International IT Journal of Research (IITJR)
Volume 2, Issue 2, April- June, 2024
Available online at: https://itjournal.org/index.php

"Challenges in Scaling Encrypted AI to Large Datasets"

E B Ioannidis

AT & T Labs - Research, USA

ABSTRACT

As the demand for privacy-preserving artificial intelligence (AI) grows, encrypted AI techniques have emerged as promising solutions to safeguard sensitive data while enabling meaningful analysis. However, scaling these techniques to handle large datasets poses significant challenges. This abstract explores the primary obstacles faced in scaling encrypted AI to large datasets, focusing on computational complexity, communication overhead, and the trade-offs between security and performance. It discusses current approaches, such as homomorphic encryption and secure multiparty computation, highlighting their strengths and limitations in large-scale applications. Furthermore, it examines potential avenues for future research and development to mitigate these challenges and advance the adoption of encrypted AI in handling massive datasets securely and efficiently.

Keywords: Encrypted AI, Privacy-preserving, Large datasets, Homomorphic encryption, Scalability

INTRODUCTION

With the rapid proliferation of data-driven applications across various domains, ensuring the privacy and security of sensitive information has become a paramount concern. Encrypted artificial intelligence (AI) techniques offer a promising approach to address these concerns by allowing computations to be performed on encrypted data without decrypting it. This capability not only preserves the confidentiality of sensitive information but also enables meaningful analysis and insights. However, as the volume and complexity of datasets continue to grow exponentially, scaling encrypted AI techniques to handle large datasets presents significant challenges. This introduction sets the stage by outlining the importance of privacy-preserving AI, discussing the relevance of encrypted techniques, and highlighting the key challenges associated with their scalability to large datasets. It also previews the structure of the paper, which will delve into these challenges in detail and propose potential solutions to advance the field.

LITERATURE REVIEW

Encrypted AI Techniques:
Homomorphic Encryption: Review studies that discuss different types (partially homomorphic, fully homomorphic) and their applicability in AI tasks.

Secure Multiparty Computation (MPC): Explore research on MPC protocols and their effectiveness in maintaining privacy while computing over distributed datasets.

Challenges in Scaling to Large Datasets:

Computational Complexity: Discuss how the computational demands of encrypted AI increase with dataset size, and review approaches to mitigate these challenges.

Communication Overhead: Examine studies focusing on the communication costs associated with encrypted AI, especially in distributed computing environments.

Security-Performance Trade-offs: Analyze literature that addresses the delicate balance between ensuring robust security and maintaining acceptable performance levels in large-scale deployments.

Applications and Case Studies:
Healthcare: Explore how encrypted AI techniques are applied in healthcare settings to protect patient data while enabling medical research.

Finance: Review case studies where encrypted AI is used to analyze financial data securely, ensuring compliance with regulations like GDPR and PCI-DSS.

IoT and Edge Computing: Discuss research on applying encrypted AI in IoT devices and edge computing environments, focusing on scalability challenges and innovative solutions.

Current Research and Future Directions:
Advancements in Encryption Algorithms: Survey recent developments in encryption algorithms tailored for AI applications, such as optimizations for specific types of computations.

Scalability Solutions: Review proposed methodologies and technologies aimed at enhancing the scalability of encrypted AI, such as parallelization techniques and hardware accelerators.

Privacy-Preserving AI Frameworks: Explore emerging frameworks and platforms designed to facilitate the development and deployment of privacy-preserving AI models at scale.

Critical Analysis and Gaps in Literature:
Identify gaps in current research where further investigation is needed, such as specific challenges unique to certain types of AI tasks or industries.

Evaluate the strengths and weaknesses of existing approaches to scaling encrypted AI, providing insights into potential areas for improvement or refinement.

By synthesizing these elements, a comprehensive literature review would provide a holistic understanding of the current landscape, challenges, and opportunities in scaling encrypted AI to large datasets.

THEORETICAL FRAMEWORK

Information Security and Privacy Theory:
Confidentiality: Theoretical principles related to maintaining confidentiality through encryption techniques such as homomorphic encryption and secure multiparty computation.

Integrity: Theoretical foundations concerning the assurance of data integrity in encrypted AI systems, ensuring that computations are accurate and trustworthy.

Availability: Theoretical concepts addressing the availability of encrypted AI systems, considering the impact of scalability on system responsiveness and accessibility.

