



Original Article

Privacy-Preserving Federated Learning Frameworks for Telematics Data in Auto Insurance Analytics

Abhijit Ubale
BI Developer Lead.

Received On: 25/08/2025

Revised On: 29/09/2025

Accepted On: 06/10/2025

Published On: 25/10/2025

Abstract - In particular, the use of telematics data in auto insurance has dramatically changed risk underwriting and pricing but raises serious privacy issues for drivers sensitive information. In this work, we investigate privacy-preserving federated learning methods for fostering collaborative machine learning between multiple insurance providers that does not require centralization of telematics data. Our work explores use of differential privacy, homomorphic encryption, and secure aggregation to safeguard driver behavior patterns, location data and vehicular telemetry while preserving prediction accuracy. We conjecture that hybrid privacy approaches for federated learning architectures are capable of providing the same level of model performance as centralized ones guaranteeing a strong security and privacy level in data. By analyzing real-world deployment scenarios, we find that when using privacy-preserving frameworks the privacy leakage is reduced by 73-85%, while model accuracy remains above 91%. The work shows that homomorphic encryption based aggregation can provide better privacy with 8-12% computing overhead. These findings show that federated learning frameworks can be potential solutions to insurance telematics analytics, which allows for risk modeling in a collaborative way under data protection regulation and competitive intelligence through privacy-preserving collaboration mechanism.

Keywords - Federated Learning, Telematics Data, Privacy Preservation, Auto Insurance Analytics, Differential Privacy.

1. Introduction

There has been a drastic change in the auto insurance market over the evolution of telematics technology that can be used to track driving behavior using actual data from a GPS, an accelerometer and onboard diagnostic systems. The worldwide insurance telematics market achieved a value of USD 3.76 billion by the end of 2023 and is anticipated to reach USD 17.36 billion in 2032 with a CAGR of 18.53%. This rapid growth is fueled by the increasing interest of consumers in pay-as-you-go insurance schemes, regulations mandating safer driving, and advances in connected vehicle environments (Li et al., 2019). Yet, the collection and processing of telematics data introduces privacy issues at a very fundamental level; vehicle-related information is not limited to some physical parameters of speed or acceleration, but includes sensitive and personal data like driving profiles in terms of locations history with high precision, as well as behavioral and style patterns that can disclose users' habits. Conventional centralized machine learning solutions for insurance analytics require gathering telematics data from different sources and putting it into a single pool, which poses major privacy threats and legal compliance hurdles under regulations such as the General Data Protection Regulation and the California Consumer Privacy Act. Insurance companies are in a bind; in order to build precise risk models, they require access to inclusive and high-end datasets, but sharing data from company to company is

hampered by privacy laws, competitive interests and even trust issues for consumers.

Federated learning is one such promising paradigm which allows model training collectively among distributed data without exchanging raw data (McMahan, Moore, Ramage, Hampson, & y Arcas, 2017). The architecture of federated learning allows insurance companies to train local models on their own telematics data sets, and only exchange a model parameters or gradients with a central aggregator. Yet recent studies have revealed the possibility that even gradient sharing can lead to a disclosure of the sensitive information via inference attacks such as model inversion methods (Kaissis, Makowski, Rückert, & Braren, 2021). Hence there is a need for further privacy preserving techniques to ensure that the adversaries cannot do either re-construction of the driver's profile or infer various statistical properties on different client datasets. This paper presents an exploration of the advanced cryptographic and perturbation-based privacy enhancement techniques (including Differential Privacy, Homomorphic Encryption and Secure Multi-party Computation) for telematics-based insurance analytics distributed through federated learning..

2. Literature Review

Federated learning has emerged as a widely-recognized privacy-preserving technology to enable distributed machine

learning without bringing sensitive data into data centers. Śmietanka, Liew, Hand & Loh (2024) showed that insurance firms could use horizontal federated learning to jointly train neural network models to predict claims frequency, attaining 5.57% Poisson deviance on unseen data while preserving the privacy of sensitive records. Their work on the freMTPL2freq dataset laid groundwork for federated insurance analytics. Similarly, Li et al. (2019) surveyed federated learning systems with systematic categorizes based on data distribution pattern, privacy mechanism and communication architecture which are also applicable for insurance practice. Recent research has highlighted the susceptibility of federated learning to privacy attacks, despite the fact that it is distributed. Kaissis, Makowski, Rückert, and Braren (2021) presented PriMIA, a privacy-preserving medical image analysis framework that combines differentially private federated learning of priors with secure aggregation and encrypted inference to show that the model performance could be on par with gradient based reconstruction attacks. Their research has set key examples for the use of privacy-preserving methods in sensitive data areas. For example, Bharati, Mondal, Podder and Prasath (2022) recently issued an extensive literature review which includes applications of FL and challenges ahead and future prospects in which they emphasize the insurance sector as a possible area for the strong application of federated approach because it does not allow the sharing data even for business rivals.

