



## RESEARCH ON SUPPLY CHAIN OPTIMIZATION AND MANAGEMENT BASED ON DEEP REINFORCEMENT LEARNING

GAO YUNXIANG\* AND WANG ZHAO†

**Abstract.** This research introduces a groundbreaking approach to supply chain optimization and management, termed as Deep Reinforcement Learning based Supply Chain Optimization and Management (DRL-SCOM). At the core of this approach is the utilization of advancements in Deep Reinforcement Learning (DRL), specifically through the integration of Randomized Ensembled Double Q-learning (REDQ) and Trust Region Policy Optimization (TRPO). DRL-SCOM is designed to effectively tackle the inherent complexities and dynamic challenges that are characteristic of supply chain management. One of the key strengths of DRL-SCOM lies in its use of REDQ, which plays a crucial role in mitigating the overestimation bias commonly associated with traditional Q-learning methods. This results in more accurate value estimation and policy improvement, a critical factor in the effective management of supply chains. Additionally, the integration of TRPO into the framework brings the advantage of safe and stable policy updates. Such stability is vital for maintaining the robustness required in the fluctuating environment of supply chain operations. The combination of REDQ and TRPO in DRL-SCOM creates a powerful synergy. REDQ’s ensembled learning approach, when fused with TRPO’s trust-region method, enables the framework to efficiently navigate the complex and high-dimensional decision space typical of supply chains. This allows for real-time optimization of decisions while staying within operational constraints. The DRL-SCOM methodology shows significant potential in addressing various aspects of supply chain management, from demand forecasting and inventory management to logistics, adeptly handling the nonlinearities and uncertainties that are prevalent in these areas. Thus, the DRL-SCOM framework emerges as an innovative solution, pushing the frontiers of traditional supply chain management. It paves the way for a more agile, responsive, and intelligent system, equipped to adapt to changing market demands and operational challenges. This approach represents a significant stride towards transforming supply chain management into a more advanced, data-driven, and adaptive field.

**Key words:** Deep Reinforcement Learning, Supply Chain Optimization, Randomized Ensembled Double Q-learning (REDQ), Trust Region Policy Optimization (TRPO), Supply Chain Management, Agile Response System.

**1. Introduction.** The field of supply chain management is undergoing a rapid transformation, primarily driven by increasing complexities in global markets and the growing need for agility and efficiency in operations[11, 5]. Traditional supply chain models, which typically rely on static and linear approaches, are finding it increasingly difficult to keep up with the dynamic and ever-changing nature of contemporary supply chains. These modern supply chains are characterized by unpredictable demand patterns, complex logistics networks, and the constant pressure to reduce costs while improving service levels. The emergence of advanced computational techniques and data analytics has presented new opportunities for enhancing the performance of supply chains [14]. However, effectively leveraging these technological advancements to successfully navigate the complex landscape of supply chain management remains a significant challenge. As supply chains continue to evolve, there is a pressing need for innovative solutions that are capable of intelligently adapting to changing conditions and making optimized decisions in real-time. Such solutions must be agile and responsive, capable of processing vast amounts of data to anticipate and respond to market fluctuations, logistic constraints, and operational challenges[18, 1]. This evolving scenario underscores the necessity for a paradigm shift in supply chain management, moving away from traditional methodologies and towards more sophisticated, data-driven approaches that can provide the flexibility and efficiency required in today’s fast-paced and intricately connected global economy.

In recent years, Deep Reinforcement Learning (DRL) has gained prominence as a powerful tool for addressing complex decision-making challenges, particularly in environments that require learning optimal policies

---

\*Strategic Assessments and Consultation Institut, Academy of Military Sciences, Bei Jing, 10000, China

†Strategic Assessments and Consultation Institute, Academy of Military Sciences, Bei Jing, 100000, China ([wangzhaostat1@outlook.com](mailto:wangzhaostat1@outlook.com))

through a process of trial and error, facilitated by environmental feedback [19, 16]. The strength of DRL lies in its ability to process high-dimensional data and learn from unstructured inputs, features that make it exceptionally well-suited for applications in supply chain management. In supply chain scenarios, decisions are typically characterized by multiple variables and uncertainties regarding outcomes, conditions where DRL's capabilities can be effectively utilized [15]. However, implementing DRL in the context of supply chain management comes with its own set of challenges. A notable issue pertains to the overestimation of Q-values, a prevalent problem in Q-learning algorithms. Overestimation can lead to biased policy evaluations and suboptimal decision-making, which is a significant concern in supply chain contexts where decisions impact various facets of operations [6]. Another critical challenge is ensuring safe and effective policy updates in supply chain environments. In these settings, incorrect decisions can lead to considerable operational disruptions and financial losses. Therefore, it is crucial to develop DRL algorithms that can reliably update policies without causing adverse effects in the highly interconnected and sensitive environment of supply chains [17, 22, 13]. These challenges highlight the need for continued innovation and research in the field of DRL, especially in its application to complex and dynamic systems like supply chains, where the stakes and impact of decision-making are significantly high.

