International Journal of AI, Big Data, Computational and Management Studies
Noble Scholar Research Group | Volume 5, Issue 3, 155-160, 2024
ISSN: 3050-9416 | https://doi.org/10.63282/3050-9416.IJAIBDCMS-V5I3P116

Leveraging Predictive Analytics and Redis-Backed Caching to Optimize Specialty Medication Fulfillment and Pharmacy Inventory Management

Srikanth Reddy Gudi
Cigna Evernorth Health Services Inc., Charlotte, North Carolina, USA.

Abstract: Specialty medications have long shelf lives, high costs, and demand patterns that are often unpredictable, all of which makes specialty medication inventory management a big challenge for the pharmaceutical industry. This work explores the use of predictive analytics in combination with Redis backed caching systems, to streamline specialty medication dispensing and cost effective management of pharmacy inventory. The key aim is to conduct regression analysis for ML-based demand forecasting performance, evaluate real-time inventory systems with Redis caching performance, and evaluate gains in Order Fulfillment Rate performance measurement. This study explored quantitative research design using secondary data from pharmacy management systems and health care published databases. Specifically, the hypothesis was that integrated predictive analytics along with Redis caching has a huge impact on reducing stockouts and enhancing fulfillment efficiency. And tests showed that pharmacies that use these technologies had 23% lower inventory holding costs and 31% better order fulfillment accuracy. These predictive models achieved 87% accuracy in forecasting conclusions for specialty medications. Post discussion, implementation of Redis cache showed 78% of latency of database queries were vastly reduced, and this helped in making inventory real-time decisions. Integrating technology to improve pharmacy efficiency and clinical outcomes for specialty medications: An updated review of the literature Abstract Background The impact of technology integration into specialty pharmacy practice is documented in the literature; however, these studies are relatively scarce and often outdated.

Keywords: Predictive Analytics, Redis Caching, Pharmacy Inventory Management, Specialty Medications, Machine Learning.

1. Introduction

As specialty medications have become more central to healthcare delivery, there has been a fundamental shift in the pharmaceutical supply chain. While accounting for only 2 percent of the patient population, specialty medications comprise nearly half (50%) of all drug spending because they are typically more expensive, require more stringent storage conditions and are optimized for select therapeutic targets (Aitken & Kleinrock, 2019). Such a gulf poses unique challenges for the ability of pharmacy inventory management systems to balance patient access to needed medications while ensuring the financial health of the pharmacy. Demand for specialty medications is inherently volatile and any demand forecast provides little insight into actual demand, making traditional inventory management approaches using historical consumption patterns and safety stock calculations ineffective when managing specialty inventory. Not only does a single stockout event disrupt therapy and subsequently impact patient outcomes, overstocking causes huge economic losses through expiration and storage fees (Shah, 2004). Predictive analytics has emerged as a combination of predictive analytics leveraging machine learning algorithms to predict demand patterns more accurately than traditional methods.

Redis, an open-source, in-memory data structure store with outstanding performance characteristics, has also seen widespread adoption in healthcare information systems. Redis serves as a caching layer that allows for sub-millisecond response for frequently accessed data, making it ideal for real-time inventory management use cases (Carlson, 2013). Personalized predictive analytics meet redis-backed caching model, a new paradigm to deal with unpredictable demand and responsive systems challenges in pharmacy operations - this is real science, hot off the press. Introduction: The Indian pharmaceutical market which is valued around $42 billion in 2021 (Pillai & Kumar, 2022), has its own complexities for the inventory optimization bringing to light the diverse nature of the patients and frequent changes in the regulation (Basak & Sathyanarayana, 2010) In response to managing portfolios of specialty medications and complying with the regulatory practices, Indian pharmacies are more inclined towards adopting technology solutions. This study investigates the potential of predictive analytics and Redis caching in the optimization of specialty medication fulfillment in this setting. This research has important implications, not only for operational efficiency, but also for patient safety and health care accessibility. Note that full treatment for cancer, rheumatoid arthritis, and multiple sclerosis (Tichy et al., 2019) depends on the timely availability of some specialty medications. This will aid pharmacies in catering better patient care inseparable from financially balancing out stock distribution by creating stronger technological backbones for preserving optimal stock against demand.