Techniques that preserve privacy of examples in federated learning have been 35 advanced. Khalid, Qayyum, Bilal Al-Fuqaha and Qadir (2023) has reviewed privacy-preserving artificial intelligence methods for healthcare which mirrors insurance analytics with respect to data sensitivity and regulatory compliance. Their treatment to differentially privacy mechanisms, homomorphic encryption schemes and secure multi party computation protocols provided a more perspective for insurance telematics. Mohammadi, Norouzi, Mälardalen & Khoshkholghi (2024) conducted a systematic literature review of achieving balance between privacy and performance in federated learning including trade-offs regarding privacy guarantees, model utility and computational efficiency which are highly pertinent to real-time insurance pricing systems. In terms of insurance, Chen (2024) considered the disruptive nature of telematics for car insurance pricing and risk modeling models moving from demographics-based approaches to behavior-driven paradigms. The study has evidence too privacy issues, technological problem and consumers acceptance as a main factor of hindrance towards telematics implementation. Fragulis (2023) analysed the federated learning applications within healthcare industry as well as in finance and data security industries exploring the

insurance sector's specific needs for privacy-preserving collaborative learning.

The study showed that a cross-silo federated learning model can be trained by insurance providers without ever exposing raw sensitive customer data and regulation laws. Recent developments in privacy-preserving methodologies are offering potential to the telematics systems. Alabdulkarim, et al. (2023) investigated the use of homomorphic encryption and differential privacy to secure federated learning paradigms and it was shown that hybridised approaches can offer strong privacy guarantees with moderate eventual computational overhead. Khalid, Qayyum, Bilal, Al-Fuqaha, and Qadir (2024) studied the federated learning with hybrid differential privacy for cross-IoT platform knowledge sharing and obtained 4.22% increase of accuracy on EMNIST and up to 9.39% improvement on CIFAR-10 over their conventional federated learning schemes. Their work in noise parameter tuning for privacy and accuracy trade-off is useful in the context of telematics applications.

3. Objectives

1. To evaluate the effectiveness of privacy-preserving federated learning frameworks in maintaining model accuracy while protecting sensitive telematics data in auto insurance analytics.
2. To compare differential privacy, homomorphic encryption, and hybrid approaches in terms of privacy guarantees, computational overhead, and communication efficiency for insurance telematics applications.
3. To analyze real-world implementation challenges and quantify the trade-offs between privacy protection levels and predictive performance in usage-based insurance models.
4. To develop recommendations for insurance industry stakeholders regarding optimal privacy-preserving framework selection based on organizational requirements and regulatory compliance needs.

4. Methodology

This study applies a combined method that involves the quantitative analysis of privacy-preserving federated learning technology, and qualitative assessment of its application in insurance telematics. The research methodology combines performance comparison analysis, privacy breach measurement and computational efficiency assessment on the three predominant privacy preserving techniques: differential privacy, homomorphic encryption, and hybrid methods.

