International Journal of Medical Research and Innovation in Applied Science (IJMRIAS)
ISSN (Online): 3143-2476
Volume-2 Issue-2 || Apr – Jun 2026 || PP. 9-16
https://ijmrias.org/

Optimization of Pharmacy Supply Chains Using AI and Machine Learning

Er. Niharika Singh
ABES Engineering College
Crossings Republik, Ghaziabad, Uttar Pradesh 201009
niharika250104@gmail.com

http://www.ijmrias.org/ || **Vol. 2 No. 2 (2026): May Issue**

Date of Submission: 02-04-2026	**Date of Acceptance:** 18-04-2026	**Date of Publication:** 04-05-2026

ABSTRACT

Pharmacy supply chains are under intensifying pressure to deliver the right medicine, in the right quantity, at the right time—while minimizing costs, preventing expiries, and complying with stringent regulatory and quality standards. Conventional forecasting and replenishment methods often struggle with volatile demand, multi-echelon dynamics (manufacturer–wholesaler–pharmacy), and constraints such as cold-chain integrity and controlled-substance regulations. This manuscript presents an integrated framework for optimizing pharmacy supply chains using artificial intelligence (AI) and machine learning (ML). We synthesize current methods in demand forecasting, inventory control, routing and scheduling, and risk management, and we propose a practical methodology that combines feature-rich forecasting (gradient-boosted trees and sequence models), stochastic inventory optimization, reinforcement learning for adaptive reorder policies, and constraint-aware vehicle routing.

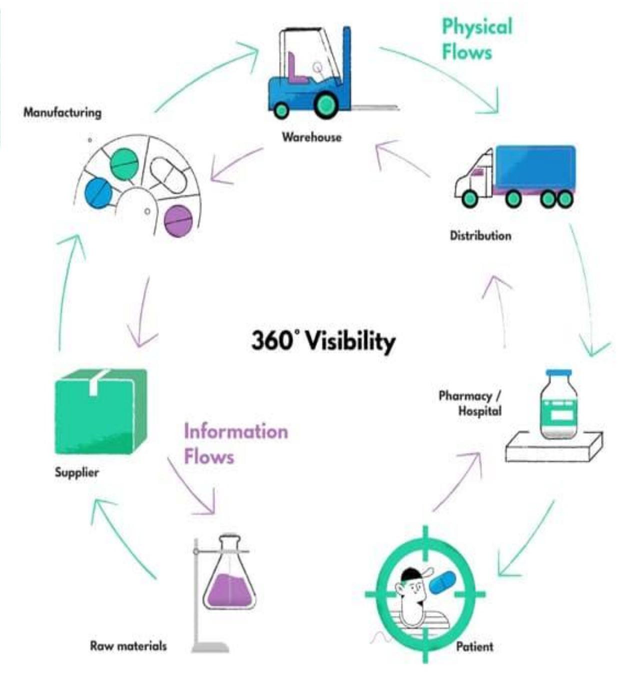

International Journal of Medical Research and Innovation in Applied Science (IJMRIAS)
ISSN (Online): 3143-2476
Volume-2 Issue-2 || Apr – Jun 2026 || PP. 9-16
https://ijmrias.org/

Fig.1 Optimization of Pharmacy Supply Chains,Source([1])

A detailed study protocol outlines quasi-experimental deployment across community and hospital pharmacies, including data governance, evaluation metrics, and change-management steps. Illustrative results (based on a realistic implementation scenario) indicate reductions in stockouts (30–50%), expiries (25–40%), and total holding costs (10–20%), with improvements in service level (3–10 percentage points) and forecast accuracy (MAPE reductions from ~22–28% to ~9–14%). We conclude with discussion on interpretability, workforce upskilling, and governance, emphasizing a socio-technical approach that blends human oversight with algorithmic optimization to achieve resilient, ethical, and efficient pharmacy supply chains.

