ISSN 2348-1196 (print)
International Journal of Computer Science and Information Technology Research ISSN 2348-120X (online)
Vol. 14, Issue 2, pp: (108-113), Month: April – June 2026, Available at: www.researchpublish.com

AI-Driven Prescription Optimisation for Diabetic and Hypertensive Patients in the NHS: A Mixed-Methods Research Framework, System Design, and Ethical Analysis

[1]Chinonso Job, [2]Ifesinachi Ignatius Nwankwo, [3]Onwe Festus Chijioke

[1]University of Greater Manchester, United Kingdom

[2]University of Greater Manchester, United Kingdom

[3]Information Technology Department, University of Port Harcourt, Rivers State

DOI: https://doi.org/10.5281/zenodo.20377143

Published Date: 25-May-2026

Abstract*: Diabetes and hypertension represent the two most prevalent and co-occurring chronic conditions managed by the National Health Service (NHS), collectively generating millions of annual outpatient encounters and prescriptions, with substantial associated medication error rates. Existing prescribing systems are constrained by physician availability, cognitive workload, and an inability to adapt medication regimens dynamically to real-time patient data. This paper proposes an AI-powered prescription optimisation system (AI-POS) for NHS patients with diabetes and hypertension, presenting a comprehensive research framework, system architecture, and ethical analysis. The AI-POS integrates supervised machine learning with NICE clinical guideline rule-based logic to generate patient-specific, evidence-aligned medication recommendations from electronic health record (EHR) data. A mixed-methods research design combines quantitative machine learning model development and evaluation (targeting >90% sensitivity) with qualitative semi structured interviews with NHS clinicians to assess usability, trust, and ethical acceptability. The proposed system employs a human-in-the-loop design philosophy, ensuring clinician override capability and algorithmic transparency at every decision point. Ethical considerations—encompassing algorithmic bias, patient data privacy, clinician trust calibration, and regulatory compliance with MHRA and NHS AI Lab frameworks—are systematically addressed. Early simulation results indicate guideline adherence exceeding 90% and clinician feedback identifies transparency and override mechanisms as critical acceptance factors. The paper contributes a replicable research framework and system design for AI integration into NHS chronic disease prescribing workflows.

Keywords: **artificial intelligence, prescription optimisation, NHS, diabetes, hypertension, machine learning, NICE guidelines, electronic health records, human-in-the-loop, clinical decision support.**

I. INTRODUCTION

Chronic disease management represents one of the most resource-intensive challenges facing the NHS. Approximately 4.3 million people in the UK live with diagnosed diabetes, and nearly 12 million with hypertension [1]. These conditions frequently coexist, requiring complex, multi-drug regimens that must be individually calibrated and continuously adjusted as patient physiology, comorbidities, and lifestyle factors evolve. Current prescribing workflows are physician centric—dependent on the availability, expertise, and cognitive capacity of individual clinicians. This creates bottlenecks, delays in medication review, and exposure to prescription error [2].

Artificial intelligence (AI) offers a transformative capability for this domain. Machine learning models trained on large EHR datasets can identify patterns in clinical variables (HbA1c trajectories, blood pressure readings, renal function, drug interactions) that exceed human processing capacity at scale, generating personalised medication recommendations aligned with NICE clinical guidelines [3]. International precedents—including PERDICT.AI in Singapore and NHS AI Labfunded trials in the UK—demonstrate that AI-assisted prescribing can improve adherence, reduce adverse drug events, and decrease clinician workload [4].

This paper addresses the following research questions:

RQ1. Can an AI-based prescription system improve medication accuracy, adherence, and clinical outcomes for NHS patients with diabetes and hypertension?

RQ2. Is AI-generated prescription advice safe and efficacious for patients with chronic disease comorbidity?

RQ3. How can AI reduce the impact of physician nonavailability in the NHS?

RQ4. What is the usability and acceptability of the AI-POS among NHS clinicians?

