

A Data-Driven Predictive Framework for Reducing Medication Non-Adherence Cost in the U.S. Healthcare System

Author Name & Affiliations

Akash Kumar
MBA Business Analytics, Barney
School of Business, University of
Hartford, USA
akumar@hartford.edu

Maham Ali Siddiqui
Pharm-D, Ziauddin University, PK,
MBA University of Hartford, USA
msiddiqui@hartford.edu

Syed Jafar
Pharm-D, MHA Healthcare
Management, School of Health
Sciences, University of New Haven,
USA
sjafa1@unh.newhaven.edu

Kishor Kumar Somasekhar Nair
Masters in Computer Science,
College of Engineering Technology
and Architecture (CETA), University
of Hartford, USA
knair@hartford.edu

Article Processing History:

Received: 02 Jul 2026
Revised: 28 Aug 2026
Published: 30 Sep 2026

Index Terms:

Medication adherence, predictive analytics, machine learning, ensemble learning, healthcare economics, health informatics, risk stratification, explainable AI, chronic disease management, U.S. healthcare.

Abstract

Recent estimates from the healthcare industry peg the annual cost burden due to non-adherence in the United States at USD 290 billion to USD 500 billion [1], [2] of direct costs, which is one of the largest avoidable cost drivers. This paper presents a 5-step data-driven pipeline that takes as input patient demographic, clinical, behavioral and access-related data, and processes it through feature engineering for 18 variables, followed by 3 complementary base learners (logistic regression, random forest and gradient boosting), a soft voting ensemble of the 3 base learners and a final output that includes both a tiered risk classification per patient and an estimated avoidable annual cost under a targeted intervention policy. The framework was tested on a sample of $N = 215$ adult patients who were on a chronic-disease pharmacotherapy at U.S. ambulatory clinical sites, where 20.0 % had a proportion-of-days-covered (PDC) less than the 0.80 threshold established by the Pharmacy Quality Alliance (PQA) for HEDIS. The soft-voting ensemble had a receiver operating characteristic (ROC) area under the curve (AUC) of 0.839, an accuracy of 0.831, and an F1 score of 0.560, and the five-fold cross validated ROC-AUC was 0.776 on a 30 % stratified hold out set. The five features with the greatest predictive power in both feature importances were health literacy, out-of-pocket copayment, patient self-efficacy, depression burden, and side-effect burden in the gradient boosting and random forest ranking. Under a conservative intervention-effectiveness assumption of 40 % and intervention cost of USD 320 per patient per year, the cost-scenario analysis of the ensemble-flagged non-adherent patients suggests a net annual saving of USD 2,564 per captured patient – a return of the targeted-intervention policy of an order of magnitude. It is interpretable in terms of the ranking of features in Gradient-Boost, relatively easy to implement on top of a standard electronic-health-record pipeline, and it can be audited by its risk tier per patient. The framework will be expanded to a prospective multi-site evaluation and to the fairness-across-subgroups analysis that is becoming a more routine requirement of current U.S. regulatory guidance for future work.

INTRODUCTION

Medication non-adherence (defined operationally as a Proportion of Days Covered (PDC) below the 0.80 threshold set by Pharmacy Quality Alliance (PQA) and the U.S. Healthcare Effectiveness Data and Information Set (HEDIS) [3]) is among the biggest avoidable cost drivers in the U.S. healthcare system. The direct economic cost burden amounting to approximately USD 290-500 billion annually when the costs of hospitalization, visiting the emergency department, repeat outpatient visits and the cost of preventable long-term complications are combined [1], [2]. Within the first year of treatment, about 50% of patients taking chronic medications discontinue them, miss or take them incorrectly [4], [5]. This leaves older adults with multiple chronic conditions, Medicare / Medicaid enrollees, and patients from racial / ethnic minorities with compounding issues of cost-related non-adherence, limited health literacy and decreased primary care access [6], [7].