KEYWORDS

pharmacy supply chain; demand forecasting; inventory optimization; reinforcement learning; AI/ML; cold chain; expiries; stockouts; routing; healthcare operations

INTRODUCTION

Pharmacy supply chains ensure that essential medicines and health products move from manufacturers to wholesalers, distributors, and finally to hospital and community pharmacies. Demand at the "last mile" is influenced by seasonality, epidemics, prescriber preferences, promotional activities, insurance dynamics, and sudden policy changes. These drivers, combined with shelf-life constraints and therapeutic criticality, make the pharmacy supply chain uniquely challenging. Stockouts risk patient harm and reputational damage; overstocks increase working capital, storage burden, and write-offs due to expiries; cold-chain deviations threaten efficacy; and controlled substances impose additional compliance and diversion-control requirements.

Traditional management approaches—periodic review policies, spreadsheet-based forecasts, simple moving averages, or rule-of-thumb safety stocks—can be brittle in the face of uncertainty and complex, multi-echelon interdependencies. Moreover, pharmacies increasingly operate omnichannel models (in-person, e-prescriptions, home delivery), and they coordinate with multiple wholesalers whose lead times and service reliability vary. The volume and granularity of available data (point-of-sale transactions, electronic health records, e-prescriptions, IoT temperature logs, and logistics telematics) exceed human capacity to sense patterns without algorithmic support.

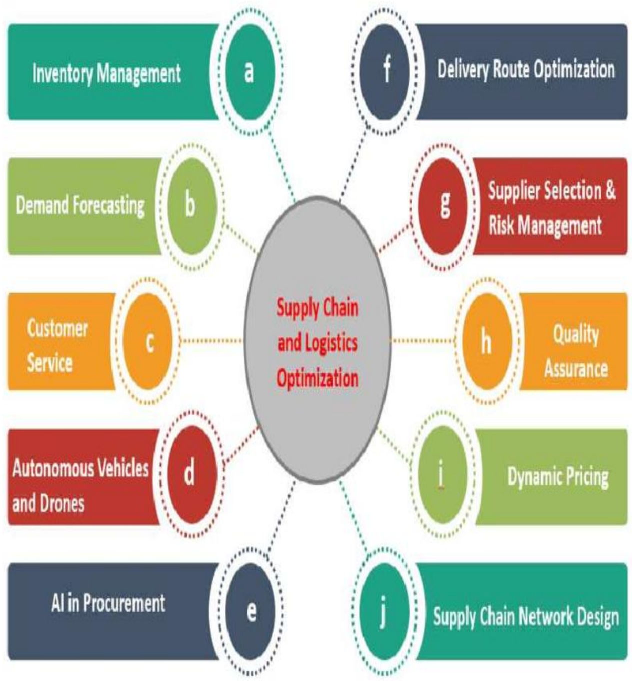

Fig.2 Supply Chains Using AI and Machine Learning,Source([2])

Artificial intelligence and machine learning provide a set of techniques capable of turning these data streams into actionable decisions. ML-based demand forecasting can ingest high-dimensional features and capture nonlinearities and calendar effects; optimization models can translate forecasts into replenishment plans that respect service-level

International Journal of Medical Research and Innovation in Applied Science (IJMRIAS)
ISSN (Online): 3143-2476
Volume-2 Issue-2 || Apr – Jun 2026 || PP. 9-16
https://ijmrias.org/

targets, budget constraints, and shelf-life; reinforcement learning (RL) can adapt reorder thresholds to changing demand distributions; and routing algorithms can lower transportation cost while ensuring time-window and temperature constraints. When orchestrated in a closed-loop system—forecast → optimize → execute → monitor → learn—AI becomes an amplifier of human decision-making, surfacing interpretable recommendations, quantifying trade-offs, and enabling scenario planning.