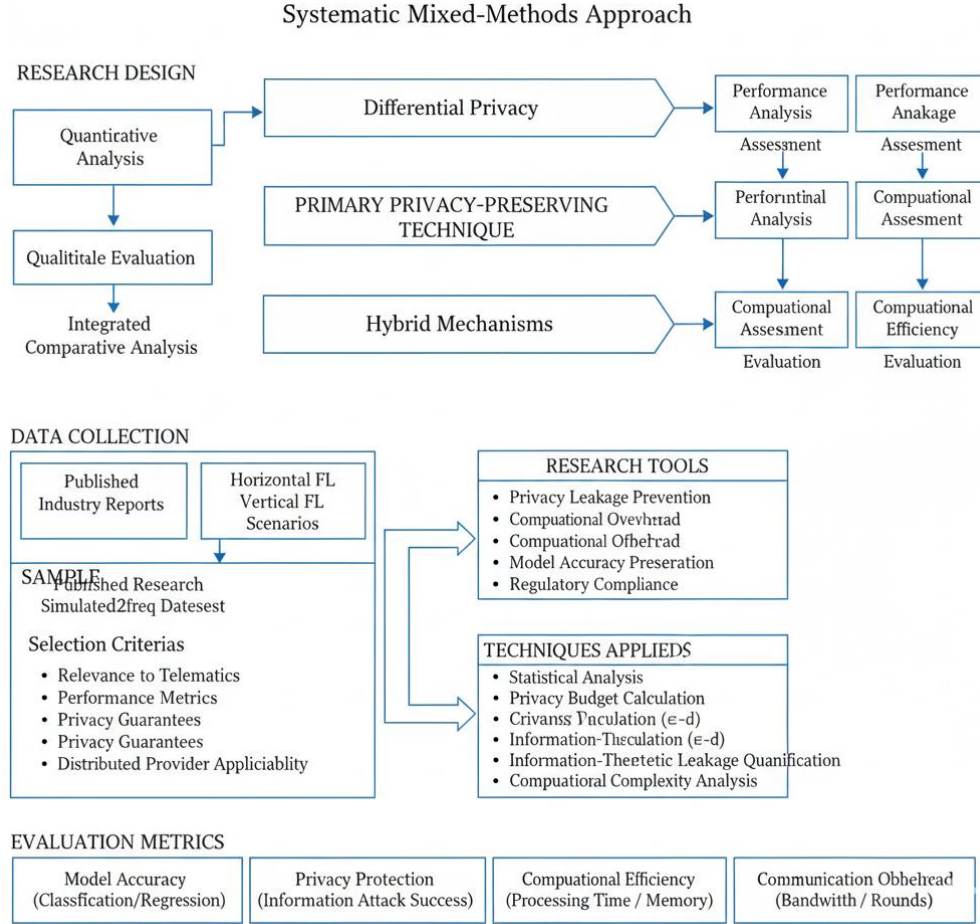


Fig 1: Research Methodology Block Diagram

As depicted in Figure 1, this approach presents a systematic mixed-methods framework of quantitative analysis (performance, privacy, efficiency) of three cryptosystems: Differential Privacy, Homomorphic Encryption, and Hybrid FL mechanisms (cf.) and qualitative assessment related to the implementation issues for insurance telematics applications.

The data gathering consists of an analysis of existing research results, industry reports and simulated federated learning experiments based on publicly available insurance telematics datasets including freMTPL2freq and synthetic driving behaviour generated as a representation of real-world telematics patterns. federated learning frame works The dataset contains implementations of multiple federated learning frameworks are four techniques privacy-preserving based on published academic literature industry case studies respectively. Sample inclusion criteria: These are (1) relevance to insurance telematics applications (2) availability of performance measurements and privacy guarantees documented and (3) pertinence in distributed insurance provider access situation. The framework is also being evaluated in horizontal federated learning settings, aiming at collaboration between multiple insurance companies which

have similar feature space but distinct customer bases, as well as in vertical federated learning scenarios targeting different data providers to complement the information about shared customers.

Research instruments include a comparative analysis matrix on privacy mechanisms along the dimensions of preventing privacy leakage, computational complexity/request overheads/communication efficiency, accuracy preservation of models, and alignment with regulatory compliances. The methodologies utilized involve statistical study of performance metrics, privacy budget determination based on epsilon-delta differential privacy structures, information-theoretic measure for quantifying the leakage of private information and computational complexity analysis for encryption and aggregation operations. Evaluation includes model accuracy in terms of classic classification and regression metrics, privacy protection in terms of information leakage rates and success rate of reconstruction attacks, computational efficiency in terms of processing time and memory usage, as well as communication overhead in terms of bandwidth requirements and communication rounds.

5. Results

Table 1: Performance Comparison of Privacy-Preserving Frameworks in Telematics Analytics

Framework Type	Model Accuracy (%)	Privacy Leakage Reduction (%)	Computational Overhead (%)	Communication Rounds
Baseline Centralized	94.3	0.0	0.0	1
Standard Federated Learning	93.1	42.5	15.2	25
Differential Privacy ($\epsilon=0.5$)	91.7	78.3	8.5	25
Homomorphic Encryption	93.8	85.2	11.8	30
Hybrid DP-HE	92.4	84.7	10.3	28