2. Literature Review

Literature Review The use of predictive analytics in the healthcare supply chain management has attracted significant attention in the literature. Kilimci et al. For example, Cormack et al.(2019) found that machine learning algorithms (gradient boosting and random forest models) were more accurate than time series methods in forecasting demand for a pharmaceutical. It was found that their research achieved between 15–25% higher accuracy than exponential smoothing techniques. Similarly, Merkuryeva et al. In, a predictive analytics enabled demand-driven material requirements planning (MPR) is examined and with the predictive knowledge substantial decreases in carrying cost of inventory such as grocery and electronics etc. Studies of specialty pharmacy operations reveal particular challenges in managing inventory. Haque et al. Identified that unique management strategies are needed for specialty medications based on high unit costs, limited distribution networks and dispensing needs unique to individual patients (2017). They found that classic economic order quantity models failed to account for the demand fluctuation that characterizes how specialty medications are consumed.

The technical literature discusses this aspect of Redis implementation in healthcare, suggesting these attributes lend it well to being an application for high-performance applications. Grolinger et al. External Link: NoSQL Databases in Health Care: Review of Current Applications and Future Directions (2013)(External Link: NoSQL Databases in Health Care: Review of Current Applications and Future Directions (2013), Externally linked, open in a new window) reviewed the utility of NoSQL databases for managing health data and found significant enhancements to query performance for time-relevant queries through the use of in-memory caching in databases such as Redis. Complementing this work, Nguyen et al. The work of Dave, et al. (2020) found and described the use of caching strategies for pharmacy management systems with 60–80% reductions in the database load using Redis caching levels. Applications of machine learning in inventory optimization have yielded promising results across various industries. Syntetos et al. A detailed analysis on forecasting and stock control methods is presented by (2016) and emphasizes the need to select an algorithm adapted to demand specificities. Based on the demand patterns of specialty medications, their framework guides the choice of the appropriate predictive models. Chen et al. And (2020) also focused on an artificial intelligence use case that tackled the pharmaceutical inventory problem, yielding significant gains in forecast accuracy and inventory turnover. An emerging research direction is the coupling of real-time analytics with caching technologies. Kang et al. From year 2021, (2021) studied machine learning and cache-aware intelligent inventory management systems and demonstrated complementary advantages when combined with machine learning and cache. According to their findings, combined methods were more effective than using either technology independently.

3. Objectives

The study pursued four specific objectives to comprehensively evaluate the proposed technological integration:

- To assess the accuracy and effectiveness of machine learning-based predictive analytics models in forecasting specialty medication demand across different therapeutic categories.
- To evaluate the performance improvements achieved through Redis-backed caching implementation in pharmacy inventory management systems, specifically measuring query latency reduction and system throughput enhancement.
- To measure the impact of integrated predictive analytics and caching systems on key operational metrics including order fulfillment rates, stockout frequency, and inventory holding costs.
- To develop recommendations for optimal technology configuration and implementation strategies for specialty pharmacy operations in the Indian healthcare context.

4. Methodology

The present study used secondary data analysis within a quantitative research design to assess the impact of predictive analytics and Redis caching on pharmacy inventory management. Qualitative research in design comparing between traditional inventory management system and technological based systems using different inventory management systems. We collected data from published healthcare databases, reports from pharmacy management systems, and peer-reviewed literature publications with implementation outcomes. The study investigated operational data from 847 specialty pharmacy locations in various geographic areas based on healthcare industry reports published between 2018 and 2023. Inclusion criteria for sample included pharmacies with 500 or more specialty medication prescriptions filled per month with some prior technology implementation records. Participants were recruited using a stratified sampling approach to ensure representation from hospital-based, retail chain, and independent specialty pharmacies.

Research tools involved the use of statistical analysis software to develop the reader and processing to evaluate the machine learning model. Performance measurement information, including Redis cache hit rates, query response times from system logs, and inventory transaction information. Calculations against actual consumption data were used to quantify the performance of the approaches through Mean Absolute Percentage Error and Root Mean Square Error of forecasts. Methods The data originating from pharmacy benefit manager databases, inventory management system reports and published implementation case studies were systematically extracted. Statistical analysis Utilizing descriptive statistics, inferential statistics (i.e. t-tests and ANOVA, for group comparisons), and regression analysis to examine the relationship between technology implementation and operational outcomes. Given the nature of the data used in this study, we anonymized and aggregated all data to prevent compromising any proprietary information while maintaining value in our analysis.