There is a large pool of recent literature that argues for a proactive approach in which limited resources for intervention – such as outreach by pharmacists, mobile reminders, synchronized refill programmes, co-payment relief – are targeted first to the patients at highest risk of non-adherence before the clinical and cost consequences have become apparent [8], [9]. This task of proactive identification fits well with machine learning (ML): The problem is a supervised binary classification problem with a large and expanding feature space derived from the electronic health records, insurance claims, pharmacy refill data, and patient-reported outcomes [10], [11]. A number of studies published between 2020 and 2024 have shown that ensemble ML models can predict medication adherence with ROC-AUC values ranging from 0.75 to 0.85 on claims based cohorts [12], [14], and that the incorporation of behavioral and access variables (copayment, pharmacy distance, self-efficacy, health literacy) further improves performance [15], [16].

It is less developed in the literature, however, an end-to-end framework that (i) codifies the 5-step data-driven adherence-prediction pipeline in a reusable way, (ii) quantifies the gradient-boosting / random-forest / logistic-regression ensemble on a cohort of size small enough to be representative of a single-site ambulatory practice (N on the order of 200-300 patients), (iii) ranks the clinical, behavioral and access variables together on explainability, and (iv) translates the classification output into a dollar-denominated cost-savings estimate under an explicit targeted-intervention policy. These four gaps are taken up in the present paper.

This paper makes three contributions. First, it outlines a five-step data-driven predictive model, which can be implemented on top of a standard EHR pipeline: ingestion, feature engineering, base-learner training, soft-voting ensemble, and risk/cost output. Second, it tests the framework on a data set of 215 U.S. adult patients undergoing chronic-disease pharmacotherapy and provides classification performance metrics (ROC-AUC, PR-AUC, accuracy, precision, recall, F1), a ranking of features from the gradient boosting classifier and a cost scenario analysis. Third, it translates the framework into the context of the HEDIS reporting system and considers implications for deployment in an ambulatory setting in the United States.

The rest of the paper is organized as follows. Section II summarizes the work related to machine learning approaches to medication-adherence prediction and cost-of-nonadherence estimates in the U.S. context. The proposed five-stage framework is provided in Section III. The description of the cohort and features and pre-processing procedure are described

in Section IV. The modelling methodology is detailed in Section V. The results are reported in Section VI for the four classifiers presented and the cost-scenario analysis. Section VII covers implications, limitations and future work. Section VIII concludes.

RELATED WORK

COST OF MEDICATION NON-ADHERENCE IN THE U.S.

It is estimated that approximately USD 290-500 billion per year are spent on direct healthcare costs due to medication non-adherence in the United States [1], [2]. A significant portion of that burden reflects cost-related non-adherence (CRNA) – the portion of patients who miss, split, or delay doses due to out-of-pocket costs – that has been shown in the largest national representative surveys of cardiovascular-risk populations in the USA [6]. The issue of polypharmacy among older people with multiple chronic illnesses exacerbates the situation – older people with 5 or more chronic medicines are over-represented in adherence-related hospital admission and ED visits [17], [18]. The results justify the use of a cost denominated output in the suggested output structure, and the selection of insurance type, copayment amount and the number of polypharmacy as feature variables.

MACHINE LEARNING FOR ADHERENCE PREDICTION

The adherence-prediction problem has been the subject of many recent studies using machine-learning techniques. Using a large U.S. health-plan dataset with 1.2 million patients, Lu et al. [12] developed an ensemble of deep-learning and classical-learning classifiers and obtained ROC-AUCs near 0.80 for the classification of the twelve-month PDC. Mason et al. [13] used gradient boosting and random forest models to train on a cohort of people with atherosclerotic cardiovascular disease and concluded that non-clinical features (depression, loneliness, socioeconomic position) were meaningfully added to the clinical baseline. In the study by Zakeri et al. [14] the methodological selections were systematized for sixteen studies on cardiovascular-disease adherence prediction and the following recommendations were proposed: to use an ensemble of classifiers, to perform cross-validated hyperparameter search and to treat the imbalanced non-adherence outcome explicitly through class-weight treatment. Koesmahargyo et al. [19] revealed that claims-based feature and visual markers captured using a smartphone can complement each other in psychiatric cohorts. Each one of these studies justifies the particular selection of components in the present paper: logistic regression, random forest, gradient boosting, class weight balancing, AUC (cross valid), clinical, behavioral and access variables.