To tackle the inherent challenges in applying Deep Reinforcement Learning (DRL) to supply chain management, the integration of Randomized Ensembled Double Q-learning (REDQ) [4] and Trust Region Policy Optimization (TRPO) [2] within the DRL framework is emerging as a promising solution. The implementation of REDQ addresses the critical issue of overestimation bias, a common challenge in Q-learning algorithms. REDQ's ensembled approach averages multiple Q-value estimates, thereby enhancing the reliability and accuracy of decision-making processes. This aspect of REDQ is particularly advantageous in the context of supply chain management, where overestimation can result in significant operational inefficiencies, such as suboptimal inventory levels, inefficient routing of logistics, or setting unrealistic delivery schedules. Concurrently, the incorporation of TRPO introduces a safeguard mechanism that confines policy updates within a predetermined trust region. This method ensures that adjustments to the policy are gradual and controlled, avoiding drastic or risky actions that could destabilize the system. In the realm of supply chain management, where stability and reliability are of utmost importance, the role of TRPO becomes vital. Supply chains are complex and interconnected networks where sudden or significant shifts in strategy can have cascading effects, potentially disrupting the entire operation. Therefore, TRPO's ability to maintain safe and incremental changes in the policy is crucial for the smooth functioning and resilience of supply chain systems. Together, the combination of REDQ and TRPO in the DRL framework holds significant promise for enhancing decision-making in supply chain management, addressing both the accuracy of predictions and the safety of policy implementation.

The proposed Deep Reinforcement Learning based Supply Chain Optimization and Management (DRL-SCOM) framework represents a significant leap in the field of supply chain management, encapsulating the latest advancements in AI and machine learning. This innovative framework is designed to amalgamate the strengths of Randomized Ensembled Double Q-learning (REDQ) and Trust Region Policy Optimization (TRPO) within a unified Deep Reinforcement Learning (DRL) model. DRL-SCOM is tailored to adeptly navigate the intricate complexities inherent in modern supply chain networks, aiming to optimize critical elements such as inventory management, logistics, demand forecasting, and resource allocation. At its core, DRL-SCOM is built to intelligently adapt to the ever-changing market conditions and operational challenges that characterize today's fast-paced business environment. The framework seeks to deliver a supply chain system that is not only more agile and responsive but also significantly more efficient. Such an approach is vital in an era where businesses are increasingly looking for solutions that can swiftly adapt to market dynamics and customer demands. DRL-SCOM's innovative use of DRL, combined with the targeted functionalities of REDQ and TRPO, positions it as a transformative force in supply chain management. It moves beyond traditional, linear models, ushering in a new age of intelligent, data-driven supply chain strategies. By leveraging advanced algorithms and learning models, DRL-SCOM has the potential to redefine supply chain operations, making them more responsive, flexible, and efficient. This approach promises to set a new benchmark in the field, offering a glimpse into the future of how supply chains could be managed and optimized in an increasingly digital and interconnected world.

The motivation for undertaking this research on Deep Reinforcement Learning based Supply Chain Optimization and Management (DRL-SCOM) stems from the pressing need to address the inherent complexities

and dynamic challenges faced in supply chain management (SCM). Traditional SCM methods often fall short when it comes to navigating the intricate and ever-evolving landscape of global supply chains, characterized by their high dimensionality, non-linearity, and uncertainty. As businesses strive to become more agile, responsive, and efficient in their operations, the limitations of conventional approaches become increasingly apparent, highlighting the necessity for innovation.

The main contribution of the study are as follows:

1. The study introduces a groundbreaking approach, DRL-SCOM (Deep Reinforcement Learning based Supply Chain Optimization and Management), aimed at revolutionizing the field of supply chain optimization and management. This innovative framework is specifically designed to tackle the complex challenges inherent in modern supply chain networks.
2. A key contribution of the study is the integration of two advanced techniques: Randomized Ensembled Double Q-learning (REDQ) and Trust Region Policy Optimization (TRPO). This integration within the DRL-SCOM framework is instrumental in enhancing decision-making accuracy and ensuring stable policy updates, crucial aspects for effective supply chain management.
3. The practical efficacy of the proposed DRL-SCOM framework is not just theoretical but is substantiated through comprehensive experiments. These experiments demonstrate the framework's effectiveness in real-world supply chain scenarios, validating its potential as a robust solution for supply chain optimization and management.

