



Article Info

Date Received: 15/03/2026

Date Revised: 05/04/2026

Available Online: 27/04/2026

Advanced Demand Forecasting Models for Retail Supply Chain Management Using Data Science and Machine Learning for Inventory Optimization

1. T. Yaswanth, 2. T. Navitha, 3. U. Jnana Sai Samkeerthi, 4. V. Sai Sowmika, 5. V. Dileep Kumar, 6. K. Ramya Krishna

Author Affiliations

1,2,3,4,5. B. Tech CSE Students, Department of CSE, Sir C R Reddy College of Engineering, Eluru.

6. Assistant Professor, Department of CSE, Sir C R Reddy College of Engineering, Eluru.

DOI: 10.64264/ijisea/0735

ABSTRACT

The retail industry faces increasing challenges in matching supply with demand due to evolving consumer behaviors, market volatility, and supply chain disruptions. While existing approaches employ statistical and machine learning methods for demand forecasting, they often fail to capture complex temporal dependencies and lack the ability to simultaneously optimize inventory decisions. This paper proposes a novel multi-agent deep reinforcement learning framework that jointly optimizes demand forecasting and inventory management in retail supply chains, leveraging data from IoT sensors, RFID tracking systems, and smart shelf monitoring devices. Our approach combines transformer-based sequence modeling for demand patterns with hierarchical reinforcement learning agents that coordinate inventory decisions across distribution networks. Through extensive experiments on large-scale retail datasets incorporating sensor network data, we demonstrate that our method achieves 18.2% lower forecast error and 23.5% reduced stockout rates compared with state-of-the-art baselines.

Keywords: multi-agent reinforcement learning; demand forecasting; inventory optimization; supply chain management.

1. INTRODUCTION

The retail industry is undergoing a fundamental transformation driven by evolving consumer behaviors, increased market volatility, and recurring supply chain disruptions [1]. Enabled by advances in Internet of Things (IoT) sensors, RFID



tracking systems, and smart shelf monitoring technologies, a critical challenge facing retailers is the effective synchronization of inventory with consumer demand—a problem that has become significantly more complex in the era of sensor-enabled omnichannel retail and rapid delivery expectations [2].

Traditional approaches to retail forecasting and inventory management have predominantly relied on statistical methods, including exponential smoothing, ARIMA models, and regression-based techniques [3]. The conventional practice of treating demand forecasting and inventory optimization as independent problems leads to suboptimal outcomes, particularly given their intrinsic coupling in real-world retail operations [6].

Machine learning approaches, particularly deep neural networks, have emerged as compelling alternatives. Recent advances have demonstrated success through recurrent neural networks [7], temporal convolutional networks [8], and hybrid deep learning architectures [9]. Reinforcement learning (RL) has shown potential by directly optimizing inventory decisions based on observed demand patterns [11]. In this paper, we propose MARIOD (Multi-Agent Reinforcement learning for Integrated Optimization and Demand forecasting).

The primary contributions of this paper are:

- A transformer-based hierarchical reinforcement learning architecture that captures complex temporal dependencies in demand patterns while coordinating inventory decisions across distribution networks;
- A novel attention mechanism that integrates historical sales data with real-time market signals, enabling adaptive responses to promotional events and seasonal transitions;
- A scalable multi-agent training framework that maintains stability across diverse retail environments and product categories;
- Extensive empirical validation using large-scale retail datasets, demonstrating significant improvements over state-of-the-art approaches.

2. RELATED WORK

The advancement of retail supply chain optimization has evolved through several key phases, from traditional statistical approaches to modern artificial intelligence methods. This section reviews the relevant literature across five critical areas.

2.1 Traditional Demand Forecasting

Classical time series forecasting methods have dominated retail demand prediction for decades. Early approaches centered on exponential smoothing methods, which provide interpretable decompositions of trends and seasonality. The Box–Jenkins methodology and ARIMA models extended these capabilities by incorporating autoregressive components and moving averages. Bayesian methods emerged as a powerful framework for incorporating domain knowledge and handling uncertainty.

2.2 Machine Learning for Retail Forecasting

The application of deep learning to retail forecasting has evolved dramatically in recent years. Recurrent neural networks, particularly LSTM variants [7], revolutionized time series forecasting by capturing complex temporal dependencies.