EXPLAINABLE AND DEPLOYABLE ML IN HEALTHCARE

It has been argued that clinically viable applications of healthcare ML systems must be explainable and auditable [20], [21]. Post-hoc interpretability techniques such as model-agnostic SHAP and LIME and permutation importance have become widely adopted [22]. The specific patient's reason for being identified as high-risk is crucial when designing interventions in the medication-adherence context: If the reason is copayment, a prescription-benefits intervention is appropriate; if the reason is health literacy, a patient-education intervention is appropriate [23]. In this paper, the top feature-important output from the gradient boosting classifier is reported as the main explainability output.

DIGITAL-HEALTH INTERVENTIONS AND HEDIS REPORTING

There are several different modalities for which interventions targeting the high-risk subset flagged by the ML have been evaluated: pharmacist-led outreach, mobile health reminders, synchronized refill programs and bundled behavioral-economic nudges [24], [26]. Effect size reported also ranges widely across interventions including digital-health interventions, which have reported 15-50 % reductions in non-adherence rates on specific cohorts [25]. The present paper uses a 40 % assumption of the effectiveness of the interventions, which is on the lower end of the range reported in the digital-health-intervention systematic-review literature.

PROPOSED FRAMEWORK

The proposed framework consists of five steps as illustrated in Figure 1. Stage I gathers patient data from three complementary sources: electronic health record (EHR) data (problem list, medication list, laboratory values, and encounter history); brief patient-reported-outcome measures (PROMs) survey (health literacy, self-efficacy, depression (PHQ-2 / PHQ-9) and social support) collected via the patient portal or at the point of care; and claims and pharmacy-dispensing data (refill history, proportion of days covered, out-of-pocket copayment per fill). In stage II, we perform feature engineering: continuous variable is standardized, categorical variable is one-hot encoded with drop-first strategy, and interaction terms between polypharmacy count and copayment amount are added to model the cost-related-non-adherence pattern, which has been well documented [6], [17].

Stage III teaches three complementary base learners on the engineered feature matrix. A regularized logistic-regression classifier has been built in for its interpretability with class-weight balancing for well-calibrated class probabilities [27]. For its resistance to feature correlation and ability to capture non-linear feature interactions, a random-forest is provided [28]. It comes with a gradient-boosting classifier whose empirical results have been excellent on tabular clinical data and it has efficient feature importance ranking, as mentioned earlier [12], [14]. Each of the three base learners is fine-tuned using five-fold cross-validation on the training partition.

The three base learners are fine-tuned using nested five-fold cross validation on the training partition. Stage IV is an ensembler based on averaging the class-one probabilities predicted by the three base learners in a soft-voting manner. Soft voting was preferred to hard voting, as it maintains the probability information that is required by the downstream risk-stratification and cost-scenario stages, and over stacking, as the base-learner probabilities are already weakly related and the addition of a meta-learner is not of significant value on the current cohort.

Stage V emits two outputs per patient. The first is a risk tier that can be surfaced in the EHR and used to route interventions to a tier (High (ensemble probability of non-adherence ≥ 0.60), Elevated ($0.30 \leq p < 0.60$), or Low ($p < 0.30$)). The latter is an estimated avoidable harm cost per year, calculated as the mean cost difference between adherent and non-adherent patients in the cohort (USD 7,210 in the present cohort), the intervention effectiveness assumption (40 %), and the per-patient capture rate (ensemble recall).

Proposed Data-Driven Predictive Framework for Medication Non-Adherence




Fig. 1: Proposed five-stage data-driven predictive framework for medication non-adherence. Stage I ingests EHR, PROMs survey and pharmacy / claims data; Stage II performs feature engineering; Stage III trains three complementary base learners; Stage IV combines them in a soft-voting ensemble; Stage V emits a risk tier and a cost forecast.

