IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510109

MediCycle AI: A Smart Framework for Safe and Sustainable Redistribution of Surplus Medicines

Mrs.Anjana H S[1], Atharva K R[2], Gnanasagar G N[3], Jeevan P[4], Siddeshwar B V[5]

East West Institute of Technology, Department of Information Science and Engineering Bengaluru, India[1-5]

Abstract: Unused and surplus medicines accumulate across households, hospitals, and non-governmental organizations, creating environmental and public-health burdens while access to essential medicines remains uneven. Take-back and reverse-logistics programmes concentrate on collection and disposal governance, demand forecasting targets procurement, and verification, expiry recognition, and AI decision support are demonstrated in isolation. This paper proposes MediCycle AI, an integrated, auditable framework for the safe, expiry-aware, demand-driven redistribution of eligible surplus medicines.It combines image-based verification with OCR for label and expiry extraction; demand forecasting using random forest and XGBoost, with LSTM only where the dataset justifies it; inventory management with near-expiry detection; a transparent multi-factor recipient-matching and priority-scoring model; and controlled generative and agentic AI layers under retrieval grounding and mandatory human approval. Hard safety constraints—expired or unverifiable medicines are never redistributed—are separated explicitly from soft prioritization factors such as demand, urgency, shelf life, quantity, distance, and recipient priority. The paper presents the architecture, database design, protocol, and evaluation metrics. As of submission, implementation and quantitative results are pending; the manuscript reports an evaluation plan and expected outcomes rather than empirical results, and positions the framework strictly as decision support requiring professional and regulatory validation.

Index Terms: Surplus Medicine Redistribution, Pharmaceutical Reverse Logistics, Demand Forecasting, XGBoost, OCR, Computer Vision, Medicine Verification, Recipient Matching, Generative AI, Agentic AI, Healthcare Supply Chain, SDG.

I. INTRODUCTION

Medicines are wasted at every point of the pharmaceutical lifecycle, yet scarcity persists in the same communities. Review evidence on take-back programmes reports wide variation in structure, facilities, ownership, legislation, and cost, and links accessible collection points, regular schedules, clear ownership, and defined financial responsibilities to better outcomes. Reverse-logistics research frames medication waste as an environmental and public-health issue but stops at collection and disposal rather than redistribution . The enabling technologies exist independently: drug-label identification reaches up to 88% precision using combined image and text embeddings ; expiry recognition reaches 97.74% accuracy with fully convolutional networks; hospital demand forecasting has progressed to hybrid optimized machine learning , knowledge-graph deep learning , and forecast-driven replenishment support ; and AI is increasingly positioned within healthcare supply-chain resilience . Evidence on healthcare LLMs documents summarization, prediction, and administration alongside reliability, bias, and privacy concerns , while agentic-AI surveys emphasize memory, tool use, planning, and alignment . Each capability is demonstrated in isolation, and no study in the verified corpus integrates verification, expiry awareness, forecasting, inventory state, recipient priority, proximity, and matching in one auditable framework. The principal contributions are: C1: an integrated, auditable framework combining verification, forecasting, inventory, matching, and recommendation in one pipeline. C2: a computer-vision-plus-OCR verification and expiry-extraction pipeline that separates package recognition from pharmaceutical authentication. C3: demand forecasting integrated into allocation, converting predictions into expiry-aware, priority-sensitive recommendations. C4: a transparent multi-factor donor–recipient scoring model with documented weight selection, sensitivity analysis, and baseline comparison. C5: a controlled generative and agentic AI layer under retrieval grounding, hard safety constraints, and mandatory human approval.

II. LITERATURE SURVEY

A. Surplus Collection and Pharmaceutical Reverse Logistics

A PRISMA-based review of take-back programmes (27 studies, 15 countries) reports wide variation in structure, facilities, ownership, legislation, and cost [1]. A review of medications reverse logistics with a mixed-method proposal for Brazil identifies medicine waste as an environmental and public-health issue and emphasizes information management and stakeholder roles [2]. Neither provides a demand-aware redistribution mechanism.

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510109

B. Pharmaceutical Demand Forecasting

A hybrid Grey Wolf Optimizer plus XGBoost model for hospital demand forecasting, evaluated on weekly dispensing data from two Thai hospitals, captures temporal and environmental utilization factors [3]; a knowledge-graph approach with Graph Convolutional Networks and LSTM captures inter-drug relationships missed by univariate models [4]. A 2026 preprint extends forecasting into decision support: XGBoost forecasts of 14- and 30-day demand with a VEN volatility buffer reduced stock-out days by 81.5% and ordering days by 71.6%, with MAPE of 7.66%–8.06% at 14 days and 3.97% 4.61% at 30 days [8].

