

Synchronous Field Optimization for Pharmaceutical Logistics: A Governed Enterprise-To-Edge Ai Architecture for Predictive Field Intelligence

¹Dr. Sunita Arjun Yadav and ^{2*}Dr. Marielia Assumption

^{1,2} Somaiya School of Basic and Applied Sciences , Somaiya Vidyavihar university

Received: 14th June, 2026; Revised: 20th July 2026; Accepted: 20th August, 2026; Available Online: 05th September 2026

ABSTRACT

However, as pharmaceutical logistics is increasingly connected with sophisticated computation systems, traditional push-type linear models are limited by the challenges associated with heterogeneity of data, delay in feedback, fragmentation in work flows, and growing demands for traceability and governance. The present paper introduces the Synchronous Field Optimization (SFO) model, which involves a distributed closed loop between an enterprise predictive intelligence core and the field edge nodes. In turn, the proposed framework combines structured operation data with unstructured field data using natural language processing (NLP), retrieval-augmented generation (RAG), predictive modeling, and constrained Next-Best-Action (NBA) optimization. The utility function formally represents the expected benefit, operation costs, regulatory risks, privacy/security risks, and uncertainty in data quality. For testability, the proof-of-concept framework experiment consisting of 20,000 records was designed and evaluated based on the predictive discrimination, calibration, and constrained action-utility criteria. In the simulated environment, the increasing richness of SFO features resulted in better predictive discrimination compared to structured data, and constrained NBA generated greater mean utility than rule-based push. These findings are representative of architecture behaviors in an ideal synthetic setting, but do not imply anything about the real-world efficacy of pharmaceutical applications. The paper also describes a scenario-driven 30-year roadmap that covers multimodal and edge AI (2026-2035), synthetic patient digital twins (2036-2045), and decentralized digital manufacturing scenarios (2046-2056). The paper's contribution is methodological, and pertains to architectural contributions in the form of SFO's governed enterprise-to-edge orchestration model.

Keywords: pharmaceutical logistics; artificial intelligence; Next-Best-Action; natural language processing; retrieval-augmented generation; edge computing; digital twins; supply-chain intelligence; responsible AI; design science.

How to cite this article: Sunita Arjun Yadav, Marielia Assumption, Synchronous Field Optimization for Pharmaceutical Logistics: A Governed Enterprise-To-Edge Ai Architecture for Predictive Field Intelligence *Int J Drug Deliv Technol.* 2026;16(84s): 1986-1995; DOI: 10.25258/ijddt.16.84s.201

Source of support: Nil.

Conflict of interest: None

1. INTRODUCTION

The pharmaceutical distribution and field activity is increasingly becoming a data-driven phenomenon. There is a need for coordination of issues like demand, stock levels, product quality, shortages, limitations in access, evidence dissemination, regulations, and decision-making on multiple levels. Artificial intelligence (AI) in the realm of supply chain management has been analyzed in many papers to reveal its application to areas like forecasting, logistics, production, and others, while studies devoted to the health care supply chain confirm the comparative immaturity of AI implementation. The integration of artificial intelligence into organizational operations can support predictive decision-making, supply-chain coordination, operational-risk reduction, and more effective technology-enabled management (Singh & Kaur, 2025). Artificial intelligence and machine-learning techniques are increasingly being applied in healthcare to support automated disease detection, classification, and early diagnosis from medical images (Pandey, 2025). This suggests the research potential in that area because the

field needs not only predictions but also execution and feedback.

Within the context of the study, the notion of medicine circulation describes a process wherein pharmaceuticals, operational data, evidence, field intelligence, and decisional cues are circulated within the delivery system of pharmaceuticals. The notion is wider than the idea of movement within the supply chain. It involves communication channels wherein enterprise-level systems can exchange information with their field-based counterparts and vice versa.

The original SFO framework involved an enterprise-level AI at its core, field-level nodes, a synthesis loop using NLP technology, and a 30-year-long computational trajectory. The current manuscript preserves the core of the SFO but reinterprets some of the claims for better scientific justification. Particularly, the current model does not make assumptions regarding autonomous clinical decision-making, unrestricted ambient surveillance, automatic model mutation, and replacement of physical

**Author for Correspondence: Dr. Marielia Assumption*

logistics with future technologies. These factors are presented only as design constraints and potential scenarios.