5. Results

Pharmacy operations realized significant performance gains in the analysis of predictive analytics and the implementation of Redis caching within specialty pharmacy operations. Technology effectiveness was assessed using six analytical dimensions across incorporated data.

Table 1: Predictive Analytics Model Performance Comparison

Model Type	MAPE (%)	RMSE	Accuracy (%)	Training Time (sec)
Random Forest	12.3	45.6	87.7	234
Gradient Boosting	11.8	42.3	88.2	312
LSTM Neural Network	13.1	48.9	86.9	567
ARIMA (Traditional)	18.7	67.4	81.3	45
Exponential Smoothing	21.2	78.2	78.8	23

Comparative performance metrics for different demand forecasting models on specialty medication inventory data appears in Table 1. It perform best with an accuracy of 88.2% and with the lowest RMSE of 42.3, which also suggests that demand predictions are only slightly deviating from actuals (best of all models). Conventional statistical methods such as ARIMA and exponential smoothing performed far worsely, recording MAPE scores over 18%. Despite training times taking up to 567 seconds, the accuracy of the LSTM neural networks could not better a simple ensemble method. Conclusion These results validate that modern machine learning methods are considerably superior in predicting specialty medication demand over traditional prediction methods.

Table 2: Redis Cache Performance Metrics

Metric	Before Implementation	After Implementation	Improvement (%)
Average Query Latency (ms)	245	54	78.0
Peak Query Latency (ms)	1,234	187	84.8
Database Load (queries/sec)	12,450	3,120	74.9
Cache Hit Rate (%)	N/A	92.3	—
System Throughput (transactions/hr)	8,760	15,430	76.1

As shown in table 2, with the Redis caching added, there is significant improvement in the performance of pharmacy inventory systems. Query latency average went down from 245 milliseconds to 54 milliseconds a 78% improvement for near real time inventory visibility. A 92.3% cache hit rate implies a well-designed caching strategy, where the vast majority of queries for inventory were served from memory instead of the database. We enhanced system throughput by 76.1%, permitting pharmacies to process vastly increased transaction volumes during peak operational cycles. This increase in performance not only means a better experience for pharmacy staff, but also a quicker order processing functionality.

Table 3: Inventory Management Outcome Metrics

Metric	Traditional System	Technology-Enhanced	Statistical Significance
Stockout Rate (%)	8.7	3.2	$p < 0.001$
Order Fulfillment Rate (%)	76.4	94.8	$p < 0.001$
Inventory Turnover Ratio	6.2	9.4	$p < 0.01$
Days of Inventory	58.9	38.8	$p < 0.01$
Expired Medication Rate (%)	4.3	1.1	$p < 0.001$

As we can see in Table 3, post technology implementation, we observe large values of the specified inventory management indicators with respect to the baseline. This led to significant reductions in stockout rates, from 8.7% down to 3.2%, thereby minimizing the chances of therapy interruption for patients who need specialty medications. The percentage of orders fulfilled on first attempt (i.e., those without any kind of intervention required) increased from around 76.4 percent to around 94.8 percent, suggesting that technology-enhanced pharmacies completed many more prescription orders on the first attempt with little or no further help required. Inventory turnover ratio increased from 6.2 to 9.4, reflecting improved utilization of capital while days of inventory decreased from 58.9 to 38.8, indicating lean inventory positions. These differences were statistically significant at $p < 0.01$ or better) so it is improbable that the substantial improvements would be due to chance.

Table 4: Financial Impact Analysis (Annual, Per Pharmacy)

Cost Category	Traditional ($)	Technology-Enhanced ($)	Savings ($)	Reduction (%)
Inventory Holding Costs	234,560	180,611	53,949	23.0
Stockout Costs	87,430	32,149	55,281	63.2
Expired Medication Losses	156,780	40,763	116,017	74.0
Technology Investment	0	45,000	(45,000)	—
Net Annual Benefit	—	—	180,247	—

Financial Benefit of Predictive Analytics & Redis Caching Implementation Each of these benefits is quantified in Table 4. Due to a combination of lower inventory levels and improved demand forecasting, annual inventory holding costs fell by 23%, from $234,560 to $180,611. This effort resulted in the largest portion of savings in the area of expired medication losses, which fell by 74% from $156,780 to $40,763 per year as inaccurate forecasting was avoided resulting in less over-purchasing of costly specialty products. Following a $45,000 per annum investment in technology, average net annual benefit was $180,247 per pharmacy site. This ROI highlights clear financial justification for technology adoption extending well beyond operational efficiency factors.

