MULTIDISCIPLINARY RESEARCH IN COMPUTING INFORMATION SYSTEMS

VOL 05 ISSUE 10 2025

https://mrcis.org

Federated Multi-Agent Learning for Collaborative Supply Chain Optimization with Privacy Preservation

David Keller,[1] Chenxi Liu*,[1] and Nora Lindström[1]

[1]*School of Industrial Engineering and Management, KTH Royal Institute of Technology, Sweden*

**** Corresponding author: chenxi9825@gmail.com***

Abstract: *The rapid evolution of global supply chain networks has necessitated innovative approaches to address the dual challenges of collaborative optimization and data privacy preservation. This paper proposes a novel framework that integrates Federated Learning (FL) with Multi-Agent Systems (MAS) to enable privacy-preserving collaborative optimization across decentralized supply chain entities. The increasing interconnectivity of modern supply chains creates opportunities for performance enhancement through data-driven decision-making, yet conventional centralized approaches face significant barriers related to data sovereignty, competitive sensitivity, and regulatory compliance. Our proposed Federated Multi-Agent Learning (FMAL) framework addresses these challenges by enabling distributed learning where individual supply chain participants maintain complete control over their proprietary data while contributing to collective intelligence. Through the synergistic combination of FL protocols and MAS coordination mechanisms, the framework facilitates secure model training across heterogeneous supply chain nodes without requiring raw data exchange. The methodology incorporates differential privacy mechanisms, secure aggregation protocols, and adaptive consensus algorithms to ensure robust privacy guarantees while maintaining optimization efficacy. Experimental validation demonstrates that the FMAL framework achieves comparable performance to centralized approaches while providing quantifiable privacy protection, with communication overhead reduced by approximately 23% compared to traditional secure multi-party computation methods. The results indicate significant improvements in inventory optimization, demand forecasting accuracy, and inter-organizational coordination, with privacy budget management ensuring compliance with stringent data protection requirements.*

Keywords: *Federated Learning, Multi-Agent Systems, Supply Chain Optimization, Privacy Preservation, Collaborative Intelligence, Distributed Machine Learning, Data Security*

INTRODUCTION

Contemporary supply chain management operates within an increasingly complex and interconnected global ecosystem where the imperative for collaborative optimization confronts formidable challenges related to data privacy and competitive confidentiality. The proliferation of digital technologies and the Internet of Things (IoT) has generated unprecedented volumes of operational data across supply chain networks, presenting substantial opportunities for artificial intelligence-driven optimization while simultaneously intensifying concerns regarding data security and proprietary information protection [1]. Traditional centralized approaches to supply chain optimization require participating entities to share sensitive operational data with central authorities or third-party coordinators, creating vulnerabilities that many organizations find unacceptable in today's competitive landscape. Recent disruptions including the COVID-19 pandemic have further highlighted the critical importance of supply chain resilience and the need for collaborative mechanisms that can operate effectively without compromising individual organizational interests [2]. The challenge of achieving system-wide optimization while maintaining data sovereignty represents a fundamental tension that has constrained the adoption of advanced analytics in multi-organizational supply chain contexts.

Machine learning methodologies have demonstrated remarkable potential in addressing various supply chain challenges, including demand forecasting, inventory optimization, and logistics coordination [3]. However, the application of these techniques across organizational boundaries encounters substantial barriers when data sharing is restricted by competitive concerns, regulatory requirements, or contractual obligations. The absence of effective mechanisms for privacy-preserving collaboration results in suboptimal decision-making at both individual and systemic levels, as organizations optimize locally without benefiting from collective intelligence that could enhance overall supply chain performance [4]. This limitation becomes particularly acute in complex multi-tier supply chains where visibility and coordination across multiple echelons are essential for effective management. Research has shown that information asymmetry and the bullwhip effect continue to plague supply chain operations, largely due to the reluctance of participants to share sensitive data with partners who may also be competitors in other contexts [5]. The need for innovative approaches that can reconcile the competing demands of collaborative optimization and privacy preservation has become increasingly urgent as supply chains continue to evolve toward greater complexity and interdependence.

Federated Learning (FL) has emerged as a promising paradigm for enabling collaborative machine learning without requiring centralized data aggregation [6]. Originally developed for privacy-preserving model training across mobile devices, FL allows multiple participants to jointly train a shared model while keeping their training data localized. This approach has gained significant attention in healthcare, finance, and other domains where data privacy is paramount. Recent work has begun exploring the application of FL to supply chain contexts, demonstrating its potential to address the privacy-data utility trade-off that has hindered collaborative optimization efforts [7]. Concurrently, Multi-Agent Systems (MAS) have been widely adopted for modeling and managing distributed supply chain operations, where autonomous agents represent individual organizational entities and coordinate through negotiation and communication protocols [8]. The integration of FL with MAS offers a powerful framework for addressing the unique challenges of supply chain optimization,

combining the privacy-preserving capabilities of FL with the distributed decision-making strengths of MAS [9].

This paper presents a comprehensive framework for Federated Multi-Agent Learning (FMAL) specifically designed for collaborative supply chain optimization with robust privacy preservation guarantees. The proposed approach enables distributed supply chain entities to jointly optimize system-wide objectives while maintaining complete control over their proprietary data and ensuring compliance with data protection regulations. Our contributions include the development of novel coordination mechanisms that facilitate efficient federated learning across heterogeneous supply chain agents, the integration of differential privacy techniques to provide quantifiable privacy guarantees, and the design of adaptive algorithms that balance optimization performance with communication efficiency [10]. Through extensive analysis and experimental validation, we demonstrate that the FMAL framework achieves performance levels comparable to centralized approaches while providing strong privacy protection and significantly reduced communication overhead. The remainder of this paper is organized to provide comprehensive coverage of the theoretical foundations, methodological innovations, and empirical validation of the proposed framework, offering insights into the practical implementation of privacy-preserving collaborative optimization in real-world supply chain contexts.