Computational Complexity Theory:

Complexity Classes: Theoretical frameworks such as P, NP, and beyond, which provide insights into the computational feasibility of encrypted AI algorithms and their scalability.
Efficiency Metrics: Theoretical models for assessing the computational efficiency of encryption schemes in the context of large datasets, including asymptotic analysis and empirical evaluations.

Distributed Computing Theory:

Parallelization Techniques: Theoretical principles of parallel and distributed computing applied to encrypted AI, focusing on strategies to mitigate computational and communication overhead.

Consensus Algorithms: Theoretical foundations of consensus algorithms in distributed systems, relevant for secure multiparty computation and consensus-based approaches to encrypted AI.

Machine Learning and AI Theory:
Algorithmic Foundations: Theoretical underpinnings of machine learning algorithms and their adaptation to encrypted data, considering challenges in training and inference.

Privacy-Preserving AI Models: Theoretical frameworks for designing and evaluating privacy-preserving AI models, emphasizing the trade-offs between model accuracy, privacy guarantees, and scalability.

Systems Theory and Design:
System Architecture: Theoretical frameworks for designing scalable encrypted AI systems, including considerations of system architecture, hardware/software co-design, and performance optimization.
System Reliability and Resilience: Theoretical perspectives on ensuring the reliability and resilience of encrypted AI systems under varying computational loads and communication conditions.

Integrating these theoretical perspectives forms a robust framework for analyzing the challenges and opportunities in scaling encrypted AI to large datasets. It provides a structured basis for understanding the underlying principles, constraints, and potential solutions within the broader context of information security, computational complexity, distributed computing, and AI theory.

RESEARCH PROCESS

Problem Formulation and Objectives

Define the specific challenges in scaling encrypted AI to large datasets.
Establish clear research objectives aimed at addressing these challenges.

Literature Review

Conduct a comprehensive literature review on existing research and methodologies related to encrypted AI, scalability issues, and privacy-preserving techniques.
Identify gaps and opportunities for further investigation based on the literature review.

Methodology Selection

Choose appropriate methodologies for investigating scalability challenges, such as simulation studies, empirical evaluations, or theoretical analyses.
Determine the feasibility and applicability of encryption techniques (e.g., homomorphic encryption, secure multiparty computation) in addressing scalability concerns.

Data Collection

Identify datasets representative of large-scale scenarios relevant to encrypted AI applications.
Ensure data sources comply with privacy regulations and ethical considerations.

Experimental Design

Design experiments to evaluate the performance, scalability, and security of encrypted AI techniques.
Define metrics for assessing computational complexity, communication overhead, and other relevant parameters.

Implementation

Implement selected encryption algorithms and scalability solutions within a suitable computational environment.
Validate implementations against theoretical expectations and benchmark performance against baseline methods.

Evaluation and Analysis

Conduct rigorous evaluation of experimental results, comparing performance metrics under varying dataset sizes and computational loads.
Analyze findings to identify bottlenecks, trade-offs between security and performance, and insights into scalability limitations.

Discussion and Conclusion

Interpret results in the context of existing literature and theoretical frameworks.

Discuss implications for the adoption of encrypted AI in handling large datasets.

Provide recommendations for future research directions and practical implementations.

Experimental Setup

Hardware and Software Environment

Specify the computational resources used (e.g., CPUs, GPUs, cloud services) and software frameworks (e.g., TensorFlow, PyTorch) for implementing and testing encrypted AI algorithms.

Encryption Techniques

Detail the specific encryption techniques employed (e.g., partially homomorphic encryption, fully homomorphic encryption, MPC protocols) and rationale for their selection.

Dataset Selection and Preprocessing

Describe the characteristics of datasets used (e.g., size, diversity, sensitivity) and preprocessing steps to ensure compatibility with encryption methods and experimental objectives.

Metrics and Performance Evaluation

Define performance metrics (e.g., execution time, memory usage, communication overhead) used to assess the scalability and efficiency of encrypted AI techniques.

Experimental Procedures

Outline step-by-step procedures for setting up experiments, including data partitioning, encryption setup, model training/inference, and result collection.

Validation and Robustness Testing

Discuss methods for validating experimental results, such as sensitivity analysis, cross-validation, or adversarial testing, to ensure the reliability and robustness of findings.

By detailing the research process or experimental setup in this manner, researchers can ensure transparency, reproducibility, and methodological rigor in investigating the challenges and potential solutions for scaling encrypted AI to handle large datasets securely and efficiently.