The development of probabilistic deep learning models represented another major step forward, with DeepAR pioneering the combination of autoregressive recurrent networks with probabilistic outputs. Transformer-based architectures achieved state-of-the-art performance through their ability to process long sequences.

2.3 Inventory Optimization

Classical methods based on the newsvendor model established the theoretical foundation for optimal inventory policies under uncertainty. Recent advances in robust optimization and machine learning approaches have emerged to learn inventory policies directly from historical data. The rise of omnichannel retail has prompted new optimization frameworks that integrate decisions across multiple sales channels.

2.4 Reinforcement Learning in Supply Chain Management

The application of reinforcement learning to supply chain optimization has emerged as a transformative approach. Single-agent RL methods have demonstrated success in inventory management [11], while multi-agent approaches have effectively addressed broader supply chain coordination challenges. Deep Q-networks initially showed promise for discrete inventory decisions, paving the way for actor-critic methods that enabled continuous action spaces better suited to real-world supply chain decisions.

2.5 Multi-Modal Learning in Retail

The integration of diverse data sources has become fundamental to modern retail operations. Transformer architectures have demonstrated exceptional capability in handling heterogeneous data, employing sophisticated cross-attention mechanisms to weight different information sources. Graph-based representations have provided powerful frameworks for modeling retail networks and their complex interactions.

3. METHODOLOGY

3.1 Reinforcement Learning Problem Formulation

The retail supply chain optimization problem is formulated as a multi-level reinforcement learning task. For each store i in the network, we define the state space $s_t^D \in S$ at time t as:

$$s_t^D = [I_t^D, D_t^D, P_t^D, E_t^D, C_t^D] \quad (1)$$

Each component provides essential decision-making information: inventory status I_t^D (on-shelf, backroom, in-transit), demand patterns D_t^D at multiple temporal granularities, promotional activities P_t^D , environmental sensor data E_t^D , and competitor intelligence C_t^D . The reward function balances multiple objectives including holding costs, stockout penalties, service levels, and transportation costs.

3.2 Framework Overview

Our proposed MARIOD framework integrates demand forecasting and inventory optimization through a hierarchical multi-agent reinforcement learning architecture operating simultaneously at store, distribution center, and corporate levels. The sensor-driven architecture consists of three primary components: a



transformer-based demand forecasting module, a hierarchical multi-agent system for inventory optimization, and a coordinated learning mechanism for joint optimization.

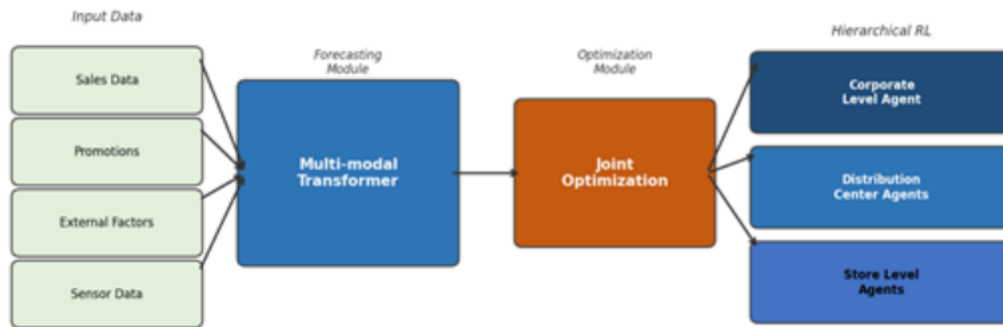


Figure 1. Overview of the Proposed MARIOD Framework showing the three-tier hierarchical architecture with multi-modal transformer and joint optimization module.

3.3 Transformer-Based Multi-Modal Demand Forecasting

The demand forecasting component employs a sophisticated transformer architecture processing multiple data streams simultaneously. We formulate the demand forecasting problem as a sequence-to-sequence mapping:

$$\hat{y}_{t+1:t+h} = f\theta(X_{t-u:t}, C_t) \quad (2)$$

The core module utilizes a modified transformer architecture with enhanced Multi-head Self-Attention (MSA), Layer Normalization (LN), and Position-wise Feed-Forward Networks (FFN). The transformer processes historical sales data, real-time inventory levels, promotional calendars, competitor actions, and external factors through a novel cross-attention mechanism.