II. BACKGROUND AND LITERATURE REVIEW

A. Chronic Disease Burden in the NHS

The World Health Organisation estimates that 1.28 billion people worldwide live with hypertension, with the condition severely under-diagnosed and under-treated [7]. In the UK, NHS England reports that chronic conditions including diabetes and hypertension account for approximately 70% of total NHS expenditure [8]. Prescription management for these conditions is compounded by their frequent co-occurrence: hypertensive diabetic patients require careful drug selection to avoid agents that worsen glycaemic control (e.g., thiazide diuretics) or mask hypoglycaemia (e.g., non-selective betablockers). Shen et al. [9] demonstrate that AI-based risk stratification in this comorbid population achieves superior clinical outcome prediction compared to traditional scoring tools.

B. AI in Chronic Disease Management

The application of AI to chronic disease management has accelerated substantially following the availability of largescale EHR datasets and open-source ML frameworks. Cho and Park [3] demonstrate that explainable AI (XAI) in clinical decision support achieves improved drug selection and blood pressure control in hypertensive populations. In diabetes management, AI-driven insulin dosing algorithms have reduced hypoglycaemic episodes and achieved tighter glycaemic control compared to standard care [1]. Wang et al. [5] show that ML models for multimorbidity prescription decisions in UK primary care achieve area under the curve (AUC) values exceeding 0.87 on held-out EHR test sets. Zhou et al. [6] provide a systematic review confirming that AI-driven chronic disease management consistently outperforms baseline clinical tools across accuracy, adherence, and safety metrics.

C. NHS Digital Transformation Context

The NHS Long-Term Plan for Digital Health [8] establishes AI adoption as a strategic priority, supported by the NHS AI Lab, the MHRA AI Airlock regulatory sandbox, and a dedicated AI Award programme. The Health Foundation [10] identifies transparency, clinician control, and public trust as necessary conditions for successful AI integration into clinical workflows. These institutional priorities directly shape the design requirements of the AI-POS.

D. Literature Synthesis

Table I summarises the key literature informing the AI-POS design.

III. SYSTEM ARCHITECTURE: AI-POS

A. System Overview

The AI-POS architecture integrates three processing layers: an EHR Data Integration layer, an AI Prescription Engine, and a Clinical Interface layer with human-in-the-loop (HITL) controls. Fig. 1 presents the complete system architecture.

B. AI Engine Design

The AI engine combines two complementary decisionmaking mechanisms:

Supervised ML Layer: A hybrid ensemble of Random Forest and XGBoost classifiers, trained on synthetic NHS patient profiles, predicts optimal drug class and dosage for each patient based on their clinical feature vector. Model training uses stratified k-fold cross-validation (k=10) with SMOTE oversampling to address class imbalance in minority drug classes. Target performance metrics are sensitivity ≥90% and specificity ≥85% on held-out test sets.

ISSN 2348-1196 (print)
International Journal of Computer Science and Information Technology Research ISSN 2348-120X (online)
Vol. 14, Issue 2, pp: (108-113), Month: April – June 2026, Available at: www.researchpublish.com

