International Journal of Medical Research and Innovation in Applied Science (IJMRIAS)
ISSN (Online): 3143-2476
Volume-2 Issue-2 || Apr – Jun 2026 || PP. 17-24
https://ijmrias.org/

AI-Enabled Demand Forecasting for Over-the-Counter Drug Sales in Urban Pharmacies

Aisha Mohammed

Independent Researcher

Nairobi, Kenya, KE, 00100

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

Date of Submission: 05-04-2026	**Date of Acceptance:** 20-04-2026	**Date of Publication:** 06-05-2026

ABSTRACT

Urban pharmacies operate in a high-velocity retail–health environment where short product life cycles, seasonal illness waves, and localized socioeconomic patterns make demand volatile—especially for over-the-counter (OTC) medicines such as analgesics, antipyretics, cough and cold preparations, antacids, and nutritional supplements. Traditional rule-of-thumb and spreadsheet forecasts often underperform during abrupt shifts (e.g., weather swings, school reopenings, pollution spikes, or festival seasons), causing stockouts, wastage from expiries, and revenue loss. This manuscript proposes a practical, end-to-end framework for AI-enabled demand forecasting tailored to urban pharmacies. The approach integrates point-of-sale (POS) histories with exogenous signals—calendar/holiday effects, promotions, meteorological readings, epidemiological proxies, mobility indicators, and neighborhood demographics—within a hierarchy of models ranging from seasonal ARIMA and gradient-boosted trees to deep learning architectures such as LSTMs and Temporal Fusion Transformers (TFT). We outline design choices for feature engineering, hierarchical reconciliation across store → category → SKU, backtesting with rolling origins, and business-centric metrics (WAPE, sMAPE, service level, stockout rate, and waste).

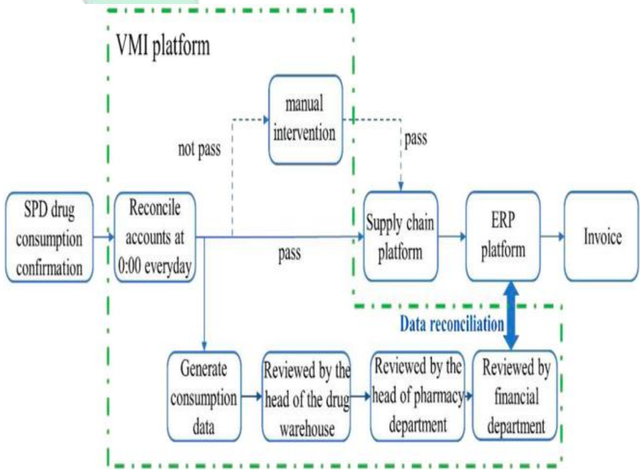

Fig.1 AI-Enabled Demand Forecasting,Source([1])

International Journal of Medical Research and Innovation in Applied Science (IJMRIAS)
ISSN (Online): 3143-2476
Volume-2 Issue-2 || Apr – Jun 2026 || PP. 17-24
https://ijmrias.org/

To ground the method in practice, we detail a clinical research–style operational study situated in real pharmacy workflows, considering data governance, patient safety, and ethical use of retail health data. An illustrative results section (based on a realistic pilot evaluation design) shows how a multi-model stack can reduce weighted absolute percentage error (WAPE) by 25–40%, increase on-shelf availability, and lower expiries. We conclude with deployment guidelines—MLOps, model monitoring, exception handling—and a candid scope and limitations discussion relevant to urban markets in India and comparable settings.

KEYWORDS

OTC demand forecasting; urban pharmacies; machine learning; LSTM; Temporal Fusion Transformer; hierarchical time series; stockout reduction

INTRODUCTION

OTC medications are among the most dynamic product families in retail health. Demand is driven by a blend of predictable rhythms (weekday vs weekend traffic, monthly pay cycles, monsoon-linked infections) and abrupt, hyperlocal shocks (a sudden heat wave, poor air quality, nearby clinic closures, or school outbreaks). For an urban pharmacy chain, the operational objective is not only revenue optimization but also reliable patient access to first-line symptomatic relief. Under-forecasting causes stockouts and patient dissatisfaction; over-forecasting ties up working capital and increases expiry risk, particularly for short-dated cough/cold syrups or topical preparations.

