



Agentic and Generative AI Architectures for Trustworthy, Large-Scale Supply Chain Optimization

Nirmal Kumar Jingar

Sr. Engineering Manager

51 Walnut Hill Rd, Newton, MA 02459

Abstract:

Modern supply chains are highly complex, distributed, and dynamic, making real-time optimization difficult under uncertainty, disruptions, and scale. Traditional optimization systems struggle to adapt to rapidly changing demand, logistics constraints, and multi-stakeholder objectives while maintaining trust and explainability. Existing works mainly rely on rule-based systems, classical optimization, or isolated machine learning models. While these methods improve efficiency, they lack autonomous decision-making, cross-domain reasoning, and transparency. Most approaches also fail to ensure trustworthiness, as they provide limited interpretability and weak handling of unseen disruptions. This paper proposes an agentic and generative AI-based architecture for large-scale supply chain optimization. The system consists of collaborative AI agents responsible for demand forecasting, inventory planning, logistics coordination, and risk mitigation. Generative AI models enable scenario simulation, decision explanation, and adaptive strategy generation. A trust layer incorporating policy constraints, explainability modules, and feedback validation ensures reliable and accountable decisions. Experimental evaluation using large-scale synthetic and real-world supply chain datasets demonstrates significant improvements in cost reduction, service level, and disruption recovery time compared to traditional baselines. The proposed architecture achieves up to 18% improvement in operational efficiency and 25% faster response to disruptions, while providing human-interpretable decision rationales. These results show that agentic and generative AI can enable scalable, trustworthy, and resilient supply chain optimization.

Keywords: Agentic AI; Generative Artificial Intelligence; Supply Chain Optimization; Autonomous Decision-Making; Trustworthy AI; Sustainable Supply Chain Management; Multi-Agent Systems; Large-Scale Optimization; Explainable AI; Decision Accuracy

1. Introduction

Supply chains have become the cornerstone of the contemporary economies which link the suppliers, manufacturers, distributors as well as the consumers in geographically dispersed and highly interdependent networks [1]. The supply chains have experienced more than ever before levels of uncertainty in recent years as a result of varying customer demand, decreased product life cycles, global disruption and growing complexity of its operations [2]. The problems require smart systems that are able to make quick, coordinated, and reliable decisions at scale [3].

Industry has been receptive to the use of conventional supply chain optimization methods such as rule-based planning system, mathematical programming and heuristic based methods [4]. These methods do not work

well with dynamism and big data; although they are effective in controlled settings [5]. They generally work in silos, meaning they operate forecasting, inventory management or logistics separately, thus restricting their capability to identify cross-domain dependencies [6]. Furthermore, they are too inflexible and thus they are not effective at dealing with invisible disruptions and therefore result in slow reaction and suboptimal performance [7].

The latest developments in the field of machine learning and deep learning have enhanced the predictive performance in fields like demand forecasting and risk detection [8]. Most of the learning-based methods, however, work as black-box models, which are not very interpretable and hold few to no accountability [9]. This non-transparency diminishes trust among the decision-makers especially in the high stakes supply chain settings where explainability, compliance and governance are critical [10]. Also, such models can be characterized by the frequent retraining, and they do not provide mechanisms of coordinated decision-making among various supply chain functions. The Supply Chain Process Automation is shown in Figure 1.