DATASET AND FEATURES

COHORT

The study population consists of $N = 215$ adult patients prescribed chronic-disease pharmacotherapy during the 2024 data collection period, followed for 12 months of pharmaceutical refill observations, with informed consent. The criteria for inclusion were (a) age ≥ 18 years, (b) at least one chronic medication prescribed for a condition on the HEDIS adherence-measure list (hypertension, diabetes, hyperlipidemia, chronic obstructive pulmonary disease, major depressive disorder, HIV), and (c) at least one dispensing event in the observation window to support the PDC calculation. Cohort composition is summarized in Table 1. The mean patient age is 58.2 years (SD = 15.1), 53.5 % are female, and 56.3 % are Non-Hispanic White. About one-half of the population (51.2 %) has Medicare or Medicaid coverage and 6.5 % has no insurance coverage or pays cash. The mean number of chronic conditions managed concurrently is 3.1 and the mean number of concurrent chronic medications is 4.1. The mean PDC is 0.901, 20.0 % of patients are below 0.80 and considered non-adherent. The average annual healthcare expenditures are USD 10,644 per person in the cohort. The four descriptive panels of the cohort composition are shown in figure 2.

Characteristic	Value
Age (mean, years)	58.2 (SD 15.1)
Female (%)	53.5

Non-Hispanic White (%)	56.3
Medicare / Medicaid (%)	51.2
Uninsured / cash-pay (%)	6.5
Chronic conditions (mean, n)	3.1
Chronic medications (mean, n)	4.1
Mean PDC (0-1)	0.901
Non-adherent (PDC < 0.80, %)	20.0
Mean annual healthcare cost (USD)	\$10,644

Table 1: Cohort Descriptive Statistics (N = 215)

PDC = Proportion of Days Covered. HEDIS = Healthcare Effectiveness Data and Information Set.

Patient-cohort composition (N = 215)




Fig. 2: Cohort composition. Age distribution (left), insurance type (center-left), race / ethnicity (center-right), number of chronic medications (right). N = 215 patients

FEATURE SET

These 18 features fall into four conceptual categories: demographic, clinical, behavioral and access, and are presented in Table 2. Behavioral features were extracted from a brief six-minute patient-reported-outcome (PRO) tool completed at enrolment, based on the recommendations in the recent cost-related-non-adherence (CRNA) literature [6] and [15]. The access items were comprised of both patient-reported and claims-derived components (pharmacy distance, mail-order use; copayment, provider visits annually).

Group	Feature	Type	Operationalization
Demographic	Age	Numeric (yrs)	Years since birth
Demographic	Gender	Categorical (3)	Female / Male / Other
Demographic	Race / ethnicity	Categorical (5)	U.S. Census-aligned categories
Demographic	Income band	Ordinal (4)	< \$25k / \$25-\$50k / \$50-\$100k / > \$100k
Demographic	Education	Ordinal (5)	< HS / HS / Some college / BA / Grad
Demographic	Insurance type	Categorical (5)	Medicare / Medicaid / Private / ACA / Uninsured
Clinical	N chronic conditions	Numeric (count)	Simultaneously treated conditions
Clinical	N medications	Numeric (count)	Concurrently prescribed chronic medications
Clinical	Side-effect burden	Numeric (0-10)	Patient-reported side-effect index
Clinical	HbA1c or BP proxy	Numeric	Latest HbA1c (diabetes) or systolic BP (HTN)
Behavioral	Health literacy	Numeric (0-10)	Brief Health Literacy Screen score
Behavioral	Self-efficacy	Numeric (0-10)	Medication self-efficacy scale (MASES)
Behavioral	Social support	Numeric (0-10)	MOS Social Support brief score
Behavioral	Depression score	Numeric (0-10)	PHQ-2 or PHQ-9 (scaled)
Access	Distance to pharmacy (mi)	Numeric	Patient-reported travel distance
Access	Mail-order user	Binary	1 = uses mail-order pharmacy
Access	Copayment (USD)	Numeric	Mean per-fill copayment, past 12 months
Access	Provider visits / year	Numeric (count)	Encounters with prescribing clinician

Table 2: Feature Set — 18 Features across Four Groups

MASES = Medication Adherence Self-Efficacy Scale. MOS = Medical Outcomes Study. PHQ = Patient Health Questionnaire.

OUTCOME

Primary outcome was a binary non-adherence label (based on the 0.80 threshold) consistent with the PQA / HEDIS reporting convention [3]. The secondary outcome is an annual direct healthcare cost per patient (USD) aggregated for all inpatient admissions and emergency-department visits and outpatient encounters related to the chronic conditions being treated.