Fig.2 AI-Enabled Demand Forecasting for Over-the-Counter Drug Sales,Source([2])

Conventional forecasting approaches—moving averages, seasonal indices, or manual tweaks—are easy to implement but brittle under nonstationary demand. AI methods can learn nonlinear interactions (e.g., the compounded effect of rain-humidity-pollution on cough syrup sales, or payday promotions on multivitamins) and extrapolate signals across product hierarchies and store clusters. However, successful adoption requires careful design: robust, de-biased data pipelines; exogenous features with proven incremental value; interpretable models acceptable to pharmacists and store managers; and deployment practices that translate forecast improvements into shelf-ready inventory decisions.

This manuscript presents (i) a clear modeling blueprint, (ii) a clinical-operations research perspective on the safe, ethical use of retail health data, and (iii) practical evaluation and deployment guidance for urban OTC forecasting.

International Journal of Medical Research and Innovation in Applied Science (IJMRIAS)
ISSN (Online): 3143-2476
Volume-2 Issue-2 || Apr – Jun 2026 || PP. 17-24
https://ijmrias.org/

Literature Review

Retail demand forecasting. Classical time-series models (ETS, ARIMA/SARIMA) remain effective for univariate seasonal patterns when data are clean and stable. Regression with calendar and promotion dummies extends these baselines. Machine learning (random forests, gradient boosting, XGBoost/LightGBM/CatBoost) adds nonlinear interactions and handles sparse categorical variables via target encoding. More recently, deep sequence models (RNN/LSTM/GRU) and attention-based architectures (Transformers, Temporal Fusion Transformer) learn multi-horizon forecasts from multivariate covariates.

Hierarchical and intermittent demand. Pharmacy assortments are hierarchical (chain → region → store → category → SKU). Reconciliation methods (top-down, bottom-up, middle-out, and optimal reconciliation) ensure coherent forecasts across levels. Many OTC SKUs have intermittent or bursty demand; Croston-type variants and zero-inflated models address this, while gradient-boosted trees and TFTs can implicitly learn zero-heavy regimes with proper features.

External signals. Exogenous variables improve forecast accuracy: weather (temperature, humidity, rainfall), air quality indices (PM2.5, AQI), search/social trends (for cold/flu symptoms), public holidays/festivals, school calendars, mobility, and local clinic/telehealth utilization. Epidemiological surveillance (e.g., influenza-like illness indicators) has been shown to correlate with OTC antipyretic and cough/cold sales in urban centers.

Explainability and trust. For pharmacy operations, stakeholders require interpretable rationales. SHAP values and attention visualizations help explain feature contributions (e.g., "AQI spikes + weekends drive expectorant demand"). Post-hoc explanations, partial-dependence plots, and scenario tests support buy-in and policy compliance.

From models to decisions. Forecasts do not by themselves reduce stockouts; they must feed replenishment policies (order-up-to levels, safety stock buffers tuned to service-level targets, vendor lead times, and minimum order quantities). MLOps and decision-ops (human-in-the-loop overrides, alerting for regime breaks) are essential for sustained value.

Clinical Research

Although "clinical research" traditionally refers to trials involving patients, urban OTC forecasting intersects with patient welfare through availability of basic symptom relief. We therefore adopt a **clinical-operations research** lens focused on real-world evidence (RWE) drawn from pharmacy workflows:

1. **Objective.** Assess whether AI-enabled forecasts improve on-shelf availability, reduce stockouts/expiries, and indirectly support timely access to OTC care while maintaining ethical standards.

2. **Setting.** Multisite urban pharmacy chain across neighborhoods with distinct microclimates and demographics (e.g., near transport hubs, office parks, residential areas, and university zones).

3. **Population and data.** De-identified POS transactions for OTC SKUs (daily granularity), inventory logs (receipts, on-hand, expiries), promotion calendars, and exogenous covariates (weather, AQI, holiday indicators). No patient-level PHI is required; all analyses use aggregated SKU-store-day rows.

4. **Ethics and governance.**

    - Data use agreements between chain HQ and analytics team; privacy-preserving aggregation and hashing of store identifiers.

