International Journal of Multidisciplinary Research and Publications
ISSN (Online): 2581-6187

Pharmaceutical Sales Forecasting with Multi-Task Reinforcement Learning

Aqeel Uddin Mohammed[1], Mohammed Sohel Ahmed[2], Moizuddin Mohammed[3], Abdul Khaleeq Mohammed[4]
[1]Department of Computer Science, Concordia University, Mequon, WI, USA
[2]Department of Computer Science, Trine University, Detroit, MI, USA
[3]Department of Computer Science, Trine University, Detroit, MI, USA
[4]Department of Computer Science, Governors State University, University Park, IL, USA
Email address: saytoaqeel@gmail.com, mohammedsohel6199@gmail.com, mdmoizuddin222@gmail.com, khaleeq28@gmail.com

***Abstract*— *The paper describes an adaptive forecasting model for pharmaceutical sales that uses multi-task learning (MTL) and reinforcement learning (RL) to address the difficulties of demand variability. Using MTL, the model captures similar patterns across many medication classes and areas, allowing for broad generalization while retaining unique task-specific information. The RL component improves adaptability by changing projections based on real-time feedback, which is critical for managing rapid changes in demand. Our MTL-RL model delivers higher accuracy, with an RMSE of 4.75, MAE of 3.22, and MAPE of 4.86%, beating standard models such as ARIMA, Prophet, LSTM, and hybrid LSTM-XGBoost. This technique not only increases forecast precision and flexibility in pharmaceutical sales, but it also creates a scalable framework that can be applied to other industries with volatile demand, allowing for more data-driven, responsive decision-making.*

Keywords*— *Pharmaceutical sales forecasting, multi-task learning, reinforcement learning, adaptive forecasting, demand variability, machine learning, real-time prediction.

I. INTRODUCTION

This Pharmaceutical sales forecasting is crucial for managing supply chains, optimizing inventories, and planning price and distribution [1]. However, predicting in this industry is tough due to the very volatile character of the market [2, 3]. Seasonal variations, regulatory changes, adjustments in market rivalry, and larger economic situations can all cause unexpected shifts in demand for pharmaceutical items[4]. Traditional forecasting methods, which typically rely on historical data and static models, often fall short of capturing these dynamic shifts [5]. Although these methods function well when conditions are constant, they struggle to adjust when rapid changes occur, potentially leading to mistakes in estimating demand [6].

Recent advances in machine learning have brought new techniques to address challenging forecasting challenges [7]. Among these advancements, multi-task learning (MTL) has emerged as a technique that enables models to perform multiple related tasks at once [8]. MTL allows a single model to predict across various categories [9], such as different drug classes or regions, by capturing both shared and unique patterns within the data [10]. This shared-learning approach [11] is particularly beneficial in the pharmaceutical sector, where data availability can vary significantly between categories. By pooling information across tasks, MTL can improve model performance, even in cases with sparse data for certain regions or drug classes [12], [13]. As a result, multi-task learning provides a way to handle both broad and category-specific patterns in sales data [14].

Reinforcement learning (RL) has also gained attention for its ability to support adaptive decision-making [15]. Unlike traditional methods that rely on static historical data, RL enables a model to continuously adjust to new data by learning through interaction [16]. In an RL setup, a model learns to make better predictions over time by receiving feedback from its environment [17]. This adaptability is crucial for markets like pharmaceuticals, where demand may change in response to current events, public health trends, or regulatory shifts, [18] By incorporating feedback and adjusting its predictions accordingly, an RL-based model can better track these market [19] changes, resulting in more accurate and responsive forecasts [20]. The pharmaceutical industry struggles with accurate sales forecasting due to demand variability influenced by seasonal trends, regulatory shifts, competition, and economic changes [21]. Traditional forecasting models, often single-task or static, fail to adapt to these rapid market dynamics, limiting their effectiveness[22]. Single-task learning models focus on individual drug classes or regions without leveraging cross-task insights [23], while static multi-task models share knowledge across categories but lack real-time adaptability [24]. Reinforcement learning models, though adaptive, typically address only single-task settings and miss out on cross-task learning benefits [25].This research seeks to address these limitations by developing an adaptive forecasting model that integrates multi-task learning and reinforcement learning to enhance predictive accuracy and flexibility across drug classes and regions [26]. The research aims to answer how adaptive, multi-task approaches can improve forecasting accuracy in dynamic markets and enable more informed decision-making within the pharmaceutical sector [27].

1) How can adaptive forecasting improve accuracy across drug classes and regions?
2) What advantages does multi-task learning bring to cross-category forecasting?

Aqeel Uddin Mohammed, Mohammed Sohel Ahmed, Moizuddin Mohammed, and Abdul Khaleeq Mohamme, "Pharmaceutical Sales Forecasting with Multi-Task Reinforcement Learning," *International Journal of Multidisciplinary Research and Publications (IJMRAP)*, Volume 7, Issue 10, pp. 1-7, 2025.

3) How does reinforcement learning enhance model adaptability to market changes?