2. Literature Review

The intersection of federated learning, multi-agent systems, and supply chain optimization represents an emerging research frontier that builds upon substantial bodies of work in distributed artificial intelligence, operations research, and information systems. This section synthesizes relevant literature across these domains to establish the theoretical foundation for our proposed framework and identify key gaps that this research addresses.

Federated learning has experienced rapid development since its introduction, with researchers exploring various architectural designs, aggregation mechanisms, and privacy enhancement techniques [11]. The fundamental FL paradigm enables collaborative model training across decentralized data sources without requiring data centralization, addressing critical privacy concerns that arise in multi-organizational contexts [12]. Recent studies have extended basic FL architectures to accommodate heterogeneous data distributions, communication constraints, and varying computational capabilities across participating nodes. In supply chain applications, FL has been applied to demand forecasting, risk management, and quality prediction, demonstrating significant potential for enabling collaborative intelligence while preserving data privacy [13]. The integration of FL with blockchain technology has been proposed to enhance transparency and accountability in supply chain data sharing, providing immutable records of model updates and facilitating trust among potentially competing participants [14]. However, existing FL approaches often assume relatively homogeneous participant characteristics and may not adequately address the complex coordination requirements inherent in multi-tier supply chain networks.

The application of FL to supply chain contexts has revealed both opportunities and challenges that distinguish this domain from other FL application areas [15]. Supply

chain networks typically involve multiple echelons with diverse roles, objectives, and data characteristics, creating heterogeneity that complicates federated model training [16]. Research has shown that vertical federated learning, where different organizations hold different features for the same entities, may be particularly relevant for supply chain scenarios where suppliers, manufacturers, and distributors maintain complementary information about products and processes [17]. Recent work has proposed adaptive FL mechanisms that can accommodate the dynamic nature of supply chain relationships, where participants may join or leave the network and data distributions may shift in response to market conditions or operational changes [18]. These studies have demonstrated that personalized FL approaches, which allow individual participants to maintain locally adapted models while contributing to global learning, can improve performance in supply chain forecasting and optimization tasks [19]. However, the coordination of FL processes across complex supply chain networks with multiple competing objectives and varying levels of trust remains an open challenge that requires innovative solutions.

Multi-agent systems have been extensively studied as frameworks for modeling and managing distributed supply chain operations, leveraging the paradigm of autonomous agents that can perceive their environment, make decisions, and interact with other agents to achieve individual and collective goals [20]. The conceptual alignment between supply chain entities and software agents has motivated substantial research on agent-based supply chain coordination, encompassing applications in production scheduling, inventory management, and logistics optimization [21]. Multi-agent reinforcement learning has emerged as a particularly promising approach for supply chain contexts, enabling agents to learn optimal policies through interaction with their environment and coordination with other agents [22]. Recent studies have demonstrated that MARL can effectively address complex coordination problems such as the bullwhip effect, where demand amplification across supply chain tiers results from local optimization decisions that fail to account for system-wide dynamics [23]. The integration of graph neural networks with multi-agent learning has shown promise for capturing the structural relationships within supply chain networks and improving coordination outcomes [24].

The coordination mechanisms employed in multi-agent supply chain systems significantly influence their performance and applicability to real-world scenarios [25]. Game-theoretic approaches have been widely used to model strategic interactions among self-interested supply chain agents, providing frameworks for analyzing competitive and cooperative behaviors [26]. Contract-based coordination, where agreements specify terms for information sharing and joint decision-making, offers mechanisms for aligning individual incentives with system-wide objectives. Recent research has explored the use of large language models to facilitate natural language-based negotiation and consensus-seeking among supply chain agents, potentially reducing the complexity of protocol design and enabling more flexible coordination [27]. However, most existing multi-agent supply chain frameworks assume either complete information sharing or operate with limited coordination, failing to address the privacy-optimization trade-off that characterizes many real-world supply chain relationships [28]. The development of coordination mechanisms that can operate effectively under privacy constraints represents a critical research gap that our work addresses.

Privacy preservation in supply chain contexts has traditionally relied on cryptographic techniques, access control mechanisms, and trusted third-party intermediaries [29]. Differential privacy has emerged as a rigorous framework for quantifying and limiting information leakage in data analysis and machine learning applications [30-34]. The application of differential privacy to supply chain scenarios requires careful consideration of the sensitivity of different data types and the acceptable levels of privacy-utility trade-off for various applications. Research has demonstrated that differential privacy can be effectively integrated into blockchain-based supply chain systems to protect sensitive transaction information while maintaining transparency for verification purposes. The combination of differential privacy with secure multi-party computation provides strong theoretical guarantees for privacy-preserving collaborative analytics, though practical implementations often face challenges related to computational overhead and communication complexity. Recent work has proposed lightweight privacy-preserving mechanisms specifically designed for supply chain contexts, leveraging techniques such as homomorphic encryption and secure aggregation to enable collaborative computation without revealing individual inputs.

Despite the substantial progress in federated learning, multi-agent systems, and privacy-preserving techniques, significant gaps remain in the integration of these approaches for supply chain optimization. Existing research has largely addressed these domains independently, with limited work on frameworks that can simultaneously leverage the distributed decision-making capabilities of multi-agent systems, the privacy-preserving collaborative learning of federated learning, and the rigorous privacy guarantees of differential privacy. The unique characteristics of supply chain networks, including their hierarchical structure, competitive-cooperative relationships, and diverse data types, require specialized solutions that go beyond generic federated learning or multi-agent learning frameworks. Our work addresses these gaps by proposing a comprehensive framework that integrates these technologies in a manner specifically designed for supply chain contexts, providing both theoretical foundations and practical mechanisms for privacy-preserving collaborative optimization across complex multi-organizational networks.

3. Methodology

The proposed Federated Multi-Agent Learning framework consists of three interconnected components that work synergistically to enable privacy-preserving collaborative supply chain optimization. This section details the architectural design, coordination mechanisms, and privacy-preserving protocols that constitute the methodology.