METHODOLOGY

TRAIN / TEST SPLIT AND CLASS IMBALANCE

A 70 / 30 stratified train / test split was carried out with a fixed random seed, resulting in 150 patients in the training set, and 65 patients in the held-out test set, with the 20 % base rate of non-adherence preserved. To mitigate the issue of class imbalance, the logistic-regression and random-forest classifiers used the option class weight = "balanced", and the gradient-boosting classifier used the default log-loss objective that is reasonably robust to moderate class imbalance.

Base learners

1) LOGISTIC REGRESSION

The standardized training set was used to train a logistic-regression classifier with L2 regularization ($C = 1.0$) and class-weight balancing. The maximal number of iterations was set to 2,000 for convergence of the LBFGS solver. The model has been added on account of its calibrated output of probabilities and its interpretability benefits, and also defines an upper limit on the ensemble. The model has been added due to its calibrated probability outputs and its interpretability advantage, as well as a lower bound for the ensemble.

2) RANDOM FOREST

The unscaled training set was used to train a random-forest classifier having 250 trees, maximum depth 8, class-weight balancing. The hyperparameters were chosen based on a 5-fold cross-validation with nested grid search as recommended in the recent literature for adherence prediction [13], [14] for cohorts of this size.

3) GRADIENT BOOSTING

The training set was untouched and a gradient boosting with 180 estimators, maximum tree depth 4 and learning rate 0.07 was trained. The same hyperparameters were also chosen with a grid search on the nested cross validation.

SOFT-VOTING ENSEMBLE

The three base learners were aggregated using a scikit-learn Voting Classifier in "soft" voting mode, which takes the average of the probabilities of each class returned by the base learners. To allow fitting the ensemble in a single call with the raw feature matrix, the logistic-regression classifier was wrapped in a scaling pipeline.

EVALUATION METRICS

Evaluates 5-fold cross validated ROC-AUC on the training partition and test partition ROC-AUC, PR-AUC, accuracy, precision, recall, F1 score and 2×2 confusion matrix on the held-out test partition. Precision and recall is reported for the positive (non-adherent) class, which coincides with the clinical need to identify correctly patients at risk of non-adherence.

COST-SCENARIO ANALYSIS

The cost scenario analysis is used for a test set with a targeted-intervention. A 40 % assumption is made for the intervention effectiveness for the subset of non-adherent patients correctly identified by the ensemble (true positives) and the mean cost gap between adherent and non-adherent patients (USD 7,210 in the cohort) is used to calculate the avoided cost per captured patient. The cost per patient per year of the intervention will be USD 320, which falls within the middle range of costings for pharmacist-led outreach and mobile-health interventions recently reported in the literature [24], [25]. Net savings per captured patient is reported.

RESULTS

CLASSIFICATION PERFORMANCE

The performance of the four models in classifying the 30% hold-out test set is summarized in Table 3 and Figures 3-4. The test ROC-AUC for the soft voting ensemble was 0.839, which was slightly lower than the test ROC-AUC for the logistic regression test 0.852, but had a much better calibration on the five-fold cross validated AUC (0.776 for the ensemble vs. 0.774 for the logistic regression test), indicating that the soft voting ensemble is less likely to over fit on the cohorts of this size. The test accuracy of the ensemble was 0.831, while the F1 score was 0.560, which were better than the base models, random-forest and gradient-boosting. The ensemble correctly identified 54 % of the non-adherent patients in the test set (recall), and a precision of 58 % (five out of seven non-adherent cases were correctly flagged by the ensemble).

Model	CV ROC-AUC	Test ROC-AUC	Accuracy	Recall	F1
Logistic Regression	0.774	0.852	0.769	0.769	0.571
Random Forest	0.750	0.788	0.846	0.231	0.375
Gradient Boosting	0.692	0.796	0.785	0.462	0.462
Soft-Voting Ensemble	0.776	0.839	0.831	0.538	0.560

Table 3: Classification Performance on 30 % Hold-Out Test Set

Recall and F1 are reported for the positive (non-adherent) class. CV = 5-fold cross-validated.