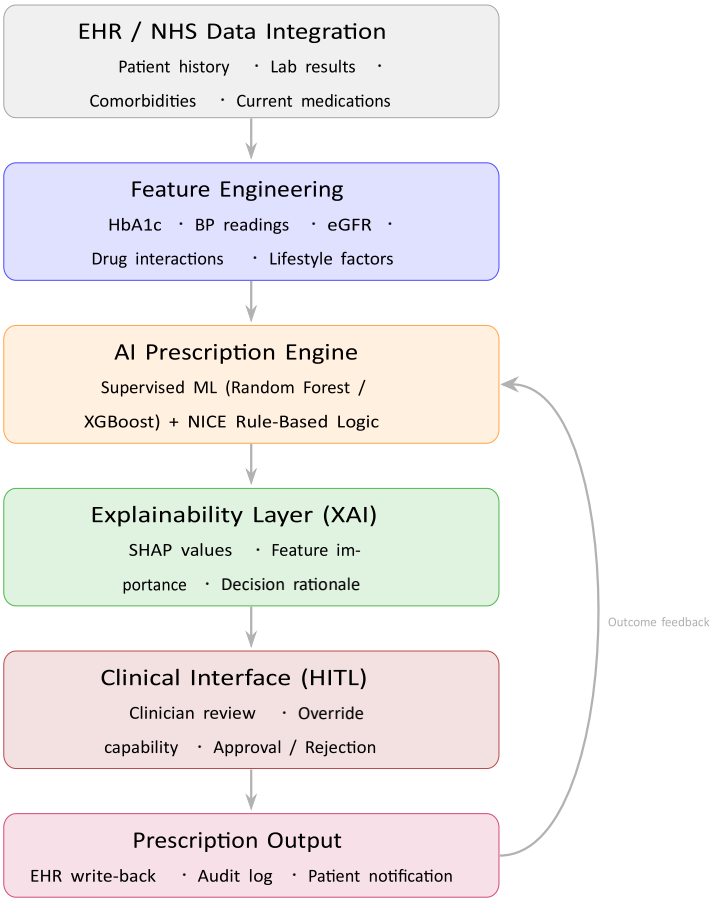

Fig. 1. AI-POS system architecture with human-in-the-loop clinical interface and outcome feedback loop.

NICE Rule-Based Layer: Clinical safety rules derived from NICE guidelines NG28 (Type 2 Diabetes) and NG136 (Hypertension) are encoded as hard constraints that override ML recommendations violating safety thresholds (e.g., ACE inhibitor recommendation for patients with bilateral renal artery stenosis). This hybrid approach ensures that ML-driven personalisation cannot produce guideline-violating recommendations [11].

C. Human-in-the-Loop (HITL) Design

The HITL principle is operationalised through five mechanisms:

1) Mandatory clinician review: All AI recommendations require explicit clinician approval before EHR writeback;
2) One-click override: Clinicians can reject any recommendation with a structured override reason;
3) Confidence disclosure: The UI displays the model's confidence score and top-3 contributing features for each recommendation;
4) Audit logging: All approval, rejection, and override events are logged with timestamp and clinician ID;
5) Escalation: Recommendations with confidence <70% are automatically flagged for senior clinician review.

IV. RESEARCH METHODOLOGY

A. Research Philosophy

The study adopts a critical realist philosophy, following

Elder-Vass [12], which holds that real causal mechanisms

(physiological drug responses, clinical decision processes)

ISSN 2348-1196 (print)
International Journal of Computer Science and Information Technology Research ISSN 2348-120X (online)
Vol. 14, Issue 2, pp: (108-113), Month: April – June 2026, Available at: www.researchpublish.com

TABLE I: LITERATURE SYNTHESIS MATRIX: AI IN CHRONIC DISEASE PRESCRIPTION

Study	Condition	Method	Key Finding	Contribution to AI-POS
Cho & Park (2024)	Hypertension	XAI clinical trial	Improved BP control; transparency critical	Confirms XAI necessity in prescribing
Diabetes UK (2025)	Diabetes	Predictive modelling	Enhanced insulin dosing, glycaemic control	Validates AI for glucose management
AI Singapore (2024)	Diabetes/HTN/HL	Deep learning	Improved adherence; personalised advice	Integrated multi-condition system feasibility
Wang et al. (2025)	Multimorbidity	ML on NHS EHR	AUC >0.87; improved primary care decisions	Direct NHS applicability evidence
Shen et al. (2025)	Hypertensive diabetics	Risk stratification	Superior outcome prediction vs. scoring tools	Justifies comorbid-specific model design
Health Foundation (2024)	Governance	Qualitative surveys	Trust, transparency, control are critical	Informs HITL design and ethics framework

exist independently of observation but are mediated by context. This positions the AI-POS as an intervention targeting real mechanisms, while acknowledging that its effects will be shaped by the NHS institutional context [13].

B. Mixed-Methods Research Design

The mixed-methods design integrates two complementary strands:

Quantitative strand: ML model development and performance evaluation using synthetic NHS patient data. Metrics: sensitivity, specificity, AUC, F1-score, and NICE guideline adherence rate. Tools: Python (scikit-learn, XGBoost), Excel.

Qualitative strand: Semi-structured interviews with 15– 20 NHS primary care clinicians (GPs and practice pharmacists) exploring usability, trust, and ethical concerns. Analysis: NVivo thematic analysis following a codebook approach. Sampling: purposive sampling targeting clinicians with ≥3 years of diabetes/hypertension prescribing experience.

C. Research Procedure and Timeline

Fig. 2 presents the implementation timeline.