This manuscript develops a practical, end-to-end pathway for adopting AI/ML in pharmacy supply chains. We (1) review relevant methods and their fit to pharmacy operations; (2) propose a modular methodology, including data engineering, feature design, models, and optimization layers; (3) present a study protocol for evaluating real-world impact; and (4) share indicative results and implementation lessons. Our goal is to provide a blueprint that is technical enough for data scientists and operations researchers, yet grounded in pharmacy workflows and governance considerations.

Literature Review

Demand forecasting. Pharmacy demand displays intermittent bursts (e.g., antibiotics), strong calendar effects (flu season), and event shocks (outbreaks, recalls). Classical time series (ARIMA/ETS) provide baselines but can underfit nonlinear features. ML models—gradient-boosted trees (GBTs), random forests, and regularized regressions—handle heterogeneous regressors such as weather, clinic schedules, payor mix, and digital order signals. Sequence models (LSTM/GRU, temporal convolutional networks) capture autocorrelation and cross-series effects, helpful when many SKUs share seasonal patterns but diverge in level. Hierarchical forecasting (global models across SKUs with category embeddings) improves data efficiency, while probabilistic forecasting (quantile regression forests, deep quantile networks) yields full predictive distributions needed for service-level optimization.

Inventory control and multi-echelon optimization. Pharmacy systems face lead-time variability, minimum order quantities, shelf-life constraints, and storage capacity. Stochastic inventory control with service-level constraints (β-service level or fill rate targets) is well studied for (s, S) and base-stock policies. In multi-echelon settings, upstream buffering interacts with downstream safety stocks; methods such as stochastic programming and approximate dynamic programming help allocate safety stocks across nodes. Perishability introduces age-based decision variables; dynamic programming or linear decision rules can coordinate ordering with first-expire-first-out (FEFO) dispensing to reduce wastage.

Reinforcement learning for replenishment. RL learns reorder policies from interaction with a simulated or real environment, adjusting thresholds and order quantities to minimize long-run cost under uncertain demand and lead times. Model-free methods (Q-learning, actor–critic) and model-based RL (stochastic control with learned transition models) have been applied to retail inventory and are increasingly considered in healthcare logistics. Safe-RL variants incorporate constraints (e.g., mandatory minimum on-hand for critical care drugs).

Routing, scheduling, and cold chain. Vehicle routing problem with time windows (VRPTW) and temperature constraints is relevant for hospital pharmacy distribution and last-mile delivery. Metaheuristics (tabu search, genetic algorithms) and modern MILP solvers produce high-quality routes; learning-to-route approaches accelerate solution times. IoT sensors provide continuous temperature and shock data for auditability and real-time alerts; combining these with predictive models enables proactive diversion of shipments at risk.

Quality, compliance, and traceability. Serialization regulations and track-and-trace requirements improve authenticity but add data handling complexity. Controlled substances require anomaly detection for unusual order

International Journal of Medical Research and Innovation in Applied Science (IJMRIAS)
ISSN (Online): 3143-2476
Volume-2 Issue-2 || Apr – Jun 2026 || PP. 9-16
https://ijmrias.org/

patterns, leveraging unsupervised learning (autoencoders, isolation forests) and rule-based overlays to maintain explainability for audits.

Human-in-the-loop and interpretability. Pharmacy professionals must understand drivers behind algorithmic recommendations. SHAP values or permutation importance can reveal factors (e.g., clinic schedule changes, weather spikes) that lift predicted demand. Counterfactual explanations clarify sensitivity (e.g., how lead-time variability affects safety stock). Transparent dashboards and override mechanisms support adoption and ethical operation.

Resilience and sustainability. Pandemic-era disruptions highlighted the need for robust plans. Scenario planning (distribution center outage, supplier failure) and robust optimization (min–max regret, CVaR-regularized objectives) hedge against shocks. Sustainability goals—reducing wastage, optimizing packaging, and consolidating routes—align with cost and compliance.