C. Medicine Identification, OCR, and Expiry Recognition

DLI-IT crops label images into text regions with a Connectionist Text Proposal Network, recognizes text with Tesseract OCR, embeds it with sentence embeddings, and retrieves reference labels by cosine similarity; trained on 1,749 opioid and 2,365 non-opioid images, it reached up to 88% precision on 300 external test images [9]. Fully convolutional date detection and parser networks report 97.74% expiry accuracy, with the ExpDate dataset released because no public dataset existed [7]. DLI-IT identifies labels rather than authenticating products, and the expiry framework targets generic packaging—both motivate the pipeline proposed here.

D. Generative and Agentic AI in Healthcare

A PRISMA-based review of conversational LLMs in healthcare (65 of 820 records) found applications in summarization, knowledge inquiry, prediction, and administration, and concluded that the evidence does not establish suitability for unsupervised clinical decisions [5]. A survey of over 200 healthcare-agent studies proposes a lifecycle taxonomy emphasizing safety, trust, and alignment [6]. These justify the retrieval-grounded, human-in-the-loop design adopted here.

E. Healthcare Supply-Chain

AI Peer-reviewed 2026 literature situates AI within healthcare supply-chain management, efficiency, and resilience [10], confirming that redistribution-specific integration remains underexplored. Table I summarizes the verified corpus.

TABLE I SUMMARY OF THE VERIFIED LITERATURE CORPUS

No.	Author(s) [Ref.] / Year	Method	Key Results	Limitations / Relevance
1	Wang, Aziz, Wang, Chik [1] / 2024	Systematic review (PRISMA; 27 studies)	Structural variation; access, schedules, ownership, financing matter	No demand-aware redistribution
2	Lima, Delgado, Santos, Florentino [2] / 2022	SLR + mixed-method Brazilian case	Waste as environmental/public-health issue	Brazil-specific; no integrated platform
3	Samniang et al. [3] / 2026	GWO + XGBoost; weekly dispensing, 2 hospitals	Captures temporal/environmental utilization factors	Two hospitals, one region; forecasting only
4	Chen, Lu, Zhang, Wan [4] / 2026	Knowledge graph + GCN + LSTM	Captures inter-drug relationships and temporal dynamics	High complexity; data-dependent
5	Wang, Yin et al. [5] / 2024	Systematic review; 65 of 820 records	Summarization, inquiry, prediction; bias and privacy concerns	Pre-Sep-2023 evidence; no unsupervised clinical use
6	Xu et al. (preprint) [6] / 2025	Survey; lifecycle taxonomy; 200+ studies	Perception, memory, tool use, planning, collaboration	Preprint; agentic workflow reference
7	Seker, Ahn [7] / 2022	FCN date detection + parser; ExpDate dataset	97.74% recognition accuracy	Generic packaging; pharma testing required
8	İşli, Aydın (preprint) [8] / 2026	XGBoost 14/30-day forecasts; VEN buffer	Stock-out days −81.5%; ordering days −71.6%	Not peer-reviewed; replenishment, not redistribution
9	Liu, Meehan, Tong et al. [9] / 2020	CTPN + Tesseract + sentence embedding	Up to 88% precision; up to 35% gain over single-modality	Identification, not authentication; CV+OCR precedent

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510109

10	Espinosa et al. [10] / 2026	Systematic review (Cochrane guidelines)	AI adoption in healthcare SCM; heterogeneous maturity	Context for AI-enabled health logistics

III. RESEARCH GAP AND MOTIVATION

Collection-focused research stops short of redistribution [1], [2]; forecasting stops short of allocation, targeting procurement and replenishment [3], [4], [8]; and enabling components are proven individually [7], [9], while generative and agentic AI surveys stress that safety requires grounding, reliability, and human oversight [5], [6]. Gap statement: comparatively limited work integrates these capabilities into a unified AI-enabled decision-support framework for safe, expiry-aware, demand-driven redistribution of eligible surplus medicines [1]–[10], a claim supported by the verified corpus without asserting absolute novelty.