Table 5: Specialty Medication Category Analysis

Therapeutic Category	Forecast Accuracy (%)	Stockout Reduction (%)	Fulfillment Improvement (%)
Oncology	85.4	58.7	21.3
Rheumatology	89.2	67.4	24.8
Multiple Sclerosis	91.3	72.1	28.6
Hepatitis C	86.7	61.2	22.4
HIV/AIDS	92.8	78.3	31.2
Rare Diseases	78.9	45.6	15.7

As shown in Table 5, the effectiveness of the technology varies among specialty medication therapeutic classes. Forecast accuracy was highest for HIV/AIDS medications at 92.8% and stockout was lowest for HIV/AIDS medications at 78.3%, suggesting more stable patient populations with return refill patterns following medication refills. For drugs used in multiple sclerosis, the prediction accuracy was also high (91.3%). The lowest performance was observed for rare disease medications with an F-measure accuracy of 78.9% and a stockout reduction of 45.6% due to the inherent unpredictability of demand for low-volume specialty products. These results at the level of the individual drug category provide actionable information that can develop implementation strategies focused on particular types of pharmacies that would benefit most from the technology as well as identifying target therapeutic areas where to focus on technology deployment augmentation.

Table 6: Implementation Timeline and Performance Progression

Implementation Phase	Duration (months)	Query Latency (ms)	Forecast Accuracy (%)	Fulfillment Rate (%)
Baseline	0	245	78.8	76.4
Initial Deployment	1-3	178	81.2	79.8
Optimization	4-6	98	84.6	86.3
Mature Operation	7-12	54	88.2	94.8
Steady State	12+	52	89.1	95.2

Performance trajectory improvement after technology implementation (Table 6) Phase 1 of Deployment (query latency reduction to 178 milliseconds and 81.2% accuracy on forecasting minor improvement, pre-built features play a role here) Months 4–6 saw orders of magnitude refinement in performance thanks to caching strategies and frequent retraining of ML models using operational data accrued over the previous 3 months of traffic. Mature operation delivered on-target performance levels with continuous step-wise improvements into steady state post twelve months. This pattern of variability also implies that full realization of benefits requires an ongoing optimization effort rather than the simple deployment of a solution, which is an important input into realistic expectation management as implementation plans are developed.

6. Discussion

These results from the application of integrated predictive analytics and Redis caching to fill specialty medications, and manage pharmacy inventory, provide strong evidence for the efficacy of the presented approach. The operational, financial, and patient-care gains identify the direct benefits hypothesized to result from technology integration, as well as important contextual evidence to guide development of implementation strategies (Aitken & Kleinrock, 2019). The better performance of the gradient boosting and random forest models over the traditional statistical models is consistent with recent literature in demand forecasting that suggests the superiority of these tree-based models over other traditional statistical methods. Because of these reasons, captured complex, non-linear demand patterns associated with the specialty medications could be effectively modelled via these ensemble methods, which are difficult to encapsulate through the existing ARIMA and exponential smoothing approaches (Kilimci et al, 2019). An accuracy of 88.2% is a significant improvement over baseline performance; however, the result of rare disease drugs attaining lower accuracy indicates natural limits to the ability to forecast highly intermittent demand patterns.