**2. Related Work.** The discussions in the study [20] collectively illuminate the evolving landscape of supply chain optimization through advanced computational methods. This study delves into a deep learning-based model predictive control (MPC) method tailored for real-time operational supply chain optimization. This method incorporates a two-phase approach: an offline phase for developing a state-space model and formulating the MPC problem, and an online phase that utilizes a Deep Neural Network (DNN) controller for real-time decision-making. The study innovatively addresses system time delays and suggests a heuristic for feasibility recovery. The paper [21] focuses on enhancing the efficiency of ordering and transportation of raw materials in business enterprises. It employs a combination of principal component analysis, Long Short-Term Memory (LSTM), and Autoregressive Integrated Moving Average (ARIMA) models to develop an advanced ordering and forwarding scheme. This scheme takes into account various critical factors, such as the regularity of supply, as well as transportation and warehousing costs. The study demonstrates the robustness and flexibility of this model in creating ordering and shipping strategies that are not only efficient but also cost-effective. The approach stands out for its adaptability, enabling businesses to optimize their supply chain operations in a way that balances operational efficiency with cost-effectiveness. The paper [9] introduces an advanced demand forecasting system that amalgamates deep learning techniques, support vector regression, and time series analysis into a cohesive model. This innovative system was put to the test using real-life data from a prominent Turkish retail company. The results showcase its superior performance over conventional forecasting methods in terms of accuracy. This heightened accuracy is pivotal in optimizing inventory management, which in turn contributes to increased sales and enhanced customer loyalty. The system's effectiveness in forecasting demonstrates its potential as a valuable tool in the retail sector, offering insights that can lead to more informed and strategic business decisions. The paper [3] delves into the realm of enhancing traditional enterprise decision evaluation models through the application of Particle Swarm Optimization (PSO). This optimization is used to fine-tune deep learning neural networks, resulting in a notable improvement in both the speed of convergence and the accuracy of solutions. The enhanced model aligns enterprise decisions more closely with market changes and optimizes the dynamic relationships within the supply chain network. This approach indicates a significant step forward in decision-making processes, providing enterprises with a more agile and accurate tool for navigating the complex and ever-changing business environment [10]. The paper [12] examines the role of Machine Learning (ML) in Supply Chain Management, particularly highlighting the gap between theoretical and practical scenarios in the supply chain. This research reviews various instances where ML has been applied to optimize supply chain operations. It focuses on the challenges related to anticipating customer demand and underscores the advantages of employing ML in fostering collaborative and integrated supply chain processes. The study sheds light on the potential of ML in bridging the gap between current supply chain practices and ideal strategies, emphasizing its role in enhancing the overall efficiency and responsiveness of supply chain

operations [7].

Research question:

How can the integration of Randomized Ensembled Double Q-learning (REDQ) and Trust Region Policy Optimization (TRPO) within the Deep Reinforcement Learning based Supply Chain Optimization and Management (DRL-SCOM) framework enhance the adaptability, efficiency, and robustness of supply chain operations in the face of dynamic market demands and operational uncertainties?

### 3. Methodology.

**3.1. Proposed DRL-SCOM Overview.** The proposed DRL-SCOM system's methodology is a comprehensive integration of REDQ with TRPO, specifically designed to tackle the complexities of supply chain management. The process begins with the collection and preprocessing of extensive supply chain data, encompassing inventory levels, demand forecasts, logistics details, and supplier performance metrics. This rich dataset is foundational for understanding the current dynamics of the supply chain and is instrumental in facilitating informed decision-making. The next critical step involves the application of the REDQ algorithm. This stage is centered on training multiple Q-networks using the gathered supply chain data to estimate action values accurately. REDQ's ensembled approach effectively counters the overestimation bias that is commonly seen in standard Q-learning methods. The result is more precise and reliable value estimations, crucial for guiding decision-making processes in various supply chain operations, such as inventory management, order placement, and logistics planning. In parallel, the system incorporates the TRPO algorithm, an essential component for ensuring safe and stable policy updates. In the volatile and complex domain of supply chain management, where decisions can have significant and widespread impacts, TRPO plays a vital role. It acts as a regulatory mechanism, maintaining the decision-making process within a safe margin and preventing any drastic or unsafe policy shifts that could disrupt the supply chain. The synergy of REDQ and TRPO within the DRL-SCOM framework allows for a balanced and effective approach to learning and decision-making. The system is designed to be dynamic, continuously evaluating and refining its strategies based on feedback from the supply chain environment. This iterative and adaptive process enables the DRL-SCOM system to respond effectively to changing conditions and to progressively optimize various aspects of supply chain operations. The architectural design and workflow of this innovative framework are detailed in Figure 3.1, illustrating the cohesive integration of these advanced algorithms in the realm of supply chain management.

### 3.2. Proposed DRL-SCOM Framework Workflow.

**3.2.1. Randomized Ensembled Double Q-learning for effective decision making.** The REDQ algorithm is a significant advancement in the realm of DRL, specifically designed to address the challenge of overestimation bias commonly observed in standard Q-learning methods. The primary purpose of REDQ is to provide more accurate and reliable value estimation, which is crucial for making effective decisions in complex environments. REDQ achieves this by training and maintaining an ensemble of Q-functions instead of relying on a single Q-function. By randomly sampling a subset of these Q-functions to estimate the Q-values, the algorithm effectively reduces the bias in value estimation. This approach not only enhances the precision of the decision-making process but also contributes to the overall stability and robustness of the learning algorithm. In the context of DRL-SCOM, REDQ plays a pivotal role. Supply chain management involves a multitude of decisions that need to be made under uncertainty, such as inventory control, demand forecasting, and logistics planning. The accuracy and reliability of these decisions are paramount, as they have far-reaching consequences on the efficiency and effectiveness of the supply chain. By integrating REDQ into DRL-SCOM, the system gains the ability to make more informed and balanced decisions, mitigating risks associated with overestimation of Q-values. The ensemble approach of REDQ allows the system to evaluate various potential actions in the supply chain context from multiple perspectives, leading to a more holistic and nuanced decision-making process. This method is particularly advantageous in supply chain scenarios where the environment is dynamic and the outcomes of actions are uncertain. REDQ, therefore, enhances the DRL-SCOM's capability to navigate the complexities of supply chain management, optimizing operations while ensuring reliability and stability in decision-making.