The case for SFO is also reinforced by regulatory guidelines in recent times. In January 2026, the U.S. Food and Drug Administration (FDA) and the European Medicines Agency (EMA) issued ten guiding principles for good artificial intelligence practices in drug development, focusing on human-centered design, risk-based approach, standards, context of use, multi-disciplinary skills, data management, model development, risk-based performance assessment, life cycle management, and essential information. These principles apply directly to an AI system for the pharmaceutical industry that goes beyond retroactive analysis.

1.1 Research Problem

The problem in research is the lack of an explicitly formulated and managed architecture for linking heterogeneous data within enterprises with field information, unstructured language processing, prediction, decision-making, and feedback within one pharmaceutical operation cycle. The problem is not merely the amount of data. Data is scattered between systems of different temporal resolution, structure, ownership, and governance. SFO solves the problem by formulating the data-to-action-to-feedback process with well-defined representations, control boundaries, and evaluation criteria for each transition within the process.

1.2 Research Questions

- **RQ1.** How can an enterprise-to-edge architecture improve the timeliness and contextualization of pharmaceutical field intelligence?
- **RQ2.** How can constrained Next-Best-Action classification be incorporated into field workflows while preserving human oversight and regulatory boundaries?
- **RQ3.** How can NLP and evidence retrieval transform permitted unstructured field information into auditable predictive signals?
- **RQ4.** What measurable performance and governance criteria are required to evaluate SFO before operational deployment?
- **RQ5.** How might the SFO architecture evolve toward digital-twin and decentralized-manufacturing scenarios over a 30-year horizon?

1.3 Contributions

- An enterprise-to-edge SFO architecture for pharmaceutical field intelligence.
- A constrained mathematical formulation for NBA selection.
- A governed NLP/evidence-retrieval pipeline for converting permitted unstructured information into model-ready signals.

- A lifecycle feedback model incorporating monitoring, drift detection, validation, and controlled model updating.
- A synthetic proof-of-concept validation demonstrating how progressively richer information layers can be evaluated without exposing patient or proprietary data.
- A scenario-based roadmap distinguishing emerging technologies from speculative long-horizon possibilities.

1.4 Novelty Positioning

NLP, RAG, predictive analytics, edge computing, NBA systems, and digital twins are individually established or actively researched technologies. The novelty claimed here is therefore not the invention of an individual algorithm. Rather, the contribution is the system-level orchestration of these components for pharmaceutical field operations through a governed bidirectional feedback architecture, together with a formal decision objective and a reproducible validation protocol.

2. BACKGROUND AND RELATED WORK

2.1 AI in Healthcare and Pharmaceutical Supply Chains

AI has been explored in the realms of forecasting, inventory control, routing, anomaly detection, supply chain risk, and production planning. According to an overview of AI applications in healthcare supply chain management published in 2026, the field has much potential, but at present is still underdeveloped in terms of implementation research. In the case of pharmaceutical supply chain, there are additional issues associated with product integrity, expiration, shortages, temperature control, and compliance. This makes the use of a pure prediction-based approach inadequate.

2.2 NLP and Unstructured Information

Notes written in free text form may include context-dependent information that cannot be included in a structured note. A systematic literature review on clinical note natural language processing revealed a heavy reliance on machine learning methods while noting difficulties in relation extraction, temporal reasoning, interpretability, and access to data. Thus, in SFO, unstructured information is considered a probabilistic signal rather than the ground truth.

2.3 Retrieval-Augmented Generation

RAG can connect a language model to external evidence repositories and thereby reduce dependence on static model parameters. A 2025 systematic review of healthcare RAG found a wide range of retrieval and generation methods but also a lack of standardized evaluation frameworks and limited assessment of ethical considerations. SFO consequently treats retrieval as a governed evidence layer with source attribution, versioning, retrieval timestamps, and validation.

2.4 Next-Best-Action Optimization

NBA systems rank candidate actions according to expected utility. Within SFO, candidate actions are operational rather than autonomous therapeutic decisions.

Examples include prioritizing an authorized follow-up, retrieving an approved evidence package, escalating an access issue, or coordinating inventory information. The framework explicitly excludes autonomous prescribing or modification of approved clinical instructions.

2.5 Edge Computing and Distributed Operations

Edge computing can place selected processing near the point of data generation. In SFO, the edge node is a constrained execution layer supporting local preprocessing, approved retrieval, lightweight inference, caching, and synchronization. It is not an independent clinical authority. Device identity, encryption, least-privilege access, secure software updates, model-version control, and audit logging are required.