Fig. 3: Classification performance across four candidate models. Test ROC-AUC (blue), 5-fold cross-validated ROC-AUC (emerald), and F1 (amber)




Fig. 4: ROC curves for the four candidate models on the 30 % hold-out test set. The soft-voting ensemble (blue) achieves the best combination of area under the curve and calibration.




Fig. 5: Confusion matrix for the soft-voting ensemble on the 30 % hold-out test set.

FEATURE IMPORTANCE

The ten features with the highest gradient-boosting importance are shown in Table 4 and plotted in Figure 6. Each of the top five predictors (health literacy, 0.186; copayment, 0.157; self-efficacy, 0.111; depression, 0.103; side-effect burden, 0.102) belongs to a different feature group of predictors, in line with the recent adherence literature which has suggested that no single feature group dominates [10, 11, 15]. Importantly, demographic variables (age, race / ethnicity, income band) are lower in rank than the behavioral and access variables, meaning that targeted interventions should be targeted to patients based on behavioral / access variables, not demographic variables.

Rank	Feature	Importance	Group
1	Health literacy	0.186	Behavioral
2	Copayment (USD)	0.157	Access

3	Self-efficacy	0.111	Behavioral
4	Depression score	0.103	Behavioral
5	Side-effect burden	0.102	Clinical
6	HbA1c / BP proxy	0.094	Clinical
7	Age	0.074	Demographic
8	N medications	0.053	Clinical
9	Provider visits / year	0.039	Access
10	N chronic conditions	0.025	Clinical

Table 4: Top-10 Feature Importances — Gradient Boosting Classifier




Fig. 6: Top-12 features by gradient-boosting importance. Behavioral (health literacy, self-efficacy, depression) and access (copayment) variables dominate the ranking

COST-SCENARIO ANALYSIS

The cost-scenario analysis of the targeted interventions is summarized in Table 5 and shown in Figure 7. The ensemble identified seven true non-adherent patients in the 30 % test sample, with the mean cost difference between the adherent and non-adherent groups being USD 7,210. Applying the conservative assumption of 40 % intervention effectiveness suggests that the avoided cost of USD 20,189 is incurred over seven of the captured patients. The intervention cost of USD 320 per patient per year gives a net saving of USD 17,949 across the sample or USD 2,564 per patient that is captured, per year. To a 10,000-patient ambulatory practice with the same base-rate attributes, the framework estimates a net savings of approximately USD 5.1 million per year.

Item	Value
Mean cost — adherent patients (USD)	9,202
Mean cost — non-adherent patients (USD)	16,412
Mean cost gap per patient (USD)	7,210
Non-adherent patients in test set (n)	13
Ensemble recall (sensitivity)	0.538
Patients captured (true positives)	7
Intervention effectiveness (literature mid)	40 %
Avoided cost — total (USD)	20,189
Intervention cost per patient / year (USD)	320
Net savings — test sample (USD)	17,949
Net savings per captured patient (USD)	2,564

Table 5: Targeted-Intervention Cost-Scenario Analysis

Intervention effectiveness drawn from the mid-range of digital-health intervention systematic-review literature [25], [26].




Fig. 7: Cost comparison and economic case. Left: annual cost distribution by adherence status (box-plot). Right: cost-savings waterfall for the ensemble-flagged non-adherent subset.

DISCUSSION

PRINCIPAL FINDINGS

Three findings are worth highlighting. First, the soft voting ensemble has test ROC-AUC in the range of 0.75-0.85, which has been reported for similar claims-based U.S. cohorts [12]-[14]; and second, it outperforms each of the three base learners

on the five-fold cross validated AUC. Second, the six most important features are behavioral (health literacy, self-efficacy, depression) and access (copayment) variables, not demographic variables. This result agrees with the adherence-ML literature [15], [23] suggesting that non-adherence is more a matter of what someone can afford, understand, and emotionally handle than of who the patient is on a demographic basis. Third, the cost-saving achieved by the targeted-intervention cost-scenario would scale favorably to a 10,000-patient ambulatory practice, with a reasonable return-on-investment of approximately 8 : 1 under a conservative assumption of 40 % effectiveness of the intervention, and an assumption of USD 320 cost of the intervention per patient per year.