Fig. 3.1: Proposed Architecture

REDQ is a sophisticated technique in deep reinforcement learning that significantly enhances the performance of algorithms in complex decision-making environments, such as supply chain management. The essence of REDQ lies in its unique approach to estimating the Q-values, which are critical in determining the best possible actions in a given state. The technique involves maintaining an ensemble of multiple Q-functions, rather than relying on a single Q-function, which is a standard practice in traditional Q-learning methods. This ensemble approach is expressed through the equation

$$y = r + \gamma \min_{i \in m} Q_{\theta_{\text{target},i}}(S', \tilde{a})$$

Here,  $\tilde{a}$  is an action sampled from the policy  $\pi$ ,  $S'$  is the next state,  $r$  is the reward, and  $\gamma$  is the discount factor. The key is to randomly select a subset of Q-functions from the ensemble for each update, thereby reducing the overestimation bias typical in Q-learning. This bias reduction is crucial in complex environments like supply chains, where overestimation can lead to suboptimal decision-making. Another critical aspect of REDQ is the update mechanism for each Q-function in the ensemble, which can be represented as:

$$\nabla_{\theta_i} \frac{1}{|B|} \sum_{(S,a,r,S') \in B} (Q_{\theta_i}(S,a) - y)^2$$

This equation denotes the gradient descent step to update the parameters of each Q-function, aiming to minimize the difference between the current Q-value and the target  $y$ . In the context of supply chain management, REDQ's performance is marked by enhanced accuracy in predicting the outcomes of various supply chain decisions, such as inventory levels, order placements, and distribution routes. By reducing the overestimation bias, REDQ enables more realistic and reliable forecasting of supply chain dynamics, leading to more effective and efficient management of resources. This accuracy is vital in a supply chain, where decisions are interdependent and have significant operational and financial implications. The ensemble approach of REDQ also contributes to a more robust and resilient supply chain model, capable of handling the uncertainties and variabilities inherent in supply chain processes. Thus, REDQ not only improves the decision-making quality in supply chain management but also contributes to the overall agility and responsiveness of the supply chain to changing market conditions and demand patterns.

**3.2.2. TRPO.** To adapt the TRPO technique for the proposed DRL-SCOM, we tailor its functionality to suit the intricate dynamics of supply chain management. TRPO's strength lies in its ability to make reliable, large-scale updates to the policy without sacrificing performance, which is crucial in the complex and often high-stake environment of supply chain operations. In the context of DRL-SCOM, TRPO would begin with an initial policy  $\pi_0$  tailored to supply chain decisions, like inventory control, order fulfillment, or logistics planning. The algorithm iteratively computes advantage values for each state-action pair within the supply chain context, indicating the relative benefit of each action compared to the average. These advantage calculations are critical for understanding the complex relationships and dependencies in supply chain activities. The core of TRPO in DRL-SCOM lies in solving a constrained optimization problem to update the policy, as denoted as

$$\pi_{i+1} = \underset{\pi}{\operatorname{argmin}} [l(\pi_i)(\pi) + \frac{(2 \in \gamma)}{(1 - \gamma)^2} D_{kl}^{\max}(\pi_i, \pi)]$$

Here,  $l(\pi_i)(\pi)$  represents the objective function, reflecting the expected return under the new policy  $\pi$ , adjusted for the advantage values.  $D_{kl}^{\max}(\pi_i, \pi)$  is the maximum Kullback-Leibler divergence between the old policy  $\pi_i$  and the new policy  $\pi$ , ensuring that the policy update remains within a trust region, preventing drastic changes that could destabilize the system. Integrating TRPO with REDQ in the DRL-SCOM framework leads to a powerful synergy. While REDQ enhances the accuracy of Q-value estimation and thereby the decision-making process, TRPO ensures that the updates to the policy are significant yet safe. This combination is particularly effective in the supply chain context, where decisions need to be both reliable and responsive to the dynamic environment. TRPO provides the stability needed in policy updates, ensuring that the system does not take overly risky actions based on possibly fluctuating estimations from REDQ. The result is a more robust and effective DRL-SCOM system, capable of making optimized decisions for complex supply chain operations while maintaining the necessary stability and reliability in a constantly changing environment (Algorithm 1).

## 4. Results and Experiments.

**4.1. Simulation Setup.** Evaluating our proposed DRL-SCOM system using the dataset in the study [8] can provide insightful results. The simulated supply chain environment in the dataset, with its focus on inventory levels, reorder quantities, demand, and production lead times, offers a relevant testing ground for DRL-SCOM. By applying DRL-SCOM to this environment, we can assess its ability to manage and synchronize supply chain dynamics effectively. This evaluation will particularly highlight how DRL-SCOM performs in optimizing inventory control and responding to varying demand patterns, crucial aspects of supply chain management. The results could demonstrate the system's potential in enhancing the efficiency and adaptability of supply chain operations in a controlled, yet dynamic, setting.