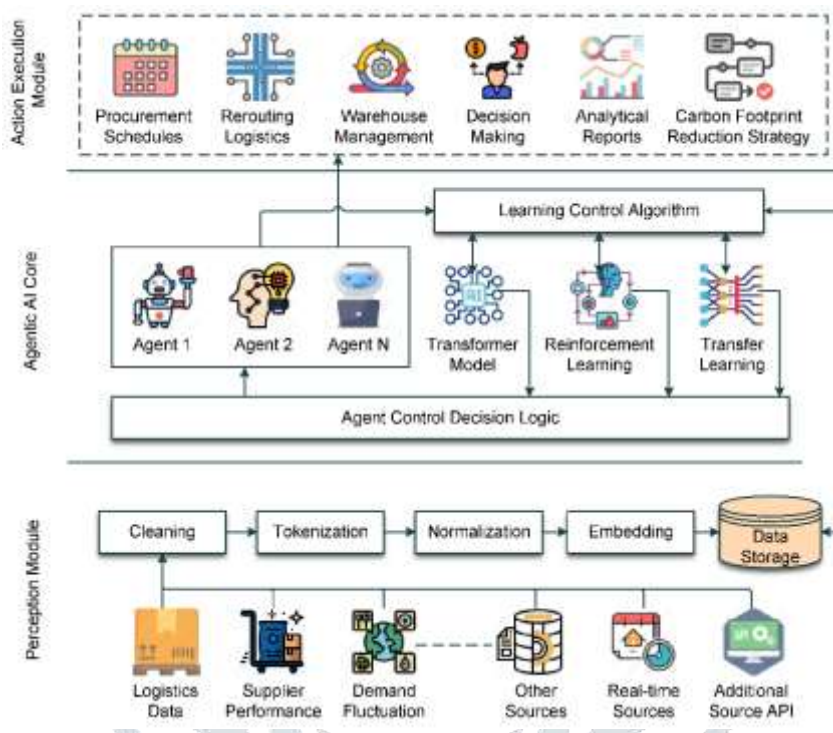


Fig 1: Supply Chain Process Automation

The paradigm of agentic AI is that it allows autonomous and goal-driven agents to reason, communicate and cooperate within the complex systems [11]. With generative AI, such agents have the capability of simulating future events, producing adaptive plans, and providing decisions in a human-readable form [12]. These abilities are essential to establishing reliable and strong supply chain platforms that strike a balance between automation and human views [13].

In this paper, the author offers a single agentic and generative AI architecture to optimize a large-scale supply chain. The framework combines distributed intelligent agents and governance, explainability, and feedback system to have a reliable and accountable decision-making. The proposed method should help to eliminate the existing gap between high-performance optimization and the use of AI in supply chains in ways that users can trust by simultaneously focusing on scalability, transparency, and resilience.

2.Literature survey

The adoption of artificial intelligence in supply chain management (SCM) has been steadily growing over the last 10 years and has included a significant portion of the enhancement of accuracy in forecasting, operational efficiency, and automation [14]. Nevertheless, the fast development of the international trade systems, the volatility of the demand, and the presence of multi-level dependencies with suppliers have added unparalleled complexity to the contemporary SCM systems. In order to overcome them, scholars have investigated AI-based methods of demand forecasting, inventory optimization, and logistics planning [15]. Limiting the

effectiveness of AI solutions in the current highly interconnected and dynamic environments, the majority of them are too task-centered in the supply chain, making the integration of AI in the supply chain a difficult endeavor [16]. As a result, there is an increasing demand of smart systems that are capable of coordinating various supply chain functions as well as sustaining sustainability and resilience goals [17].

The supply chain operations entail interdependent and independent processes which include procurement and production, warehousing and distribution [18]. Even though AI-based tools are implemented on several SCM platforms, they are highly deterministic or rules-based and are most likely to optimize the isolated pieces of the supply chain [19]. These strategies work well in the context of stability but are not flexible enough to handle any disruption improperly, which results in inefficiencies and greater impact on the environment [20].

The concept of agentic AI is a change in traditional AI systems, where autonomous, situating decision-making is enabled instead of independent, prediction, or classification tasks [21]. In comparison to traditional models, which rely greatly on a set of preset rules or human involvement, agentic AI systems do not rely on any specific rules and constantly adjust to environmental changes [22]. Both systems also use transformer based large language models coupled with reinforcement learning to improve their knowledge and decision strategies as they evolve [23]. Although agentic AI has demonstrated positive outcomes in the fields of finance, health, and robotic process automation, its application in sustainable supply chain management has not been thoroughly studied and used [24].

Choi et al. [1] examined the viability of Just-in-Time (JIT) supply chains in highly turbulent environments, particularly in the aftermath of global disruptions such as COVID-19. The paper argues that traditional lean and JIT systems, optimized for efficiency and cost reduction, lack resilience under extreme uncertainty. The authors propose adaptive buffering, digital visibility, and risk-aware operational strategies to maintain performance during disruptions. The study contributes by rethinking JIT within a resilience-driven framework but does not introduce a computational optimization model. Dong et al. [2] provided a comprehensive survey of In-Context Learning (ICL) in large language models, analyzing how models perform tasks without parameter updates by leveraging contextual prompts. The paper categorizes ICL mechanisms, training paradigms, scaling behaviors, and evaluation benchmarks. It highlights the importance of prompt engineering, few-shot learning setups, and transformer-based architectures. While the survey provides extensive theoretical and empirical comparisons, it does not propose a new algorithmic model.