Redis caching performed remarkably well in fulfilling system performance requirements based on real-time inventory management. Pharmacy staff are able to obtain up-to-date information on available inventory instantly with a resultant 78% reduction in average query latency, which facilitates rapid decision-making for prescription fulfilment (Grolinger et al., 2013). A cache hit rate of 92.3% suggests this was the right move in the architecture design as inventory data that was accessed most frequently was kept in memory. According to Carlson, quote bigger always is better: The performance level here is a great deal greater than standard usage only implementations and validates Redis as ideal technological innovation for healthcare source of information applications. This study illustrates operational improvements that have important implications for the quality of patient care. Cutting stockouts from 8.7% to 3.2% means patients relying on specialty drugs for maintaining therapy for their chronic conditions can expect reduced interruptions in their therapy (Tichy et al., 2019). Because specialty medications often address critical illness such as cancer and autoimmune diseases, continuous medication access is not just a matter of operational efficiency but is a fundamental issue of safety for the patient (Haque et al., 2017).

An analysis of financials shows a strong return on investment for the adoption of the technology. The net annual benefit of 180,247 dollars per pharmacy site far exceeds implementation costs, with payback period often within six months. By enhancing demand forecasting (through the work of the SAV DHE team), the largest savings was also in expired medication losses (74% reduction from baseline) (Chen et al., 2020). Because, for specialty medications, the unit costs of many individual items typically exceed $1,000, even small reductions of expirations lead to large savings that are multiplied across by portfolio of inventory. Such category-level analysis can gain actionable insights to help with the prioritization of implementations. Predictive analytics provided the most benefit for therapeutic categories with relatively predictable demand patterns, which were typically the categories related to HIV/AIDS and multiple sclerosis medications, while rare disease medications showed more modest improvements (Merkuryeva et al., 2019). These results support diversified inventory strategies at the pharmacy level, with predictive analytics applied most aggressively for high-volume specialty categories and higher safety stocks used for the more unpredictable rare disease products (Nguyen et al., 2020).

Analysis of implementation timelines shows that realizing the associated benefits necessitates ongoing optimization effort long after initial deployment. We (Kang et al., 2021) went from a baseline forecast accuracy of 78.8% to a steady-state accuracy of 89.1% over twelve months of using and improving the model in production. This finding is crucial for expectation management and resource allocation, as organizations need to realize that implementation of optimization should not be seen as a one-time project, but rather an exercise that entails continuous investments (Shah, 2004). Technology adoption is a mix of opportunities and challenges in the Indian pharmaceutical context. While the specialty medication market is expanding quickly (Palumbo & Vitaliano, 2008; Basak & Sathyanarayana, 2010), making the need for improved inventory management capabilities quite compelling (Zhang, Manickam, & Varma, 2013), pharmacy infrastructure variability across locations may influence feasibility of implementation. Organizations looking to implement technology must weigh infrastructure needs and plan roll out in stages, appropriate to the environments in which they operate.

7. Conclusion

This paper illustrates that predictive analytics plus Redis backed caching systems are able to improve fulfilment of specialty medication orders and also leads to better pharmacy stock handling. Demand forecast accuracy of 88.2% using machine learning algorithms which far outperformed traditional statistical methods. With Redis caching, we slashed query latency by 78% and sustained cache hit rates of 92.3%, making inventory visibility in real-time possible and key to operational decision-making. But some operational improvements also happened: 63.2% reduction in stock-out rates, 24% improvement in order fulfilment and 74% decrease in expired medication losses. The key financial analysis was calculating the net annual benefit of $180,247 per pharmacy location achieved on an annual basis that provides clear evidence of return on technology investment. Findings by category suggest largest gains found in predictable therapeutic categories and indicate that innovative inventory strategies are required for rare disease medications. Realizing the full benefit during implementation can take up to twelve months of continued effort to optimize the process. Our findings lend support to technology adoption recommendations for specialty pharmacies looking to improve their operations and clinical care outcomes as pharmaceutical supply chains become increasingly complex.

References

1. Aitken, M., & Kleinrock, M. (2019). Medicine use and spending in the U.S.: A review of 2018 and outlook to 2023. *IQVIA Institute for Human Data Science*, 1-49.
2. Basak, S. C., & Sathyanarayana, D. (2010). Pharmacy education in India. *American Journal of Pharmaceutical Education*, 74(4), 68.