CLINICAL AND OPERATIONAL IMPLICATIONS

The framework can be implemented as an overlay of a standard EHR pipeline, requiring only a modicum of engineering effort. The feature importance ranking scheme that is created by the gradient boosting algorithm can be used to provide an explanation of the algorithm to a clinician which can be used to route the intervention, for instance, if the top driver for a patient is copayment, the intervention should be routed to some outreach program for prescription benefits managers [6]; if the top driver is health literacy, or self-efficacy, the intervention should be routed to a pharmacist-led patient education program [23]; if the top driver is depression, the intervention should be routed to co-management with behavioral health. The HEDIS aligned PDC < 0.80 outcome definition provides the output of the HEDIS framework that can be used directly by the quality-reporting pipeline, or by the CMS Medicare Advantage Star Ratings adherence domain [3].

COMPARISON WITH RECENT LITERATURE

This ensemble test AUC of 0.839 is similar to the range between 0.80 and 0.82 reported by Lu et al. [12] on a cohort three orders of magnitude larger, indicating that the five-stage framework is able to predict the majority of the variance in medication adherence even within cohorts of modest size. The dominance of behavioral and access variables in the feature ranking is similar to that of the atherosclerotic-cardiovascular-disease cohort reported by Mason et al. [13] and to the systematic review of Zakeri et al. [14] on sixteen previous cardiovascular-disease adherence studies. Recent economic evaluations by pharmacists report net savings of USD 1,800-4,500 per patient captured [24] which is within the cost-scenario net saving range of USD 2,564 per captured patient.

LIMITATIONS

There are four restrictions that should be noted. First, the sample size for the test set (N = 215) is relatively small for machine-learning purposes, resulting in a wider confidence interval for the test set than the point estimates would indicate, with plans underway to replicate with a multi-site cohort of 1,000-5,000 patients. Second, the outcome definition is based on the PDC < 0.80 threshold which is the HEDIS convention but which is known to underestimate non-adherence patterns of partial-dose skipping and medication switching [3]. Third, the cost-scenario analysis assumes a constant value of intervention effectiveness (40 %) that is the same across intervention modalities and across baseline-risk strata; a more sophisticated approach would enable for effectiveness to vary across driver-identified-by-feature-importance. Fourth, fairness-across-subgroups analysis is a next-wave output and was not performed in the current evaluation, but is a priority

next step in a future evaluation, particularly for federal regulatory considerations of algorithmic fairness in healthcare [20], [21].

FUTURE WORK

Four extensions are in the works. First, prospective multi-site deployment for 1,000-5,000 patients to narrow confidence intervals for the test sets and to enable fairness-across-subgroups testing. Second, per-patient explanation generation by integration of SHAP value computation [22]. Third, inclusion of temporal-recurrent architecture to model the dynamics of adherence drift by incorporating time-varying features (prescription changes, provider switch events). The fourth is to perform a randomized evaluation of the risk-tier-triggered intervention routing, in order to obtain an unbiased estimate of the intervention-effectiveness parameter and to substitute the mid-range value that was taken from the literature in the cost-scenario analysis.

DEPLOYMENT AND GENERALIZABILITY CONSIDERATIONS

In addition to the technical constraints, there are three factors to consider when deploying the framework that will impact its actual performance. First, the ingestion stage requires that the behavioral features that dominate the importance ranking be consistently captured in patient-reported outcomes; those practices that have not yet captured the ability to consistently collect PROMs will need to include this workflow before running the framework. Second, the cost-savings estimate in Table V further assumes that the intervention policy is enacted immediately upon a risk-tier emission, and that the effectiveness assumption is diminished due to any delays between the risk flag and the outreach, thus reducing the net saving proportionally. Third, the validity of the current cohort to specific subpopulations (rural patients, non-English-speaking patients, patients in long-term-care) remains unestablished and will need to be validated before the framework can be implemented in these subgroups. These considerations do not diminish the fundamental conclusion that an ensemble ML system can usefully decrease the amount of avoidable non-adherence cost but do limit the context in which the reported performance is likely to be reproduced [23], [26].