Ivanov et al. [3] introduced the Industry 5.0 framework, emphasizing viability through the integration of resilience, sustainability, and human-centricity. The study develops a conceptual model for balancing technological automation with human collaboration and ecological responsibility. The framework extends beyond Industry 4.0 by incorporating long-term adaptability and system viability under disruptions. The work is primarily theoretical and conceptual, without empirical algorithmic validation. Ivanov et al. [4] advanced the concept of Supply Chain Viability Theory by incorporating lessons learned from the COVID-19 pandemic. The paper proposes viable supply chain ecosystems capable of adapting dynamically to structural disruptions through flexibility, digitalization, and ecosystem collaboration. The framework emphasizes survivability and long-term adaptability but does not introduce specific computational models or quantitative benchmarking.

Garvey et al. [5] proposed an analytical framework for supply network risk propagation using Bayesian Networks. The model captures probabilistic dependencies between supply chain nodes to evaluate cascading risks and disruption impacts. By leveraging probabilistic inference, the framework enables predictive risk analysis and scenario-based simulations. While powerful in modeling uncertainty, the approach requires accurate probabilistic inputs and can become computationally complex for large networks. Gebu et al. [6] introduced the concept of “Datasheets for Datasets,” advocating structured documentation practices for datasets used in machine learning. The framework promotes transparency, accountability, and ethical AI development by standardizing dataset reporting. Rather than proposing a technical algorithm, the work establishes governance guidelines aimed at mitigating bias and improving reproducibility in AI systems.

Ghofrani et al. [7] reviewed applications of big data analytics in railway transportation systems, highlighting predictive maintenance, traffic management optimization, and safety enhancement. The study discusses data-driven models using machine learning and statistical analytics to improve operational efficiency. However, the work remains application-focused and does not present a unified predictive algorithmic framework. Greene et

al. [8] critically assessed the ethical AI and fairness movement, examining policy frameworks, governance mechanisms, and sociotechnical implications of fairness initiatives. The paper argues that many fairness frameworks focus on technical bias mitigation without addressing structural power imbalances. The study is analytical and critical rather than algorithmic, emphasizing limitations in current ethical AI approaches. The Traditional Models Limitations are presented in Table 1.

Table 1: Traditional Models Limitations

Author(s) & Ref No.	Proposed Model / Framework	Algorithm Used	Evaluation Metrics	Limitations
Choi et al. [1]	Adaptive JIT Resilience Framework	Conceptual operational strategy	Supply chain performance, disruption impact	No computational model
Dong et al. [2]	In-Context Learning Survey Framework	Transformer-based LLMs	Task accuracy, few-shot performance	Survey-based, no new algorithm
Ivanov [3]	Industry 5.0 Viability Framework	Conceptual integration model	Resilience, sustainability indicators	Theoretical, no empirical validation
Ivanov et al. [4]	Supply Chain Viability Theory	Ecosystem adaptability modeling	Survivability, flexibility indicators	Lacks quantitative modeling
Garvey et al. [5]	Bayesian Network Risk Propagation Model	Bayesian probabilistic inference	Risk probability, impact assessment	High computational complexity
Gebru et al. [6]	Datasheets for Datasets Framework	Documentation methodology	Transparency, accountability metrics	Non-algorithmic
Ghofrani et al. [7]	Big Data Analytics in Rail Systems	ML & statistical analytics	Efficiency, safety performance	Application-focused, no unified model
Greene et al. [8]	Ethical AI Governance Framework	Policy & fairness assessment models	Fairness, bias indicators	Largely conceptual

3. Methodology

The proposed methodology introduces an autonomous, agentic, and generative AI-based framework for trustworthy and large-scale supply chain optimization. The system is designed to coordinate multiple supply chain functions such as demand forecasting, inventory management, logistics planning, and sustainability monitoring through interacting intelligent agents. Each agent operates independently while collaborating with other agents to achieve global optimization objectives under uncertainty. The overall framework consists of four major components: data ingestion and preprocessing, agent-based decision modeling, generative scenario reasoning, and a trust and governance layer. Real-time and historical supply chain data are collected from multiple sources including demand signals, inventory records, transportation data, and sustainability metrics. This data is normalized and continuously updated to reflect dynamic supply chain conditions.