2.6 Digital Twins

Use of digital twins has become more evident within the process of drug discovery, development, manufacturing and precision pharmacotherapy. Reviews from 2026 have highlighted possibilities but have also noted data integration, model robustness, privacy concerns and regulatory acceptance as important roadblocks. Hence, the SFO roadmap sees the concept of digital twins as an evolving one and not a ready replacement for clinical trials or physical medicine distribution.

3. METHODOLOGY

3.1 Research Design

This research employs a design science approach. The process involves problem identification, objective identification, architectural design, formal decision modeling, experimentation, evaluation, and communication. The experimentation process here is a proof of concept designed to test if the suggested architecture is computable and if there are any differences produced by the architectural components designed.

3.2 System Boundary

The SFO framework comprises enterprise operational data, authorized field information, edge computing, NLP and evidence retrieval, predictive models, NBA optimization, outcome tracking, and model life cycle management. It does not involve self-diagnosis, treatment, labeling modification of authorized products, or independent clinical decision making.

3.3 Data Representation

Let x_t denote structured enterprise data, $e_{i,t}$ localized edge information, $u_{i,t}$ permitted unstructured

observations, $z_{i,t}$ an NLP-derived representation, $p_{i,t}$ a predicted probability, $a_{i,t}$ a candidate action, and $y_{i,t}$ an observed outcome. The NLP transformation is represented as:

$$z_{i,t} = \text{NLP}(u_{i,t})$$

$$p_{i,t} = f_{\theta}(x_t, e_{i,t}, z_{i,t})$$

The predictive function may be implemented using logistic regression, gradient-boosted trees, neural networks, transformer encoders, or another validated method. Model selection should depend on context of use, calibration, interpretability, latency, and governance requirements.

3.4 Constrained Next-Best-Action Objective

The recommended action is defined as:

$$a^*_{i,t} = \arg\max_{a \in A_t} E[U(a | x_t, e_{i,t}, z_{i,t})]$$

A constrained utility formulation is:

$$a^*_{i,t} = \arg\max_a [U(a) - \lambda_1 C(a) - \lambda_2 R(a) - \lambda_3 D(a) - \lambda_4 Q(a)]$$

$U(a)$ denotes expected operational benefit; $C(a)$ operational cost; $R(a)$ regulatory or policy risk; $D(a)$ privacy/security risk; and $Q(a)$ uncertainty or data-quality penalty. λ parameters represent approved operating weights. Hard constraints can be represented as $g_j(a, x) \leq 0$, with infeasible actions removed before ranking.

3.5 Governed Feedback and Model Updating

Following an authorized action:

$$D_{t+1} = D_t \cup \{(x_t, e_{t,z_t}, a_{t,y_t})\}$$

A governed update can be represented as:


$$\theta_{t+1} = \text{Update}(\theta_t, D_{t+1})$$

This formulation does not imply unrestricted self-modification. Data-quality checks, drift detection, offline validation, subgroup evaluation, approval gates, version control, rollback procedures, and post-deployment monitoring are required before a new model is released.

4. SFO ARCHITECTURE

SFO architectural model comprises five interconnected layers that include enterprise AI core, field edge, NLP/evidence layer, NBA layer under constraints, and outcome monitoring. The architecture transforms field activities from the reporting layer into the controlled learning signal source.

Figure 1. Synchronous Field Optimization architecture



4.1 Enterprise AI Core

The AI core of the enterprise is the coordination hub for validated structured data, model definitions, approved evidence stores, policies, monitoring, and model evaluation. The AI core must ensure lineage from raw data to features, prediction, action, and result.

4.2 Field Edge Node

The edge node executes the following tasks locally: text processing, authorized search, transformation of features, inference, caching, and synchronization. The proposed architecture makes the assumption that local processing can help increase efficiency and avoid transfer of unnecessary data over the network, yet at the same time, it acknowledges the additional burden of cybersecurity and versioning.

4.3 NLP and Evidence Layer

The NLP/evidence layer converts validated unstructured observations to structured form and finds relevant approved information. The system implementation should have the following components: data minimization, entity extraction, intent classification, temporal processing, embeddings, retrieval, sourcing, and output validation. RAG should not be considered an unrestricted source of clinical expertise.

4.4 Human-in-the-Loop NBA

SFO makes a distinction among recommendations to be shown, recommendations to be confirmed, and actions needing specialized escalation. Overrides should be recorded as repeat overrides would suggest a problem with model specification or data.