3. Gunda, S. K. (2023). Software defect prediction using advanced ensemble techniques: A focus on boosting and voting methods. 2023 International Conference on Electronic Systems and Intelligent Computing (ICESIC), Chennai, India, 157–161. https://doi.org/10.1109/ICESIC61777.2023.10846550
4. Carlson, J. L. (2013). *Redis in action*. Manning Publications.
5. Gunda SK, Yettapu SDR, Bodakunti S, Bikki SB. Decision Intelligence Methodology for AI-Driven Agile Software Lifecycle Governance and Architecture-Centered Project Management, 2023 Mar. 30;4(1):102-8. https://doi.org/10.63282/3050-9262.IJAIDSML-V4I1P112
6. Chen, Y., Li, K., Kawamoto, K., & Kohane, I. S. (2020). Machine learning-enabled predictive analytics for pharmaceutical inventory management. *Journal of Biomedical Informatics*, 104, 103395.
7. Gunda, S. K. (2023). Enhancing software fault prediction with machine learning: A comparative study on the PC1 dataset. 2023 Global Conference on Communications and Information Technologies (GCCIT), Bangalore, India, 1–4. https://doi.org/10.1109/GCCIT63234.2023.10862351
8. Grolinger, K., Higashino, W. A., Tiwari, A., & Capretz, M. A. (2013). Data management in cloud environments: NoSQL and NewSQL data stores. *Journal of Cloud Computing: Advances, Systems and Applications*, 2(1), 22.
9. Haque, S. N., Gellad, W. F., & Engel, K. G. (2017). Specialty medication management: Key considerations for hospital-based pharmacies. *American Journal of Health-System Pharmacy*, 74(7), 482-488.
10. Gunda, S. K. (2023). Comparative analysis of machine learning models for software defect prediction. 2023 International Conference on Power, Energy, Control and Transmission Systems (ICPECTS), Chennai, India, 1–6. https://doi.org/10.1109/ICPECTS62210.2023.10780167
11. Kang, J., Lee, S., & Park, S. (2021). Intelligent inventory management system using machine learning and caching mechanisms. *Computers & Industrial Engineering*, 156, 107274.
12. Gunda, S. K. (2023). Machine learning approaches for software fault diagnosis: Evaluating decision tree and KNN models. 2023 Global Conference on Communications and Information Technologies (GCCIT), Bangalore, India, 1–5. https://doi.org/10.1109/GCCIT63234.2023.10861953
13. Kilimci, Z. H., Akyuz, A. O., Uysal, M., Akyokus, S., Uysal, M. O., Bulbul, B. A., & Ekmis, M. A. (2019). An improved demand forecasting model using deep learning approach and proposed decision integration strategy for supply chain. *Complexity*, 2019, 9067367.
14. Gunda, S. K. (2023). Analyzing machine learning techniques for software defect prediction: A comprehensive performance comparison. 2023 Asian Conference on Intelligent Technologies (ACOIT), Kolar, India, 1–5. https://doi.org/10.1109/ACOIT62457.2023.10939610
15. Merkuryeva, G., Valberga, A., & Smirnov, A. (2019). Demand forecasting in pharmaceutical supply chains: A case study. *Procedia Computer Science*, 149, 3-10.
16. Nguyen, T., Li, Z., & Spiegler, V. (2020). Caching strategies for pharmaceutical inventory systems: A performance evaluation. *International Journal of Production Economics*, 228, 107743.
17. Shah, N. (2004). Pharmaceutical supply chains: Key issues and strategies for optimization. *Computers & Chemical Engineering*, 28(6-7), 929-941.
18. Syntetos, A. A., Babai, Z., Boylan, J. E., Kolassa, S., & Nikolopoulos, K. (2016). Supply chain forecasting: Theory, practice, their gap and the future. *European Journal of Operational Research*, 252(1), 1-26.
19. Gunda, S. K. G. (2023). The Future of Software Development and the Expanding Role of ML Models. International Journal of Emerging Research in Engineering and Technology, 4(2), 126-129. https://doi.org/10.63282/3050-922X.IJERET-V4I2P113
20. Tichy, E. M., Hoffman, J. M., Suda, K. J., Rim, M. H., Tadrous, M., Cuellar, S., & Schumock, G. T. (2019). National trends in prescription drug expenditures and projections for 2019. *American Journal of Health-System Pharmacy*, 76(15), 1105-1121.
21. Gunda, S. K. (2023). Device for continuous software testing and validation (UK Registered Design No. 6400738). Registered with the UK Intellectual Property Office, Class 14-02, granted in November 2023.