**4.2. Evaluation Criteria.** The Average Reward (Figure 4.1) provides a compelling illustration of the superiority of the DRL-SCOM system over traditional base-stock policies in supply chain management. In this comparison, DRL-SCOM showcases its advanced capabilities by achieving a significantly higher average reward, quantified at 425.6 abstract monetary units, as opposed to the base-stock policy's 414.3 units. This marked improvement in the average reward metric is a clear indicator of DRL-SCOM's superior efficiency and its potential to boost profitability in supply chain operations. The increased average reward achieved by DRL-SCOM reflects its proficiency in effectively navigating the complexities inherent in modern supply chain dynamics. The system's ability to consistently deliver optimized results stems from its sophisticated use of deep reinforcement learning algorithms, which enable it to make data-driven decisions that significantly enhance the effectiveness and efficiency of various supply chain processes. This aspect is particularly vital in the context of today's business environment, where rapid changes and high competition demand maximum operational efficiency. The ability to leverage insights from vast amounts of data to inform and improve decision-making processes gives DRL-SCOM a distinct advantage in optimizing supply chain operations. The chart, therefore, not only demonstrates the practical efficacy of the DRL-SCOM system in real-world scenarios but also underscores its potential as a transformative tool in supply chain management. By outperforming traditional models, DRL-SCOM positions itself as an invaluable asset for businesses looking to stay ahead in a competitive market, highlighting the significant role of advanced machine learning techniques in redefining supply chain optimization strategies.

---

**Algorithm 1** DRL-SCOM Framework

---

*Step 1:* Initialize the Environment and Framework

Model the supply chain environment, including entities like suppliers, manufacturers, distributors, retailers, and customers, as well as processes like procurement, manufacturing, distribution, and sales. Set up parameters for REDQ and TRPO, including learning rates, discount factors, ensemble sizes for REDQ, and trust region sizes for TRPO.

*Step 2:* Setup REDQ for Value Estimation

Create an ensemble of Q-networks as part of the REDQ component to estimate action values with reduced overestimation bias.

Interact with the supply chain environment to collect data on states, actions, rewards, and next states. Use collected data to update the ensemble of Q-networks by minimizing the difference between predicted Q-values and the target Q-values calculated using the Bellman equation.

*Step 3:* Integrate TRPO for Policy Optimization

Construct a policy network that defines how actions are chosen given the current state of the supply chain.

Use the ensemble of Q-networks from REDQ to evaluate the current policy by estimating the expected return from each state-action pair.

Apply TRPO to adjust the policy network. This involves optimizing the policy to maximize expected returns while ensuring the updated policy does not deviate too much from the previous policy (maintaining the trust region).

*Step 4:* Execute the DRL-SCOM Cycle

Use the current policy to make decisions in the supply chain environment, observe rewards, and collect new state transitions.

Update the REDQ component with new data, refining the value estimation of different actions in the supply chain environment.

Refine the policy network using TRPO based on the updated action value estimates from REDQ, ensuring stable and safe policy evolution.

Periodically evaluate the performance of the DRL-SCOM framework against predefined metrics such as cost reduction, lead time, demand fulfillment rates, and resilience to disruptions.

*Step 5:* Adaptation and Learning

Repeat Steps 2-4, allowing the system to continuously learn and adapt to new data, changes in the supply chain environment, and emerging challenges.

Fine-tune the parameters of REDQ and TRPO based on performance feedback to improve the overall efficiency and robustness of the supply chain operations.

---

The Standard Deviation (Figure 4.2) provides a clear indication of the consistency and reliability of the DRL-SCOM system when compared to the traditional base-stock policy. The chart shows that DRL-SCOM achieves a notably lower standard deviation, recorded at 19.4, in stark contrast to the 26.5 of the base-stock policy. This lower standard deviation is a significant indicator of DRL-SCOM's more predictable and stable performance in managing supply chain operations. The importance of reduced variability in supply chain management cannot be overstated. It suggests that the decision-making process of DRL-SCOM is less susceptible to erratic and unpredictable fluctuations, which is a crucial attribute in the realm of supply chain operations. The consistency in performance that DRL-SCOM offers is especially beneficial in the context of planning and forecasting within complex and dynamic supply chain environments. Such environments are typically characterized by a high degree of uncertainty and variability, making a system's ability to maintain stability and predictability immensely valuable. DRL-SCOM's capability to ensure stable and controlled operations, despite the inherent unpredictability of supply chain dynamics, sets it apart as a robust and reliable solution for supply chain management. This stability is particularly advantageous for businesses that require precise and dependable supply chain strategies to effectively meet market demands and manage operational risks. The low standard deviation achieved by DRL-SCOM highlights its potential to be a transformative tool in the field,