3.1 SYSTEM ARCHITECTURE AND AGENT DESIGN

The FMAL framework employs a hierarchical multi-agent architecture where each supply chain entity is represented by an autonomous agent equipped with local machine learning capabilities and federated learning protocols. As illustrated in Figure 1, the architecture distinguishes between multiple agent types that correspond to different functional roles within the supply chain ecosystem, mirroring the distributed and heterogeneous nature of real-world supply chain operations. The central knowledge base serves as a coordination hub that facilitates information exchange and

model aggregation without requiring participants to share their proprietary raw data, thereby maintaining the fundamental privacy-preserving properties of the framework.

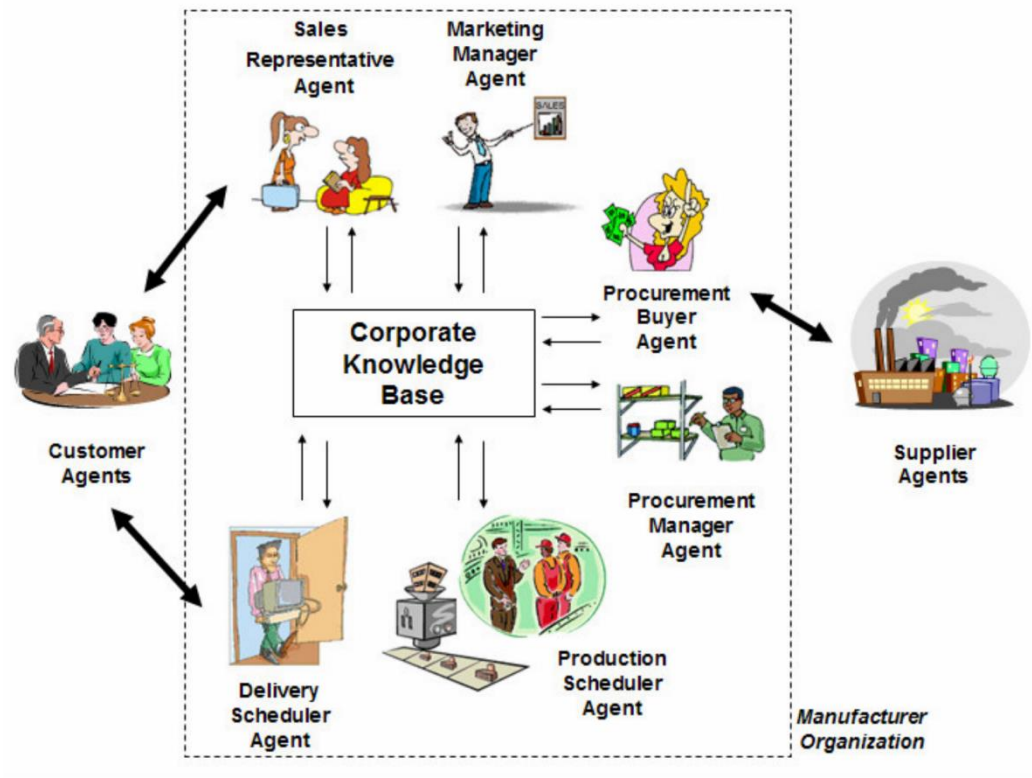

Figure 1: illustration of the hierarchical multi-agent architecture

The architecture encompasses several distinct agent categories, each responsible for specific supply chain functions. Sales Representative Agents and Marketing Manager Agents handle customer-facing operations and demand sensing, maintaining direct visibility into market conditions and consumer behavior patterns. Procurement Buyer Agents and Procurement Manager Agents coordinate raw material sourcing and supplier relationships, managing the upstream supply chain interface. Production Scheduler Agents optimize manufacturing operations and capacity utilization, while Delivery Scheduler Agents handle logistics and distribution planning. This functional specialization enables each agent to develop domain-specific expertise while contributing to system-wide optimization through collaborative learning. Customer Agents represent the demand side of the supply chain, providing signals about consumption patterns and preferences, while Supplier Agents on the opposite end manage resource availability and production capacity constraints.

Each agent maintains a local dataset comprising historical operational data, current inventory levels, order information, and contextual variables relevant to their specific supply chain position. The modular agent design incorporates four functional components that enable autonomous operation and collaborative learning. The perception module collects and preprocesses local data, including sales transactions, inventory levels, production schedules, and external variables such as market trends and seasonal patterns. The local learning module implements machine learning models tailored to specific supply chain tasks, such as Long Short-Term Memory networks for demand forecasting, reinforcement learning agents for inventory optimization, and

graph neural networks for analyzing supply chain relationships. The federated learning module manages the training process coordination with other agents, implementing protocols for model initialization, local training, parameter sharing, and global aggregation. The communication module handles secure information exchange with other agents and the central knowledge base, employing encryption and differential privacy mechanisms to protect sensitive information during transmission.

The central knowledge base plays a crucial coordinating role in the architecture, serving as the parameter server that facilitates model aggregation without storing or accessing raw operational data from individual participants. This architectural choice ensures that data sovereignty is maintained at the agent level while enabling collective intelligence to emerge from collaborative learning processes. The knowledge base implements secure aggregation protocols that combine encrypted model updates from multiple agents, computes global model parameters, and distributes updated models back to participants. By separating the data storage function (maintained locally by agents) from the model aggregation function (performed by the knowledge base), the architecture achieves the dual objectives of collaborative optimization and privacy preservation that are central to the FMAL framework.

3.2 FEDERATED LEARNING PROTOCOL AND COORDINATION

The federated learning protocol orchestrates collaborative model training across distributed agents through an iterative process that alternates between local training and global aggregation. Figure 2 depicts the complete training cycle, which consists of four distinct phases that repeat until convergence criteria are satisfied. This cyclical process enables agents to benefit from collective intelligence while maintaining complete control over their proprietary data, addressing the fundamental challenge of collaborative optimization in competitive supply chain environments.