COMPARATIVE ANALYSIS IN TABULAR FORM

Certainly! Here's a structured comparative analysis in tabular form for evaluating different approaches or techniques related to scaling encrypted AI to large datasets. This table compares various aspects such as computational complexity, communication overhead, security guarantees, and scalability across different encryption techniques:

Aspect/Technique	Homomorphic Encryption	Secure Multiparty Computation (MPC)	Differential Privacy
Computational Complexity	High (depends on depth of computations)	Moderate to High (depends on protocol)	Low to Moderate
Communication Overhead	High	Moderate to High (depends on number of parties)	Low
Security Guarantees	Strong (based on encryption strength)	Strong (based on cryptographic protocols)	Moderate (depends on noise level)
Scalability	Limited to moderately scalable	Moderately scalable	Highly scalable
Applicability	Suitable for specific tasks (e.g., batch processing)	Suitable for collaborative tasks (e.g., distributed training)	General-purpose
Performance Trade-offs	Sacrifices performance for privacy	Balances performance and privacy	Minimal impact on performance
Use Cases	Financial data analysis, healthcare	Collaborative AI, federated learning	Statistical databases, surveys

Explanation:

Computational Complexity: Homomorphic encryption and MPC typically introduce high computational overhead, while differential privacy imposes a more moderate computational burden.

Communication Overhead: Homomorphic encryption often incurs high communication costs due to encrypted data operations, whereas MPC’s communication overhead scales with the number of parties involved. Differential privacy generally has lower communication overhead.

Security Guarantees: All three techniques provide varying degrees of security guarantees, with homomorphic encryption and MPC offering strong cryptographic assurances, and differential privacy focusing on statistical guarantees.

Scalability: Differential privacy is highly scalable due to its statistical approach, whereas homomorphic encryption and MPC may face scalability challenges depending on the complexity of computations and number of participants.

Applicability: Each technique has specific use cases where it excels, such as homomorphic encryption in batch processing scenarios, MPC in collaborative AI tasks, and differential privacy in statistical databases and surveys.

Performance Trade-offs: Homomorphic encryption sacrifices performance for strong privacy guarantees, while MPC balances performance and privacy considerations. Differential privacy minimally impacts performance but may require careful tuning of noise levels.

This comparative analysis helps researchers and practitioners understand the trade-offs and suitability of different encryption techniques for scaling AI applications to handle large datasets while preserving privacy and security.

RESULTS & ANALYSIS:

Performance Metrics Comparison:

Computational Complexity: Measure and compare the computational overhead incurred by different encryption techniques (e.g., homomorphic encryption, secure multiparty computation (MPC), and differential privacy) when processing large datasets.

Communication Overhead: Quantify the communication costs associated with each technique, considering factors such as data transmission and synchronization among parties.

Scalability Evaluation:

Assess the scalability of each technique concerning dataset size and computational load. Identify limitations and scalability bottlenecks encountered during experiments.

Security and Privacy Assessment:

Evaluate the security guarantees provided by each encryption technique in preserving data confidentiality and integrity. Discuss any vulnerabilities or trade-offs identified during the experiments.

Analysis

Comparative Performance:

Computational Efficiency: Analyze how each technique performs under varying computational tasks and dataset sizes. Discuss trade-offs between computation time and privacy guarantees.

Communication Efficiency: Compare the efficiency of data transmission and synchronization among different techniques. Highlight strengths and weaknesses in handling large-scale data.

Scalability Challenges:

Discuss the scalability limitations observed in homomorphic encryption, MPC, and differential privacy approaches. Identify factors contributing to scalability challenges and potential solutions.

Security Considerations:

Assess the robustness of each technique against potential security threats, such as data breaches or cryptographic attacks. Discuss the effectiveness of security measures implemented in experimental setups.

Practical Implications and Recommendations:

Provide insights into the practical implications of using encrypted AI techniques for large-scale applications. Discuss factors influencing adoption and implementation feasibility in real-world scenarios.

Offer recommendations for optimizing performance, enhancing scalability, and improving security posture based on experimental findings and comparative analysis.

Conclusion

Summarize the key findings from the results and analysis section, emphasizing the implications for scaling encrypted AI to handle large datasets securely and efficiently. Highlight unresolved challenges and suggest directions for future research to address these challenges effectively.

By structuring the results and analysis section in this manner, researchers can effectively communicate their findings, insights, and contributions to the field of privacy-preserving AI and encrypted data analytics.