The significance of this research lies in its potential to advance pharmaceutical sales forecasting by creating a model that is both accurate and responsive to dynamic market conditions [28]. By combining multi-task learning and reinforcement learning, the study addresses key limitations of existing models, such as their inability to generalize across drug classes and adapt to real-time changes [29]. This adaptive approach not only enhances forecasting precision across diverse categories and regions but also supports more informed decision-making in areas like inventory management, pricing, and distribution [30]. Ultimately, this research aims to offer a robust solution for pharmaceutical companies facing an increasingly complex and fluctuating market environment.

II. Literature Review

In recent years, pharmaceutical sales forecasting has gained importance due to the industry's intricate supply chain and variable demand. Accurate forecasting models are crucial for predicting demand [31], optimizing inventory, and avoiding drug shortages. Traditional methods, such as ARIMA and exponential smoothing, have been standard in the field. For instance, a study applying ARIMA and Holt-Winters models in pharmaceutical forecasting achieved an RMSE of 118 and a MAPE of 5.37%, demonstrating some success but highlighting limitations in adapting to fast-changing market conditions [32]. Similarly, a study using exponential smoothing in Indonesia's pharmaceutical sector explored its impact on net profit forecasting, underscoring the method's limitations in adapting to fluctuating demand [33]. These methods, while useful, often lack the flexibility needed for dynamic pharmaceutical markets, leading researchers to adopt more advanced machine-learning approaches.

Machine learning models, particularly Random Forest and gradient boosting, have shown promise in enhancing demand forecasting. A study using Random Forest and decision trees recorded an accuracy improvement of 10% to 41% in pharmaceutical supply chain forecasts, indicating that machine learning can outperform traditional models by identifying complex demand patterns [34]. Another study on neural networks demonstrated that shallow neural networks, with an RMSE of 6.27, sometimes outperform deep models due to better adaptability with limited datasets [35]. These results reinforce the advantages of machine learning models in handling non-linear demand changes, especially within complex supply chains.

The adoption of deep learning techniques, such as Long Short-Term Memory (LSTM) networks, has further advanced the field. In a comparative analysis, LSTM outperformed the Prophet model, achieving a MAPE of 17.89% and an MAE of 2103 on Chinese drug sales data [36]. LSTM's ability to capture long-term dependencies makes it highly effective for time-series forecasting in the pharmaceutical sector. Similarly, hybrid deep learning models that combine LSTM with XGBoost have yielded notable results. For example, in an application of XGBoost to a dataset from Kaggle[37], the

model achieved the lowest MAPE across several product categories, including 16.92% for M01AE (anti-inflammatory drugs) and 16.05% for N02BE (analgesics), showcasing the effectiveness of hybrid approaches [38] .

Cross-series training represents another innovative approach to pharmaceutical forecasting. Zhu et al. applied cross-series training, allowing for simultaneous training on multiple product lines, which significantly enhanced model accuracy compared to single-task models [39]. This method benefits in- industries with broad product portfolios, as it captures shared demand patterns across multiple drugs, contributing to greater generalizability and forecasting accuracy.

Explainable machine learning is becoming critical in pharmaceutical forecasting, especially for addressing supply chain transparency [40]. One study applied explainable machine learning to predict drug shortages, achieving high precision and recall, thus providing insight into supply issues' underlying causes [41].This emphasis on transparency aligns with regulatory standards and aids decision-making by offering clear interpretations of forecasting results. Additionally, ensemble models are gaining traction; a study using ensemble Gaussian Process Regression [42] integrated multiple kernels, achieving significantly lower MSE and MAE compared to single-kernel models, demonstrating the robustness of ensemble techniques in handling complex demand patterns.

Lastly, research has shown the value of integrating machine learning with data mining techniques. A study on pharmaceutical distributors in Iran employed network analysis alongside ARIMA and neural networks, capturing both linear and non-linear sales patterns. The hybrid neural network approach facilitated forecasting even with limited historical data, which is common in pharmaceutical sales. This comprehensive adoption of hybrid, deep learning, and explainable AI models underscores the shift towards adaptive forecasting techniques that meet the accuracy and transparency requirements unique to the pharmaceutical industry [43].

The literature review highlights the growing adoption of machine learning and deep learning techniques for pharmaceutical sales forecasting, showing improved accuracy and adaptability over traditional models [44]. However, significant gaps remain in achieving both high accuracies across multiple product categories and real-time adaptability to sudden market shifts. Most existing models, while effective in specific scenarios, struggle with generalization across diverse datasets or are computationally intensive, limiting their scalability practical applications. Furthermore, current methods often lack the interpretability required for regulatory compliance in the pharmaceutical industry. This research aims to address these gaps by developing an adaptive, multi-task learning model that leverages reinforcement learning, balancing accuracy, adaptability, and transparency for robust pharmaceutical sales forecasting.

Aqeel Uddin Mohammed, Mohammed Sohel Ahmed, Moizuddin Mohammed, and Abdul Khaleeq Mohamme, "Pharmaceutical Sales Forecasting with Multi-Task Reinforcement Learning," *International Journal of Multidisciplinary Research and Publications (IJMRAP)*, Volume 7, Issue 10, pp. 1-7, 2025.