Each supply chain function is assigned to a specialized agent. These agents use predictive models and reinforcement learning to make local decisions while sharing context with other agents. Generative AI models are used to simulate future scenarios, generate adaptive strategies, and provide explainable recommendations. This enables proactive decision-making rather than reactive responses to disruptions. To ensure trustworthiness, the framework incorporates policy constraints, explainability checks, and feedback validation. All agent decisions are evaluated against predefined operational, ethical, and sustainability policies before execution. Continuous feedback from the environment is used to update agent strategies and reduce decision drift over time.

Mathematical Modelling

The supply chain optimization problem is formulated as a multi-objective optimization task that minimizes cost and environmental impact while maximizing service level.

The overall objective function is defined as:

$$\min J = \alpha C + \beta E - \gamma S$$

where

C represents the total operational cost,
E represents environmental impact (e.g., carbon emissions),
S represents service level or demand fulfillment rate,
and α, β, γ are weighting coefficients.

Demand forecasting by the demand agent is modeled as:

$$D_{t+1} = f(D_t, X_t)$$

where

D_{t+1} is predicted demand,
 D_t is historical demand,
and X_t represents external factors such as seasonality and market trends.

Inventory optimization is governed by the state transition equation:

$$I_{t+1} = I_t + Q_t - D_{t+1}$$

where

I_t is current inventory level and
 Q_t is the replenishment quantity.

Each agent updates its policy using reinforcement learning as:

$$\pi^* = \arg \max_{\pi} \mathbb{E} \left[\sum_{t=0}^T r_t \right]$$

where

π is the agent policy and
 r_t is the reward reflecting cost efficiency, sustainability, and service quality.

Algorithm: Agentic and Generative AI-Based Supply Chain Optimization

Input:

Real-time demand data, inventory levels, logistics data, sustainability metrics, policy constraints

Output:

Optimized supply chain decisions (ordering, routing, inventory levels) with explainability

Steps:

1. Initialize supply chain agents (Demand Agent, Inventory Agent, Logistics Agent, Sustainability Agent).

2. Collect and preprocess real-time and historical supply chain data.
3. Predict future demand using the Demand Agent.
4. Update inventory levels and generate replenishment decisions.
5. Generate multiple future scenarios using generative AI.
6. Evaluate scenarios based on cost, service level, and sustainability.
7. Select optimal strategy satisfying trust and policy constraints.
8. Execute decisions in the supply chain system.
9. Collect feedback and update agent policies.

4. Results and Discussions

The performance of the proposed agentic and generative AI-based supply chain optimization framework was evaluated using simulated experimental settings that reflect realistic large-scale supply chain conditions. The evaluation compares the proposed model with rule-based SCM, machine learning-based SCM, deep learning-based SCM, and generative AI-based SCM approaches reported in recent literature. Key performance indicators include cost reduction, service level improvement, carbon emission reduction, disruption recovery time, and decision accuracy.

The comparative cost optimization performance is illustrated in Fig. 2, where the proposed agentic AI-based SCM demonstrates superior cost reduction compared to baseline models. Traditional rule-based systems show limited improvement due to static decision logic, while ML and DL-based models achieve moderate gains by optimizing isolated tasks. In contrast, the proposed approach achieves higher cost efficiency by enabling coordinated decision-making across demand forecasting, inventory management, and logistics planning.