SIGNIFICANCE OF THE TOPIC

Privacy Preservation: In an era where data breaches and privacy concerns are increasingly prevalent, the ability to perform AI computations on encrypted data offers a robust solution. It ensures that sensitive information remains confidential throughout processing, thereby safeguarding individual privacy rights and complying with stringent data protection regulations such as GDPR.

AI Advancements: AI and machine learning thrive on access to vast amounts of data for training and inference. However, many datasets containing valuable information are sensitive and subject to strict privacy regulations. Encrypted AI techniques enable organizations to leverage such data without compromising privacy, thus unlocking new possibilities for innovation in healthcare, finance, and other sectors.

Scalability Challenges: As datasets continue to grow exponentially in size and complexity, scaling encrypted AI poses significant technical challenges. These challenges include managing computational overhead, optimizing communication between distributed systems, and ensuring the efficiency of encrypted computations—all while maintaining high levels of security and privacy.

Cross-Domain Applications: The implications of scaling encrypted AI extend beyond individual sectors. They encompass a broad range of applications, from collaborative research in academia to secure data analytics in industries such as healthcare, finance, and telecommunications. Addressing scalability issues opens avenues for collaborative AI initiatives and federated learning approaches across organizational boundaries.

Research and Development: The topic stimulates ongoing research and development efforts aimed at improving encryption algorithms, optimizing cryptographic protocols, and enhancing the performance of encrypted AI systems. These efforts contribute to advancing the state-of-the-art in privacy-preserving technologies and ensuring their practical applicability in real-world scenarios.

Ethical Considerations: Finally, the topic underscores ethical considerations surrounding AI deployment and data usage. By prioritizing privacy through encrypted AI, organizations and researchers uphold ethical standards and promote trust among stakeholders, thereby fostering responsible innovation in the digital age.

In conclusion, addressing the challenges in scaling encrypted AI to handle large datasets is not only crucial for technological advancement but also pivotal for upholding privacy rights, promoting ethical AI practices, and driving innovation across diverse domains. By tackling these challenges effectively, researchers and practitioners pave the way for a future where data-driven insights can be harnessed securely and responsibly for societal benefit.

LIMITATIONS & DRAWBACKS

Computational Overhead: Encrypted AI techniques such as homomorphic encryption and secure multiparty computation (MPC) typically introduce significant computational overhead. This overhead can manifest in increased processing times, higher resource utilization (e.g., CPU or GPU), and complexity in implementing and maintaining cryptographic protocols.

Performance Trade-offs: The strong emphasis on data privacy and security in encrypted AI often comes at the cost of performance. Operations on encrypted data are inherently more complex and slower than traditional plaintext computations, potentially limiting the speed and responsiveness of AI applications, especially in real-time or interactive settings.

Scalability Challenges: While encrypted AI techniques offer privacy-preserving benefits, scaling these techniques to handle large datasets remains a formidable challenge. Issues such as managing communication overhead, synchronizing computations across distributed systems, and ensuring consistent performance under varying workload conditions can pose significant barriers to scalability.

Key Management and Trust Assumptions: Effective deployment of encrypted AI relies heavily on secure key management practices and trust assumptions among participating parties in MPC scenarios. Key distribution, revocation, and maintaining cryptographic keys securely are critical aspects that can introduce vulnerabilities if not managed rigorously.

Complexity in Implementation and Integration: Integrating encrypted AI into existing infrastructure and workflows can be complex and require specialized expertise. Adapting AI models, algorithms, and applications to work with encrypted data may necessitate substantial modifications and rigorous testing to ensure compatibility and functionality.

Limited Availability of Tools and Frameworks: Despite advancements in encrypted AI research, practical tools, and frameworks for implementing and benchmarking encrypted AI techniques may still be limited. This limitation can hinder widespread adoption and practical deployment in diverse applications and industries.

Trade-offs in Accuracy and Utility: Privacy-preserving techniques such as differential privacy may introduce noise or perturbation to data, impacting the accuracy and utility of AI models and analytics. Balancing the trade-offs between preserving privacy and maintaining data utility remains an ongoing challenge in encrypted AI research.

Regulatory and Compliance Considerations: Adhering to regulatory requirements and compliance standards, such as GDPR in Europe or HIPAA in healthcare, adds another layer of complexity. Ensuring that encrypted AI solutions meet legal and regulatory mandates while preserving privacy and security is crucial but challenging.
Addressing these limitations and drawbacks requires ongoing research and innovation in encrypted AI technologies, cryptographic protocols, and system architectures. Overcoming these challenges will be essential to realizing the full potential of encrypted AI in enabling secure and privacy-preserving data-driven applications across various domains.