IV. PROBLEM STATEMENT AND RESEARCH QUESTIONS

RQ1: how can verification, demand forecasting, and multi-factor matching be integrated into a single auditable redistribution framework? RQ2: how accurately can an object-detection-plus-OCR pipeline identify medicine packages and extract expiry dates? RQ3: can machine-learning forecasting beat naive and moving-average baselines under timebased validation, and does that translate into demand-aware allocation? RQ4: does transparent multi-factor scoring rank recipients better than rule-based, nearest-recipient, expiry-only, and demand-only baselines? RQ5: does integration reduce wastage and improve redistribution efficiency relative to single-criterion processes? RQ6: how can grounded generative and human-constrained agentic AI support decisions without replacing professionals or overriding safety rules?

Problem statement (formal): despite advances in forecasting, computer-vision identification, OCR-based expiry extraction, and healthcare AI decision support, no demonstrated integrated framework safely converts verified surplus medicines into demand-driven redistributions—without substituting for professional or regulatory judgment.

V. OBJECTIVES

O1—evaluate an image-detection-plus-OCR pipeline; report precision, recall, F1, and date-extraction accuracy against an OCR-only baseline. O2—train naïve, moving-average, random-forest, and XGBoost models; report MAE, RMSE, MAPE, and R^2 under time-based validation with leakage controls. O3—implement normalized scoring over demand, urgency, shelf life, quantity, distance, and priority; document weights and sensitivity; compare with four baselines using Precision@K, Recall@K, NDCG@K, and match success rate. O4—deliver the end-to-end system and measure API latency, query time, matching time, throughput, and resource utilization. O5—integrate a retrieval-grounded layer and a monitoring→recommendation→human-approval workflow with audit logging. O6—measure medicines redistributed, waste avoided, redistribution time, distance, and completion rate in controlled backtests, reporting expected results as expected.

VI. PROPOSED METHODOLOGY

A. System Specification

Input: surplus-medicine information (category, batch, manufacturer, quantity, manufacturing and expiry dates, storage condition, location), package images, OCR text, donation records, recipient requests, historical demand, and inventory state. Processing: identification; OCR extraction; expiry validation; eligibility assessment; inventory registration; forecasting; recipient filtering; matching; priority scoring; recommendation; notification; and redistribution tracking. Output: a ranked list of eligible redistribution opportunities requiring human approval.

B. Medicine Eligibility Model

Hard constraints encode safety and legality; any violation renders a medicine ineligible regardless of soft scores: expired medicines must not be redistributed; unverifiable identity is not automatically eligible; and storage-condition and regulatory restrictions must be respected. Soft factors—demand, urgency, distance, shelf life, quantity, and recipient priority—rank candidates and never override hard constraints. The rule set requires qualified professional and regulatory validation; formal WHO/CDSCO citation will be added only after the instruments are verified.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510109

C. Matching and Priority Formulation

Recipient ranking uses a normalized, transparent scoring function:

$$PS(r) = w_1 \cdot D_r + w_2 \cdot U_r + w_3 \cdot S_r + w_4 \cdot Q_r + w_5 \cdot L_r + w_6 \cdot P_r \quad \ldots (1)$$

where D = demand, U = urgency, S = remaining shelf life, Q = quantity suitability, L = distance, and P = recipient priority, each scaled to [0,1]. Weights w_1–w_6 are expert-defined initially and subjected to sensitivity analysis. Arbitrary assignment is avoided by reporting weight values, their rationale, and ranking sensitivity to perturbations; the model is compared against rule-based, nearest-recipient, expiry-only, and demand-only baselines.

VII. SYSTEM ARCHITECTURE

The architecture comprises nine layers: User, Presentation, API, Application Logic, AI/ML, Database, External API, Security, and Deployment. Users (donors, hospitals, NGOs, administrators) interact through the presentation layer; the API layer exposes authenticated endpoints; application logic orchestrates collection, verification, inventory, matching, recommendation, and notification; the AI/ML layer provides vision/OCR, forecasting, matching, and generative/agentic services; state lives in MongoDB Atlas and PostgreSQL; and deployment is containerized on cloud infrastructure. The stack (React/Next.js, FastAPI, MongoDB/PostgreSQL, PyTorch/XGBoost/OpenCV/EasyOCR/YOLOv8, Docker/AWS) is proposed, not claimed as implemented.

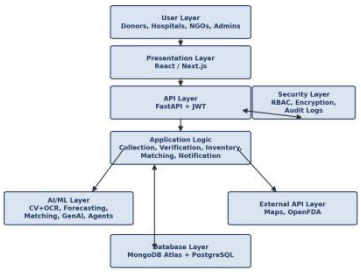


Fig. 1. Overall MediCycle AI architecture and data flow between layers.

