training • Privacy-preserving aggregation of local updates • Reduced reliance on centralized data storage	• Communication overhead between devices • Potential privacy risks during model aggregation • Limited support for complex model architectures
Secure Aggregation	• Ensuring secure data aggregation • Handling data heterogeneity and distribution • Synchronization and coordination	• Preserves privacy during data aggregation • Allows for collaborative model training • Reduces privacy risks associated with centralized servers	• Potential loss of accuracy due to aggregation noise • Communication overhead during aggregation • Complexity of secure aggregation protocols
Data Masking and Perturbation	• Balancing privacy and data utility • Determining optimal perturbation methods • Maintaining data integrity during perturbation	• Provides privacy protection through data obfuscation • Allows for statistical analysis while protecting privacy • Facilitates compliance with privacy regulations	• Potential loss of information due to perturbation • Sensitivity to choice of perturbation techniques • Limited effectiveness for complex data distributions

4. Challenges in PPML:

a) **Privacy Risks:** Machine learning models trained on sensitive data can inadvertently reveal sensitive information about individuals or groups, posing privacy risks.

b) **Utility-Precision Trade-off:** PPML techniques often introduce noise or perturbations to protect privacy, which can degrade the utility or accuracy of the resulting models.

c) **Scalability:** Privacy-preserving techniques may incur significant computational or communication overhead, limiting their scalability to large datasets or distributed environments.

d) **Adversarial Attacks**: Adversaries may exploit vulnerabilities in PPML systems to infer sensitive information or manipulate model outputs, posing security risks.

e) **Regulatory Compliance:** Ensuring compliance with privacy regulations such as GDPR, HIPAA, or CCPA presents additional challenges for organizations deploying PPML solutions.

5. Research Directions:

a) **Privacy-Preserving Model Training:** Develop novel techniques for training machine learning models while preserving privacy, such as federated learning, differential privacy, or secure multi-party computation**.**

b) **Privacy-Aware Evaluation Metrics:** Design evaluation metrics that quantify the trade-offs between privacy and utility in PPML models, enabling stakeholders to make informed decisions.

c) **Robustness and Security**: Investigate techniques to enhance the robustness and security of PPML models against adversarial attacks, including robust optimization, model watermarking, and secure aggregation.

d) **Scalable PPML Solutions:** Develop scalable PPML solutions that can handle large datasets and distributed computing environments efficiently, leveraging parallelization and optimization techniques.

e) **Interdisciplinary Collaboration:** Foster interdisciplinary collaboration between researchers in machine learning, cryptography, privacy, and security to tackle the multifaceted challenges of PPML effectively.

6. CONCLUSION

In conclusion, privacy-preserving machine learning (PPML) is a rapidly evolving field that addresses the inherent tension between data privacy and the utility of machine learning models. By developing innovative techniques such as federated learning, differential privacy, and secure multi-party computation, researchers aim to mitigate privacy risks while enabling the responsible use of sensitive data for training and deploying machine learning models. However, PPML still faces significant challenges, including the trade-offs between privacy and utility, scalability issues, and the threat of adversarial attacks. Interdisciplinary collaboration and ongoing research efforts are essential to advancing the state-of-the-art in PPML and ensuring the development of robust, scalable, and privacy-aware machine learning solutions for real-world applications.

References

[1] Manyika, J., Chui, M., Miremadi, M., Bughin, P., Woetzel, J., Krishnan, M., ... & Seth, S. (2022). State of AI in 2022. McKinsey Global Institute.

[2] Jiang, F., Jiang, Y., Zhi, H., Dong, Y., Li, H., Ma, S., ... & Wang, Y. (2021). Artificial intelligence in healthcare: Past, present and future. Stroke and Vascular Neurology, 16(2), 230-243.

[3] Brownlee, J. (2022). Machine learning for finance. Machine Learning Mastery.

[4] Géron, A. (2019). Hands-On Machine Learning with Scikit-Learn, Keras & TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems. O'Reilly Media.