4.5 Monitoring

The monitoring layer collects the information regarding acceptance/rejection of recommendations, reason codes when available, outcome measures, latency, quality of

data, data drift, policy exceptions, and security incidents. It is important to monitor both accepted and rejected recommendations because otherwise there will be a case of selection bias.

5. SYNTHETIC PROOF-OF-CONCEPT VALIDATION

In order to transform SFO from a mere theoretical construct into something that can be tested empirically, we created a dataset of 20,000 records which are artificial, not patient-specific or drug-specific but operational observations. Some of the variables used are region, HCP category, days in inventory, trend of demand, access barrier, engagement, evidence relevance, artificial NLP score, data quality, outcome, and action taken.

5.1 Synthetic Data-Generation Assumptions

The synthetic outcome probability was generated from a nonlinear logistic process combining operational, engagement, access, text-derived, evidence, and data-quality factors. Candidate actions were assigned utilities with penalties for cost, risk, and data-quality uncertainty. The simulation was seeded for reproducibility. Because the data-generating process is known by construction, the results should be interpreted as a controlled demonstration of model behavior, not as an estimate of real-world pharmaceutical performance.

5.2 Predictive Models

Four feature configurations were evaluated using logistic regression: (1) Baseline ML using structured operational features; (2) ML + NLP adding the synthetic text-derived signal; (3) ML + NLP + Evidence adding evidence relevance; and (4) Full SFO adding data-quality information. A stratified 70:30 train-test split was used with a fixed random seed.

5.3 Predictive Performance

Table 1. Predictive performance on the synthetic held-out test set.

Model	AUROC	AUPRC	F1	Brier
Baseline ML	0.689	0.677	0.633	0.223
ML + NLP	0.694	0.681	0.642	0.222
ML + NLP + Evidence	0.696	0.682	0.641	0.221
Full SFO	0.696	0.682	0.642	0.221




Figure 2. AUROC comparison across synthetic SFO feature configurations.

5.4 Next-Best-Action Evaluation

A rule-based push strategy was compared with the full constrained SFO action-selection mechanism. The rule-based strategy prioritized inventory coordination when inventory was below ten days, access escalation when an

access barrier was present, evidence support when evidence relevance exceeded a threshold, and otherwise selected follow-up. The SFO strategy ranked five candidate actions using the constrained utility function.

Table 2. Synthetic action-selection comparison.

Strategy	Mean constrained utility	Agreement with oracle
Rule-based push	0.584	0.775
Full SFO (synthetic simulation)	0.602	0.991





Figure 3. Mean constrained utility in the synthetic action-selection experiment.

5.5 Calibration

Calibration was assessed with the Brier score. Lower values indicate better probabilistic calibration. The comparison is useful because an operational NBA system

should not only rank observations correctly but also produce probabilities that can be interpreted as uncertainty estimates.

**Figure 4.** Brier-score comparison; lower values indicate lower probabilistic error.

5.6 Interpretation of the Synthetic Results

The synthetic experiment demonstrates the intended directional behavior of the architecture: adding contextual signals to structured features changes predictive performance, and the constrained action layer can be evaluated separately from the prediction layer. The best AUROC among the simulated configurations was 0.696 for Full SFO. The full SFO action-selection simulation achieved a mean constrained utility of 0.602, compared with 0.584 for the rule-based strategy. These numerical results are generated from synthetic assumptions and therefore should not be generalized to real pharmaceutical settings.

6. DISCUSSION

The most important finding of this work is a methodological one: SFO can be separated into individually testable levels of operation, instead of being seen as a black box AI model. Prediction performance can be measured with discrimination and calibration measures; NLP and evidence retrieval can be evaluated using extraction, retrieval, attribution, and factuality measures; NBA can be evaluated using action utility and acceptability; and the whole architecture can be evaluated using operational and governance measures.

Finally, the synthetic experiment shows why feedback-based architecture cannot be defined only in terms of accuracy of prediction. In the end, a deployed system chooses actions with constraints. Choosing an action that gives high benefit according to prediction but that has high cost, high policy risk, high privacy risk, or uncertainty might be incorrect.

The boundaries of RAG are carefully defined. Although RAG for healthcare has been shown to have a number of promising use cases, there is a lack of standardization in data and assessment, as well as large variability between datasets and their evaluation. SFO must therefore ensure provenance and verification. An operational version of the tool will need to evaluate not only generation quality but also retrieval recall, relevancy of source documents, attribution accuracy, contradictions, and output faithfulness.