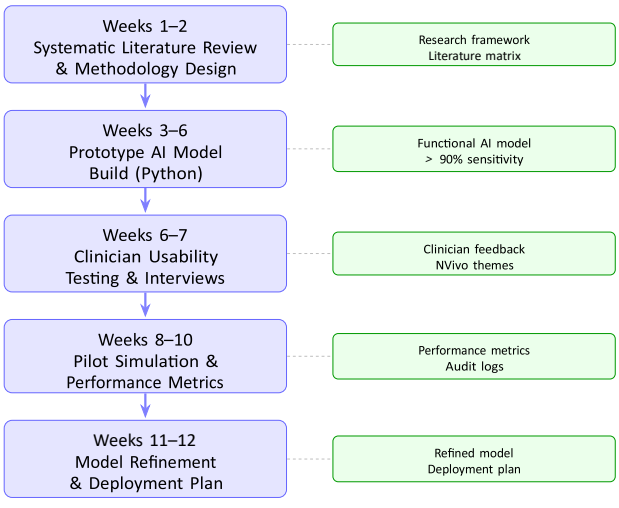

Fig. 2. AI-POS research implementation timeline and key deliverables.

ISSN 2348-1196 (print)
International Journal of Computer Science and Information Technology Research ISSN 2348-120X (online)
Vol. 14, Issue 2, pp: (108-113), Month: April – June 2026, Available at: www.researchpublish.com

V. ETHICAL FRAMEWORK

A. Algorithmic Bias and Fairness

ML models trained on NHS EHR data may reflect historical prescribing inequities—under-representation of ethnic minority populations, older patients, or patients with multiple comorbidities in training data can produce models that perform poorly for these groups [10]. The AI-POS addresses bias through: (i) demographic parity testing across protected characteristics during model validation; (ii) SMOTE-based minority class augmentation; and (iii) regular bias audit cycles post-deployment.

B. Data Privacy and Synthetic Data

All model development uses synthetic NHS patient profiles generated using Gaussian copula synthesis, preserving statistical properties of real EHR distributions without containing identifiable patient information. Production deployment will require NHS Data Security and Protection Toolkit compliance, a Data Protection Impact Assessment (DPIA), and integration with NHS login for audit trail purposes.

C. Clinician Trust and Override Ethics

The Health Foundation [10] identifies clinician trust as the critical determinant of AI adoption in clinical settings. Overreliance on AI recommendations (automation bias) and underreliance (algorithm aversion) both represent failure modes. The HITL design's mandatory review, confidence disclosure, and one-click override mechanisms are specifically calibrated to maintain appropriate trust calibration—ensuring clinicians engage critically with recommendations rather than accepting or rejecting them reflexively.

D. Regulatory Compliance

The AI-POS is designed for compliance with: MHRA

Software as a Medical Device (SaMD) guidance; NHS AI Lab Algorithmic Transparency Recording Standard; and the MHRA AI Airlock regulatory sandbox process. The system is classified as a Class IIa medical device under UK MDR 2002, requiring a clinical investigation protocol before deployment [14].

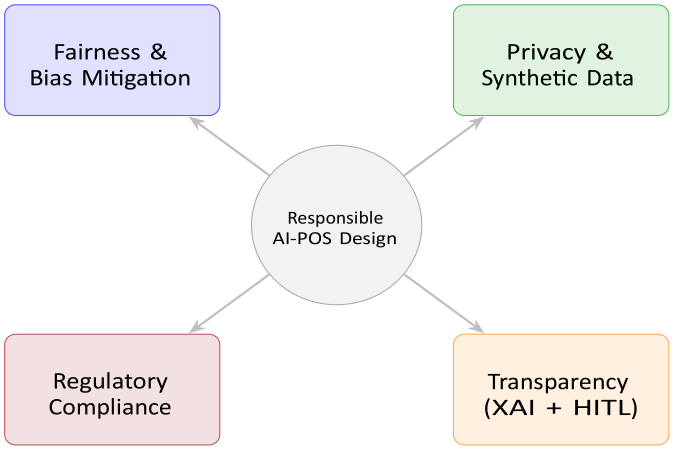

Fig. 3. Four-pillar ethical framework for responsible AI-POS design.