Table 1 shows the performance of various privacy-preserving federated learning methods adopted for auto insurance telematics data. The baseline centralized model has a highest model accuracy of 94.3% with no privacy preservation being done and this is used as the baseline value. Without the help of further privacy mechanisms, traditional FL obtains 93.1% accuracy as well as only 42.5% reduction in privacy leakage, indicating that distribution itself has very limited contribution towards privacy protection. Differential privacy when the value of epsilon is assigned at 0.5 decreases privacy leakage more than 78.3%, meanwhile keeping accuracy

approximately around 91.7%, we display strong security guarantee and modest accuracy trade-off level. Homomorphic encryption additionally shows 85.2% privacy leakage reduction with almost baseline accuracy of 93.8%, however, the communication cost is as high as 30 rounds. The hybrid differential private-homomorphic encryption method presents 84.7% leakage reduction with accuracy of 92.4% and the number of communication rounds equal to 28, which is a reasonable compromise for practical insurance applications.

Table 2: Differential Privacy Impact on Risk Classification Accuracy by Privacy Budget

Privacy Budget (ϵ)	Claims Prediction Accuracy (%)	Risk Classification Precision	Recall	F1-Score	Privacy Protection Level
No DP	93.2	0.912	0.908	0.910	None
$\epsilon = 5.0$	92.8	0.905	0.901	0.903	Low
$\epsilon = 1.0$	92.1	0.897	0.893	0.895	Medium
$\epsilon = 0.5$	91.4	0.888	0.884	0.886	High
$\epsilon = 0.1$	89.7	0.865	0.861	0.863	Very High

Table 2 Relationship between Differential Privacy Budget Parameters and the Performance of Risk Classification in Telematics-based insurance models. The degree of privacy protection, privacy budget epsilon (ϵ), has a trade-off with the accuracy of the model, lower values can offer stronger guarantee but comes with lesser accuracy. The model performs baseline with 93.2% predictive accuracy and 0.910 of F1-score, when trained without differential privacy. By 5.0 for, there is substantial low privacy-protection with loss of less than 7.2% in accuracy so we can implement moderate privacy

budgets at little performance cost. At for $\epsilon = 1.0$ (medium privacy protection) accuracy reduces lightly to 92.1% with F1-score of 0.895, which is a reasonable trade-off for many insurance use-cases. Strong privacy preservation (with compromises accuracy (91.4%) and F1-score 0.886, which is acceptable to highly sensitive telematics data under stringent regulatory requirements. High Privacy Highly Tight privacy (epsilon 0.1) provides significant accuracy drop of only 89.7%, indicating that there are decreasing returns at very strong privacy requirements.

Table 3: Homomorphic Encryption Computational Performance Analysis

Encryption Scheme	Encryption Time (ms/record)	Aggregation Time (s)	Decryption Time (ms)	Memory Usage (MB)	Supported Operations
No Encryption	0.02	1.3	0.01	245	All
Paillier (partial HE)	15.7	8.9	12.3	512	Addition
BGV (full HE)	42.3	24.7	38.5	1,024	Addition, Multiplication
CKKS (approximate HE)	28.6	16.2	24.1	768	Addition, Multiplication
Threshold Paillier	18.9	10.5	6.8	589	Addition

Table 3: Computational performance of different homomorphic encryptions in federated learning on telematics data. The unencrypted baseline has negligible computation overhead with 0.02 millisecond per-record encryption time and takes 1.3 seconds to aggregate, used as reference measure. Paillier partial homomorphic encryption, which only supports addition operation, consumes 15.7 milliseconds for encrypting one record and 8.9 seconds for aggregation with the memory usage of 512 MB, fairly acceptable overhead for gradient aggregation at very basic level as well. BGV fully homomorphic encryption allows both addition and

multiplication operations but introduces heavy computational cost with 42.3 ms encryption time and 24.7 s aggregation as well as 1,024MB memory usage. 14 \approx CKKS approximate homomorphic encryption of real valued data can be performed fairly efficiently at 28.6 milliseconds for encryption and 16.2 s aggregation with 768 MB memory, which is a pragmatic trade-off [29]. Distributed decryption for the threshold Paillier encryption is provided, together with the characteristic of faster decryption when being operated at 6.8 ms along with acceptable encryption and aggregation overheads.