VIII. MEDICINE VERIFICATION AND OCR

The pipeline follows a fixed sequence: image capture → preprocessing → detection/classification → package identification → OCR → name/batch/expiry extraction → database cross-check (e.g., OpenFDA) → eligibility decision, extending demonstrated precedents [7], [9]. Evaluation uses precision, recall, F1, OCR word accuracy, expiry-extraction accuracy, and identification accuracy. A critical distinction is maintained: the pipeline recognizes printed packaging, not a medicine's physical integrity or provenance, and image recognition is never presented as proof of authenticity [9]. Matches below the confidence threshold route to human review.

IX. DEMAND FORECASTING

Baselines are naive and moving average; machine-learning models are linear regression, random forest, and XGBoost; deep learning is used only if dataset size and structure justify it [3], [4]. Validation uses strict time-based splits, no future information in training, tuning on validation folds only, and walk-forward validation where series length permits. Features include historical and lagged demand, month, season, location, category, stock level, and disease-season indicators where valid. Metrics are MAE, RMSE, MAPE, and R^2; leakage controls restrict predictors to information available at decision time [8].

X. AI-BASED MEDICINE MATCHING

Matching embeds the scoring function of Section VI-C in a filtering pipeline: candidates must satisfy hard eligibility filters, after which soft scores rank the survivors. Scores, weights, and rationale are exposed to administrators, consistent with evidence that auditable, priority-aware rules improve hospital outcomes [8]. Evaluation uses Precision@K, Recall@K, and NDCG@K [9], plus match success and completion rates in backtests.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510109

XI. RECOMMENDATION ENGINE

The engine generates operational suggestions—suitable recipient, distribution priority, medicine priority, route, donation priority, and, where medically and legally appropriate, alternative options—framed as informational suggestions requiring professional validation. No recommendation implies clinical interchangeability without authoritative medical evidence; the engine is a decision-support instrument, not a prescriber.

XII. GENERATIVE AI COMPONENT

The generative layer is a retrieval-grounded module with bounded functions: inventory summaries, explanations of recommendations, redistribution reports, demand-trend summaries, and query assistance, following evidence that healthcare LLMs help with summarization and administration but raise reliability, bias, and privacy concerns requiring oversight [5]. Guardrails: the model must not prescribe, diagnose, declare medicines interchangeable, or override expiry, safety, or regulatory rules; groundedness, factuality, privacy, and security are evaluated; and every output is logged.

XIII. AGENTIC AI COMPONENT

A multi-agent workflow provides decision support under human control: Inventory → Expiry Monitoring → Demand Forecasting → Recipient Matching → Priority → Recommendation → Notification Agent → Human Approval → Redistribution. Each agent observes state, reasons over predefined rules and data, calls authorized tools, and produces recommendations only; no agent holds autonomous clinical or high-risk authority, following evidence that healthcare agents require perception, memory, tool use, planning, and safety alignment [6]. Agent goals, tools, guardrails, and auditability are documented per agent.

Dataset	Source	Access	Records (Verified)	Limitations
DLI-IT labels	DailyMed [9]	Public download	1,749 opioid; 2,365 nonopioid; 300 external test	Identification, not authentication
ExpDate	Seker & Ahn [7]	Released by authors	Not reported	Generic packaging; pharma testing required
İşli & Aydın panel	Retrospective hospital records (preprint) [8]	Not publicly released	35,801 unit–item–day records; 1,269 series	Preprint; replenishment context
MediCycle project dataset	To be created	Project-specific	Not yet collected	Requires ethical and regulatory vetting

TABLE II DATASETS AND THEIR VERIFIED STATUS

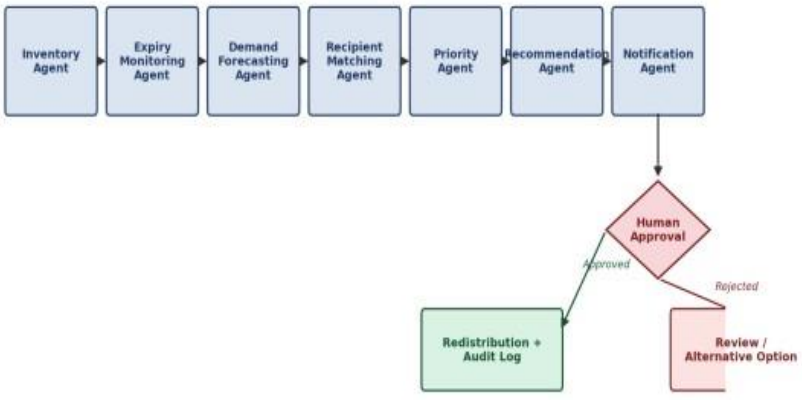


Fig. 2. Agentic AI workflow with mandatory human approval before redistribution.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510109