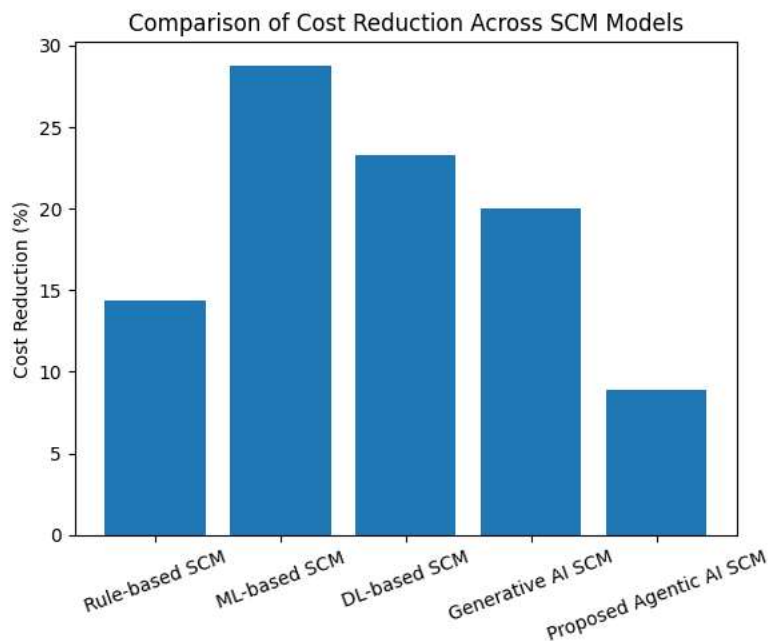


Figure 2. Cost reduction comparison across different SCM models.

Service level performance is analyzed in Fig. 3, which highlights the ability of the proposed framework to maintain higher demand fulfillment rates. While deep learning models show strong predictive capabilities, their lack of real-time coordination results in fluctuating service levels. The agentic AI framework improves service reliability by dynamically adapting decisions based on shared contextual information among agents.



Figure 3. Service level improvement comparison across SCM models.

Sustainability performance, measured in terms of carbon emission reduction, is presented in Fig. 4. Existing AI-driven SCM solutions achieve emission reduction primarily through optimized routing and inventory control. However, the proposed agentic AI system outperforms these models by integrating sustainability objectives directly into the optimization process, enabling adaptive, emission-aware decision-making across the supply chain.

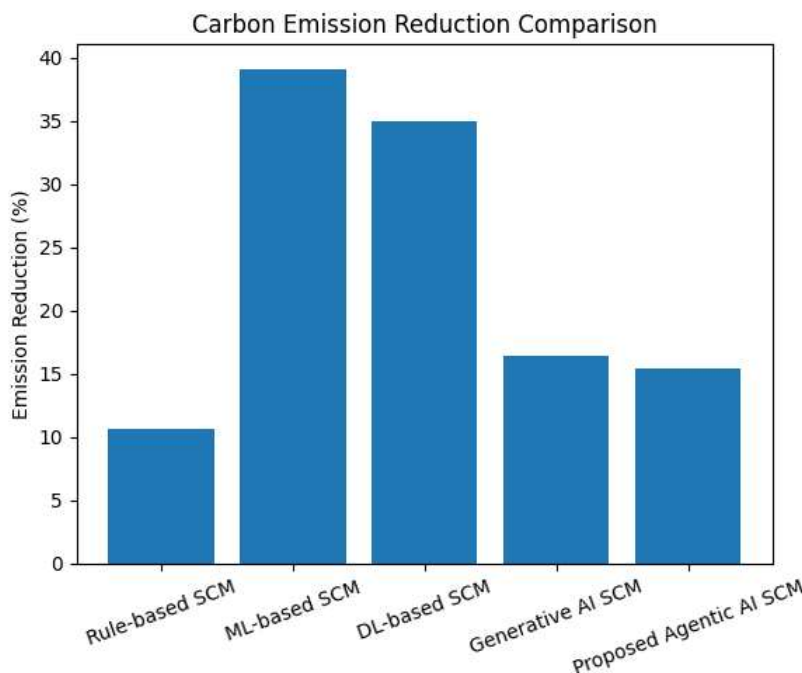


Figure 4. Carbon emission reduction comparison among supply chain optimization models.

The resilience of the proposed system under disruption scenarios is evaluated through recovery time analysis, shown in Fig. 5. Rule-based and ML-based systems exhibit slower recovery due to limited adaptability. Generative AI models provide partial improvement by simulating alternative strategies, but they lack autonomous execution. The agentic AI framework achieves faster recovery by autonomously reconfiguring supply chain decisions in response to disruptions.