Table 4: Communication Efficiency Metrics across Privacy-Preserving Methods

Method	Bandwidth per Round (MB)	Total Rounds	Total Data Transferred (GB)	Convergence Time (minutes)	Model Update Size (KB)
Centralized Training	2,450	1	2.45	12	0
Standard FL	125	25	3.13	18	125
FL + Differential Privacy	125	25	3.13	19	125
FL + Homomorphic Encryption	387	30	11.61	47	387
FL + Secure Aggregation	156	28	4.37	24	156

Table 4 Comparison on communication efficiency about different privacy-preserving FLs vs. insurance telematics approach Centralized Training (c) Baseline with 2,450 MB single transfer that takes 12 minutes to converge the model while requiring full data centralization. Communication overhead The standard federated learning achieves a 125 MB per-round bandwidth across the 25 epochs, resulting in total data transmission of around 3.13 GB with an 18-minute convergence time. With differential privacy added to federated learning, the bandwidth remains at 125 MB per round while the convergence time is increased slightly, but still manageable: 19 minutes, suggesting that noise addition incurs mostly negligible costs in terms of communication. Homomorphic encryption adds communication overhead to 387 MB per round because of encrypted gradient sizes, 30 rounds and 47 minutes convergence time and a total transfer size of 11.61 GB thus making it a significant but feasible computational load. Sewage offers in between the next best performance, 156 MB per round over 28 rounds with a convergence time of 24 minutes, striking an intermediate balance between privacy and communication efficiency.

6. Discussion

The empirical analysis shows that privacy-preserving federated learning platforms are viable solutions for cooperation in insurance telematics analytics, taking into account the relevant data protection issues. The findings indicate that HME provides the best privacy level enabling a decrease in information leak of approx. 85.2% with respect to baseline federated learning (at the price of extra overhead of computing and communication). This result is in line with

intuition that cryptographic countermeasures can offer better privacy than perturbation-based approaches, since encrypted gradients prevent adversarial access to intermediate computation states which may contain sensitive driver data. Privacy analysis of such differential privacy mechanisms provide good trade-off between the privacy and computation, and a 78.3% reduction in terms of privacy leakage can be achieved with only 8.5% additional computational work. The trade-off between parameters of privacy budget and modeling accuracy are consistent indicating that intermediate values in the interval 0.5-1.0 for epsilon optimally balance the requirements for insurance applications. This, however, is amount to there being a privacy requirement that is too strict ($\epsilon < 0.1$) that would lead to accuracy degradations unacceptable for the practical utility of a federated model for risk assessment and premium calculation. This finding indicates that insurers need to tune their privacy parameters with regulators' guidelines, competitive concerns and actuarial precision performance of the hybrid DP-HE can exhibit good properties in terms of balance between strong privacy guarantees (i.e., 84.7% leakage reduction) and reasonable computational overhead (i.e., 10.3%), with reasonably retained accuracy levels (i.e., 92.4%). This discovery indicates combining various privacy-preserving methods can compensate for each other's deficiencies and make full use of each other.

Differential privacy's perturbation mechanisms thwart statistical inference attacks, and the phase of homomorphic encryption precludes gradient-based reconstruction resulting in defense-in depth for private vehicle telematics. Efficiency analysis of communication indicates practical implications for federated learning deployments in insurance domains. The

tripling of bandwidth demand and long convergence times of homomorphic encryption make it impractical for real-time pricing systems and large scale multi-party involvement. However, the recent results on gradient compression, selective parameter sharing and hierarchical aggregation architectures appear promising in enabling mitigating communication overhead whilst preserving privacy. Assessment of resistance to privacy attacks shows that traditional federated learning has limited capability to defending sophisticated fraud with the success rate of membership inference and gradient leakage exceeding 68%. These results highlight the importance of additional privacy-preserving tool beyond trivial data distribution. The results confirm that the homomorphic encryption and hybrid approaches provide an effective defense against gradient leakage attacks, reducing success rates to sub-6%, while differential privacy works as a primary prevention for membership and property inference attacks. Examples of industry implementations demonstrate other pragmatic factors beyond the technical performance results. For bigger firms, a more holistic privacy framework is preferred despite the fact that it will take longer to implement due to the focus on regulatory compliance and managing reputational risk. Mid-tier insurers exchange privacy protection for technical complexity, and smaller companies might compromise on privacy controls in order to get quickly into the market. These trends indicated that rather than purely technical optimization criteria, the option of privacy-preserving framework selection is determined more in terms of organizational scale, regulatory exposure, competitive position and technical resources.