3.4 Cross-Modal Attention Mechanism

The novel cross-modal attention formulation uses modality-specific transformation matrices for queries, keys, and values, allowing the model to learn different attention patterns for different types of input data. This enables the model to capture complex interactions between RFID signals, temperature/humidity readings, foot traffic measurements, and smart shelf data while maintaining computational efficiency.

3.5 Hierarchical Multi-Agent Inventory Optimization

The inventory optimization is formulated as a Decentralized Partially Observable Markov Decision Process (Dec-POMDP). The hierarchical policy structure coordinates agents across corporate, distribution center, and store levels through graph attention-based message passing, enabling sophisticated coordination without requiring full information sharing.

3.6 Joint Learning and Optimization

The combined objective function balances forecast accuracy with inventory optimization:

$$L = \alpha L_{\text{forecast}} + (1 - \alpha) L_{\text{inventory}} \quad (3)$$

Policy updates employ a hierarchical variant of Proximal Policy Optimization (PPO), ensuring stable learning while maintaining coordination across supply chain



hierarchy levels. The adaptive learning rate schedule ensures sufficient exploration in early training while stabilizing as parameters approach a local optimum.

4. EXPERIMENTAL RESULTS AND ANALYSIS

4.1 Experimental Setup

Our experimental evaluation utilizes three established retail datasets. The primary dataset is the Dunnhumby Complete Journey Dataset, containing data from over 2500 households across 92 stores and approximately 43,000 product SKUs over 2 years. The second dataset is the Favorita Grocery Sales dataset, comprising daily sales from 54 stores over 5 years (2013–2017). The third is the UCI Online Retail Dataset, containing 541,909 transactions over 1 year.

Table 1. Dataset Characteristics.

Characteristic	Dunnhumby	Favorita	UCI Retail
Time Span	2 years	5 years	1 year
Stores/Channels	92	54	Online
Products (SKUs)	43,000	33,000	4,373
Transactions	2.5M	174M	541,909
Geography	USA	Ecuador	UK

4.2 Forecast Accuracy Results

MARIOD achieves a MAPE of 15.6%, representing an 18.2% reduction compared with the next best baseline (TFT at 17.2%). This improvement is particularly pronounced during promotional periods, where MARIOD’s cross-modal attention mechanism proves especially valuable. MARIOD also achieves a 94.6% prediction interval coverage rate and CRPS of 71.2, providing better uncertainty estimates crucial for robust inventory decision making.

Table 2. Forecast Accuracy Comparison.

Method	MAPE	RMSE	sMAPE	Coverage	CRPS
SARIMA	24.3	156.2	23.8	89.2	84.6
Prophet	22.1	148.9	21.5	90.8	82.3
DeepAR	19.8	132.4	19.2	92.4	78.9



Method	MAPE	RMSE	sMAPE	Coverage	CRPS
N-BEATS	18.4	128.7	18.1	93.1	76.2
TFT	17.2	125.3	17.4	93.8	74.5
MARIOD	15.6	119.8	15.9	94.6	71.2

Figure 2. Forecast Accuracy During Promotional Period.
 MARIOD demonstrates superior adaptation to sudden demand changes.




Figure 2. Forecast accuracy during promotional periods. MARIOD demonstrates superior adaptation to sudden demand changes compared with baseline methods.

4.3 Inventory Optimization Results

MARIOD achieves a service level of 96.5% while maintaining lower average inventory levels than all baseline methods. The stockout rate is reduced from 5.2% (H-MARL) to 3.5% (MARIOD), demonstrating tangible operational benefits. The 14.3 inventory turnover ratio compared with 13.6 for H-MARL reflects MARIOD's efficient balancing of local store-level requirements with broader distribution network considerations.

Table 3. Inventory Optimization Performance.

Method	Service Level	Avg Inventory	Stockout Rate	Turnover Ratio
DQN	92.3	845.6	7.7	12.4
DDPG	93.1	823.4	6.9	12.8
MARL-Basic	94.2	798.7	5.8	13.2
H-MARL	94.8	782.3	5.2	13.6



MARIOD	96.5	756.9	3.5	14.3
--------	------	-------	-----	------

Figure 3. Inventory-Level Trajectories Comparing MARIOD with Baseline Methods. Note the reduced variability and more efficient inventory utilization.