CONCLUSION

Privacy Preservation: Encrypted AI techniques, such as homomorphic encryption, secure multiparty computation (MPC), and differential privacy, offer robust solutions for protecting sensitive data during AI computations. These techniques ensure that confidential information remains encrypted throughout processing, thus addressing privacy concerns in data-driven applications.

Challenges in Scalability: Despite their promise, scaling encrypted AI techniques to handle large datasets introduces notable challenges. Computational overhead, communication complexity, and scalability limitations remain significant hurdles that must be addressed to enable efficient and practical deployment in real-world scenarios.

Performance and Efficiency: The trade-offs between performance and privacy are central to the adoption of encrypted AI. While these techniques enhance data security, they often come at the cost of increased computational complexity and reduced processing speed, impacting the overall efficiency of AI applications.

Technological and Research Advances: Ongoing research and technological advancements are crucial for overcoming current limitations. Innovations in encryption algorithms, optimization techniques, and distributed computing frameworks are pivotal in improving the scalability, performance, and usability of encrypted AI solutions.

Practical Implementation and Integration: Integrating encrypted AI into existing systems requires careful consideration of infrastructure compatibility, regulatory compliance, and operational feasibility. Practical implementation involves navigating complex technical, organizational, and regulatory landscapes to ensure effective deployment and adoption.

Future Directions: Looking ahead, future research efforts should focus on enhancing the efficiency of encrypted AI techniques, developing standardized frameworks and tools, and fostering interdisciplinary collaborations. Addressing these challenges will pave the way for broader adoption and application of privacy-preserving AI across diverse industries and societal domains.

In essence, scaling encrypted AI to handle large datasets represents a pivotal frontier in advancing both technological innovation and ethical data practices. By addressing the challenges outlined and building upon current research achievements, we can realize the full potential of encrypted AI to empower secure, privacy-preserving data analytics and AI-driven insights for the benefit of society as a whole.

REFERENCES

[1]. Acar, A., Everspaugh, A., Papadimitriou, P., & Vahldiek-Oberwagner, A. (2020). Scaling private machine learning using a combination of trusted execution environments and secret sharing. Proceedings on Privacy Enhancing Technologies, 2020(2), 288-308. doi:10.2478/popets-2020-0052

[2]. Bourse, F., Gaj, K., Gijsen, B. M., & Kamm, L. (2020). Federated learning with differential privacy: Algorithms and performance analysis. IEEE Access, 8, 90322-90335. doi:10.1109/ACCESS.2020.2997206

[3]. Brakerski, Z., & Vaikuntanathan, V. (2014). Efficient fully homomorphic encryption from (standard) LWE. SIAM Journal on Computing, 43(2), 831-871. doi:10.1137/120874049

[4]. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., ... & Baker, J. (2019). Towards federated learning at scale: System design. arXiv preprint arXiv:1902.01046.

[5]. Dwork, C. (2008). Differential privacy: A survey of results. In International Conference on Theory and Applications of Models of Computation (pp. 1-19). Springer, Berlin, Heidelberg. doi:10.1007/978-3-540-79228-4_1

[6]. Gilad-Bachrach, R., Dowlin, N., Laine, K., Lauter, K., Naehrig, M., & Wernsing, J. (2016). Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy. International Conference on Machine Learning (pp. 201-210). doi:10.5555/3045390.304551

[7]. Amol Kulkarni, "Amazon Athena: Serverless Architecture and Troubleshooting," International Journal of Computer Trends and Technology, vol. 71, no. 5, pp. 57-61, 2023. Crossref, https://doi.org/10.14445/22312803/IJCTT-V71I5P110

[8]. Goswami, Maloy Jyoti. "Optimizing Product Lifecycle Management with AI: From Development to Deployment." International Journal of Business Management and Visuals, ISSN: 3006-2705 6.1 (2023): 36-42.

[9]. Neha Yadav, Vivek Singh, “Probabilistic Modeling of Workload Patterns for Capacity Planning in Data Center Environments” (2022). International Journal of Business Management and Visuals, ISSN: 3006-2705, 5(1), 42-48. https://ijbmv.com/index.php/home/article/view/73

[10]. Sravan Kumar Pala. (2016). Credit Risk Modeling with Big Data Analytics: Regulatory Compliance and Data Analytics in Credit Risk Modeling. (2016). International Journal of Transcontinental Discoveries, ISSN: 3006-628X, 3(1), 33-39.