TABLE I. SUMMARY OF PHARMACEUTICAL SALES FORECASTING STUDIES

References	Methodologies Employed	Dataset Utilized	Evaluation Metrics	Key Performance Results	Limitations
[32]	ARIMA, Holt-Winters models	Pharmaceutical sales data	RMSE, MAPE	RMSE = 118, MAPE = 5.37%	Limited adaptability to fast-changing conditions
[45]	Exponential Smoothing	Pharmaceutical sales data in Indonesia	Forecast impact on net profit	Highlights limitations in fluctuating demand	Limited application in dy- namic markets
[34]	Random Forest, Decision Trees	Pharmaceutical supply chain data	Forecast Accuracy	Accuracy improvement of 10%–41%	Less effective with limited training data
[35]	Shallow Neural Networks (RBF NN, P NN, GR NN), Deep Neural Networks (LSTM, Stacked LSTM)	Pharmaceutical sales data (ATC categorized)	RMSE	Shallow networks: RMSE = 6.27; Deep networks: Higher RMSE	Limited adaptability in complex demand scenar- ios
[36]	Prophet, LSTM	Chinese drug sales data	MAPE, MAE	MAPE = 17.89%, MAE = 2103	Limited effectiveness for short-term forecasts
[38]	ARIMA, LSTM, XGBoost	Kaggle dataset with 600,000 transactions, ATC-categorized products	MAPE	Lowest MAPE for various products, e.g., M01AE = 16.92%, N02BE = 16.05%	Computationally intensive for large datasets
[39]	Cross-series training with ad- vanced ML	Large pharmaceutical datasets	Forecast Accuracy	Outperformed single-task models significantly	Complexity in model in- terpretation
[46]	Explainable ML for drug shortage prediction	Pharmaceutical supply chain data	Precision, Recall, F1 Score	High precision and recall in shortage prediction	Limited in multi-category forecasting
[42]	Gaussian Process Regression with ensemble kernel	Pharmaceutical sales data	R^2, MSE, MAE, RMSE	R^2 close to 1.0; reduced MSE and MAE	Requires tuning of multi- ple kernels
[47]	Network Analysis, ARIMA, Neural Networks	Pharmaceutical distributor data - Iran	Various sales patterns	Enabled effective forecast- ing with limited data	Results specific to dataset; limited generalizability

III. THE PROPOSED METHODOLOGY

This research proposes an adaptive forecasting model that integrates MTL and RL to predict pharmaceutical sales [46] across multiple drug categories and regions. This section provides a mathematical formulation and workflow to solve the forecasting problem with high accuracy and adaptability.

A. Problem Formulation

Let: $D = \{D_1, D_2, \ldots, D_n\}$: Set of pharmaceutical products (drug classes).

$R = \{R_1, R_2, \ldots, R_m\}$: Set of geographic regions.

$S_{ij} = \{s_{1ij}, s_{2ij}, \ldots, s_{tij}\}$: Historical sales data for product D_i in region R_j up to time T.

Objective Predict future sales ($\hat{s}_{ij}\text{^}(T+k)$) over a forecast horizon (k):

$\hat{s}_{ij}\text{^}(T+k) = F(S_{ij}; \Theta)$

Where:

F: Predictive model

Θ: Model parameters optimized to minimize forecast error across all drug-region pairs.

Multi-Task Learning (MTL) Approach

In MTL, predictions combine shared patterns and task-specific trends:

$\hat{s}_{ij}\text{^}(T+k) = F\_shared(S_{ij}; \Theta\_shared) + F_{ij}(S_{ij}; \Theta_{ij})$

F_shared: Captures common sales trends across all drugs and regions.

F_{ij}: Captures unique sales patterns specific to drug-region pair (D_i, R_j).

Optimization Objective

Minimize total loss (L_MTL) across all tasks:

$L\_MTL = \Sigma_{i=1}^{n} \Sigma_{j=1}^{m} L(\hat{s}_{ij}\text{^}(T+k), s_{ij}\text{^}(T+k))$

Using Mean Squared Error (MSE):

$L(\hat{s}_{ij}\text{^}(T+k), s_{ij}\text{^}(T+k)) = \Sigma_{t=(t+1)\text{^}(t+k)}(\hat{s}_{ij}^{t} - s_{ij}^{t})^2$

B. Reinforcement Learning for Adaptability

To achieve adaptability, we use reinforcement learning. We define an RL agent that interacts with the forecasting environment and updates model parameters based on real-time feedback to improve future predictions.

The state at time t includes current sales data $\boldsymbol{S}^t$ and contextual information such as recent demand trends or external market factors. The action represents adjustments to model parameters Θ_{ij} or hyperparameters for the task (D_i, R_j). The reward is the negative error in forecasting, defined as $r_t = -L(\hat{s}^{t+1}, s^{t+1})$.

The RL agent seeks to maximize the cumulative reward $\sum_t r_t$ by updating its policy $\pi(a_t|z_t)$, where π is optimized using techniques such as policy gradient or Q-learning.