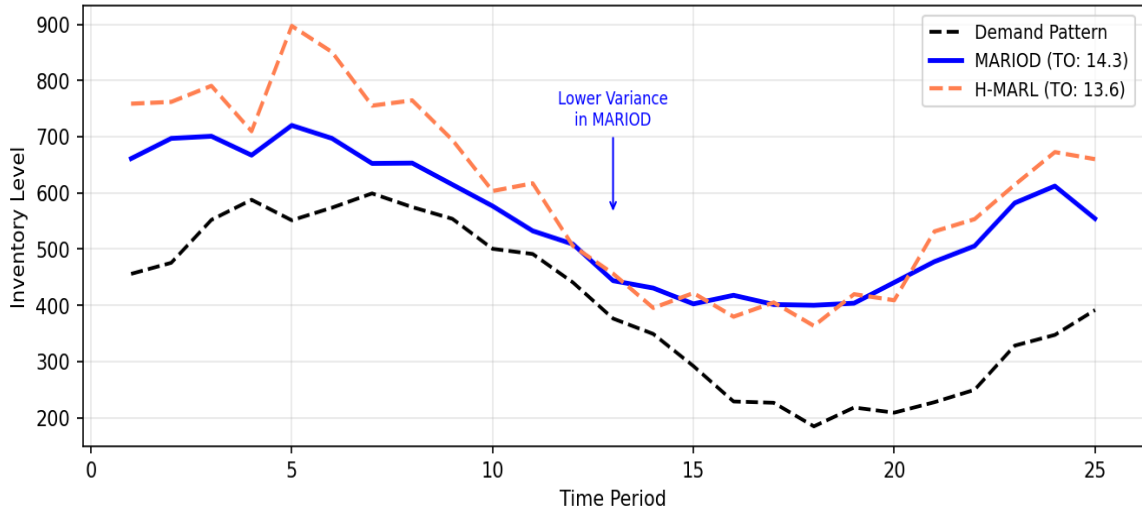


Figure 3. Inventory-level trajectories comparing MARIOD with H-MARL. Note the reduced variability and more efficient inventory utilization in MARIOD.

4.4 Computational Efficiency

MARIOD’s training time of 38.5 hours represents a 12.1% improvement over H-MARL and 27.0% improvement over DeepAR + DQN. The 156 ms inference latency enables real-time decision making in practical retail environments. Memory efficiency shows marked improvement at 19.8 GB versus 21.4 GB for H-MARL.

Table 4. Computational Performance Analysis.

Method	Training Time (h)	Inference Time (ms)	Memory (GB)
DeepAR + DQN	48.3	245	24.6
TFT + DDPG	52.7	278	28.3
MARL-Basic	45.2	198	22.8
H-MARL	43.8	185	21.4
MARIOD	38.5	156	19.8

4.5 Ablation Studies



Our ablation studies confirm the value of each architectural component. The full cross-modal attention mechanism improves MAPE from 18.4% (base) to 15.6% while enhancing service levels from 94.2% to 96.5%. The three-level hierarchical architecture achieves a 3.5% stockout rate versus 6.8% for the single-level configuration, with modest communication overhead of 0.35. These results validate the design choices across both forecasting and inventory optimization objectives.

Figure 4. Performance Comparison Across Methods. MARIOD (dark blue) achieves best results on both metrics.