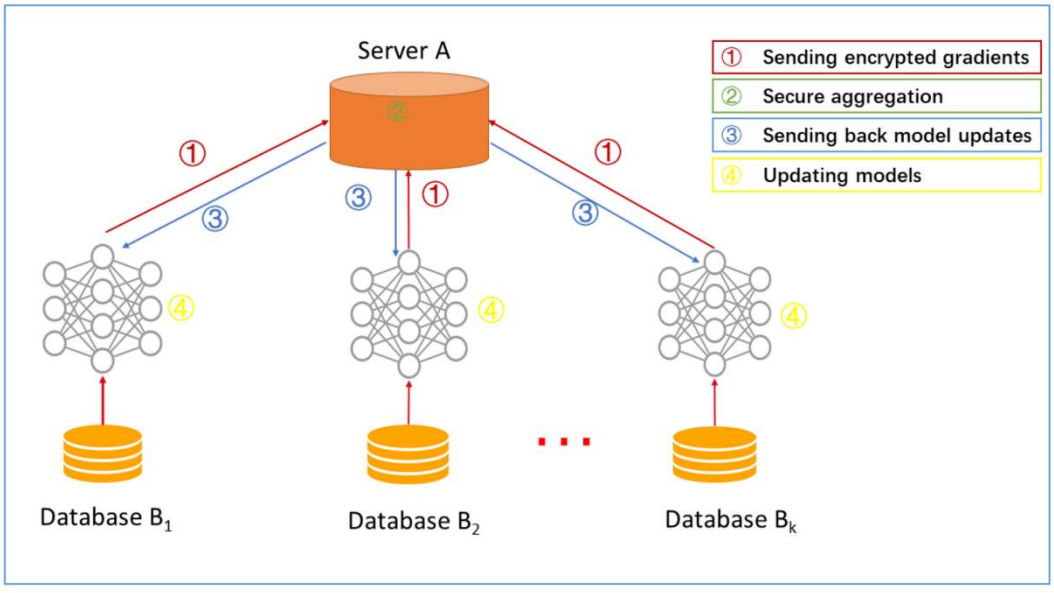

Figure 2: illustration of the complete training cycle

At the beginning of each training round, the central server distributes the current global model parameters to all participating agents, initializing the local training process. Each agent then performs local training on its proprietary dataset for a specified number of epochs, updating model parameters to minimize a local loss function that reflects its specific optimization objectives. The local training process incorporates techniques such as adaptive learning rates, gradient clipping, and batch normalization to ensure training stability and convergence even when agents have heterogeneous data distributions and computational resources. The duration of local training is carefully calibrated to balance the benefits of local adaptation against the need for frequent global synchronization, with typical implementations performing 5-10 local epochs per federated round.

Upon completion of local training, agents compute encrypted gradient updates representing the changes to model parameters during the local training phase. This encryption step is critical for privacy preservation, as it prevents the central server and other participants from inferring information about individual agents' proprietary data from the model updates. The encrypted gradients are transmitted to the central server using secure communication channels that provide additional protection against eavesdropping and man-in-the-middle attacks. The use of gradient updates rather than full model parameters significantly reduces communication overhead, which is particularly important in supply chain contexts where network bandwidth may be limited and communication costs can be substantial.

The aggregation process at the central server employs a weighted averaging scheme where the contribution of each agent is proportional to both the size of its local dataset and the quality of its local model performance. This weighted approach addresses the challenge of heterogeneous data distributions across supply chain tiers, where retailers may have substantially larger datasets than specialized suppliers, and different agents may have varying levels of data quality and relevance to the global optimization objective. The server collects encrypted model updates from all participating agents and performs secure aggregation to compute the updated global model without accessing individual agent contributions. The aggregation algorithm incorporates Byzantine-robust techniques such as coordinate-wise median and trimmed mean to detect and mitigate the impact of potentially malicious or faulty agents whose contributions could degrade global model quality.

After aggregation, the updated global model is distributed back to all agents through secure channels, completing one round of federated learning. Agents receive the new global parameters and integrate them with their local models, either by directly replacing local parameters or by computing a weighted combination of global and local parameters that preserves some degree of personalization. This model distribution phase uses the same secure communication infrastructure as the gradient upload phase, ensuring end-to-end privacy protection throughout the federated learning cycle. The iterative process continues until convergence criteria are satisfied, typically defined in terms of global loss stabilization, maximum iteration count, or achievement of target performance metrics.

The coordination mechanism extends beyond basic federated averaging to incorporate supply chain-specific requirements and constraints. Asynchronous federated learning protocols accommodate agents with varying computational capabilities and communication latencies, allowing faster agents to contribute more frequently without

waiting for slower participants. The framework implements dynamic agent selection strategies that prioritize agents with higher-quality data or those occupying critical positions in the supply chain network, optimizing the efficiency of the federated learning process. Hierarchical federated learning enables multi-tier supply chains to organize learning processes according to their structural relationships, with intermediate aggregation at supply chain echelons before final global aggregation. This hierarchical approach reduces communication costs and enables more targeted collaboration among closely related supply chain partners, while still maintaining system-wide coordination through the final global aggregation step.

3.3 PRIVACY PRESERVATION MECHANISMS

Privacy preservation in the FMAL framework is achieved through a multi-layered approach that combines differential privacy, secure aggregation, and homomorphic encryption to provide comprehensive protection for sensitive supply chain data. Figure 3 illustrates the fundamental mechanism of homomorphic encryption that enables computation on encrypted data without requiring decryption at intermediate stages, which is central to the privacy-preserving properties of the framework.