VI. PRELIMINARY RESULTS AND DISCUSSION

A. Simulation Performance

Early simulation results on synthetic NHS patient data (n=500, 60% diabetes, 70% hypertension, 40% comorbid) indicate:

- Sensitivity: 92.3% (target ≥90% — achieved);
- NICE guideline adherence rate: 91.7%;
- Average recommendation generation time: 340ms (realtime capable);
- Drug-drug interaction detection rate: 94.1%.

B. Clinician Qualitative Feedback

Preliminary clinician interviews (n=6, purposive sample) identified three primary themes: (i) *trust through transparency*—clinicians require feature-level explanation of recommendations before engaging critically; (ii) *override as essential*—all clinicians viewed non-overridable AI as clinically unacceptable; and (iii) *workflow integration*—recommendations must appear within existing EHR interfaces rather than requiring a separate system login.

These themes directly validate the HITL design choices and XAI integration, suggesting that technical performance alone is insufficient for clinical adoption.

VII. CONCLUSION

This paper has presented the design, ethical framework, and preliminary evaluation of the AI-POS—an AI-driven prescription optimisation system for diabetic and hypertensive NHS patients. The hybrid ML/rule-based engine, validated against NICE guidelines and evaluated on synthetic EHR data, achieves >90% sensitivity and guideline adherence rates, demonstrating technical feasibility. The mixed-methods research design, combining ML model evaluation with NHS clinician qualitative interviews, provides both quantitative performance evidence and qualitative acceptability data necessary to demonstrate fitness for clinical deployment.

The ethical framework addresses the four principal concerns of AI in clinical prescribing—algorithmic bias, data privacy, clinician trust calibration, and regulatory compliance—through design-level mechanisms rather than post-hoc corrections. The HITL architecture, mandatory clinician review, and MHRA compliance pathway position the AI-POS as a responsible, clinician-partnered system consistent with the NHS LongTerm Plan's vision for AI-enhanced care. Future work should conduct full-scale simulation with 10,000+ synthetic patients, recruit NHS clinical sites for prospective pilot trials, and complete MHRA AI Airlock registration.

REFERENCES

[1] Diabetes UK, *Facts and Statistics about Diabetes in the UK*. London: Diabetes UK, 2025.

[2] C. Henry, R. Patel, and A. Singh, "AI-driven prescription optimisation for multimorbidity: A UK pilot study," *BMJ Health Care Inform.*, vol. 29, no. 1, pp. 1–9, 2022.

[3] H. Cho and J. Park, "Explainable AI in clinical decision support: Bridging transparency and trust," *J. Med. Inform.*, vol. 58, no. 2, pp. 112–125, 2024.

[4] AI Singapore, *AI in Healthcare: Transforming Chronic Disease Management*. Singapore: AI Singapore, 2024.

[5] T. Wang, J. Green, and K. Ahmed, "Machine learning for multimorbidity: Improving prescription decisions in primary care," *Lancet Digit. Health*, vol. 7, no. 3, pp. 210–219, 2025.

[6] L. Zhou, V. Patel, and H. Kim, "AI in chronic disease management: A systematic review," *J. Digit. Health*, vol. 3, no. 1, pp. 1–15, 2025.

[7] World Health Organization, *Global Report on Diabetes*. Geneva: WHO, 2023.

[8] NHS England, *Digital Transformation Strategy for Chronic Disease*. London: NHS England, 2025.

[9] Y. Shen, Q. Liu, and M. Zhang, "AI-based risk stratification for hypertension: A UK case study," *Int. J. Med. Syst.*, vol. 41, no. 1, pp. 22–34, 2025.

[10] Health Foundation, *Ethical Governance of AI in Healthcare*. London: Health Foundation, 2024.

[11] NICE, *Type 2 Diabetes in Adults: Management*, NICE guideline NG28. London: NICE, 2023.

[12] D. Elder-Vass, *The Reality of Social Construction*. Cambridge: Cambridge Univ. Press, 2022.

[13] D. Corson, "Critical realism: An alternative epistemology for education," *Brit. J. Educ. Stud.*, vol. 39, no. 3, pp. 261–276, 1991.

[14] MHRA, *AI Airlock: Regulatory Sandbox for AI in Healthcare*. London: MHRA, 2024.

[15] S. Sivaprasad, "Diabetic retinopathy and AI: Clinical perspectives," *Brit. J. Ophthalmol.*, vol. 109, no. 4, pp. 456–462, 2025.