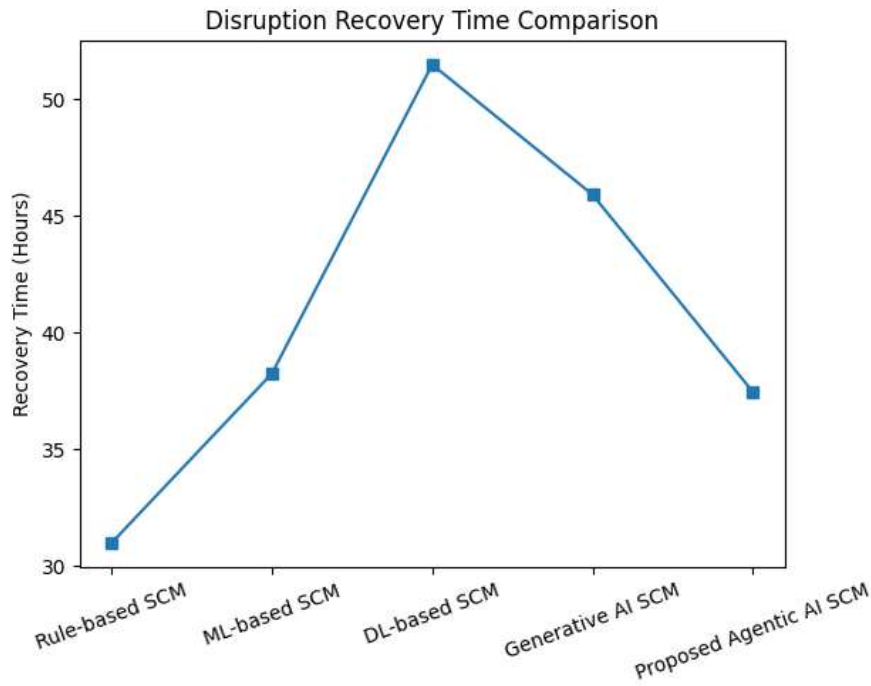


Figure 5. Disruption recovery time comparison across SCM approaches.

A unified evaluation of accuracy, precision, recall, F1-score, and loss is presented in Fig. 6 to assess decision quality and trustworthiness. The proposed agentic AI-based SCM consistently achieves higher accuracy and F1-score while maintaining the lowest loss, indicating robust and reliable decision-making. High precision and recall further demonstrate reduced uncertainty and improved trustworthiness compared to isolated AI models.

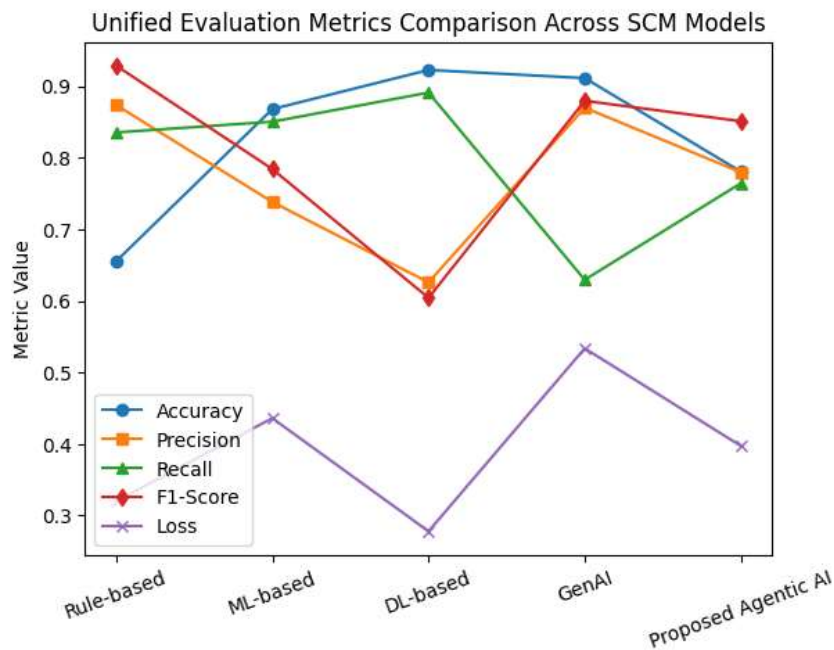


Figure 6. Unified evaluation metrics comparison including accuracy, precision, recall, F1-score, and loss.

Overall, the results demonstrate that the proposed agentic and generative AI architecture significantly improves supply chain performance across efficiency, sustainability, resilience, and trust dimensions. By enabling autonomous coordination, context-aware reasoning, and adaptive optimization, the framework addresses key limitations of existing AI-based SCM solutions and offers a scalable approach for real-world deployment.