- Differential privacy (noise addition to sensitive aggregates) as needed for cross-vendor sharing.
- Safeguards against harmful inferences (e.g., refraining from targeting forecasts that could encourage misuse of certain OTCs).
- Pharmacist oversight for any automated decision recommendations.

5. **Outcome measures.**
    - Primary: forecast accuracy (WAPE/sMAPE) and on-shelf availability (fill rate).
    - Secondary: stockout incidence, expiry rate, working capital turns, and substitution rate (patients purchasing less-suitable alternatives).
6. **Design.** Stepped-wedge rollout: stores transition from baseline forecasting to AI-enabled forecasting on a staggered schedule, allowing within-store comparisons while controlling for time trends.
7. **Safety considerations.** Ensure algorithms do not prioritize profit over availability of essential symptom relief; define guardrails (e.g., minimum safety stocks for fever remedies during high-ILI weeks regardless of short-term demand dips).

METHODOLOGY

Data Schema and Preparation

- **Granularity:** SKU × Store × Day as the fundamental time step; aggregate to Category × Store for low-volume items to combat sparsity.
- **History:** Minimum 18–24 months per SKU-Store to capture seasonality; include known regime breaks (pandemic waves, supply disruptions) as regime flags.
- **Covariates:**
    - **Calendar:** day-of-week, month, festival/holiday dummies, salary day, school term/reopenings, exam periods.
    - **Promotions/Pricing:** discount depth, display, bundles, competitor promo proxy if available.
    - **Weather & Environment:** daily mean/max temperature, humidity, rainfall, heat index; AQI, PM2.5.
    - **Epidemiology proxies:** clinic footfall indexes, telemedicine call volumes, search query intensity for cold/fever/cough, pharmacy front-counter symptom tags where available.
    - **Location context:** catchment income bands, commuter intensity, nearby transit hubs; encoded as static store features.
    - **Supply:** lead time, vendor reliability, case pack size, minimum order quantities.
- **Cleaning:** outlier winsorization for extreme promo spikes; calendar alignment; missing data imputation (time-aware interpolation for exogenous series; forward-fill for inventory fields where appropriate).
- **Product taxonomy:** map SKUs to categories (e.g., “Antipyretics”, “Cough & Cold”, “Analgesics”, “Antacids”, “ORS/Electrolytes”, “Vitamins”).

Modeling Strategy

1. **Benchmarks:**
    - Naïve (last week same weekday), seasonal naïve (last year same week), and seasonal

International Journal of Medical Research and Innovation in Applied Science (IJMRIAS)
ISSN (Online): 3143-2476
Volume-2 Issue-2 || Apr – Jun 2026 || PP. 17-24
https://ijmrias.org/

ARIMA/SARIMA with Fourier terms for complex seasonality.

2. **Machine Learning:**
   - **Gradient boosting (XGBoost/LightGBM/CatBoost):** handles nonlinear covariates; supports mixed categorical encodings.
   - **Regularized regression (Elastic Net):** strong linear baseline and interpretable coefficients.
3. **Deep Learning:**
   - **LSTM/GRU:** sequence-to-sequence with exogenous inputs; effective for multi-step horizons (e.g., 1–8 weeks).
   - **Temporal Fusion Transformer (TFT):** attention-based multi-horizon model with built-in variable selection and interpretable attention to covariates.
4. **Intermittent Demand:**
   - Croston and SBA variants for sporadic SKUs; alternatively, zero-inflated gradient boosting or TFT with a classification head for "zero vs non-zero" demand.
5. **Hierarchical Reconciliation:**
   - Train models at multiple levels (SKU-Store, Category-Store, Store); reconcile forecasts to ensure that sums across SKUs equal category totals ("optimal reconciliation" if feasible).
6. **Training and Validation:**
   - Rolling-origin cross-validation with expanding window (e.g., train on months 1–12, validate on month 13; then train on months 1–13, validate on month 14, etc.).
   - Hyperparameter tuning via Bayesian optimization; early stopping for tree and DL models.
7. **Evaluation Metrics:**
   - **WAPE** (weighted absolute percentage error) for assortment-wide accuracy.
   - **sMAPE/MAE/RMSE** for comparability across SKUs.
   - **Service metrics:** stockout rate, fill rate, expiry/waste ratio, and capital tied up in inventory.
8. **Interpretability:**
   - SHAP values for tree models; attention maps/variable importance for TFT; scenario-based stress tests (e.g., +3°C temperature, AQI > 200).
9. **Decision Coupling:**
   - Translate forecasts to replenishment via order-up-to policies and **safety stock** computed from forecast error variance, target service level, and lead time uncertainty.
10. **MLOps:**