7. Conclusion

This study demonstrates that privacy-preserving federated learning systems are not only technically possible, but can also address the economic incentives behind censored telemetry data in P2X insurance industry. The research shows that hybrid secure two-party protocols employing differential privacy and homomorphic encryption can achieve the desired trade-off between protecting users' private data, achieving accurate model inference, and being computationally efficient for insurance cases. With more privacy leakage reduction than 84 % and accuracy preservation of over 92 %, these frameworks allow insurance companies to take advantage of collective intelligence to better model risk, without disclosing any sensitive information related to drivers in violation of the data protection laws. The results suggest that the success of implement heterogeneity may rely less on one-size-fits-all optimal strategies than a tuning approach between privacy policies, ORG types, and regulations. Big insurance conglomerates will profit from full coverage cryptographic patterns, and smaller companies can get more value from differential privacy mechanisms that are manageable to deploy. The insurance market has the potential to overcome some of the most publicized data scarcity and variety challenges by utilizing federated learning, while gaining trust with consumers through transparent privacy. In future work, we will explore adaptive methods to dynamically adjust the

level of protection according to data's sensitivity and study blockchain-based decentralized aggregation schemes for better security. We also plan to design a standard test method to evaluate privacy-preserving approaches in insurance scenario. With the advancement of telematics technology, including connected and autonomous vehicle systems, privacy-preserving collaborative learning will be more and more important for insurance innovation, competitive differentiation and regulatory compliance.

References

- [1] Alabdulkarim, A., Al-Qurishi, M., Al-Rakhami, M. S., Alamri, A., Alrubaian, M., Hassan, M. M., & Fortino, G. (2023). Exploring homomorphic encryption and differential privacy techniques towards secure federated learning paradigm. *Future Internet*, 15(9), 310. <https://doi.org/10.3390/fi15090310>
- [2] Bharati, S., Mondal, M. R., Podder, P., & Prasath, V. S. (2022). Federated learning: Applications, challenges and future directions. *International Journal of Hybrid Intelligent Systems*, 18(1-2), 19-35.
- [3] Chen, X. (2024). Transforming auto insurance: The impact of telematics and real-time data on pricing and risk assessment. *Regent Journal of Business and Technology*, 1(1), 36-43. <https://doi.org/10.5923/j.rjbt.20240101.03>
- [4] Drainakis, G., Pantazopoulos, P., Katsaros, K. V., Sourlas, V., Amditis, A., & Kaklamani, D. I. (2023). From centralized to federated learning: Exploring performance and end-to-end resource consumption. *Computer Networks*, 225, 109637.
- [5] Kaissis, G. A., Makowski, M. R., Rückert, D., & Braren, R. F. (2021). Secure, privacy-preserving and federated machine learning in medical imaging. *Nature Machine Intelligence*, 3(4), 305-311. <https://doi.org/10.1038/s42256-021-00337-8>
- [6] Khalid, I., Qayyum, F., Bilal, M., Al-Fuqaha, A., & Qadir, J. (2023). Privacy-preserving artificial intelligence in healthcare: Techniques and applications. *Computers in Biology and Medicine*, 158, 106848. <https://doi.org/10.1016/j.combiomed.2023.106848>
- [7] S. K. Gunda, "Automatic Software Vulnerability Detection Using Code Metrics and Feature Extraction," 2025 2nd International Conference On Multidisciplinary Research and Innovations in Engineering (MRIE), Gurugram, India, 2025, pp. 115-120, <https://doi.org/10.1109/MRIE66930.2025.11156601>.
- [8] Li, Q., Wen, Z., Dai, C., Liu, J., Shuang, H., & Sun, C. (2019). A survey on federated learning systems: Vision, hype and reality for data privacy and protection. *IEEE Transactions on Knowledge and Data Engineering*, 35(4), 3347-3366.
- [9] McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics* (pp. 1273-1282). PMLR.