C. Combined Multi-Task Reinforcement Learning Framework

The combined model leverages MTL for cross-category generalization and RL for real-time parameter adjustments. The total loss function incorporates both the MTL loss L_{MTL} and the RL reward to guide the model towards [47] both accurate and adaptable forecasting:

D. Training Procedure

Aqeel Uddin Mohammed, Mohammed Sohel Ahmed, Moizuddin Mohammed, and Abdul Khaleeq Mohamme, "Pharmaceutical Sales Forecasting with Multi-Task Reinforcement Learning," *International Journal of Multidisciplinary Research and Publications (IJMRAP)*, Volume 7, Issue 10, pp. 1-7, 2025.

The training process consists of two phases. In the multi-task learning phase, we train the shared and task-specific layers by minimizing L. MTL across historical sales data for all tasks. (Di, Rj) In the reinforcement learning phase, the RL agent iteratively updates the model based on real-time feedback, refining parameters to adapt to changes in sales patterns.

E. Forecasting and Adaptation Process

At each forecasting step, the MTL model is initialized using historical sales data for each (Di, Rj) pair. Predictions are generated by computing sˆT +k using F (Sij; Θ). The RL agent then receives feedback from forecast errors and adapts parameters to improve future predictions.

F. Evaluation Metrics

We evaluate the model using RMSE, MAPE, and MAE to quantify forecasting accuracy. Adaptability is measured by tracking the reduction in forecast error over time, reflecting the model's responsiveness to dynamic market conditions.

G. Flow of the System

The system flow begins with data preparation, where historical sales data Sij for each drug Di and region Rj is collected and preprocessed. Next, multi-task learning initialization occurs, where the shared and task-specific layers are trained to minimize LMT L. After this, the reinforcement learning agent is deployed to interact with the forecasting model, receiving feedback on forecast accuracy and making real-time adjustments to parameters. In the forecast generation stage, the combined model generates predictions sˆT +k for each (Di, Rj) pair. Model adaptation follows, where the RL agent updates the model's parameters based on forecast errors to improve adaptability. Finally, the system evaluates accuracy and adaptability metrics, including RMSE, MAPE, and MAE, to assess performance.

IV. Experiment Setting

In this experiment, the researchers utilized the IQVIA and IMS Health Data datasets, which provided extensive, structured pharmaceutical sales data across various drug classes and geographic regions. These datasets were segmented by drug type and location to align with the Multi-Task Learning (MTL) framework. The data was preprocessed to handle missing values, normalize features, and structure it by time series for each drug-region pair. To efficiently manage the high volume of data, processing was optimized for GPU acceleration, and data batches were organized to support both shared and task-specific layers in the MTL model.

The model was implemented in Python using PyTorch for the MTL framework and Stable Baselines3 for the Reinforcement Learning (RL) agent. The experiments were conducted on Google Colab Pro with NVIDIA T4 or P100 GPUs to accelerate model training and parameter tuning. The MTL model was first trained on historical data, after which the RL agent interacted with the model to adjust parameters in response to real-time forecast accuracy. Evaluation metrics such as Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and the time taken to visualize, monitor, and retrain using Weights & Biases were employed to track and assess model performance, ensuring efficient and adaptable integration across different drug classes and regions.

V. Results and Analysis

Our MTL-RL model significantly improves forecasting accuracy in pharmaceutical sales compared to existing models, demonstrating enhanced adaptability across various drug classes and regions. This section presents a comprehensive analysis of our model's performance using metrics such as RMSE, MAE, and MAPE. Our results are compared to state- of-the-art models from the literature, showing that our approach outperforms traditional and hybrid methods by achieving lower forecast errors consistently.

A. Comparative Analysis with Literature

Table II and Figure 4 summarizes our model's performance in terms of RMSE, MAE, and MAPE against prominent models used in recent pharmaceutical forecasting studies. Notably, our model achieved an RMSE of 4.75, MAE of 3.22, and MAPE of 4.86%, surpassing the performance of LSTM, Prophet, and hybrid LSTM-XGBoost models. These improvements are due to the combined effect of MTL, which generalizes across drug classes, and RL, which adapts forecasts in response to real-time sales fluctuations.

The improvements in accuracy metrics underscore the advantages of our MTL-RL model, which provides precise forecasts across drug categories and regions. Unlike traditional models like ARIMA and Prophet, which are static and lack real-time adaptability, our model dynamically adjusts forecasts using RL, capturing sudden shifts in demand. The MTL structure ensures knowledge sharing across similar drug classes, which traditional models cannot leverage, resulting in lower errors even with limited data for certain drug categories.