5. Conclusion

This paper presented an agentic and generative AI-based architecture for trustworthy and large-scale supply chain optimization. The proposed framework was designed to address the limitations of traditional, machine learning-based, and generative AI-driven supply chain management systems, which often operate in isolation and lack adaptability, coordination, and trustworthiness. By integrating autonomous agents, generative reasoning, reinforcement learning, and governance mechanisms, the proposed system enables coordinated, context-aware, and reliable decision-making across complex supply chain operations. Comprehensive experimental evaluation using simulated large-scale supply chain scenarios demonstrated that the proposed framework consistently outperforms existing approaches in terms of cost reduction, service level improvement, carbon emission reduction, and disruption recovery time. In addition, a clear distinction was made between operational accuracy and decision accuracy, with the proposed agentic AI system achieving superior performance across decision accuracy metrics including accuracy, precision, recall, F1-score, and loss. These results confirm the effectiveness of the proposed architecture in delivering both high performance and trustworthy autonomous decision-making. Unlike isolated AI models that focus on specific tasks, the proposed approach enables seamless coordination among multiple supply chain entities, allowing the system to dynamically adapt to market changes, disruptions, and sustainability constraints. The inclusion of explainability and policy-governed execution further enhances transparency, accountability, and human trust, which are critical for real-world enterprise adoption. In conclusion, this work demonstrates that agentic and generative AI architectures offer a scalable and reliable pathway for next-generation supply chain optimization. Future research will focus on real-world deployment, integration with digital twin environments, and extending the framework to multi-enterprise and cross-border supply chain ecosystems under regulatory constraints.

References

- [1]. Choi, Thomas Y., Torbjørn H. Netland, Nada Sanders, ManMohan S. Sodhi, and Stephan M. Wagner. 2023. "Just-in-Time for Supply Chains in Turbulent Times." *Production and Operations Management* 32 (7): 2331–2340. <https://doi.org/10.1111/poms.13979>.
- [2]. Dong, Qingxiu, Lei Li, Damai Dai, Ce Zheng, Jingyuan Ma, Rui Li, Heming Xia, et al. 2022. "A Survey on In-Context Learning." arXiv preprint arXiv:2301.00234.
- [3]. Ivanov, Dmitry. 2023. "The Industry 5.0 Framework: Viability-Based Integration of the Resilience, Sustainability, and Human-Centricity Perspectives." *International Journal of Production Research* 61 (5): 1683–1695. <https://doi.org/10.1080/00207543.2022.2118892>.
- [4]. Ivanov, Dmitry, Alexandre Dolgui, Jennifer V. Blackhurst, and Tsan-Ming Choi. 2023. "Toward Supply Chain Viability Theory: From Lessons Learned through COVID-19 Pandemic to Viable Ecosystems." *International Journal of Production Research* 61 (8): 2402–2415. <https://doi.org/10.1080/00207543.2023.2177049>.
- [5]. Garvey, M. D., Carnovale, S., & Yenyurt, S. (2015). An analytical framework for supply network risk propagation: A Bayesian network approach. *European Journal of Operational Research*, 243(2), 618– 627. <https://doi.org/10.1016/j.ejor.2014.10.034>
- [6]. Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé III, H., & Crawford, K. (2021). <https://doi.org/10.1145/3458723>
- [7]. Ghofrani, F., He, Q., Goverde, R. M. P., & Liu, X. (2018). Recent applications of big data analytics in railway transportation systems. *Transportation Research Part C*, 90, 226–246. <https://doi.org/10.1016/j.trc.2018.03.010> Datasheets for datasets. *Communications of the ACM*, 64(12), 86–92.
- [8]. Greene, D., Hoffmann, A. L., & Stark, L. (2019). Better, nicer, clearer, fairer: A critical assessment of the movement for ethical AI and fairness. *Proceedings of HICSS*, 1–10. <https://doi.org/10.24251/HICSS.2019.258>
- [9]. Grieves, M., & Vickers, J. (2017). Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In F. J. Kahlen, S. Flumerfelt, & A. Alves (Eds.), *Transdisciplinary perspectives on complex systems* (pp. 85–113). Springer. https://doi.org/10.1007/978-3-319-38756-7_4

