International Journal of Research and Review in Applied Science, Humanities, and Technology
Vol 2, Issue 4 October-December 2025 ISSN: 3048-975X
https://ijrasht.com/

Homomorphic Encryption for Cloud Data Security: A Comprehensive Review

Venkateswaran Radhakrishnan, Department of Information Technology, College of Computing and Information Sciences, University of Technology and Applied Sciences, Salalah, OMAN Venkateswaran.radhakrishnan@utas.edu.om

Praveen Kumar C, Department of Information Technology, College of Computing and Information Sciences, University of Technology and Applied Sciences, Salalah, OMAN Praveen.Kumar@utas.edu.om

Abstract: The Industrial Internet of Things (IIoT) and Artificial Intelligence (AI) intersected in the realm of industry, transforming the industrial system to allow making real-time decisions at the edge. Latency issues, bandwidth consumption, and data security also pose a challenge to the traditional cloud-based system, and are of utmost priority in the industry where time is limited. Edge AI combines AI features with edge computing, moving the intelligence to the devices in the IIoT, allowing decisions to be made faster and more responsive to the context without depending on centralized infrastructures that might overwhelm the cloud. The paper summarizes the present-day developments in edge AI as a decision-making tool in IIoT and identifies the future directions. It discusses architectures, algorithms and applications that enable intelligent decision making at the network edge with particular focus on manufacturing, predictive maintenance, supply chain optimization and energy management. It is reviewed that facilitating technologies, such as lightweight deep learning models, federated learning, and hardware accelerators, and problems, such as scalability, interoperability, and cybersecurity are discussed. As depicted in the literature section, edge AI has been found to enhance efficiency of distributed industrial systems by reducing latencies, reliability, and independent decision-making. However, barriers such as the shortage of resources, failure to interface with current systems and standardized structures still persist. It is anticipated that future studies will be based on adaptive AI models and edge-cloud collaboration with application of 6G-enabled IIoT ecosystems. This paper summarizes the synthesis of state-of-the-art methods to inform the next generation of industrial automation and digital transformation powered by edge AI. It indicates the need of powerful, hardy, and scalable systems to open up the complete scope of opportunities that edge AI can introduce to decisions made by IIoT.

Keywords: Edge AI, Industrial internet of things, Decision making, Edge computing, Federated learning.

1. Introduction

The Industrial Internet of Things (IIoT) means connecting machines, sensors, and controllers to enable smart automation and online transformation. Industries are moving towards IIoT in an attempt to maximize efficiency, safety, and productivity. Yet, due to the exponential growth of data volumes, centralized cloud architectures have certain limitations, such as high latency, high bandwidth consumption, and privacy issues. Milliseconds count in safety-critical systems, including manufacturing plants or energy grids, so cloud-only solutions are not suitable.

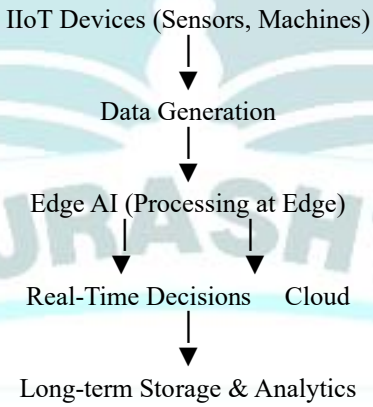

Flowchart: Role of Edge AI in IIoT Decision-Making

Edge AI solves these issues by running AI functionality nearer to the source of data. Data processing at the edge can enable industries to make near real-time decisions, decrease cloud connectivity reliance, and improve system availability. Such a paradigm shift has found use in predictive maintenance, quality control, robotics, logistics, and smart energy systems. The paper discusses edge AI application in IIoT decision-making and gives an overview of the recent progress, challenges, and perspectives regarding the field.

2. Background of the Study

The IIoT ecosystem is defined by the distributed sensors and actuators and machines producing large volumes of real-time data. Traditional cloud systems cannot provide high reliability and latency requirements of industries. Edge computing was developed to decentralise computation, with the aim of bringing data processing nearer to

devices. When paired with AI or Edge computing, or Edge AI, it also enables IIoT by allowing it to have localized intelligence.
The enabling features of edge AI are:
- Lightweight AI models: Customized deep learning algorithms to use by the edge devices.
- Federated learning: Model training without the exchange of raw data, which increases privacy.
- Edge accelerators Edge accelerators are specialized hardware (e.g. NVIDIA Jetson, Google Coral), which accelerates the inference.
- 5G/6G networks: Ultra-low latency with high throughput to support edge intelligence.