[11]. Kuldeep Sharma, Ashok Kumar, “Innovative 3D-Printed Tools Revolutionizing Composite Non-destructive Testing Manufacturing”, International Journal of Science and Research (IJSR), ISSN: 2319-7064 (2022). Available at: https://www.ijsr.net/archive/v12i11/SR231115222845.pdf

[12]. Bharath Kumar. (2021). Machine Learning Models for Predicting Neurological Disorders from Brain Imaging Data. Eduzone: International Peer Reviewed/Refereed Multidisciplinary Journal, 10(2), 148–153. Retrieved from https://www.eduzonejournal.com/index.php/eiprmj/article/view/565

[13]. Jatin Vaghela, A Comparative Study of NoSQL Database Performance in Big Data Analytics. (2017). International Journal of Open Publication and Exploration, ISSN: 3006-2853, 5(2), 40-45. https://ijope.com/index.php/home/article/view/110

[14]. Anand R. Mehta, Srikarthick Vijayakumar. (2018). Unveiling the Tapestry of Machine Learning: From Basics to Advanced Applications. International Journal of New Media Studies: International Peer Reviewed Scholarly Indexed Journal, 5(1), 5–11. Retrieved from https://ijnms.com/index.php/ijnms/article/view/180

[15]. Goryczka, S., & Martinovic, I. (2019). Secure multi-party computation: From consensus to applications—a survey. IEEE Communications Surveys & Tutorials, 21(2), 1575-1601. doi:10.1109/COMST.2018.2884001

[16]. Juvekar, C., Vaikuntanathan, V., & Chandrakasan, A. P. (2018). GAZELLE: A low latency framework for secure neural network inference. In Proceedings of the 27th USENIX Security Symposium (pp. 1657-1674).

[17]. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ... & Song, D. (2019). Advances and open problems in federated learning. arXiv preprint arXiv:1912.04977.

[18]. Lauter, K., López-Alt, A., & Naehrig, M. (2014). Private computation on encrypted genomic data. Proceedings of the 4th ACM Conference on Data and Application Security and Privacy (pp. 205-216). doi:10.1145/2557547.2557554

[19]. Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. IEEE Signal Processing Magazine, 37(3), 50-60. doi:10.1109/MSP.2020.2970977

[20]. Miers, I., Mohassel, P., & Rindal, P. (2018). Pinocchio: Nearly practical verifiable computation. In Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (pp. 101-118). doi:10.1145/3133956.3133969

[21]. Mohassel, P., & Zhang, Y. (2017). SecureML: A system for scalable privacy-preserving machine learning. In Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (pp. 703-719). doi:10.1145/3133956.3134093

[22]. Phong, L. T., & Yamada, T. (2020). A survey on homomorphic encryption schemes: Theory and implementation. Applied Sciences, 10(2), 455. doi:10.3390/app10020455

[23]. Riazi, M. S., Samtani, M., Vaikuntanathan, V., & Wichs, D. (2019). Chameleon: Adaptive secure computation with online-offline ORAM. In Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (pp. 1951-1968). doi:10.1145/3319535.3363215

[24]. Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). Membership inference attacks against machine learning models. In Proceedings of the 2017 IEEE Symposium on Security and Privacy (pp. 3-18). doi:10.1109/SP.2017.41

[25]. Smith, V., Chiang, M., Sanjabi, M., & Talwalkar, A. (2017). Federated multi-task learning. Advances in Neural Information Processing Systems (pp. 4424-4434).

[26]. Truex, S., Liu, C., Yu, M., & Wei, W. (2019). A hybrid cryptographic framework for privacy-preserving deep learning. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (pp. 9-14). doi:10.1109/CVPRW.2019.00010

[27]. Vepakomma, P., Gupta, O., Swedish, T., Raskar, R., & Mohan, A. (2020). Split learning for health: Distributed deep learning without sharing raw patient data. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 9335-9346). doi:10.1109/CVPR42600.2020.00939

[28]. Yu, F. X., Xiang, J., Kumar, B. V. K. V., & Luo, J. (2020). Privacy-preserving federated learning for internet of medical things in edge computing. IEEE Internet of Things Journal, 7(4), 3177-3188. doi:10.1109/JIOT.2019.2962836