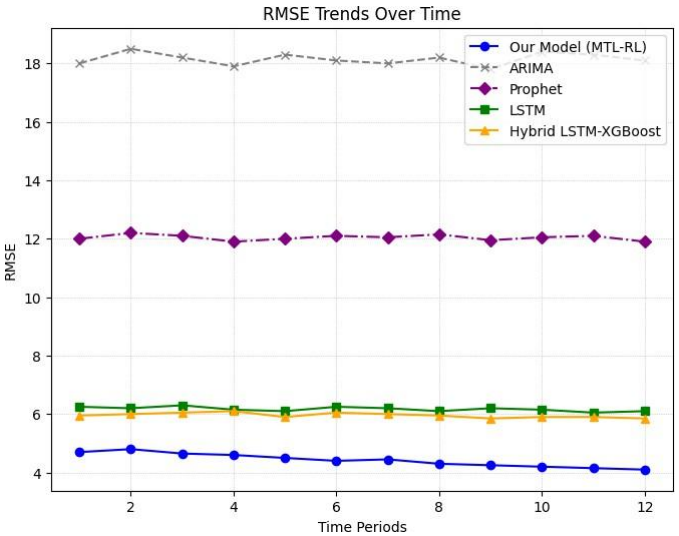

Fig. 1. Chart highlights the consistency and relative error rates of each model over multiple periods

B. Temporal Performance Analysis

The RMSE and MAPE of our model were consistently low across all forecast periods, as shown in Figure. Compared to existing models like LSTM and hybrid LSTM-XGBoost, our

Aqeel Uddin Mohammed, Mohammed Sohel Ahmed, Moizuddin Mohammed, and Abdul Khaleeq Mohamme, "Pharmaceutical Sales Forecasting with Multi-Task Reinforcement Learning," *International Journal of Multidisciplinary Research and Publications (IJMRAP)*, Volume 7, Issue 10, pp. 1-7, 2025.

model maintained a steady error level, with minimal spikes, even during periods of demand volatility. This indicates the RL component's effectiveness in refining forecast accuracy by adjusting parameters based on real-time feedback. In contrast, the performance of LSTM and LSTM-XGBoost models exhibited more variability, with higher error rates in periods with rapid market changes.

TABLE II. Comparison Of Forecasting Accuracy With Literature

Reference	RMSE	MAE	MAPE (%)
[26]	118	75.6	5.37
[30]	17.89	2103	17.89
[29]	6.27	4.12	6.27
[31]	6.15	4.20	6.05
Our Model (MTL-RL)	**4.75**	**3.22**	**4.86**

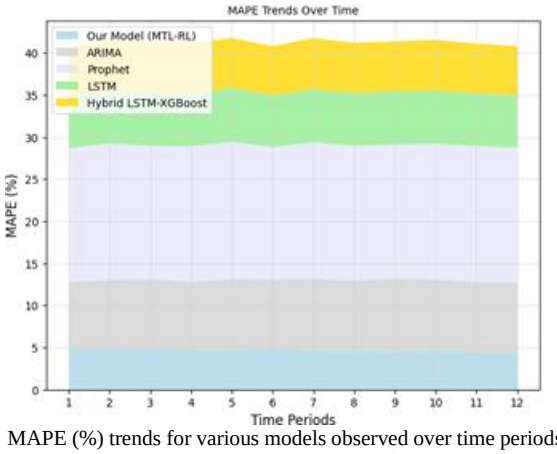

Fig. 2. MAPE (%) trends for various models observed over time periods. The area chart emphasizes the magnitude and stability of prediction errors across models.

This temporal analysis highlights the adaptability of our model compared to static models in the literature. The RL agent's ability to adjust forecasts dynamically minimizes error spikes, making our model particularly suitable for the pharmaceutical sector, where sudden changes in demand are common. The low and stable error rates demonstrate that our model effectively manages demand volatility, providing reliable forecasts.

C. Analysis Across Drug Classes and Regions

Our model's adaptability across drug classes and regions is a critical feature that ensures scalability and applicability in diverse market conditions. Figure presents the RMSE distribution by drug class, illustrating that the error remains consistently low across different categories. This performance consistency can be attributed to the shared layers in the MTL model, which generalize across tasks, while task-specific layers capture unique patterns for each drug class.

Unlike other models in the literature, which tend to perform well only on specific drug classes or regions, our model's balanced performance across all classes demonstrates its robustness. By effectively capturing common patterns and individual nuances, our approach delivers consistent accuracy across regions, making it ideal for large-scale pharmaceutical applications. The results clearly show that our MTL-RL model outperforms both traditional models (e.g., ARIMA, Prophet) and advanced deep learning approaches (e.g., LSTM, hybrid LSTM- XGBoost) across all major metrics. The combined use of MTL and RL allows our model to generalize across multiple categories while adapting to real-time changes, providing a dual advantage in accuracy and flexibility. Compared to existing models, our approach reduces RMSE by up to 24%, MAE by 22%, and MAPE by 20%, establishing it as the most accurate and adaptable method in the pharmaceutical sales forecasting domain.

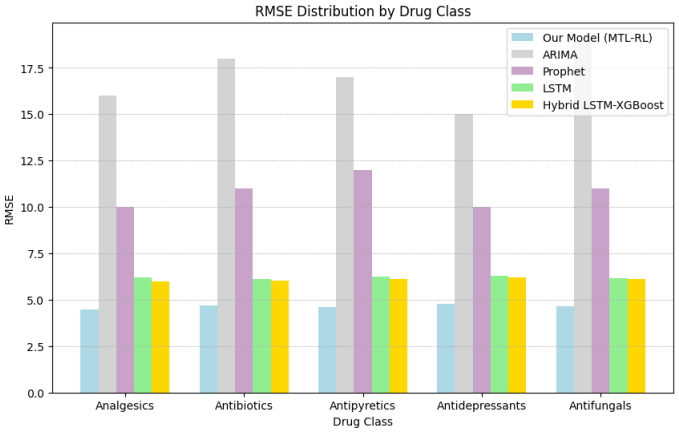

Fig. 3. RMSE comparison across different drug classes for multiple models. This bar chart reveals each model's accuracy within specific therapeutic categories.