Organization is yet another factor. Even a robust technical implementation of the recommendation system can be hindered by high field workload, reduced autonomy, and surveillance of performance. SFO must therefore assess such human factors as workload, trust, usability, overrides, and perceived usefulness.

7. COMPARISON WITH CONVENTIONAL AND PARTIAL ARCHITECTURES

Table 3. Conceptual comparison of architectural capabilities. The table describes functional scope, not empirical superiority.

Architecture	Prediction	NLP	Edge	NBA	Feedback	Governance
Traditional push/CRM	Limited	Limited	No	Rule-based	Delayed	External
Supply-chain ML	Yes	Usually no	Optional	Partial	Partial	Variable
RAG assistant	No	Yes	Optional	No	No	Variable
Field analytics	Yes	Yes	Sometimes	Partial	Partial	Variable
SFO	Yes	Yes	Yes	Yes	Yes	Native

8. GOVERNANCE, ETHICS, SECURITY, AND REGULATORY CONSIDERATIONS

The reason why the design of governance is required in SFO is that it involves sensitive clinical, professional, and operational information. WHO guidelines on the use of AI in health care point out the importance of human autonomy, privacy, transparency, accountability, equity, and good governance. In order to build SFO, the design should minimize data collection; separate identifiers from feature analysis when feasible; ensure access control by roles; secure data transmission and storage; and maintain logs.

The regulatory context of an AI component will depend on the context of use. The January 2026 guiding principles of FDA/EMA point out the importance of context of use, risk-based approach, data governance, model development, performance evaluation, and lifecycle management. The final ICH M15 of FDA of 2026 on model-informed drug development points out the importance of planning, evaluation of model, documentation, and evidence evaluation. All this highlights the importance of specifying the purpose, evidence, and performance of an SFO component.

8.1 Privacy and Ambient Data

Ambient data may suggest total, passive surveillance and is not recommended for use in a production specification where the method of data gathering and its legal basis

have yet to be defined. SFO utilizes the more limited term “authorized field signals.” Capture of voice or text must comply with appropriate consents and restrictions on privacy, employment, etc.

8.2 Bias and Fairness

Predictive algorithms may recreate past inequalities. SFO should measure performance on all the important subgroups and see whether there is a systematic difference in recommendations based on any proxies or missingness patterns. Fairness assessment should be formulated based on the use case.

8.3 Cybersecurity

The distributed nature of edge nodes creates more attack surfaces. Controls that are necessary include the identity of the device, configuration, communications security, least privilege access, software and model signing, updates, anomaly detection, and fast revocation. Blockchain, if applicable, should never be used to replace traditional cybersecurity controls.

9. 30-YEAR SCENARIO ROADMAP

The initial manuscript suggested three phases based on time. The updated one still sticks to the same format but views it as scenario analysis, not a prediction. The usefulness of the roadmap lies in highlighting the factors and questions which might become relevant as the technology advances.

SFO 30-year scenario roadmap (2026-2056)

Era I Multimodal & Edge AI	Era II Digital Twins	Era III Decentralized Manufacturing Scenario
-------------------------------	-------------------------	---

2026-2035

2036-2045

2046-2056

Figure 5. Scenario-based SFO roadmap from 2026 to 2056.

9.1 Era I — Autonomous Multi-Modal and Edge AI (2026–2035)

The topics of this case study include multimodal AI, RAG under control, edge intelligence, inventory prediction, and field co-pilots. The realistic path for research should be

incremental in nature: first, automating information gathering, then reducing administrative work, enhancing predictions, and finally offering actions that can be ranked in an understandable manner.

9.2 Era II — Synthetic Patient Digital Twins (2036–2045)

This scenario transitions from operational prediction to personalized simulation. Digital twin models could utilize longitudinal clinical data, pharmacokinetics/pharmacodynamics modeling, genomics, and monitoring data. Current research shows the way forward but highlights integration, accuracy, security, and regulatory adoption as significant challenges. It is important not to interpret this scenario as an expectation that physical clinical trials will become extinct.