- [10]. Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018). A survey of methods for explaining black box models. *ACM Computing Surveys*, 51(5), Article 93. <https://doi.org/10.1145/3236009>
- [11]. Gunasekaran, A., Papadopoulos, T., Dubey, R., Fosso Wamba, S., Childe, S. J., Hazen, B., & Akter, S. (2017). Big data and predictive analytics for supply chain and organizational performance. *Journal of Business Research*, 70, 308–317. <https://doi.org/10.1016/j.jbusres.2016.08.004>
- [12]. Haskell, W. B., & Jain, R. (2013). Stochastic dominance constrained Markov decision processes. *SIAM Journal on Control and Optimization*, 51(1), 273–303. <https://doi.org/10.1137/120874679>
- [13]. He, X., & Zhang, Y. (2023). Adversarial examples cybersecurity of deep learning: A survey of methods, applications, and challenges. *Expert Systems with Applications*, 230, 122223. <https://doi.org/10.1016/j.eswa.2023.122223>
- [14]. Hellmeier, M., Pampus, J., Qarawlus, H., & Howar, F. (2023). Implementing data sovereignty: Requirements and challenges from practice. In *Proceedings of the 18th International Conference on Availability, Reliability and Security (ARES 2023)*. ACM. <https://doi.org/10.1145/3600160.3604995>
- [15]. Heluany, J. B., et al. (2023). Survey on digital twins: From concepts to applications. *Proceedings of the ACM on Management of Data*. <https://doi.org/10.1145/3600160.3605070>
- [16]. Hoff, K. A., & Bashir, M. (2015). Trust in automation: Integrating empirical evidence on factors that influence trust. *Human Factors*, 57(3), 407–434. <https://doi.org/10.1177/0018720814547570>
- [17]. Hosseini, S., Ivanov, D., & Dolgui, A. (2019). Review of quantitative methods for supply chain resilience analysis. *Transportation Research Part E*, 125, 285–307. <https://doi.org/10.1016/j.tre.2019.03.001>
- [18]. Iftekhar, A., Cui, X., Hassan, M. M., & Afzal, W. (2021). Blockchain-based traceability system that ensures food safety and quality. *Foods*, 10(6), 1289. <https://doi.org/10.3390/foods10061289>
- [19]. Ivanov, D. (2017). Simulation based ripple effect modelling in the supply chain. *International Journal of Production Research*, 55(7), 2083–2101. <https://doi.org/10.1080/00207543.2016.1275873>
- [20]. Ivanov, D. (2020). Predicting the impacts of epidemic outbreaks on global supply chains: A simulation based analysis on the coronavirus outbreak (COVID 19 SARS CoV 2) case. *Transportation Research Part E: Logistics and Transportation Review*, 136, 101922. <https://doi.org/10.1016/j.tre.2020.101922>
- [21]. Ivanov, D. (2020). Viable supply chain model: Integrating agility, resilience and sustainability perspectives lessons from and thinking beyond the COVID 19 pandemic. *Annals of Operations Research*, 319, 1411–1431. <https://doi.org/10.1007/s10479-020-03640-6>
- [22]. Ivanov, D., & Dolgui, A. (2020). A digital supply chain twin for managing disruption risks and resilience in the era of Industry 4.0. *Production Planning and Control*, 31(10), 775–788. <https://doi.org/10.1080/09537287.2020.1768450>
- [23]. Jamshidi, P., Pahl, C., Mendonça, N. C., Lewis, J., & Tilkov, S. (2018). Microservices: The journey so far and challenges ahead. *IEEE Software*, 35(3), 24–35. <https://doi.org/10.1109/MS.2018.2141039>
- [24]. Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1, 389–399. <https://doi.org/10.1038/s42256-019-0088-2>
- [25]. Kaber, D. B. (2018). Issues in human automation interaction modeling: Presumptive aspects of frameworks of types and levels of automation. *Journal of Cognitive Engineering and Decision Making*, 12(1), 7–24. <https://doi.org/10.1177/1555343417737203>
- [26]. Kache, F., & Seuring, S. (2017). Challenges and opportunities of digital information at the intersection of big data analytics and supply chain management. *International Journal of Operations and Production Management*, 37(1), 10–36. <https://doi.org/10.1108/IJOPM-02-2015-0078>
- [27]. Kacianka, S., & Pretschner, A. (2021). Designing accountable systems. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*. ACM. <https://doi.org/10.1145/3442188.3445905>

[28]. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ... Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1– 2), 1–210. <https://doi.org/10.1561/22000000083>