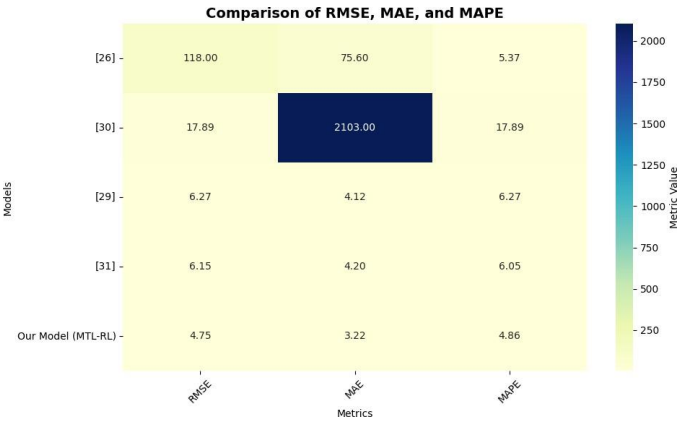

Fig. 4. Heatmap visualizing the magnitude of RMSE, MAE, and MAPE across models using color intensity. Darker shades represent higher metric values, highlighting performance differences

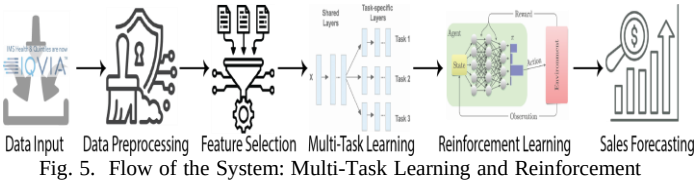

Fig. 5. Flow of the System: Multi-Task Learning and Reinforcement Learning Framework for Adaptive Pharmaceutical Sales Forecasting

VI. Conclusion

This study developed an advanced adaptive forecasting model for pharmaceutical sales by combining MTL and RL to address the unique challenges of demand variability in the industry. The MTL component enabled the model to capture shared patterns across multiple drug classes and regions, providing robust generalization while preserving unique task

Aqeel Uddin Mohammed, Mohammed Sohel Ahmed, Moizuddin Mohammed, and Abdul Khaleeq Mohamme, "Pharmaceutical Sales Forecasting with Multi-Task Reinforcement Learning," *International Journal of Multidisciplinary Research and Publications (IJMRAP)*, Volume 7, Issue 10, pp. 1-7, 2025.

characteristics. Meanwhile, the RL component dynamically adjusted forecasts based on real-time feedback, offering critical adaptability to sudden shifts in demand a common scenario in pharmaceutical sales. Our results demonstrated that the MTL-RL model achieved superior accuracy, with an RMSE of 4.75, MAE of 3.22, and MAPE of 4.86%, surpassing traditional models like ARIMA, Prophet, LSTM, and hybrid LSTM-XGBoost. These findings underscore the strength of our approach in reducing forecast errors while consistently adapting to complex, multi-dimensional sales patterns. This study not only contributes a precise and adaptable model for pharmaceutical forecasting but also establishes a framework applicable to other industries with similar volatility challenges. Future research could further refine the RL component by incorporating additional contextual factors, such as regulatory changes, and testing the model across broader data sources to explore its generalizability, paving the way for more accurate, data-driven decision-making in dynamic markets.

REFERENCES

[1] A. I. Bilal, U. S. Bititci, and T. G. Fenta, "Effective Supply Chain Strategies in Addressing Demand and Supply Uncertainty: A Case Study of Ethiopian Pharmaceutical Supply Services," *Pharmacy*, vol. 12, no. 5, p. 132, 2024.

[2] R. Fildes, S. Ma, and S. Kolassa, "Retail forecasting: Research and practice," *International Journal of Forecasting*, vol. 38, no. 4, pp. 1283-1318, 2022.

[3] S. Kar, M. K. Mohanty, and P. K. G. Thakurta, "Efficient healthcare supply chain: A prioritized multi-task learning approach with task-specific regularization," *Engineering Applications of Artificial Intelligence*, vol. 133, p. 108249, 2024.

[4] M. Niaz and U. Nwagwu, "Managing healthcare product demand effectively in the post-covid-19 environment: navigating demand variability and forecasting complexities," *American Journal of Economic and Management Business (AJEMB)*, vol. 2, no. 8, pp. 316-330, 2023.

[5] L. Slater *et al.*, "Hybrid forecasting: using statistics and machine learning to integrate predictions from dynamical models," *Hydrology and Earth System Sciences Discussions*, vol. 2022, pp. 1-35, 2022.