9.3 Era III — Decentralized Digital Manufacturing Scenario (2046–2056)

Convergence between secure digital provenance, distributed manufacturing, bioprinting, and computer-based biological design is described in the last scenario. There are enough pieces of evidence for distributed and local manufacturing to warrant scientific investigation, but the bold hypothesis that physical logistics of pharmaceuticals will be substituted with encryption of molecular codes lacks proof. This is why this scenario is more theoretical in nature and must be taken as such.

10. IMPLEMENTATION AND VALIDATION ROADMAP

- State the context of use and disallowed uses for each SFO element.
- Identify structured, semi-structured, and unstructured data sources.
- Apply data lineage, provenance, retention, and access policies.
- Design a governed NLP pipeline with quantified extraction efficacy.
- Design an approved source RAG with version control and attribution.
- Train and tune predictive models via temporally appropriate validation.
- State the NBA action candidate set and hard constraints.
- Define utility weightings via multi-disciplinary stakeholder input.
- Design human verification and escalation policies.
- Protect edge devices and model update processes.
- Detect drift, bias, latency, override, and policy exceptions.
- Advance from synthetic validation to retrospective validation and then controlled prospective testing.

11. LIMITATIONS

- The empirical experiment involves using synthetic data and hence cannot establish real-world pharmaceutical efficacy.

- The synthetic data generation process has been developed by the researchers and does not necessarily reflect real world operational distributions or dependencies.

- The utility weights used in the proof-of-concept are indicative and need to be calibrated empirically.

- Historical observational data can be biased; the predictive improvement cannot establish real-world efficacy.

- There are risks of extraction failure, retrieval failure, hallucinations, and biases in NLP and RAG.

- Stages 2036-2045 and 2046-2056 of the roadmap involve considerable uncertainty scientifically, economically, technologically, and regulation-wise.

- The framework is operational in nature; the claim about clinical efficacy requires separate validation.

12. FUTURE RESEARCH

- Create an open synthetic benchmark for pharmaceutical field intelligence workflows.

- Benchmark tree-based, neural, and transformer models versus logistic regression.

- Create a true RAG system on a closed corpus of approved material.

- Benchmark retrieval recall, attribution correctness, factuality, and contradictions.

- Consider causality or quasi-experimentation in evaluating NBA recommendations.

- Evaluate edge-to-enterprise sync, model versioning, and failure resilience.

- Conduct human factors studies of workload, trust, usability, and overrides.

- Create formal model cards, context-of-use documentation, auditing frameworks, and governance.

- Investigate interoperability with pharmaceutical and healthcare data standards.

- Research whether digital twin components can be included in SFO without conflating operational intelligence and clinical decision support.

13. CONCLUSION

Synchronous Field Optimization is the proposed architecture for governing enterprise-to-edge computing applied to field intelligence and medicine circulation. The framework integrates enterprise analysis, localized processing, natural language processing and evidence discovery, Next-Best-Action optimization with constraints, and outcome tracking in a single information loop. The principle of SFO is that pharmaceuticals AI infrastructure is to optimize not only prediction but actions themselves as well.

The proof-of-concept is synthetic in nature and designed to illustrate the possibility to test this principle without use of patient data. It does not provide evidence of effectiveness of SFO in practice; the experiment just sets up a reproducible computational pattern to compare structured prediction, contextual evidence, evidence relevancy and constrained action selection.

The timeline from 2026-2056 expands on the current framework that is evolving into multimodal and edge artificial intelligence and further on to digital twins and decentralized manufacturing. It is too early to tell whether these scenarios will become a reality; nevertheless, they should be considered as potential research hypothesis. In general, SFO allows bridging data science architecture and pharma operational architecture from theory to practice.