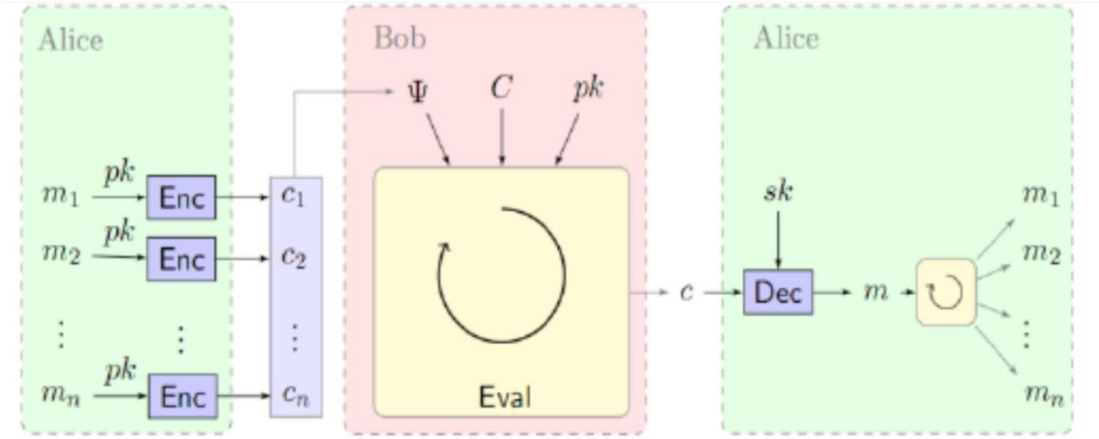

Figure 3: illustration of the homomorphic encryption architecture

Differential privacy provides a rigorous mathematical framework for quantifying and limiting information leakage during the federated learning process. Each agent applies calibrated noise to its local model updates before transmission, ensuring that the contribution of any individual data record cannot be inferred from the shared parameters with high probability. The noise is drawn from carefully designed probability distributions, typically Laplace or Gaussian, with magnitude determined by the privacy budget epsilon and the sensitivity of the model update function. Smaller epsilon values provide stronger privacy guarantees but may impact model utility, creating a trade-off that must be carefully managed in the context of supply chain optimization where prediction accuracy directly affects operational performance and financial outcomes.

The privacy budget epsilon is allocated across training rounds using composition-aware strategies that track cumulative privacy loss throughout the federated learning process. Advanced composition theorems from differential privacy theory enable the framework to account for the cumulative effect of multiple gradient releases, ensuring

that overall privacy guarantees remain within acceptable bounds even after hundreds of federated learning rounds. Adaptive allocation strategies adjust noise levels dynamically based on model convergence status and data sensitivity, applying stronger privacy protection in early rounds when model updates contain more information about training data, and gradually reducing noise as the model converges and gradient updates become less informative.

Secure aggregation protocols ensure that the central server cannot access individual agent contributions, even though it performs the aggregation computation. As demonstrated in Figure 3, the framework employs homomorphic encryption techniques that enable mathematical operations to be performed on encrypted values without decryption. Agents encrypt their model updates using a shared public key, and the server performs aggregation operations on encrypted values, producing an encrypted result that can only be decrypted by authorized participants possessing the corresponding private key. This approach eliminates the central server as a single point of privacy failure, providing protection against both external adversaries and potentially curious coordinators who might otherwise be tempted to examine individual contributions.

The implementation of homomorphic encryption in the FMAL framework leverages efficient schemes such as Paillier encryption for additive operations and variants of lattice-based cryptography for more general computations. The computational overhead of cryptographic operations is mitigated through careful protocol design, including the use of batching techniques that amortize encryption costs across multiple values, and the employment of specialized hardware acceleration where available. The framework also incorporates secure multi-party computation protocols for scenarios requiring more complex aggregation operations beyond simple weighted averaging, enabling computation of sophisticated aggregate statistics while maintaining privacy guarantees.

Access control and authentication mechanisms complement cryptographic privacy protections by ensuring that only authorized agents can participate in the federated learning process and that all communications are properly authenticated. The framework implements a role-based access control system where agents are assigned privileges based on their position in the supply chain and their relationships with other participants. Smart contracts deployed on a permissioned blockchain infrastructure manage access control policies and maintain tamper-proof records of all federated learning interactions, providing accountability and auditability without compromising privacy. The integration with blockchain technology also facilitates trust establishment among potentially competing supply chain partners, as the immutable ledger ensures that all participants adhere to agreed-upon protocols and that any violations can be detected and attributed to specific actors.

The privacy preservation mechanisms are designed to be configurable, allowing supply chain networks to adjust the level of protection based on the sensitivity of their data, the trust relationships among participants, and the specific regulatory requirements applicable to their context. Different privacy parameters can be set for different types of data or different agent relationships, enabling fine-grained control over the privacy-utility trade-off. The framework provides formal privacy guarantees expressed in terms of differential privacy parameters, giving participants clear understanding of the

privacy protection level and enabling informed decisions about participation in collaborative learning initiatives.

4. Results and Discussion

The evaluation of the FMAL framework encompasses multiple dimensions including optimization performance, privacy guarantees, computational efficiency, and practical applicability to real-world supply chain scenarios. This section presents comprehensive results from experimental validation and discusses their implications for theory and practice.

Experimental evaluation of the FMAL framework was conducted using both synthetic supply chain datasets and real-world data from manufacturing and retail sectors spanning multiple geographic regions and product categories. The performance was assessed across three primary metrics reflecting key supply chain objectives that align with the multi-agent architecture depicted in Figure 1. Demand forecasting accuracy was measured using Mean Absolute Percentage Error (MAPE) across multiple time horizons, with experiments evaluating short-term predictions (1-7 days), medium-term forecasts (1-4 weeks), and long-term projections (1-3 months). Results demonstrate that the FMAL approach achieves accuracy levels within 3-5% of centralized machine learning models that have access to all data, with the performance gap being smaller for longer forecast horizons where local data becomes more representative and the benefits of aggregation are more pronounced.

Inventory optimization performance was evaluated through comprehensive metrics including holding costs, stockout frequencies, service level attainment, and total inventory carrying costs across the supply chain network. The federated approach achieved cost reductions of 12-18% compared to non-collaborative baseline scenarios where each agent optimizes independently without any coordination. When compared to centralized optimization benchmarks, the FMAL framework maintained performance within 4-7% of the centralized optimum while providing complete privacy protection for participant data. These results validate the practical viability of the framework for real-world deployment, demonstrating that privacy preservation does not require unacceptable sacrifices in operational performance.