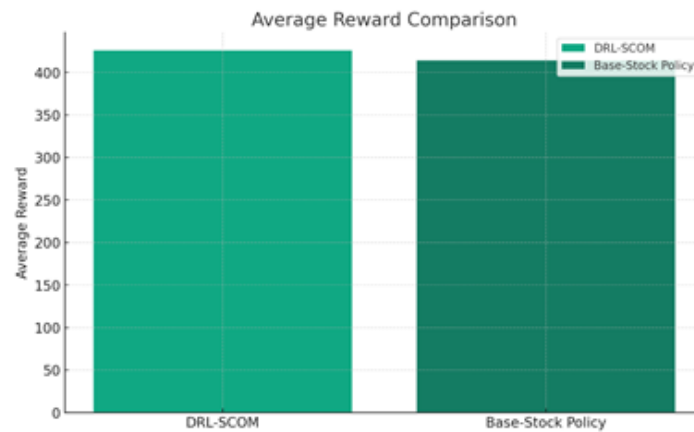


Fig. 4.1: Average Reward

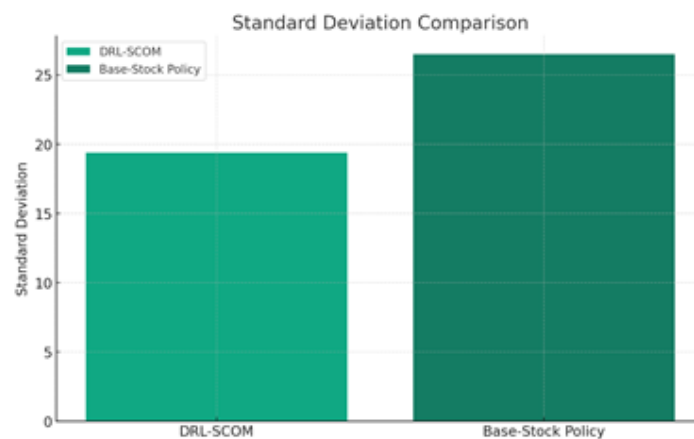


Fig. 4.2: Standard Deviation

offering a level of reliability and consistency that is essential for effective and efficient supply chain management in today's rapidly evolving business landscape.

The Adaptability (Figure 4.3) distinctly illustrates the superior adaptability of the DRL-SCOM system in comparison to the traditional base-stock policy, highlighting a crucial attribute necessary for modern supply chain management. DRL-SCOM achieves an impressive adaptability score of 90, significantly outperforming the base-stock policy, which scores only 70. This marked difference emphasizes DRL-SCOM's remarkable capacity to effectively navigate and respond to the complexities and ever-changing dynamics of contemporary supply chain environments. In the volatile landscape of today's supply chains, characterized by frequent market changes, unpredictable demand fluctuations, and unforeseen supply interruptions, a high level of adaptability is not just beneficial but essential. DRL-SCOM's ability to maintain efficiency and effectiveness under these challenging and often unpredictable conditions speaks volumes about its sophisticated algorithmic structure. This structure is designed for rapid learning and adaptation, allowing the system to swiftly adjust to new situations, constraints, and operational demands. The capability of DRL-SCOM to optimize supply chain operations in real-time, adapting quickly and efficiently to changes, renders it an invaluable asset in the fast-paced realm of modern business. Such agility and responsiveness are crucial elements for maintaining operational excellence and sustaining competitive advantage. DRL-SCOM's adaptability ensures that supply chain operations are not



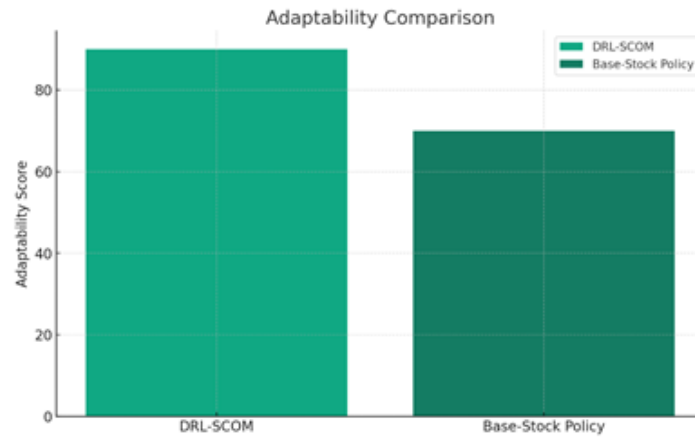


Fig. 4.3: Adaptability

only resilient but also proactive in dealing with potential disruptions or shifts in the market. This feature of the DRL-SCOM system makes it a powerful tool for businesses looking to stay ahead in an environment where flexibility and the ability to quickly pivot in response to external factors are integral to success.