REFERENCES

1. Singh, A. R., & Kaur, A. AI adoption in agriculture: A managerial framework for sustainable agribusiness development. *Journal of Artificial Intelligence in Agriculture and Applied Research*, 2025;1(1):1–12. <https://agrpjournals.com/journals/journal/index.php/jaiaar/article/view/2/2>
2. Pandey P. Artificial intelligence for early detection of dental caries from radiographic images: a machine learning. *Journal of Dentistry and Clinical Sciences*. 2025;1(1):31-39. <https://agrpjournals.com/journals/journal/index.php/jdacs/article/view/12/11>
3. Toorajipour R, Sohrabpour V, Nazarpour A, Oghazi P, Fischl M. Artificial intelligence in supply chain management: a systematic literature review. *J Bus Res*. 2021;122:502-517. doi:10.1016/j.jbusres.2020.09.009.
4. Espinosa O, Basto S, Ordóñez A, Arias ML, Sepúlveda D, Bhutta MF, et al. A review of artificial intelligence in healthcare supply chains: untapped potential? *BMC Health Serv Res*. 2026;26:742. doi:10.1186/s12913-026-14559-2.
5. Sheikhalishahi S, Miotto R, Dudley JT, Lavelli A, Rinaldi F, Osmani V. Natural Language Processing of Clinical Notes on Chronic Diseases: Systematic Review. *JMIR Med Inform*. 2019;7(2):e12239. doi:10.2196/12239.
6. Agüero D, Nelson SD. The Potential Application of Large Language Models in Pharmaceutical Supply Chain Management. *J Pediatr Pharmacol Ther*. 2024;29(2):200-205. doi:10.5863/1551-6776-29.2.200.
7. Amugongo LM, Mascheroni P, Brooks S, Doering S, Seidel J. Retrieval augmented generation for large language models in healthcare: a systematic review. *PLOS Digit Health*. 2025;4(6):e0000877. doi:10.1371/journal.pdig.0000877.
8. Liu S, McCoy AB, Wright A. Improving large language model applications in biomedicine with retrieval-augmented generation: a systematic review, meta-analysis, and clinical development guidelines. *J Am Med Inform Assoc*. 2025;32(4):605-615. doi:10.1093/jamia/ocaf008.
9. Venkatapurapu SP, Clegg L, Nowojewski A, Kimko H, Olabode D, Sawant-Basak A, Vishwanathan K. Digital twins for accelerating drug discovery and development: opportunities and challenges. *Drug Discov Today*. 2026;31(2):104617. doi:10.1016/j.drudis.2026.104617.
10. Ecker A, Szabó G, Szalma J, Fichó E, Reguly I, Csikász-Nagy A. Digital twin approaches for interpretable side effect prediction in drug discovery. *Drug Discov Today*. 2026;31(3):104637. doi:10.1016/j.drudis.2026.104637.
11. Eadie AL, Fernandez Lynch H, Scheinerman N, Parikh RB. The arrival of digital twins and in silico trials in drug development. *Nat Med*. 2026;32(6):1967-1971. doi:10.1038/s41591-026-04344-3.
12. Niarakis A, Moingeon P. Digital twins to accelerate target identification and drug development for immune-mediated disorders. *FEBS Open Bio*. 2026;16(6):1048-1059. doi:10.1002/2211-5463.70228.
13. US Food and Drug Administration; European Medicines Agency. Guiding Principles of Good AI Practice in Drug Development. January 2026.
14. US Food and Drug Administration. M15 General Principles for Model-Informed Drug Development. Final Level 1 Guidance. June 2026.
15. US Food and Drug Administration. Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products: Draft Guidance for Industry and Other Interested Parties. January 2025.
16. World Health Organization. Ethics and governance of artificial intelligence for health: WHO guidance. Geneva: WHO; 2021.
17. World Health Organization. Ethics and governance of artificial intelligence for health: guidance on large multi-modal models. Geneva: WHO; 2025.
18. US Food and Drug Administration. Distributed Manufacturing and Point-of-Care Manufacturing of Drugs. Discussion Paper. Silver Spring, MD: FDA; 2022.
19. Pournader M, Ghaderi H, Hassanzadegan A, Fahimnia B. Artificial intelligence applications in supply chain management. *Int J Prod Econ*. 2021;241:108250. doi:10.1016/j.ijpe.2021.108250.

APPENDIX A. SYNTHETIC DATASET AND REPRODUCIBILITY

A synthetic CSV dataset accompanies this manuscript. It contains 20,000 records generated with a fixed random seed (42). The dataset is intended solely for

methodological demonstration. Researchers can reproduce the proof-of-concept using the documented variable definitions, train-test split, logistic-regression models, and constrained utility formulation.

Variable	Type	Role
region	categorical	Synthetic geographic grouping
hcp_type	categorical	Synthetic provider category
inventory_days	continuous	Operational inventory proxy
demand_trend	continuous	Demand-change proxy
access_barrier	binary	Operational access signal
engagement	continuous	Synthetic interaction signal
evidence_relevance	continuous	Evidence retrieval relevance proxy
synthetic_nlp_signal	continuous	Derived unstructured-text proxy
data_quality	continuous	Data-quality confidence proxy
outcome	binary	Synthetic target
nba_action	categorical	Selected constrained action
nba_utility	continuous	Constrained action utility