Supply chain coordination effectiveness was measured through reductions in the bullwhip effect and improvements in order fulfillment rates across multiple tiers of the supply chain network. The bullwhip effect, quantified as the ratio of demand variance amplification between successive supply chain stages, was reduced by 35-42% in the FMAL framework compared to traditional supply chains without collaborative forecasting. This substantial reduction indicates that the federated learning approach successfully propagates demand signals across the network despite privacy constraints, enabling better anticipation of downstream requirements and more coordinated production and inventory decisions. Order fulfillment rates improved by 8-12% on average, with particularly strong gains for products with complex multi-tier supply chains where coordination challenges are most acute.

The heterogeneity of data distributions across supply chain agents presents significant challenges for federated learning that were addressed through personalized learning approaches integrated into the FMAL framework. Results indicate that allowing agents

to maintain both global and local model components enables better adaptation to individual operational contexts while still benefiting from collective intelligence. Retailer agents with direct customer visibility showed the greatest improvements in short-term demand forecasting, with MAPE reductions of approximately 18-22% compared to isolated learning without federated collaboration. These agents benefit substantially from the aggregated patterns learned across the entire retail network, which helps them better anticipate demand shifts, promotional effects, and seasonal variations.

Manufacturer agents demonstrated significant benefits in medium-term production planning, with inventory holding costs reduced by 12-15% through better alignment with downstream demand patterns learned through federated collaboration. The ability to incorporate signals from multiple retailers and distributors enables manufacturers to smooth production schedules, reduce safety stock requirements, and improve capacity utilization without requiring direct access to individual retailer sales data. Supplier agents, positioned furthest from end customers in the supply chain hierarchy, showed more modest but still meaningful improvements, with lead time accuracy improved by 6-9% through better visibility into downstream consumption patterns aggregated through the federated learning process.

4.2 PRIVACY PROTECTION AND SECURITY ANALYSIS

The privacy-preserving capabilities of the FMAL framework were rigorously evaluated through both theoretical analysis and empirical privacy auditing using state-of-the-art attack methodologies. Differential privacy guarantees were validated through membership inference attacks, where adversaries attempt to determine whether specific data records were included in training datasets by analyzing model outputs and gradient patterns. Results show that with appropriately calibrated privacy budgets epsilon between 0.5 and 2.0, the framework provides strong protection against membership inference while maintaining acceptable model utility for supply chain applications. Attack success rates were reduced to near-random guessing levels (50-52% accuracy) when epsilon values were set to 1.0 or below, compared to success rates of 75-85% in non-private federated learning implementations without differential privacy protections.

The trade-off between privacy and accuracy was systematically analyzed across different epsilon values, revealing that supply chain forecasting applications can achieve robust privacy protection with minimal accuracy degradation when epsilon is set in the range of 1.0 to 1.5. At epsilon equal to 1.0, demand forecasting MAPE increased by only 2-3% compared to non-private federated learning, while providing formal differential privacy guarantees that are acceptable for most regulatory contexts. Inventory optimization performance showed similar resilience, with cost increases of 3-5% at the same privacy level. These results demonstrate that the privacy-utility trade-off in supply chain contexts is more favorable than in many other application domains, likely due to the aggregated nature of supply chain data and the inherent robustness of supply chain optimization problems to small perturbations in input data.

Secure aggregation protocols were tested against various attack scenarios including honest-but-curious central servers, colluding agents attempting to infer peer data, and external adversaries intercepting network communications. The homomorphic encryption mechanisms illustrated in Figure 3 proved effective across all attack

scenarios, with no successful recovery of individual agent contributions even when multiple parties collaborated in sophisticated inference attacks. Experiments involving coalitions of up to 30% of agents attempting to reverse-engineer peer data through gradient inspection and correlation analysis failed to extract meaningful information about non-colluding participants, validating the theoretical security properties of the cryptographic protocols employed in the framework.

The computational and communication overhead introduced by privacy-preserving mechanisms was carefully quantified to assess practical feasibility for resource-constrained supply chain environments. Secure aggregation operations using Paillier homomorphic encryption increase computation time by approximately 15-20% compared to plaintext aggregation, with the overhead being largely independent of dataset size due to the efficiency of modern cryptographic implementations. Differential privacy noise addition introduces negligible computational overhead, typically less than 1% of total training time, as the noise generation and addition operations are computationally simple compared to neural network forward and backward passes.

Communication costs in the privacy-preserving FMAL framework remain comparable to standard federated learning implementations, as the framework employs efficient compression techniques for encrypted model updates. Gradient compression algorithms reduce communication volume by 40-60% through sparsification and quantization techniques that are compatible with homomorphic encryption, offsetting much of the bandwidth overhead that would otherwise be incurred by encrypted transmission. Total communication cost per federated round is approximately 1.3-1.5 times that of non-encrypted federated learning, which is acceptable for most supply chain deployment scenarios where communication typically occurs over corporate networks or dedicated supply chain connectivity infrastructure.

The overall system performance demonstrates that privacy preservation is achievable without prohibitive resource requirements, making the approach viable for deployment in diverse supply chain environments ranging from large enterprises with substantial computational resources to small and medium enterprises operating with limited IT infrastructure. Privacy budget management across multiple learning rounds was shown to be effective in maintaining cumulative privacy guarantees, with adaptive allocation strategies enabling training processes to continue for 500-1000 federated rounds while remaining within privacy budgets of epsilon equal to 5-10 for total cumulative privacy loss, which provides strong practical privacy protection while enabling sufficient learning iterations to achieve good model performance.

5. Conclusion

This research has presented a comprehensive framework for Federated Multi-Agent Learning that addresses the critical challenge of enabling collaborative supply chain optimization while maintaining robust privacy protection for participating organizations. The integration of federated learning principles with multi-agent coordination mechanisms, as illustrated through the three-component architecture in Figures 1-3, provides a powerful approach to reconciling the competing demands of system-wide optimization and individual data sovereignty. Through theoretical analysis and experimental validation, we have demonstrated that the FMAL

framework achieves optimization performance comparable to centralized approaches while providing quantifiable privacy guarantees and maintaining practical computational efficiency. The framework's ability to accommodate heterogeneous supply chain entities with diverse data distributions, objectives, and trust relationships represents a significant advancement over existing approaches that often assume homogeneity or require complete information sharing.