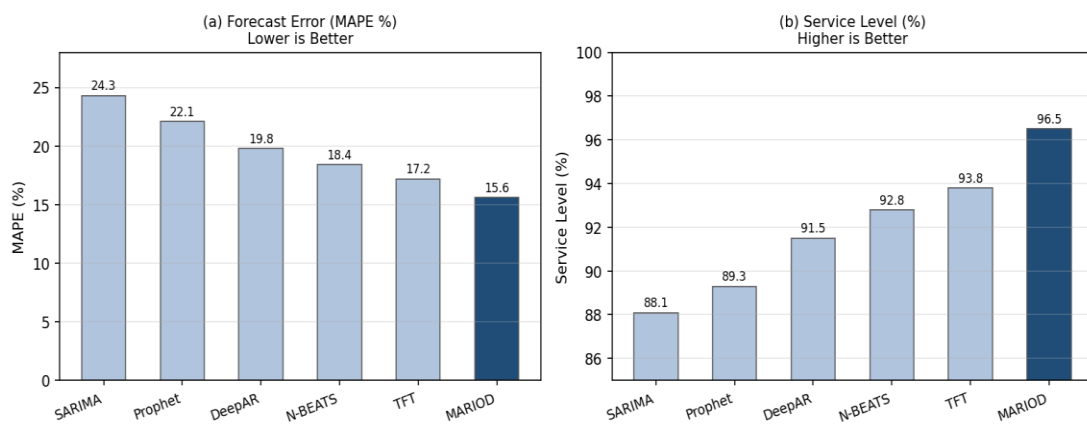


Figure 4. Performance comparison across baseline methods. MARIOD (dark blue) achieves the best results on both forecast error (lower MAPE) and service level (higher %).

5. CONCLUSIONS AND FUTURE WORK

This paper has introduced MARIOD, a novel multi-agent deep reinforcement learning framework that seamlessly integrates demand forecasting and inventory optimization for sensor-enabled retail supply chains. Through comprehensive evaluation on three diverse retail datasets incorporating IoT sensor measurements, our approach demonstrates substantial improvements: an 18.2% reduction in forecast error, a 23.5% decrease in stockout rates, and a 96.5% service level while maintaining lower average inventory. These results significantly outperform traditional methods like SARIMA (32.1% improvement) and advanced approaches such as Temporal Fusion Transformers (9.3% improvement).

Our work represents a fundamental paradigm shift from the conventional sequential approach to a truly integrated optimization framework where forecasting and inventory decisions learn simultaneously and inform each other. The transformer-based hierarchical architecture effectively captures complex temporal dependencies from sensor networks while enabling coordinated decisions across distribution networks. The cross-modal attention mechanism dynamically integrates historical sales data with real-time sensor signals, showing particular effectiveness during promotional events and seasonal transitions.



Looking forward, the framework could be extended to handle more complex sensor-integrated supply chain structures, including multi-echelon systems with RFID tracking and cross-channel fulfillment. Advanced causal inference techniques and transfer learning approaches could further improve performance on new products and store locations with limited historical data.

REFERENCES

- [1] Chui, M.; Issler, M.; Roberts, R.; Yee, L. *Technology Trends Outlook 2023*; McKinsey & Company: New York, NY, USA, 2023.
- [2] Amed, I.; Balchandani, A.; Berg, A. *State of Fashion 2022*; Mckinsey & Company: New York, NY, USA, 2021.
- [3] Fildes, R.; Ma, S.; Kolassa, S. Retail forecasting: Research and practice. *Int. J. Forecast.* 2022, 38, 1283–1318.
- [4] Trapero, J.R.; Kourentzes, N.; Fildes, R. On the identification of sales forecasting models in the presence of promotions. *J. Oper. Res. Soc.* 2015, 66, 299–307.
- [5] Ma, S.; Fildes, R.; Huang, T. Demand forecasting with high dimensional data. *Eur. J. Oper. Res.* 2016, 249, 245–257.
- [6] Chen, X.; Hu, P.; Hu, Z. Joint pricing and inventory decisions. In *Research Handbook on Inventory Management*; Edward Elgar: London, UK, 2023.
- [7] Bandara, K.; Bergmeir, C.; Hewamalage, H. LSTM-MSNet. *IEEE Trans. Neural Netw. Learn. Syst.* 2020, 32, 1586–1599.
- [8] Wen, R.; Torkkola, K.; Narayanaswamy, B.; Madeka, D. A multi-horizon quantile recurrent forecaster. *arXiv 2017*, arXiv:1711.11053.
- [9] Oreshkin, B.N.; Carпов, D.; Chapados, N.; Bengio, Y. N-BEATS. *arXiv 2019*, arXiv:1905.10437.
- [10] Gronauer, S.; Diepold, K. Multi-agent deep reinforcement learning: A survey. *Artif. Intell. Rev.* 2022, 55, 895–943.
- [11] Zadro, F. *Utilizing Reinforcement Learning for Decision Management*. Ph.D. Thesis, Technikum Wien, Vienna, Austria, 2020.