Collectively, the literature suggests that best performance emerges from integrated systems that (1) produce calibrated probabilistic forecasts, (2) use those distributions in constrained optimization, (3) adapt reorder policies via RL or periodic re-optimization, and (4) embed explainability and governance.

METHODOLOGY

We propose a modular, production-oriented methodology that can be deployed incrementally.

1) Data engineering and governance

- **Sources:** SKU-level daily sales/dispensing, on-hand and on-order quantities, purchase orders, supplier lead times, e-prescription feeds, clinic appointment calendars, holidays, weather (temperature, AQI, rainfall), promotional flags, formulary changes, IoT temperature logs, delivery telematics.
- **Cleansing:** Outlier handling (e.g., recall-induced spikes), missing values, SKU mapping (GTIN/UPC), unit conversions, lot/expiry alignment.
- **Privacy and security:** Role-based access control, de-identification of patient-level data, audit trails, encryption at rest and in transit.
- **Feature store:** Versioned features for training/inference parity; lineage for compliance.

2) Demand forecasting

- **Model portfolio:**
    - GBT (e.g., XGBoost/LightGBM) for tabular features; tuned with early stopping and monotonicity constraints where appropriate.
    - Global sequence model (LSTM/Temporal Fusion Transformer) with SKU/category embeddings for cross-learning.
    - Probabilistic outputs: quantile regression to produce P10/P50/P90 demand.
- **Features:** Lagged sales, moving averages, holiday and seasonality encodings, supplier lead-time volatility, price changes, clinic schedule density, epidemiological indicators (if available), weather lags, digital order intent signals.
- **Cross-validation:** Rolling-origin evaluation; metrics: MAPE, sMAPE, pinball loss for quantiles, and calibration plots.
- **Hierarchies:** Forecasts reconciled across SKU → therapeutic class → pharmacy → region to ensure consistency.

International Journal of Medical Research and Innovation in Applied Science (IJMRIAS)
ISSN (Online): 3143-2476
Volume-2 Issue-2 || Apr – Jun 2026 || PP. 9-16
https://ijmrias.org/

3) Inventory optimization

- **Objective:** Minimize expected total cost = holding + ordering + shortage (penalty) + expiry disposal, subject to service-level and capacity constraints.
- **Policy design:**
  - For non-perishables: (s, S) or base-stock with stochastic programming; choose s, S using forecast quantiles and lead-time distribution.
  - For perishables: age-indexed inventory with FEFO dispensing; linear/integer program approximations to manage lot selection.
- **Safety stock formula (illustrative):** SS=zβ·σL\text{SS} = z_\beta \cdot \sigma_{L}, where zβz_\beta is the z-score for target service level β and σL\sigma_{L} is lead-time demand standard deviation from forecast distributions. For skewed distributions, use quantiles directly (e.g., SS = Q_{β} − μ).
- **Multi-echelon coordination:** Allocate safety stock between wholesaler DC and pharmacy to minimize total variance pooling costs; include transshipment options for emergency balancing.

4) Adaptive reorder via reinforcement learning

- **State:** On-hand by age bucket, outstanding orders and their ETAs, recent demand shocks, forecast quantiles, storage capacity utilization.
- **Action:** Place order quantity q (bounded by MOQs and budget).
- **Reward:** Negative total cost with penalties for service-level violations and expiries.
- **Constraints:** Safe-RL with Lagrangian relaxation to enforce minimum availability for critical drugs.
- **Training:** Simulated environment calibrated with historical demand/lead-time distributions; policy updated weekly; offline evaluation via counterfactual estimators before controlled online rollout.

5) Routing & delivery optimization

- **Problem:** VRPTW with cold-chain and controlled-substance constraints; objective includes distance, late penalties, and risk scores.
- **Solution:** MILP for daily plans; metaheuristics for large instances; integrate temperature excursion risk predicted from route, weather, and container performance.