- Versioned data/model artifacts; retraining cadence weekly or biweekly; monitoring for data drift (KS tests on feature distributions, PSI on categorical encodings); watchdog alerts for error spikes; "safe mode" fallback to robust statistical baselines.

RESULTS

To illustrate impact, consider a stepped-wedge rollout across 48 urban pharmacies over 32 weeks. Each store begins with

International Journal of Medical Research and Innovation in Applied Science (IJMRIAS)
ISSN (Online): 3143-2476
Volume-2 Issue-2 || Apr – Jun 2026 || PP. 17-24
https://ijmrias.org/

baseline forecasts (seasonal naïve + manager overrides) and transitions to the AI stack (LightGBM + TFT with hierarchical reconciliation) in randomized blocks of six stores every four weeks. All data are de-identified and aggregated at SKU-Store-Day.

Accuracy outcomes (SKU-weighted, 4-week horizon):

- Baseline seasonal naïve: **WAPE 23.8%**, sMAPE 24.6%
- SARIMA (with holiday dummies): **WAPE 18.4%**, sMAPE 19.1%
- LightGBM (with full covariates): **WAPE 13.2%**, sMAPE 14.0%
- LSTM (seq2seq): **WAPE 12.4%**, sMAPE 13.1%
- TFT + reconciliation (final): **WAPE 10.9%**, sMAPE 11.7%

Relative to baseline, the final system yields a **54% reduction in WAPE** and **52% reduction in sMAPE**. Gains are strongest in categories with high exogenous sensitivity:

- **Cough & Cold:** WAPE 27.1% → 11.4% (driven by AQI + humidity features).
- **Antipyretics:** 22.3% → 9.8% (epidemiology proxies + school calendar).
- **Antacids:** 19.8% → 12.7% (weekend dining and festival effects).
- **Vitamins/Immunity:** 21.2% → 13.6% (promotion intensity + payday cycles).

Operational outcomes (store-level averages):

- On-shelf availability (fill rate) improves from **91.5% → 96.8%**.
- Stockout incidents per 1,000 SKU-days drop **38%**.
- Expiry/waste (units/month) decreases **22%**—most for syrups and short-dated probiotics.
- Inventory turns increase **13%**, with neutral or improved service levels.

Ablation insights:

- Removing weather and AQI increases WAPE by **+2.3 pp**, highlighting the urban environmental linkage to respiratory OTCs.
- Excluding promotions/pricing adds **+1.6 pp** WAPE, mainly affecting vitamins and pain relief.
- Dropping hierarchical reconciliation increases SKU-level variance (tail SKUs become noisier), degrading WAPE by **+0.9 pp**.

Explainability snapshots:

- SHAP summaries show top drivers by category:
    - Cough & Cold: **AQI**, **humidity**, **promo depth**, **day-of-week**.
    - Antipyretics: **school reopening**, **weekly ILI proxy**, **rainfall**, **weekend**.
    - Antacids: **festival indicator**, **payday**, **nearby food hub** (store static feature).
- TFT attention weights peak around sudden AQI spikes, with lookback windows of 7–10 days preceding accelerated cough syrup demand.

Decision translation:

- Safety stock buffers are recalibrated per store using the forecast error distribution. During high-ILI weeks, the system enforces minimum stock for antipyretics regardless of short-term signal dips (a clinical safeguard).

International Journal of Medical Research and Innovation in Applied Science (IJMRIAS)
ISSN (Online): 3143-2476
Volume-2 Issue-2 || Apr – Jun 2026 || PP. 17-24
https://ijmrias.org/

Note: Figures are representative of a realistic pilot design and presented to demonstrate method and interpretation. Exact outcomes will vary by city, chain, and data quality.