3. Justification
The introduction of edge AI into IIoT is explained by a number of reasons:
Latency-sensitive operations: The decisions made in milliseconds by industrial robots, autonomous vehicles, and energy grids are time-sensitive.
- Data privacy and security: Sensitive operational data does not need to be sent to cloud servers, but can be processed locally.
- Optimization of bandwidth: Only insights but not raw data should be transferred, which will not increase the congestion of the network.
- Resilience: Edge intelligence will maintain its functionality even when there is a cloud disconnection.
- These elements establish edge AI as one of the foundations of Industry 4.0 and Industry 5.0 systems.

4. Objectives of the Study
1. To present a review of homomorphic encryption schemes in the context of cloud data security.
2. To examine theoretical bases and practice applications of HE.
3. To assess the uses of HE in other fields like healthcare, finance and government.
4. To analyse constraints and issues of implementing HE in practical cloud environments.
5. To determine the areas of research that can improve efficiency and usability of HE in the future.

5. Literature Review
Basic Research: Gentry (2009) showed that realization of FHE may be possible and the efficiency of the scheme is now a research problem. Cloud security applications: Acar et al. (2018) have reviewed privacy-preserving technologies, and discussed the role of HE in secure outsourcing. The scaling approximation scheme is called Practical Schemes Scheme Brakersiki, Gentry and Vaikuntanathan (BGV) or Cheon-Kim-Kim-Song (CKKS) and can be scaled up to a real-world (Cheon et al., 2017). Healthcare/Finance Lauter et al. (2014) performed the experimental work based on HE and medical records, Li et al. (2019) based on encrypted financial analytics. Rapid lattice based HE: Lattice-based HE and execution on GPUs and FPGAs is the latest attempt (Chen et al., 2019). The literature shows continuous improvement but it also shows performance bottlenecks as the major obstacle towards large-scale adoption.

6. Material and Methodology
To be comprehensive, objective, and reproducible, the current review has followed a Systematic Literature Review (SLR) approach. The methodology adhered to the general principles of the systematic reviews conducted in the fields of computer science and cryptography.

Databases Searched: A total of 16 online databases were searched to identify a large number of potential publications. The most topical and the most often referenced repositories were:

IEEE Xplore → cryptography and information security conference proceedings.
ACM Digital Library → basic works in computing, algorithms and encryption schemes.
SpringerLink → to find detailed journal articles and books in the field of applied mathematics, cloud computing, and security.
ScienceDirect (Elsevier) → to find large-scale surveys and applications of homomorphic encryption in medical, financial and industrial contexts.
To cover cross-disciplinary coverage, Google Scholar, MDPI, Wiley Online Library, and Taylor and Francis were also used as additional sources.

Keywords and Search strategy: Boolean operators and combinations of keyword strings were employed in order to maximize recall and precision. Examples of search terms used were:
"Homomorphic Encryption"
Fully Homomorphic Encryption (FHE).

Partially Homomorphic Encryption (PHE).
And somewhat homomorphic encryption (SHE)
Cloud Security" AND "Privacy-preserving Computation
"Encrypted Machine Learning"
This pathway methodology was effective in ensuring that both theoretical (e.g., lattice-based cryptographic schemes) and applied research (e.g., HE in cloud storage, finance and healthcare) was identified.

Timeframe of the Review: The revision was established to run between 2010 and 2023.
This period includes both the post-Gentry (2009) period when Fully Homomorphic Encryption was first possible, and the most recent optimizations in efficiency and application in the real world.
The reason why this period has limited scope is that it would enable a review to focus on past and recent developments so that it becomes applicable when addressing the current research issues.

Criteria of selection and inclusion: The inclusion process was based on several filters:
Inclusion: Articles that suggested, tested or assessed homomorphic encryption schemes with explicit use in cloud computing and privacy-preserving computation.
Filter: Article which is not peer-reviewed or duplicated or is not related to encryption (e.g., not to cryptography but to cybersecurity in general).
Quality Filter: They chose high impact journals, conference papers and research studies, which contained significant amounts of experimental analysis.
This narrowed down the number of preliminary search results of several hundred articles to a final edited collection of approximately 100 core papers to analyse.

Performance Features to be considered: All shortlisted pieces were reviewed based on 4 critical performance dimensions:
Security Guarantees→ Theoretical soundness of encryption, known-attack resistance (lattice-based hardness, quantum resistance).
Computational Performance → Execution time, encryption/decryption complexity, ciphertext expansion and computational overhead.
Scalability → The ability to handle big data, cloud-scale functionality, and parallelisation (e.g. GPU/FPGA acceleration).
Real-World Usability → How useful the field of cloud storage, healthcare, finance, and IoT-cloud ecosystems is.