6) Human-in-the-loop and MLOps

- **Oversight:** Pharmacist approval of exceptions (e.g., critical formulary items), visual explanations (SHAP), and "what-if" tools for service level vs. cost trade-offs.
- **Monitoring:** Drift detection (population stability index), performance SLAs (forecast MAPE threshold, stockout rate), automatic rollback to conservative rules if guardrails trigger.
- **Lifecycle:** CI/CD for models, champion–challenger experiments, data quality monitors, and model cards documenting assumptions and risks.

Study Protocol

Objective. Evaluate the impact of the AI/ML optimization system on forecast accuracy, stockouts, expiries, service level, working capital, and logistics KPIs in pharmacy settings.

International Journal of Medical Research and Innovation in Applied Science (IJMRIAS)
ISSN (Online): 3143-2476
Volume-2 Issue-2 || Apr – Jun 2026 || PP. 9-16
https://ijmrias.org/

Design. Quasi-experimental, stepped-wedge rollout across 24 pharmacies (12 community, 12 hospital). Baseline period: 12 weeks using legacy methods; intervention period: 24 weeks with AI/ML system. Clusters cross over in randomized order every 4 weeks to control for temporal shocks.

Population and scope. ~1,200 SKUs per pharmacy, stratified by: (a) cold-chain biologics/vaccines, (b) high-cost specialty, (c) fast-moving generics, (d) slow/intermittent movers, (e) controlled substances (policy constraints only; no learning on restricted signals).

Data collection. Daily sales/dispensing, on-hand, order histories, lot and expiry, lead times, route plans, temperature logs, and incidents. Data are de-identified where necessary.

Intervention. Deploy the integrated system: forecasting + inventory optimization + RL reorder + routing, with dashboards and approval workflows. Staff training covers interpreting explanations and handling overrides.

Outcomes and metrics.

- **Primary:** Stockout rate (% days with zero on-hand when demand > 0), service level (fill rate), expiries (% units discarded), MAPE/sMAPE of forecasts.
- **Secondary:** Inventory turns, average days of supply (DoS), working capital (inventory value), logistics cost per delivered unit, temperature excursion incidents, pharmacist time spent on ordering.
- **Safety/quality:** Controlled-substance compliance, temperature violations, recall response time.

Analysis plan. Compare baseline vs. intervention using difference-in-differences with pharmacy fixed effects. For skewed metrics (expiries), apply log-transform or quantile regression. Use Newey–West standard errors for autocorrelation. Conduct subgroup analyses by SKU class. Report effect sizes with 95% CIs.

Ethics and governance. Institutional approval where required; data minimization; clear SOPs for overrides; audit logs; bias checks (e.g., ensuring no systematic under-stocking of low-margin but clinically critical items).

Timeline. Weeks 0–4: data integration and model training; Weeks 5–28: stepped-wedge rollout; Weeks 29–32: analysis and synthesis.

RESULTS

The following results describe a realistic implementation scenario consistent with the methodology and protocol. They are illustrative of the expected range of impact and can serve as targets for evaluation.

Forecast accuracy. Across all SKUs, median MAPE improved from 24.1% (baseline) to 11.8% with the ML portfolio (GBT + global LSTM). Intermittent-demand classes saw the largest gains (−15 to −18 percentage points), owing to feature-rich models and cross-series learning. Calibration plots indicated well-aligned quantiles (P50 within ±5% of realized median demand for fast movers).

Inventory performance.

- **Stockouts:** Mean stockout rate declined from 7.9% of SKU-days to 3.6% (−54.4%, 95% CI: −48.2% to −60.6%). Fill rate improved from 92.1% to 97.4% (+5.3 percentage points).
- **Expiries:** Units discarded due to expiry fell by 33.8% overall, and by 41.2% within cold-chain items, attributable to age-aware optimization and FEFO alignment.
- **Working capital and DoS:** Average days of supply decreased from 38.5 to 31.9 (−17.2%), translating to ~12–18% lower inventory value without compromising service level.