XIV. DATABASE AND SYSTEM IMPLEMENTATION

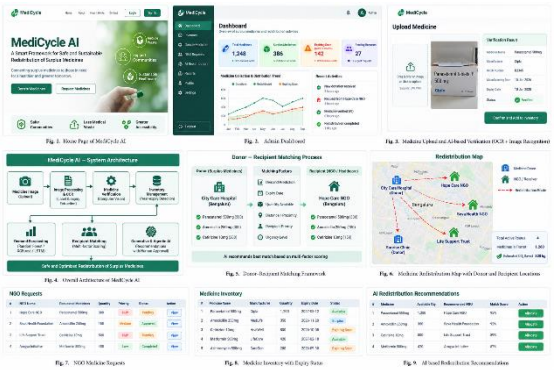
The schema comprises Users, Organizations, Medicines (name; category; manufacturer; batch), Donations (donor; medicines; quantity; dates; storage), Inventory (medicine; quantity; expiry; location; status), RecipientRequests (organization; needs; urgency; quantity), DemandHistory, Forecasts (medicine; location; horizon; value; version),Matches (donation; request; score; rank; weights), Redistributions (match; status; route;timestamps), notifications, and AuditLogs (actor; action; timestamp; outcome). Compound indexes cover (status, expiry) for inventory and (date, location, category) for demand history.Implementation is proposed with JWT authentication, role-based access control, TLS, encrypted fields, and immutable audit logging; deployment scale is not reported.

XV. EXPERIMENTAL SETUP

Experiment 1—identification/OCR: labeled package-image dataset (to be created); YOLOv8 + EasyOCR; precision, recall, F1, word accuracy; baseline OCR-only. Experiment 2—expiry extraction: date-region detection and parsing, benchmarked against [7]. Experiment 3—forecasting: naive, moving average, random forest, and XGBoost under timebased validation. Experiment 4—matching: four baselines vs. multi-factor scoring using Precision@K, Recall@K, NDCG@K, and success rate. Experiment 5—recommendation quality: expert rating of appropriateness. Experiment 6—end-to-end latency, throughput, and utilization.

XVI. RESULTS AND DISCUSSION

No experimental results have been generated as of submission; this section is therefore an evaluation plan, and no accuracy, latency, or outcome figures are reported. Expected outcomes are stated as hypotheses: (H1) the multi-factor matcher will achieve higher NDCG@K and success rate than single-criterion baselines; (H2) random-forest and XGBoost models will outperform naive and moving-average baselines in MAE/RMSE; (H3) the ablation will show shelf life and demand as the largest contributors to ranking performance; and (H4) verification latency will remain within operational bounds. Tables III–VI (forecasting accuracy, verification and OCR results, matching results, and system performance) are defined and will be populated only from executed experiments; all four remain pending.



XVII. COMPARISON WITH EXISTING APPROACHES

MediCycle AI is contrasted with four baselines: (A) manual/rule-based redistribution; (B) nearest-recipient matching; (C) expiry-only priority; and (D) demand-only matching. B ignores demand and expiry; C ignores demand and distance; D ignores expiry and urgency; the multi-factor model combines all six factors under hard constraints [8]. For forecasting, naive and moving-average baselines are compared with random forest, XGBoost, and conditional LSTM [3]; for verification, OCR-only with CV+OCR [9]. Comparisons become quantitative only after Experiments 3–6.