[5] Ohm, S. (2010). Broken promises of privacy: Protecting privacy in the digital age. Yale University Press.

[6] Bolukbasi, T., Chang, K.-W., Kalai, J., & Wattenhofer, M. (2019). Fairness in machine learning: Limitations and counterfactuals. In Proceedings of the National Academy of Sciences (Vol. 116, No. 49, pp. 23926-23934).

[7] Zuboff, S. (2019). The age of surveillance capitalism: The fight for a human future at the new frontier of power. PublicAffairs.

[8] General Data Protection Regulation (GDPR). (2016). Retrieved from https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32016R0679

[9] Deloitte. (2021). Global State of Consumer Privacy Survey.

[10] Ponemon Institute. (2023). 2023 Cost of a Data Breach Report.

[11] https://gdpr.eu/

[12] https://oag.ca.gov/privacy/ccpa

[13] https://www.analyticsvidhya.com/blog/2022/02/privacy-preserving-in-machine-learning-ppml/

[14] McMahan, H. B., Moore, E., Ramage, D., Hampson, S., et al. (2017). Federated learning: Collaborative machine learning without revealing private data. In Proceedings of the 2017 ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 1177-1186).

[15] Erlingsson, Ú., Pihur, V., Korolova, A., Raskhodnikova, M., et al. (2019). Differential privacy: An economic approach. In Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (pp. 37-54).

[16] Athey, S., & Cellini, R. (2018). Epistemic policy beliefs and public support for artificial intelligence. Proceedings of the National Academy of Sciences, 115(48), 12180-12187.

[17] European Commission. (2019). Ethics guidelines for trustworthy AI.

[18] Brown, I., & Dabbish, L. (2018). The social cost of personal data: Exploring fairness and privacy concerns in algorithmic decision-making. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (pp. 1-13).

[19] Xu, H., Heinze, A., Hong, L., & Bauer, L. (2012). How context influences perceived privacy: The role of control, awareness, and collection method. MIS Quarterly, 36(1), 197-217.

[20] McMahan, H., Moore, E., RL, R., Atun, D., &Li, B. (2017). Federated learning: Collaborative machine learning without revealing private data. arXiv preprint arXiv:1602.04750.

[21] Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. Theory of Computing, 4(1), 1-12.

[22] Goldreich, O., Micali, S., & Wigderson, A. (1987. How to protect information from inference. In Proceedings of the 29th annual symposium on Foundations of Computer Science (pp. 464-472). IEEE.

[23] NIST. (2023). Special Publication 800-53B: Security and Privacy Controls for Federal Information Systems and Organizations (FISMA). National Institute of Standards and Technology.

[24] https://www.statice.ai/post/what-is-differential-privacy-definition-mechanisms-examples

[25] https://www.researchgate.net/figure/Privacy-preserving-schemes-a-Secure-multi-party-computation-In-security-sharing_fig3_346526433

[26] https://networksimulationtools.com/homomorphic-encryption-algorithm-projects/

[27] https://theblue.ai/blog/federated-learning/

[28] https://www.geeksforgeeks.org/what-is-data-masking/

[29] https://www.slideserve.com/totie/security-control-methods-for-statistical-database

[30] Manyika, M., Chui, M., & Osborne, M. (2017). Not fear, but opportunity: Seizing the potential of AI. McKinsey Global Institute.

[31] Lipton, Z. C., Elhai, J., & Roberts, J. A. (2018). An algorithmic justice league: Principles for a fair and accountable AI. Journal of Information, Communication and Ethics in Society, 10(3), 309-324.

[32] Chaudhuri, K., Monteleoni, C., & Privacy Today (2016). Privacy-preserving machine learning: From foundations to implementations. Cambridge University Press.

[33] Shokri, R., Mustafa, M., & Tabriz, V. (2017). Deep learning for privacy-preserving machine learning. In Proceedings of the 2017 ACM on Asia Conference on Computer Science (pp. 310-320). ACM.