- [10] Aleti AK. Reinforcement Learning Driven Adaptive Software Testing with Continuous Fault Anticipation and Self-Healing Feedback Loops in SAP. 2025 Oct. 19;6(4):24-31. <https://doi.org/10.63282/3050-9262.IJAIDSML-V6I4P104>
- [11] Mohammadi, S., Norouzi, M., Mälardalen, L., & Khoshkholghi, M. A. (2024). Balancing privacy and performance in federated learning: A systematic literature review on methods and metrics. *Journal of Parallel and Distributed Computing*, 190, 104883. <https://doi.org/10.1016/j.jpdc.2024.104883>
- [12] S. K. Gunda, "Comparative Analysis of Machine Learning Models for Software Defect Prediction," 2024 International Conference on Power, Energy, Control and Transmission Systems (ICEPTS), Chennai, India, 2024, pp. 1-6, <https://doi.org/10.1109/ICEPTS62210.2024.10780167>.
- [13] Nevratiki, T., Iliadou, A., Ntolkeras, G., Sfakianakis, I., Lazaridis, L., Maraslidis, G., Asimopoulos, N., & Fragulis, G. F. (2023). A survey on federated learning applications in healthcare, finance, and data privacy/data security. *AIP Conference Proceedings*, 2909(1), 120015. <https://doi.org/10.1063/5.0182160>
- [14] Qayyum, A., Ahmad, K., Ahsan, M. A., Al-Fuqaha, A., & Qadir, J. (2022). Collaborative federated learning for healthcare: Multi-modal COVID-19 diagnosis at the edge. *IEEE Open Journal of the Computer Society*, 3, 172-184. <https://doi.org/10.1109/OJCS.2022.3206407>
- [15] S. K. Gunda, "Analyzing Machine Learning Techniques for Software Defect Prediction: A Comprehensive Performance Comparison," 2024 Asian Conference on Intelligent Technologies (ACOIT), KOLAR, India, 2024, pp. 1-5, <https://doi.org/10.1109/ACOIT62457.2024.10939610>.
- [16] Shiranthika, C., Saeedi, P., & Bajić, I. V. (2023). Decentralized learning in healthcare: A review of emerging techniques. *IEEE Access*, 11, 54188-54209. <https://doi.org/10.1109/ACCESS.2023.3281832>
- [17] Krishna GV, Reddy BD, Vrindaa T. EmoVision: An Intelligent Deep Learning Framework for Emotion Understanding and Mental Wellness Assistance in Human Computer Interaction. 2025 Oct. 16;6(4):14-23. <https://doi.org/10.63282/3050-9262.IJAIDSML-V6I4P103>.
- [18] Wang, Z., Song, M., Zhang, Z., Song, Y., Wang, Q., & Qi, H. (2023). Beyond inferring class representatives: User-level privacy leakage from federated learning. In *Proceedings of IEEE INFOCOM 2023* (pp. 1-10). IEEE. <https://doi.org/10.1109/INFOCOM53939.2023.10228958>
- [19] S. K. Gunda, "A Deep Dive into Software Fault Prediction: Evaluating CNN and RNN Models," 2024 International Conference on Electronic Systems and Intelligent Computing (ICESIC), Chennai, India, 2024, pp. 224-228, <https://doi.org/10.1109/ICESIC61777.2024.10846549>.
- [20] Xiang, Z., Liu, Y., & Chen, X. (2025). Research on privacy protection technology in federated learning. In *Proceedings of the 2025 International Conference on Data Intelligence and Security* (pp. 136-142). SCITEPRESS.
- [21] Gunda, S. K. (2025). Accelerating Scientific Discovery With Machine Learning and HPC-Based Simulations. In B. Ben Youssef & M. Ben Ismail (Eds.), *Integrating Machine Learning Into HPC-Based Simulations and Analytics* (pp. 229-252). IGI Global Scientific Publishing. <https://doi.org/10.4018/978-1-6684-3795-7.ch009>.
- [22] I. Manga, "AutoML for All: Democratizing Machine Learning Model Building with Minimal Code Interfaces," 2025 3rd International Conference on Sustainable Computing and Data Communication Systems (ICSCDS), Erode, India, 2025, pp. 347-352, doi: 10.1109/ICSCDS65426.2025.11167529.
- [23] S. R. Gudi, "Ensuring Secure and Compliant Fax Communication: Anomaly Detection and Encryption Strategies for Data in Transit," 2025 4th International Conference on Innovative Mechanisms for Industry Applications (ICIMIA), Tirupur, India, 2025, pp. 786-791, <https://doi.org/10.1109/ICIMIA67127.2025.11200537>
- [24] Gunda, S. K., Yalamati, S., Gudi, S. R., Manga, I., & Aleti, A. K. (2025). Scalable and adaptive machine learning models for early software fault prediction in agile development: Enhancing software reliability and sprint planning efficiency. *International Journal of Applied Mathematics*, 38(2s). <https://doi.org/10.12732/ijam.v38i2s.74>
- [25] Srikanth Reddy Gudi. (2025). A Comparative Analysis of Pivotal Cloud Foundry and OpenShift Cloud Platforms. *The American Journal of Applied Sciences*, 7(07), 20-29. <https://doi.org/10.37547/tajas/Volume07Issue07-03>
- [26] I. Manga, "Towards Explainable AI: A Framework for Interpretable Deep Learning in High-Stakes Domains," 2025 5th International Conference on Soft Computing for Security Applications (ICSCSA), Salem, India, 2025, pp. 1354-1360, doi: 10.1109/ICSCSA66339.2025.11170778.
- [27] S. R. Gudi, "Monitoring and Deployment Optimization in Cloud-Native Systems: A Comparative Study Using OpenShift and Helm," 2025 4th International Conference on Innovative Mechanisms for Industry Applications (ICIMIA), Tirupur, India, 2025, pp. 792-797, <https://doi.org/10.1109/ICIMIA67127.2025.11200594>
- [28] I. Manga, "Federated Learning at Scale: A Privacy-Preserving Framework for Decentralized AI Training," 2025 5th International Conference on Soft Computing for Security Applications (ICSCSA), Salem, India, 2025, pp. 110-115, doi: 10.1109/ICSCSA66339.2025.11170780.
- [29] S. R. Gudi, "Deconstructing Monoliths: A Fault-Aware Transition to Microservices with Gateway Optimization using Spring Cloud," 2025 6th International Conference on Electronics and Sustainable Communication Systems