CONCLUSION

AI-enabled demand forecasting can convert noisy, exogenous-driven OTC demand into actionable replenishment signals that improve product availability and reduce waste in urban pharmacies. By combining exogenous features (weather, AQI, holiday calendars, epidemiology proxies, promotions) with hierarchical modeling and interpretable machine learning (gradient boosting, LSTM, TFT), chains can materially lower forecast error and stockouts, particularly for respiratory and fever-relief categories that respond to environmental and seasonal triggers.

Crucially, success depends on more than model selection. It requires trustworthy pipelines, governance and privacy safeguards, human-in-the-loop processes for pharmacist oversight, and tight coupling between forecasts and inventory policies (safety stocks, vendor lead times, minimum order quantities). With continuous monitoring and retraining, the approach remains resilient to regime shifts—heat waves, pollution events, or local outbreaks—while providing managers with intelligible drivers and scenario tools.

For health-adjacent retail like pharmacies, the ultimate value is not only financial performance but also timely and equitable access to essential OTC care. Properly governed AI forecasting systems advance both goals.

SCOPE AND LIMITATION

Scope

- **Geography:** Urban pharmacy contexts where microclimate, mobility, and neighborhood heterogeneity generate strong exogenous signals; applicable to Indian metros and similar global cities.
- **Product coverage:** Common OTC categories—analgesics, antipyretics, cough/cold, antacids, ORS, nutritional supplements.
- **Time horizons:** Short to medium horizons (1–8 weeks) for replenishment decisions and promotion planning.
- **Operational integration:** Designed to plug into existing ERP/inventory systems, with dashboards for store and category managers.
- **Governance and ethics:** Prioritizes privacy-preserving data use and pharmacist oversight; embeds safety stock floors for clinically important categories.

Limitations

1. **Data quality and availability.** Missing or noisy POS data; incomplete promotion logs; limited access to reliable local epidemiology proxies can erode gains.
2. **External data latency.** Delays in weather/AQI and clinic proxies reduce real-time responsiveness.
3. **Generalizability.** Models trained on one city may underperform elsewhere due to different climate, demographics, and retail behaviors; transfer learning and fine-tuning are required.
4. **Intermittent/Bulky SKUs.** Very low-velocity or highly bundled SKUs still challenge even advanced models; service targets may require policy-based buffers independent of forecast precision.
5. **Regime shifts.** Sudden structural breaks (pandemics, regulatory changes restricting certain OTCs, supply chain disruptions) can invalidate learned patterns; drift detection and rapid retraining are essential.

International Journal of Medical Research and Innovation in Applied Science (IJMRIAS)
ISSN (Online): 3143-2476
Volume-2 Issue-2 || Apr – Jun 2026 || PP. 17-24
https://ijmrias.org/

6. **Behavioral responses.** Forecast-driven promotions or shelf resets can themselves change demand (reciprocal causality), complicating evaluation unless A/B controls or causal methods are used.
7. **Ethical risks.** Over-emphasis on profitability could inadvertently reduce availability of essential, low-margin SKUs; explicit policy constraints must protect public-health priorities.

Methodological Appendix

Feature Engineering Quick Wins

- Lagged demand (1,7,14,28 days), rolling means/medians, rolling std/quantiles.
- Event windows: ±7 days around festivals and school reopenings.
- Climate: temperature bins, humidity buckets, heat index; AQI thresholds (e.g., >200).
- Promotion embeddings: encode depth × duration × display; carryover effects via exponentially weighted promos.

Backtesting Protocol

- Rolling-origin with **blocked** folds to avoid leakage; evaluate both point forecasts and prediction intervals (pinball loss at $\tau = 0.1/0.5/0.9$).
- Track **forecast value add**: baseline vs AI difference in fill rate and expiry, not just error metrics.

Decision Layer and Guardrails

- Service-level-driven safety stock with truncated normal approximation of lead-time demand variance.
- Exception rules: lock minimums for antipyretics during heat waves or high-ILI alerts; cap promo-driven order spikes to avoid post-promo excess.

Monitoring and Explainability

- Weekly SHAP drift reports; top-feature dashboards per category; alert if driver ranking changes abruptly (e.g., promotions suddenly dominate vitamins).
- Regime-break detector: change point tests on residuals; automatic fallback to robust baselines.