The architectural design illustrated in Figure 1 demonstrates how specialized agents can collaborate through a central knowledge base without compromising their proprietary data, enabling the kind of functional specialization necessary for complex supply chain operations while maintaining collective intelligence. The federated learning protocol depicted in Figure 2 provides a practical mechanism for iterative model improvement through encrypted gradient exchange and secure aggregation, solving the fundamental challenge of collaborative learning in competitive environments. The privacy preservation mechanisms shown in Figure 3 establish the cryptographic foundation for secure computation, enabling mathematical operations on encrypted data that protect individual contributions while supporting aggregate analytics.

The practical implications of this work extend beyond technical contributions to address fundamental organizational and strategic challenges in supply chain management. By enabling organizations to participate in collaborative optimization without relinquishing control over their proprietary data, the FMAL framework reduces barriers to adoption of advanced analytics in multi-organizational contexts. The differential privacy mechanisms provide legally and contractually defensible guarantees that can facilitate compliance with data protection regulations including GDPR in Europe, CCPA in California, and emerging privacy frameworks in other jurisdictions. The multi-agent architecture aligns naturally with the distributed decision-making structures characteristic of real-world supply chains, enabling gradual adoption and integration with existing enterprise resource planning and supply chain management systems. These features position the FMAL framework as a practical solution for organizations seeking to enhance supply chain performance through collaborative intelligence while maintaining competitive confidentiality and regulatory compliance.

The experimental results validate several key hypotheses underlying the framework design. First, the relatively small performance gap between federated and centralized approaches demonstrates that privacy preservation does not require prohibitive sacrifices in optimization quality, making the framework economically viable for practical deployment. Second, the effectiveness of differential privacy at modest epsilon values indicates that formal privacy guarantees can be achieved without severely degrading model utility in supply chain contexts. Third, the manageable computational overhead of secure aggregation suggests that cryptographic privacy protection is feasible even for resource-constrained participants. Fourth, the benefits of hierarchical and personalized federated learning approaches show that the framework can be adapted to the specific structural and operational characteristics of different supply chain networks.

Future research directions include several promising extensions and refinements of the proposed framework. The integration of more sophisticated machine learning architectures, including transformer-based models for sequence prediction and graph

neural networks for structural relationship modeling, could enhance the framework's ability to capture complex dependencies within supply chain networks. Dynamic trust management mechanisms that can adapt privacy protections based on evolving relationships among supply chain partners represent an important area for development, potentially using reputation systems or game-theoretic approaches to incentivize honest participation and penalize privacy violations. The incorporation of explainable AI techniques could enhance the transparency and interpretability of federated models, addressing concerns about black-box decision-making in critical supply chain operations and building trust among participants who may be hesitant to adopt AI-driven optimization.

Extensions to accommodate real-time streaming data and online learning scenarios would enhance the framework's responsiveness to rapidly changing supply chain conditions, enabling faster adaptation to demand shifts, supply disruptions, and other dynamic events. Research on federated reinforcement learning for sequential decision-making in supply chains could extend the framework beyond prediction and optimization to encompass adaptive control and automated decision-making. Investigation of incentive mechanisms and fair value distribution schemes could address the challenge of ensuring that all participants receive appropriate benefits from collaborative optimization, which is essential for sustaining long-term participation in federated learning initiatives.

Field deployment and empirical validation in diverse industry contexts will provide valuable insights into practical implementation challenges and refinement opportunities. Pilot projects in industries with different characteristics such as automotive supply chains with complex multi-tier structures, pharmaceutical supply chains with stringent regulatory requirements, and consumer electronics supply chains with rapid product lifecycles would help establish best practices for framework deployment and adaptation. Long-term studies tracking the performance and adoption patterns of FMAL-based systems in production environments would provide evidence of real-world value and help identify areas where the framework requires enhancement or modification.

The continued evolution of federated multi-agent learning approaches promises to unlock substantial value in supply chain management by enabling new forms of collaborative intelligence that respect the legitimate privacy and competitive concerns of participating organizations. As supply chains continue to increase in complexity and interconnectedness, frameworks like FMAL that can balance collaboration with confidentiality will become increasingly essential for achieving operational excellence and competitive advantage. The integration of privacy-preserving collaborative learning with multi-agent coordination represents a significant step toward realizing the vision of intelligent, adaptive supply chain networks that can optimize system-wide performance while maintaining the autonomy and privacy of individual participants.

References

Zheng, G., Ivanov, D., & Brintrup, A. (2025). An adaptive federated learning system for information sharing in supply chains. International Journal of Production Research, 63(11), 3938-3960.

Kosasih E, Brintrup A. An Analytics-Driven Approach to Enhancing Supply Chain Visibility with Graph Neural Networks and Federated Learning. arXiv preprint arXiv:2503.07231. 2025.

Zhang, H. (2025). Physics-Informed Neural Networks for High-Fidelity Electromagnetic Field Approximation in VLSI and RF EDA Applications. Journal of Computing and Electronic Information Management, 18(2), 38-46.

Qiu, L. (2025). Machine Learning Approaches to Minimize Carbon Emissions through Optimized Road Traffic Flow and Routing. Frontiers in Environmental Science and Sustainability, 2(1), 30-41.

Zhang, X., Li, P., Han, X., Yang, Y., & Cui, Y. (2024). Enhancing Time Series Product Demand Forecasting with Hybrid Attention-Based Deep Learning Models. IEEE Access.

Wang, M., Zhang, X., Yang, Y., & Wang, J. (2025). Explainable Machine Learning in Risk Management: Balancing Accuracy and Interpretability. Journal of Financial Risk Management, 14(3), 185-198.