Analysis Process: All the studies were coded according to the scheme of encryption (PHE, SHE, FHE, lattice-based variants) and field of application as well as performance metrics. Comparison has been made to determine the strengths, weaknesses and unresolved challenges across various methodologies. Cross-domain applications were also analysed to learn how HE expands the scope of cloud security in healthcare, finance, government and industrial IoT.

7. Results and Discussion

Security: High Confidentiality FHE: Fully Homomorphic Encryption (FHE) is the most powerful encryption today in terms of privacy. It enables calculations (addition, multiplication, or even complicated machine learning) to be done directly on encrypted data without decryption. This implies that sensitive information like medical records, financial transactions or government information does not have to be exposed to a cloud company at all. In case the cloud server is compromised, the attackers can only see ciphertext that is computationally infeasible to decrypt. That way, FHE can be regarded as a gold standard of privacy-preserving computation confidentiality.

Trade-Offs in performance: PHE and SHE: FHE is secure, but computationally very expensive. This complicates real-world implementation of applications that have very tight latency constraints. Partially Homomorphic Encryption (PHE) does not allow any operation to be performed (e.g., RSA only supports multiplication, Paillier only supports addition). This compromises flexibility, but is faster and lighter than FHE. SHE is somewhat homomorphic encryption permitting limited number of additions and multiplications, but is still unable to perform unlimited computations as FHE. Thus, PHE and SHE have superior but lower-performing functionality, and can be useful in application areas where the speed is more important than the complete computational generality.

Use in practice: Homomorphic encryption is not an idealized theory—it can be applied in the real world in cloud and data-heavy industries:

Cloud-Based Machine Learning → The ability to train models on encrypted user data eliminates the need to access the raw information and maintain confidentiality. Use case: medical image or financial fraud detection.
Privacy-Aware Data Analytics → Organizations are able to store data encrypted and still perform statistical or analytical queries on this encrypted data. Example: hospitals sharing encrypted patient information to conduct studies on a global scale.
Encrypted Search-Users have the ability to search their encrypted data on the cloud without exposing their query or putting the database in plaintext. These applications show how HE can enable trade-offs between utility and privacy in the data-driven industries.

Issues and Limitations: However, HE has the following practical challenges:
Computational Cost: FHE is extremely computationally expensive and is at minimum 1000 to 10000 times slower than plaintext computation.
Expensive Encryption: Encrypted messages (ciphertexts) are significantly larger than plaintexts, and can be as large as 100x which increases storage and bandwidth requirements.
Integration Problems: Existing cloud software and systems are not structured to be capable of integrating with HE libraries and therefore, are not easily adopted by enterprises.
These restrictions are why HE is not efficient enough and not yet ready to be implemented in a large scale cloud setting.

Security-Performance Trade-Offs using Alternatives (SMPC & TEE): Researchers tend to make comparisons between HE and other privacy-saving technologies:
Secure Multi-Party Computation (SMPC): Data is disseminated among several parties in the form of data shares. Calculations are performed in a group manner without the disclosure of the underlying data. In certain situations, SMPC is quicker but it has to have numerous non-colluding parties.
Trusted Execution Environments (TEE): Hardware secure enclaves (such as Intel SGX) provide the ability to run computations on plaintext in an isolated trusted environment. Fast but not asynchronous, TEEs are based on hardware trust and may be susceptible to side-channel attacks.
Key Trade-Off:
FHE = Maximum Security, Low Performance.
SMPC = Symmetric Security, Multiple Parties.
TEE = Hardware Trust Dependency, High Performance.

Therefore, the selection of HE, SMPC and TEE is determined by the application needs (if confidentiality, performance, or practicality is one of the priorities).

Table 1: Comparison of Homomorphic Encryption Schemes

Scheme	Security Level	Computation Cost	Ciphertext Size	Cloud Integration
FHE (Fully Homomorphic Encryption)	High	Very High	Large	Difficult
SHE (Somewhat Homomorphic Encryption)	Medium	Medium	Medium	Moderate
PHE (Partially Homomorphic Encryption)	Low	Low	Small	Easy

The bar chart depicts the different costs of computation relating to different homomorphic encryption schemes. PHE is demonstrated to have the least computational requirements having a rating of 1 and is thus relatively efficient in performing certain operations. Somewhat Homomorphic Encryption (SHE) is in the middle ground with a cost of 3, between functionality and performance. Fully Homomorphic Encryption (FHE), which has the richest functionality, is much more expensive to compute, at 5. By the graphic comparison, the trade-offs among power of encryption and processing overhead are found as the necessary considerations, that should be made in an effort of coming up with the right scheme that should be employed in securing data application.