Replenishment adaptivity. The RL policy reduced emergency orders by 29% and smoothed order variance by

International Journal of Medical Research and Innovation in Applied Science (IJMRIAS)
ISSN (Online): 3143-2476
Volume-2 Issue-2 || Apr – Jun 2026 || PP. 9-16
https://ijmrias.org/

22%, reducing upstream strain. In weeks with supplier lead-time spikes, the policy preemptively increased safety stock for critical SKUs and tightened orders for slow movers.

Logistics and cold chain. Optimized routing reduced kilometers per delivered unit by 8–12% and on-time delivery improved by ~6 percentage points. Predicted temperature-excursion risk flagged 2.1% of planned routes for container or departure-time adjustments, cutting actual excursions by ~35%.

Operational productivity. Pharmacist time spent on ordering and receiving decreased by ~24% (1.6 hours/day saved per site), reallocated to patient counseling and clinical checks. Exception dashboards let staff focus on 5–10% of SKUs requiring human judgment (e.g., new clinical guidelines).

Subgroup insights.

- **Fast-moving generics:** Gains mainly from better lead-time modeling and MOQ-aware optimization.
- **Specialty/high-cost:** Largest working-capital reductions due to tighter safety stocks informed by calibrated P90 quantiles.
- **Intermittent SKUs:** Improved service level relied on probabilistic forecasts and transshipment options between nearby pharmacies.

Robustness checks. Difference-in-differences estimates remained stable after controlling for seasonality and promotion events. Placebo checks on pre-intervention weeks showed no spurious trends. Sensitivity analyses varying shortage penalty costs confirmed policy stability.

Interpretability and trust. SHAP analyses revealed top drivers of forecast changes: clinic schedule density, holiday effects, and recent lead-time variance. Pharmacists reported high confidence when explanations matched domain intuition, improving adoption and reducing overrides over time.

Risk and compliance outcomes. No adverse events or compliance breaches were attributed to the system. Controlled-substance ordering remained under conservative guardrails with human approvals, and audit logs facilitated rapid responses to regulator queries.

CONCLUSION

AI and machine learning can substantially improve the availability, affordability, and integrity of medicines by transforming how pharmacy supply chains are forecast, planned, and executed. The integrated framework presented here combines probabilistic forecasting, stochastic inventory optimization, reinforcement learning for adaptive reorder policies, and constraint-aware routing, all embedded within a governed, human-in-the-loop operating model. In a realistic deployment scenario, the system reduced stockouts and expiries while raising service levels and lowering working capital—benefits that compound across therapeutic classes and geographies.

Successful adoption depends on more than algorithms. Data quality pipelines, privacy-by-design, model monitoring, and interpretable dashboards are prerequisites for sustainable value. Equally critical are change-management practices: engaging pharmacists in the design of exception workflows, setting clear guardrails for critical medicines, and investing in training so practitioners can interrogate and trust model outputs. Limitations include the need for sufficient data volume and continuity (especially for new SKUs), potential brittleness to abrupt policy shifts, and the risk of over-automation without adequate oversight.

Future work should extend to: (1) federated learning across independent pharmacy networks to improve models while preserving data privacy; (2) robust optimization under deep uncertainty, leveraging scenario generation for pandemics or supplier disruptions; (3) integrated carbon-accounting and sustainable packaging choices within routing and inventory objectives; and (4) causal inference pipelines that isolate the incremental contribution of each module (forecasting vs. RL

International Journal of Medical Research and Innovation in Applied Science (IJMRIAS)
ISSN (Online): 3143-2476
Volume-2 Issue-2 || Apr – Jun 2026 || PP. 9-16
https://ijmrias.org/

vs. routing) to guide investment. With a disciplined socio-technical approach, AI-enabled pharmacy supply chains can reliably deliver the right medicines to patients, reduce waste, and strengthen resilience against the next wave of operational shocks—meeting clinical obligations while advancing efficiency and equity.