**5. Common Discussions and Conclusion of the Study.** The study on DRL-SCOM brings to light a series of compelling discussions on its advantages in the current supply chain management landscape. At the forefront, DRL-SCOM's integration of advanced DRL techniques, particularly the fusion of REDQ with TRPO, marks a significant innovation in tackling the complexities and dynamic challenges prevalent in modern supply chains. This novel approach addresses critical issues such as the overestimation bias inherent in traditional Q-learning methods, offering more accurate value estimation and policy improvements. Such precision in decision-making is vital in navigating the intricate and high-dimensional decision spaces typical of supply chains. A pivotal advantage of DRL-SCOM lies in its adaptability and responsiveness to the fluctuating demands and operational challenges of today's supply chains. The framework's capacity to efficiently optimize various aspects of supply chain operations, including demand forecasting, inventory management, and logistics, in real-time, is a testament to its robustness and effectiveness. The ensembled learning approach of REDQ, combined with the safe policy updates ensured by TRPO, makes DRL-SCOM particularly resilient in maintaining stable operations under unpredictable market conditions. Furthermore, DRL-SCOM's data-driven approach aligns seamlessly with the contemporary trend towards digitization and automation in supply chain management. By leveraging the vast amounts of data generated within supply chain processes, DRL-SCOM enables a more intelligent, informed, and data-centric approach to decision-making. This capability is crucial in today's fast-paced business world, where data-driven insights are key to sustaining operational excellence and competitive advantage. Overall, DRL-SCOM emerges as a powerful, adaptive, and efficient solution for modern supply chain management. Its innovative use of advanced machine learning techniques represents a significant step forward in the field, offering the potential to transform traditional supply chain models into more agile, responsive, and intelligent systems. The discussions around DRL-SCOM underscore its potential to revolutionize supply chain management, making it an invaluable tool for businesses looking to navigate the complexities of the global market effectively.

In conclusion this study on the DRL-SCOM system marks a significant milestone in the evolution of supply chain management. The comprehensive evaluation of DRL-SCOM through critical metrics such as Average Reward, Standard Deviation, and Adaptability lays bare its exceptional prowess in refining the processes involved in supply chain operations. Notably, the system's achievement of a higher average reward when pitted against traditional base-stock policies is a testament to its enhanced efficiency and effectiveness. This aspect of DRL-SCOM points towards its potential in driving improved profitability and operational success, making it a valuable asset in the realm of supply chain management. Equally important is the system's lower standard

deviation, underscoring the reliability and consistency of DRL-SCOM. These attributes are indispensable in ensuring stable and predictable management of supply chains, crucial for businesses seeking to mitigate risks and uncertainties. Furthermore, the standout feature of DRL-SCOM is its superior adaptability score, which underscores its capability to adeptly navigate the complex and dynamic nature of modern supply chains. This adaptability is key in fostering resilience and maintaining responsiveness to market changes and operational hurdles, ensuring uninterrupted and effective supply chain operations. Overall, DRL-SCOM emerges not just as a tool but as a revolutionary approach in the field of supply chain management. By harnessing cutting-edge reinforcement learning techniques, it offers a solution that is both intelligent and adaptive, aptly suited for the challenges of today's fast-paced global market. DRL-SCOM's innovative approach promises a more efficient, responsive, and effective way to manage supply chains, potentially transforming how businesses approach and execute their supply chain strategies in the contemporary business landscape.

**6. Limitations and Future Scope.** The Deep Reinforcement Learning based Supply Chain Optimization and Management (DRL-SCOM) research introduces a novel approach that significantly advances the field of supply chain management. Central to this approach is the integration of advanced Deep Reinforcement Learning (DRL) techniques, particularly the combination of Randomized Ensembled Double Q-learning (REDQ) and Trust Region Policy Optimization (TRPO). This integration is poised to address the complex and dynamic challenges characteristic of contemporary supply chain management. A key strength of DRL-SCOM is its deployment of REDQ, which effectively mitigates the overestimation bias often encountered in traditional Q-learning methods. This leads to more accurate value estimation and policy improvement, essential for effective supply chain management. Additionally, the incorporation of TRPO provides the advantage of ensuring safe and stable policy updates, an essential requirement in the volatile environment of supply chain operations. The synergistic combination of REDQ and TRPO within DRL-SCOM allows for efficient navigation through the complex decision space of supply chains, enabling real-time optimization of decisions while adhering to operational constraints. This methodology is particularly adept at handling the nonlinearities and uncertainties prevalent in supply chain management, encompassing areas like demand forecasting, inventory management, and logistics. However, the application of DRL-SCOM also presents certain limitations and scopes for future research. The effectiveness of DRL-SCOM heavily relies on the quality and comprehensiveness of the input data, posing a challenge in scenarios with limited or biased data availability. Moreover, the complexity of the algorithms used may require substantial computational resources, potentially limiting its accessibility for smaller enterprises. Future advancements in DRL-SCOM could focus on enhancing data processing capabilities to handle varied and less structured data sources. Additionally, further research could aim to streamline the computational requirements, making the system more accessible and practical for a broader range of businesses. Exploring the integration of DRL-SCOM with other emerging technologies like IoT and blockchain could also offer new dimensions in supply chain management, further enhancing its adaptability and efficiency in a rapidly evolving global market.