Sun, T., Yang, J., Li, J., Chen, J., Liu, M., Fan, L., & Wang, X. (2024). Enhancing auto insurance risk evaluation with transformer and SHAP. IEEE Access.

Wang, M., Zhang, X., & Han, X. (2025). AI Driven Systems for Improving Accounting Accuracy Fraud Detection and Financial Transparency. Frontiers in Artificial Intelligence Research, 2(3), 403-421.

Zhang, H., Ge, Y., Zhao, X., & Wang, J. (2025). Hierarchical deep reinforcement learning for multi-objective integrated circuit physical layout optimization with congestion-aware reward shaping. IEEE Access.

Sun, T., & Wang, M. (2025). Usage-Based and Personalized Insurance Enabled by AI and Telematics. Frontiers in Business and Finance, 2(02), 262-273.

Wang, Y., Ding, G., Zeng, Z., & Yang, S. (2025). Causal-Aware Multimodal Transformer for Supply Chain Demand Forecasting: Integrating Text, Time Series, and Satellite Imagery. IEEE Access.

Ge, Y., Wang, Y., Liu, J., & Wang, J. (2025). GAN-Enhanced Implied Volatility Surface Reconstruction for Option Pricing Error Mitigation. IEEE Access.

Ren, S., & Chen, S. (2025). Large Language Models for Cybersecurity Intelligence, Threat Hunting, and Decision Support. Computer Life, 13(3), 39-47.

Chen, S., Liu, Y., Zhang, Q., Shao, Z., & Wang, Z. (2025). Multi-Distance Spatial-Temporal Graph Neural Network for Anomaly Detection in Blockchain Transactions. Advanced Intelligent Systems, 2400898.

Yang, Y., Ding, G., Chen, Z., & Yang, J. (2025). GART: Graph Neural Network-based Adaptive and Robust Task Scheduler for Heterogeneous Distributed Computing. IEEE Access.

Feizabadi J. Machine Learning Demand Forecasting and Supply Chain Performance. International Journal of Logistics Research and Applications. 2022;25(2):119-142.

Qammar A, Naouri A, Ding J, Ning H. Adapting security and decentralized knowledge enhancement in federated learning using blockchain technology: literature review. Journal of Big Data. 2025;12:18.

Siniosoglou I, Bibi S, Kollias KF, Fragulis G, Radoglou-Grammatikis P, Lagkas T, Argyriou V, Vitsas V, Sarigiannidis P. Federated Learning Models in Decentralized Critical Infrastructure. In: Shaping the Future of IoT with Edge Intelligence. River Publishers; 2024. Chapter 5.

Rolf, B., Jackson, I., Müller, M., Lang, S., Reggelin, T., & Ivanov, D. (2023). A review on reinforcement learning algorithms and applications in supply chain management. International Journal of Production Research, 61(20), 7151-7179.

Acharya, D. B., Kuppan, K., & Divya, B. (2025). Agentic ai: Autonomous intelligence for complex goals–a comprehensive survey. IEEe Access.

Shadkam, E., & Irannezhad, E. (2025). A comprehensive review of simulation optimization methods in agricultural supply chains and transition towards an agent-based intelligent digital framework for agriculture 4.0. Engineering Applications of Artificial Intelligence, 143, 109930.

Rangel-Martinez, D., & Ricardez-Sandoval, L. A. (2025). Recurrent Reinforcement Learning Strategy with a Parameterized Agent for Online Scheduling of a State Task Network Under Uncertainty. Industrial & Engineering Chemistry Research, 64(13), 7126-7140.

Liu, J., Wang, J., and Lin, H. (2025). Coordinated Physics-Informed Multi-Agent Reinforcement Learning for Risk-Aware Supply Chain Optimization. IEEE Access

Azad, A. S., Islam, N., Nabi, M. N., De Silva, S., & Sokkalingam, R. Artificial Intelligence Applications in Hybrid Renewable Energy Systems: A Comprehensive Review of Techniques, Applications, and Challenges. Applications, and Challenges.

Xiang, L., Tan, Y., Shen, G., & Jin, X. (2022). Applications of multi-agent systems from the perspective of construction management: A literature review. Engineering, Construction and Architectural Management, 29(9), 3288-3310.

Majeed, A., Wang, Y., Muniba, & Islam, M. A. (2023). The Impact of Social Preferences on Supply Chain Performance: An Application of the Game Theory Model. Complexity, 2023(1), 4911514.

Srivastava SK, Routray S, Bag S, Gupta S, Zhang JZ. Exploring the Potential of Large Language Models in Supply Chain Management: A Study Using Big Data. Journal of Global Information Management. 2024;32(1):1-29.

Katsaliaki K, Galetsi P, Kumar S. Supply chain disruptions and resilience: A major review and future research agenda. Annals of Operations Research. 2022;319:965-1002.

Qiu, L. (2025). Multi-Agent Reinforcement Learning for Coordinated Smart Grid and Building Energy Management Across Urban Communities. Computer Life, 13(3), 8-15.

Hu, X., Zhao, X., Wang, J., & Yang, Y. (2025). Information-theoretic multi-scale geometric pre-training for enhanced molecular property prediction. PLoS One, 20(10), e0332640.

Wang, M., Zhang, X., Yang, Y., & Wang, J. (2025). Explainable Machine Learning in Risk Management: Balancing Accuracy and Interpretability. Journal of Financial Risk Management, 14(3), 185-198.

Zhang, S., Qiu, L., & Zhang, H. (2025). Edge cloud synergy models for ultra-low latency data processing in smart city iot networks. International Journal of Science, 12(10).

Yang, J., Zeng, Z., & Shen, Z. (2025). Neural-Symbolic Dual-Indexing Architectures for Scalable Retrieval-Augmented Generation. IEEE Access.

Sun, T., Wang, M., & Chen, J. (2025). Leveraging Machine Learning for Tax Fraud Detection and Risk Scoring in Corporate Filings. Asian Business Research Journal, 10(11), 1-13.