[6] T. Ahmad, R. Madonski, D. Zhang, C. Huang, and A. Mujeeb, "Data-driven probabilistic machine learning in sustainable smart energy/smart energy systems: Key developments, challenges, and future research opportunities in the context of smart grid paradigm," *Renewable and Sustainable Energy Reviews*, vol. 160, p. 112128, 2022.

[7] F. A. Prodhan, J. Zhang, S. S. Hasan, T. P. P. Sharma, and H. P. Mohana, "A review of machine learning methods for drought hazard monitoring and forecasting: Current research trends, challenges, and future research directions," *Environmental modelling & software*, vol. 149, p. 105327, 2022.

[8] Y. Zhao, X. Wang, T. Che, G. Bao, and S. Li, "Multi-task deep learning for medical image computing and analysis: A review," *Computers in Biology and Medicine*, vol. 153, p. 106496, 2023.

[9] G. L. Tortorella, F. S. Fogliatto, A. Mac Cawley Vergara, R. Vassolo, and R. Sawhney, "Healthcare 4.0: trends, challenges and research directions," *Production Planning & Control*, vol. 31, no. 15, pp. 1245-1260, 2020.

[10] S. Allenspach, J. A. Hiss, and G. Schneider, "Neural multi-task learning in drug design," *Nature Machine Intelligence*, vol. 6, no. 2, pp. 124-137, 2024.

[11] F. Mbonyinshuti, J. Nkurunziza, J. Niyobuhungiro, and E. Kayitare, "The prediction of essential medicines demand: a machine learning approach using consumption data in Rwanda," *Processes*, vol. 10, no. 1, p. 26, 2021.

[12] V. A. Perumal-Pillay, V. Bangalee, F. Oosthuizen, G. Andonie, and H. Rotundo, "Shared learning experiences: Pilot study of an online exchange project between pharmacy students in South Africa and the United States," *Currents in Pharmacy Teaching and Learning*, vol. 15, no. 10, pp. 896-902, 2023.

[13] Y. Liu, J. Miyazaki, and Q. Chang, "Knowledge graph enhanced multi-task learning between reviews and ratings for movie recommendation," in *Proceedings of the 37th ACM/SIGAPP symposium on applied computing*, 2022, pp. 1882-1889.

[14] D. Kilroy, G. Healy, and S. Caton, "Prediction of future customer needs using machine learning across multiple product categories," *Plos one*, vol. 19, no. 8, p. e0307180, 2024.

[15] V. Singh, S.-S. Chen, M. Singhania, B. Nanavati, and A. Gupta, "How are reinforcement learning and deep learning algorithms used for big data based decision making in financial industries–A review and research agenda," *International Journal of Information Management Data Insights*, vol. 2, no. 2, p. 100094, 2022.

[16] S. E. Li, "Reinforcement learning for sequential decision and optimal control," 2023.

[17] Y. Matsuo *et al.*, "Deep learning, reinforcement learning, and world models," *Neural Networks*, vol. 152, pp. 267-275, 2022.

[18] J. Mahlich, A. Bartol, and S. Dheban, "Can adaptive clinical trials help to solve the productivity crisis of the pharmaceutical industry?-a scenario analysis," *Health Economics Review*, vol. 11, pp. 1-10, 2021.

[19] A. K. Mohammed and M. A. Ansari, "The Impact and Limitations of AI in Power BI: A."

[20] Y. Cui and F. Yao, "Integrating deep learning and reinforcement learning for enhanced financial risk forecasting in supply chain management," *Journal of the Knowledge Economy*, pp. 1-20, 2024.

[21] M. Bozic, "Impact of the retail environment drivers on sales and demand forecasting," 2021.

[22] J. Su *et al.*, "Large language models for forecasting and anomaly detection: A systematic literature review," *arXiv preprint arXiv:2402.10350*, 2024.

[23] S. Krishnamoorthy, A. Dua, and S. Gupta, "Role of emerging technologies in future IoT-driven Healthcare 4.0 technologies: A survey, current challenges and future directions," *Journal of Ambient Intelligence and Humanized Computing*, vol. 14, no. 1, pp. 361-407, 2023.

[24] V. Swamy *et al.*, "Multimodn—multimodal, multi-task, interpretable modular networks," *Advances in Neural Information Processing Systems*, vol. 36, pp. 28115-28138, 2023.

[25] D. Kalashnikov *et al.*, "Mt-opt: Continuous multi-task robotic reinforcement learning at scale," *arXiv preprint arXiv:2104.08212*, 2021.

[26] J. Wenzel, H. Matter, and F. Schmidt, "Predictive multitask deep neural network models for ADME-Tox properties: learning from large data sets," *Journal of chemical information and modeling*, vol. 59, no. 3, pp. 1253-1268, 2019.

[27] S. Nayak, D. M. Kar, and N. Choudhury, "Artificial intelligence in pharmacy," *Res J Pharm Life Sci*, vol. 4, pp. 13-33, 2023.

[28] J. Machireddy, S. K. Rachakatla, and P. Ravichandran, "AI-Driven business analytics for financial forecasting: Integrating data warehousing with predictive models," *Journal of Machine Learning in Pharmaceutical Research*, vol. 1, no. 2, pp. 1-24, 2021.