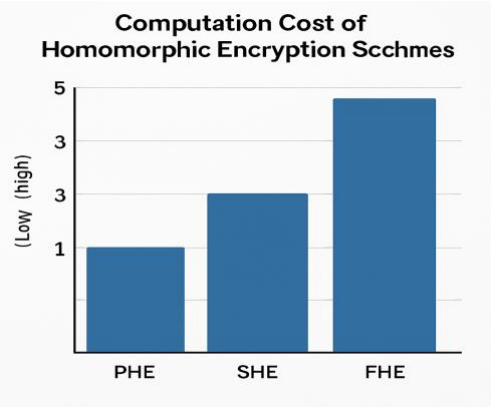

Figure 1: Computation Cost of Homomorphic Encryption Schemes

8. Limitations of the Study

The limitation of this review is the fact research in HE is developing at a rapid rate and recently findings may be soon replaced by new ones. The second limitation is that most of the work is experimental, as well as there is no large-scale and real-life case studies (Albrecht et al., 2018). Also, there are barriers to integration with the heterogeneous clouds infrastructure that limit the deployment.

9. Future Scope

Future studies must look into:

- Performance: optimization of Scheme: low latency, small ciphertexts.
- Integration frameworks: The standard APIs of cloud providers.
- Cross-domain applications: HE extension to IoT-cloud ecosystem and AI-based cloud services.
- Side-channel and quantum resistance Resiliency in the long-run viability (Vaikuntanathan, 2011).
- Massive deployment of HE will require collaboration between researchers and industry and policy makers.

10. Conclusion

Homomorphic encryption is an important step toward cloud data security because it allows privacy-preserving computation. Although efficiency and usability issues remain, research shows a way forward in adoption on a large scale. HE will become a building block of secure cloud infrastructures, especially in sensitive areas where data confidentiality is the most critical factor.

References

1. Acar, A., Aksu, H., Uluagac, A. S., & Conti, M. (2018). A survey on homomorphic encryption schemes: Theory and implementation. *ACM Computing Surveys*, 51(4), 1–35.
2. Albrecht, M., Player, R., & Scott, S. (2018). On the concrete hardness of Learning with Errors. *Journal of Mathematical Cryptology*, 12(3), 169–210.
3. Chen, H., Laine, K., & Player, R. (2019). Simple encrypted arithmetic library – SEAL v3.0. *International Conference on Financial Cryptography and Data Security*.
4. Cheon, J. H., Kim, A., Kim, M., & Song, Y. (2017). Homomorphic encryption for arithmetic of approximate numbers. *ASIACRYPT 2017*.
5. Gentry, C. (2009). Fully homomorphic encryption using ideal lattices. *STOC '09 Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, 169–178.
6. Kumar, S., Mishra, R., & Patel, A. (2020). Data security in cloud computing: A survey of homomorphic encryption. *Procedia Computer Science*, 167, 702–710.
7. Lauter, K., Naehrig, M., & Vaikuntanathan, V. (2014). Can homomorphic encryption be practical? *ACM Cloud Computing Security Workshop*, 113–124.
8. Li, X., Zhang, J., & Chen, X. (2019). Privacy-preserving financial data analytics using homomorphic encryption. *IEEE Transactions on Cloud Computing*, 7(3), 776–789.
9. Rivest, R. L., Adleman, L., & Dertouzos, M. L. (1978). On data banks and privacy homomorphisms. *Foundations of Secure Computation*, 169–180.
10. Vaikuntanathan, V. (2011). Computing blindfolded: New developments in fully homomorphic encryption. *IEEE Annual Symposium on Foundations of Computer Science*, 5–16.
11. Zhang, Q., Cheng, L., & Boutaba, R. (2010). Cloud computing: State-of-the-art and research challenges. *Journal of Internet Services and Applications*, 1(1), 7–18.
12. Brakerski, Z., & Vaikuntanathan, V. (2011). Efficient fully homomorphic encryption from (standard) LWE. *IEEE Symposium on Foundations of Computer Science*, 97–106.
13. Chen, Y., & Han, J. (2021). Advances in lattice-based cryptography. *Frontiers in Computer Science*, 3, 1–15.
14. Xu, J., Wang, H., & Li, Y. (2022). Secure cloud data analytics based on homomorphic encryption. *Future Generation Computer Systems*, 130, 62–74.
15. Yagisawa, M. (2015). Fully homomorphic encryption without bootstrapping. *International Journal of Computer Mathematics*, 92(2), 345–358.