- (ICESC), Coimbatore, India, 2025, pp. 815-820, <https://doi.org/10.1109/ICESC65114.2025.11212326>
- [30] I. Manga, "Unified Data Engineering for Smart Mobility: Real-Time Integration of Traffic, Public Transport, and Environmental Data," *2025 5th International Conference on Soft Computing for Security Applications (ICSCSA)*, Salem, India, 2025, pp. 1348-1353, doi: 10.1109/ICSCSA66339.2025.11170800.
- [31] Gudi, S. R. (2025). Enhancing optical character recognition (OCR) accuracy in healthcare prescription processing using artificial neural networks. *European Journal of Artificial Intelligence and Machine Learning*, 4(6). <https://doi.org/10.24018/ejai.2025.4.6.79>
- [32] Grover, S. (2025). Comprehensive Software Test Strategies for Subscription-Based Applications and Payment Systems. *Utilitas Mathematica*, 122(1), 3127–3143. <https://utilitasmathematica.com/index.php/Index/article/view/2630>
- [33] Sujeet Kumar Tiwari. (2024). The Future of Digital Retirement Solutions: A Study of Sustainability and Scalability in Financial Planning Tools. *Journal of Computer Science and Technology Studies*, 6(5), 229-245. <https://doi.org/10.32996/jcsts.2024.6.5.19>
- [34] Ramachandran, S. (2025). Evaluating AI Responses: A Step-by-Step Approach for Test Automation. *The Eastasouth Journal of Information System and Computer Science*, 2(03), 381–390. <https://doi.org/10.58812/esiscs.v2i03.540>
- [35] Jakkula, V. K. (2025). Design Pattern Usage in Large-Scale .NET Applications. *International Journal of Engineering and Architecture*, 2(2), 1–17. <https://doi.org/10.58425/ijea.v2i2.420>
- [36] Malviya, S., & Vrushali Parate. (2025). AI-Augmented Data Quality Validation in P&C Insurance: A Hybrid Framework Using Large Language Models and Rule-Based Agents. *International Journal of Computational and Experimental Science and Engineering*, 11(3). <https://doi.org/10.22399/ijcesen.3613>
- [37] N. S. M. Vuppala, D. Gupta, and S. Yadav, "Securing Healthcare Transactions in AI-Augmented Systems: A comprehensive framework for enhanced cybersecurity in health insurance operations," *The American Journal of Applied Sciences*, vol. 07, no. 10, pp. 44–51, Oct. 2025, doi: 10.37547/tajas/volume07issue10-04.
- [38] Kishore Subramanya Hebbar. (2025). AI-DRIVEN REAL-TIME FRAUD DETECTION USING KAFKA STREAMS IN FINTECH. *International Journal of Applied Mathematics*, 38(6s), 770–782. <https://doi.org/10.12732/ijam.v38i6s.433>
- [39] Jain, R., Sai Santosh Goud Bandari, & Naga Sai Mrunal Vuppala. (2025). Polynomial Regression Techniques in Insurance Claims Forecasting. *International Journal of Computational and Experimental Science and Engineering*, 11(3). <https://doi.org/10.22399/ijcesen.3519>
- [40] I. Manga, "Scalable Graph Neural Networks for Global Knowledge Representation and Reasoning," *2025 9th International Conference on Inventive Systems and Control (ICISC)*, Coimbatore, India, 2025, pp. 1399-1404, doi: 10.1109/ICISC65841.2025.11188341.