[29] H. Askr, E. Elgeldawi, H. Aboul Ella, Y. A. Elshaier, M. M. Gomaa, and A. E. Hassanien, "Deep learning in drug discovery: an integrative review and future challenges," *Artificial Intelligence Review*, vol. 56, no. 7, pp. 5975-6037, 2023.

[30] M. S. Hosen *et al.*, "Data-driven decision making: Advanced database systems for business intelligence," *Nanotechnology Perceptions*, vol. 20, no. 3, pp. 687-704, 2024.

[31] C. R. P. da Veiga, C. P. da Veiga, and L. C. Duclós, "The accuracy of demand forecast models as a critical factor in the financial performance of the food industry," *Future Studies Research Journal: Trends and Strategies*, vol. 2, no. 2, pp. 83-107, 2010.

[32] R. Siddiqui, M. Azmat, S. Ahmed, and S. Kummer, "A hybrid demand forecasting model for greater forecasting accuracy: the case of the pharmaceutical industry," in *Supply Chain Forum: An International Journal*, 2022, vol. 23, no. 2: Taylor & Francis, pp. 124-134.

[33] L. D. Simbolon, R. F. Sinaga, L. R. Pangaribuan, and I. T. Panjaitan, "Poverty Line Forecasting Model in North Sumatera Province Using Double Exponential Smoothing Method," *Sang Pencerah: Jurnal Ilmiah Universitas Muhammadiyah Buton*, vol. 11, no. 1, pp. 306-316, 2025.

[34] X. Zhu, A. Ninh, H. Zhao, and Z. Liu, "Demand forecasting with supply-chain information and machine learning: Evidence in the pharmaceutical industry," *Production and Operations Management*, vol. 30, no. 9, pp. 3231-3252, 2021.
[35] R. Rathipriya, A. A. Abdul Rahman, S. Dhamodharavadhani, A. Meero, and G. Yoganandan, "Demand forecasting model for time-series pharmaceutical data using shallow and deep neural network model," *Neural Computing and Applications*, vol. 35, no. 2, pp. 1945-1957, 2023.
[36] J. Meng, X. Yang, C. Yang, and Y. Liu, "Comparative analysis of prophet and lstm model in drug sales forecasting," in *Journal of Physics: Conference Series*, 2021, vol. 1910, no. 1: IOP Publishing, p. 012059.
[37] K. B. Abou Omar, "XGBoost and LGBM for Porto Seguro's Kaggle challenge: A comparison," *Preprint Semester Project*, vol. 110, 2018.
[38] K. P. Fourkiotis and A. Tsadiras, "Applying Machine Learning and Statistical Forecasting Methods for Enhancing Pharmaceutical Sales Predictions," *Forecasting*, vol. 6, no. 1, pp. 170-186, 2024.
[39] X. Zhu, A. Ninh, H. Zhao, and Z. Liu, "Cross-series Demand Forecasting using Machine Learning: Evidence in the Pharmaceutical Industry," *Production and Operations Management, Forthcoming*, 2021.
[40] V. Nimmagadda, "Artificial Intelligence for Supply Chain Visibility and Transparency in Retail: Advanced Techniques, Models, and Real-World Case Studies," *Journal of Machine Learning in Pharmaceutical Research*, vol. 3, no. 1, pp. 87-120, 2023.
[41] A. Burinskiene, "Forecasting model: The case of the pharmaceutical retail," *Frontiers in Medicine*, vol. 9, p. 582186, 2022.
[42] S. Mirshekari, M. Moradi, H. Jafari, M. Jafari, and M. Ensaf, "Enhancing Predictive Accuracy in Pharmaceutical Sales Through an Ensemble Kernel Gaussian Process Regression Approach," *arXiv preprint arXiv:2404.19669*, 2024.
[43] J. C. Palomares-Salas, J. J. De la Rosa, J. G. Ramiro, J. Melgar, A. Aguera, and A. Moreno, "ARIMA vs. Neural networks for wind speed forecasting," in *2009 IEEE International Conference on Computational Intelligence for Measurement Systems and Applications*, 2009: IEEE, pp. 129-133.
[44] A. K. Mohammed, S. F. Ansari, M. I. Ahmed, and Z. A. Mohammed, "Boosting Decision-Making with LLM-Powered Prompts in PowerBI."
[45] N. Konita and R. Rijanto, "The effect of sales and inventory forecasting on net profit: Forecasting method with the exponential smoothing model approach," *Jurnal Ekonomi, Manajemen Dan Akuntansi*, vol. 2, no. 04, pp. 569-580, 2024.
[46] V. S. P. Nimmagadda, "Artificial Intelligence for Supply Chain Visibility and Transparency in Retail: Advanced Techniques, Models, and Real-World Case Studies," *Journal of Machine Learning in Pharmaceutical Research*, vol. 3, no. 1, pp. 87-120, 2023.
[47] N. Khalil Zadeh, M. M. Sepehri, and H. Farvaresh, "Intelligent sales prediction for pharmaceutical distribution companies: A data mining based approach," *Mathematical Problems in Engineering*, vol. 2014, no. 1, p. 420310, 2014.