XVIII. SUSTAINABILITY AND SDG ALIGNMENT

SDG 3—verified surplus medicines reach populations with documented demand; SDG 12—redistribution extends medicine lifetime and reduces waste [1]; SDG 10—demand- and priority-aware allocation favors organizations serving underserved populations; SDG 9—the framework demonstrates AI-enabled healthcare logistics infrastructure [10]. Quantified SDG impact would require deployment-scale measurement and is therefore proposed, not claimed.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510109

XIX. SECURITY, ETHICAL AND REGULATORY CONSIDERATIONS

Security provisions include JWT authentication, role-based access control, password hashing, TLS, encrypted sensitive fields, data minimization, immutable audit logs, and documented provenance. Ethics and safety provisions include separation of hard constraints from soft scoring; no redistribution of expired or unverifiable medicines; explicit recognition that vision and OCR identify packaging, not authenticity [9]; grounded, non-prescriptive, human-reviewed generative outputs [5]; agentic actions gated by human approval [6]; and bias and explainability monitoring. This paper claims no regulatory approval and no clinical validation.

XX. LIMITATIONS

Acknowledged limitations: dataset availability and quality; synthetic demand data unless real records are obtained; limited image and packaging diversity; OCR errors on complex labels; no direct authenticity testing; forecasting uncertainty for intermittent demand; regulatory complexity of medicine transfer; the need for professional oversight; unknown generalization to real hospitals and NGOs; and the absence of peer review for two cited preprints ([6], [8]).

XXI. FUTURE WORK

Planned extensions: larger real-world datasets; multimodal verification; real-time demand signals; advanced time-series models; knowledge graphs [4]; explainable AI; IoT storage monitoring; blockchain or audit infrastructure where justified; route optimization; multi-agent coordination; and prospective field evaluation. Future features are described as future work, not completed functionality.

XXII. CONCLUSION

MediCycle AI is proposed as an integrated, auditable decision-support framework for the safe, expiry-aware, demanddriven redistribution of eligible surplus medicines, addressing an integration gap across take-back governance, verification, forecasting, and allocation research [1]–[4], [7]–[9]. Its design separates hard safety constraints from soft prioritization factors, grounds generative and agentic AI under human oversight [5], [6], and commits to a reproducible protocol with rank-sensitive, error-based, and operational metrics. Implementation and results are pending; claims are limited to design, with quantitative evaluation deferred to the planned experiments. The work contributes a pathway toward reducing medicine wastage and improving access, subject to professional and regulatory validation.

REFERENCES

[1]. L. S. Wang, Z. Aziz, E. S. Wang, and Z. Chik, “Unused medicine take-back programmes: a systematic review,” J. Pharm. Policy Pract., vol. 17, art. 2395535, 2024, doi: 10.1080/20523211.2024.2395535.

[2]. P. A. B. Lima, F. C. M. Delgado, T. L. dos Santos, and A. P. Florentino, “Medications reverse logistics: a systematic literature review and a method for improving the Brazilian case,” Cleaner Logist. Supply Chain, vol. 3, art. 100024, 2022.

[3]. W. Samniang, S. M. T. Shah, Y. M. Tun, A. Munawar, and K. Sarin, “A hybrid grey wolf optimized eXtreme gradient boosting-based machine learning model for hospital pharmaceutical demand forecasting,” Sci. Rep., 2026.

[4]. X. Chen, G. Lu, H. Zhang, and J. Wan, “Knowledge graph-enhanced deep learning for pharmaceutical demand forecasting,” Sci. Rep., 2026, doi: 10.1038/s41598-026-35113-4.

[5]. L. Wang, Z. Yin, et al., “Applications and concerns of ChatGPT and other conversational large language models in health care: systematic review,” J. Med. Internet Res., vol. 26, art. e22769, 2024.

[6]. G. Xu et al., “A comprehensive survey of AI agents in healthcare,” TechRxiv preprint, Nov. 14, 2025.

[7]. A. C. Seker and S. C. Ahn, “A generalized framework for recognition of expiration dates on product packages using fully convolutional networks,” Expert Syst. Appl., vol. 203, art. 117310, 2022.

[8]. [Author(s) as listed in the preprint], “Enhancing hospital drug and medical supply request processes through a clinically prioritized and demand-variability-aware adaptive decision support framework,” SSRN preprint, 2026. (Not peer-reviewed.)

[9]. X. Liu, J. Meehan, W. Tong, L. Wu, X. Xu, and J. Xu, “DLI-IT: a deep learning approach to drug label identification through image and text embedding,” BMC Med. Inform. Decis. Mak., vol. 20, art. 68, 2020.

[10]. O. Espinosa, S. Basto, A. Ordóñez, M.-L. Arias, D. Sepúlveda, and M. F. Bhutta, “A review of artificial intelligence in healthcare supply chains: untapped potential?,” BMC Health Serv. Res., vol. 26, art. 742, 2026.