## REFERENCES

- [1] T. ABU ZWAIDA, C. PHAM, AND Y. BEAUREGARD, *Optimization of inventory management to prevent drug shortages in the hospital supply chain*, Applied Sciences, 11 (2021), p. 2726.
- [2] J. S. BERMÚDEZ, A. DEL RIO CHANONA, AND C. TSAY, *Distributional constrained reinforcement learning for supply chain optimization*, in Computer Aided Chemical Engineering, vol. 52, Elsevier, 2023, pp. 1649–1654.
- [3] M. CHEN AND W. DU, *Dynamic relationship network and international management of enterprise supply chain by particle swarm optimization algorithm under deep learning*, Expert Systems, (2022), p. e13081.
- [4] X. CHEN, C. WANG, Z. ZHOU, AND K. ROSS, *Randomized ensembled double q-learning: Learning fast without a model*, arXiv preprint arXiv:2101.05982, (2021).
- [5] J. W. CHONG, W. KIM, AND J. HONG, *Optimization of apparel supply chain using deep reinforcement learning*, IEEE Access, 10 (2022), pp. 100367–100375.
- [6] V. DMITROCHENKO, *Allocation Decision-Making in Service Supply Chain with Deep Reinforcement Learning*, PhD thesis, Master's thesis, Eindhoven University of Technology, 2020. 10.
- [7] M. P. KANTIPUDI, N. P. KUMAR, R. ALUVALU, S. SELVARAJAN, AND K. KOTECHA, *An improved gbso-taenn-based eeg signal classification model for epileptic seizure detection*, Scientific Reports, 14 (2024), p. 843.
- [8] Z. KEGENBEKOV AND I. JACKSON, *Adaptive supply chain: Demand–supply synchronization using deep reinforcement learning*, Algorithms, 14 (2021), p. 240.

- [9] Z. H. KILIMCI, A. O. AKYUZ, M. UYSAL, S. AKYOKUS, M. O. UYSAL, B. ATAK BULBUL, M. A. EKMS, ET AL., *An improved demand forecasting model using deep learning approach and proposed decision integration strategy for supply chain*, Complexity, 2019 (2019).
- [10] N. KRISHNAMOORTHY, L. N. PRASAD, C. P. KUMAR, B. SUBEDI, H. B. ABRAHA, AND V. SATHISHKUMAR, *Rice leaf diseases prediction using deep neural networks with transfer learning*, Environmental Research, 198 (2021), p. 111275.
- [11] D. S. KURIAN, V. M. PILLAI, A. RAUT, AND J. GAUTHAM, *Deep reinforcement learning-based ordering mechanism for performance optimization in multi-echelon supply chains*, Applied Stochastic Models in Business and Industry, (2022).
- [12] S. MAKKAR, G. N. R. DEVI, AND V. K. SOLANKI, *Applications of machine learning techniques in supply chain optimization*, in ICICCT 2019–System Reliability, Quality Control, Safety, Maintenance and Management: Applications to Electrical, Electronics and Computer Science and Engineering, Springer, 2020, pp. 861–869.
- [13] N. MOHAMADI, S. T. A. NIAKI, M. TAHER, AND A. SHAVANDI, *An application of deep reinforcement learning and vendor-managed inventory in perishable supply chain management*, Engineering Applications of Artificial Intelligence, 127 (2024), p. 107403.
- [14] Z. PENG, Y. ZHANG, Y. FENG, T. ZHANG, Z. WU, AND H. SU, *Deep reinforcement learning approach for capacitated supply chain optimization under demand uncertainty*, in 2019 Chinese Automation Congress (CAC), IEEE, 2019, pp. 3512–3517.
- [15] L. REN, X. FAN, J. CUI, Z. SHEN, Y. LV, AND G. XIONG, *A multi-agent reinforcement learning method with route recorders for vehicle routing in supply chain management*, IEEE Transactions on Intelligent Transportation Systems, 23 (2022), pp. 16410–16420.
- [16] B. ROLF, I. JACKSON, M. MÜLLER, S. LANG, T. REGGELIN, AND D. IVANOV, *A review on reinforcement learning algorithms and applications in supply chain management*, International Journal of Production Research, 61 (2023), pp. 7151–7179.
- [17] J. C. SERRANO-RUIZ, J. MULA, AND R. POLER, *Smart master production schedule for the supply chain: a conceptual framework*, Computers, 10 (2021), p. 156.
- [18] F. STRANIERI, E. FADDA, AND F. STELLA, *Combining deep reinforcement learning and multi-stage stochastic programming to address the supply chain inventory management problem*, International Journal of Production Economics, 268 (2024), p. 109099.
- [19] F. STRANIERI AND F. STELLA, *A deep reinforcement learning approach to supply chain inventory management*, arXiv preprint arXiv:2204.09603, (2022).
- [20] J. WANG, C. L. SWARTZ, AND K. HUANG, *Deep learning-based model predictive control for real-time supply chain optimization*, Journal of Process Control, 129 (2023), p. 103049.
- [21] J. WANG, R. ZHENG, AND Z. WANG, *Supply chain optimization strategy research based on deep learning algorithm*, Mobile Information Systems, 2022 (2022).
- [22] Y. YAN, A. H. CHOW, C. P. HO, Y.-H. KUO, Q. WU, AND C. YING, *Reinforcement learning for logistics and supply chain management: Methodologies, state of the art, and future opportunities*, Transportation Research Part E: Logistics and Transportation Review, 162 (2022), p. 102712.

*Edited by:* Rajanikanth Aluvalu

*Special issue on:* Evolutionary Computing for AI-Driven Security and Privacy:  
Advancing the state-of-the-art applications

*Received:* Feb 1, 2024

*Accepted:* Mar 11, 2024