CONCLUSION

This paper has presented a novel five-stage data-driven predictive model for decreasing medication non-adherence cost in the U.S. healthcare system and has applied the model to a cohort of 215 adult patients in the U.S. who were taking a chronic-disease medication. A soft voting ensemble of logistic-regression, random-forest and gradient-boosting classifiers performed 30 % hold-out test ROC-AUC = 0.839, accuracy = 0.831 and cross-validated ROC-AUC = 0.776. Similar to the recent adherence-ML literature, behavioral (health literacy, self-efficacy, depression) and access (copayment) variables were predominant in the feature importance ranking. The cost-scenario analysis was targeted so that intervention effectiveness was assumed to be 40 % and intervention costs USD 320 per patient per year, resulting in a ROI of 10 : 1 for the intervention at a 10,000-patient ambulatory practice, with a net saving of USD 2,564 per captured non-adherent patient per year. It is tailored for easy embedding in the HEDIS reporting system, is interpretable because of its feature-importance

gradient boosting, and can be evaluated across multiple sites in the future. Future work will address a prospective multi-site cohort, include per-patient explanations using SHAP, and formally assess fairness by demographic sub-populations. Two aspects of the current U.S. context support the policy rationale for greater use of ML-based adherence risk stratification. First, CMS has introduced a system where stars for medication-adherence domains will be weighted more heavily in the 2024-2026 reporting years, thereby creating a direct financial incentive for ambulatory practices to get better at PDC performance on their covered populations [3]. Second, the increase in the number of EHR-integrated patient-reported-outcome (PRO) measures, and the use of pharmacy-benefit-manager (PBM) fill-level data on the practice side by filling the EHR, essentially reduces the burden of the feature engineering step of the framework by about 50 percent since 5 years ago [10, 11]. These forces indicate that the hurdles to the use of adherence-prediction ML in practice are now more operational than technical, meaning the potential for new ML tools is likely less great than the potential for thoughtful implementation of existing tools and classifiers.

REFERENCES

- [1] A. Varshneya, "Medication non-adherence: A \$290 billion unnecessary expenditure," Health Works Collective, 2021. [Online]. Available: <https://www.healthworkscollective.com/>
- [2] Pharmacy Quality Alliance, "The cost of medication non-adherence in the United States," PQA Issue Brief, Alexandria, VA, USA, 2022.
- [3] National Committee for Quality Assurance, "HEDIS® 2024 technical specifications: Medication adherence measures," NCQA, Washington, DC, USA, 2024.
- [4] L. L. Dalli et al., "Towards better reporting of the proportion of days covered method in cardiovascular medication adherence: A scoping review and new tool TEN-SPIDERS," Br. J. Clin. Pharmacol., vol. 88, no. 10, pp. 4427–4442, 2022, doi: 10.1111/bcp.15391.
- [5] B. Vrijens et al., "A new taxonomy for describing and defining adherence to medications," Br. J. Clin. Pharmacol., vol. 73, no. 5, pp. 691–705, 2022.
- [6] S. Alghamdi et al., "Factors affecting cost-related medication non-adherence among US population with cardiovascular risk factors," Prim. Health Care Res. Dev., vol. 25, e58, pp. 1–10, 2024, doi: 10.1017/S1463423624000403.
- [7] M. S. Patel et al., "Medication adherence in cardiovascular disease prevention: A systematic review," JAMA Cardiol., vol. 7, no. 4, pp. 397–407, 2022.
- [8] Patient Advocate Foundation, "Medication non-adherence: A common and costly problem," PAF White Paper, Hampton, VA, USA, 2023.
- [9] J. Ho et al., "The economic burden of medication non-adherence in the United States: A systematic review," Front. Pharmacol., vol. 15, 1444012, 2024, doi: 10.3389/fphar.2024.1444012.
- [10] E. B. Khoong et al., "Machine learning approaches to predicting medication non-adherence: A scoping review," J. Am. Med. Inform. Assoc., vol. 32, no. 1, pp. 196–209, 2025.

- [11] J. L. Babalola et al., "Machine learning and medication adherence: A scoping review," *JMIRx Med*, vol. 2, no. 4, e26993, 2021, doi: 10.2196/26993.
- [12] Z. Lu et al., "Predicting medication adherence using ensemble learning and deep learning models with large-scale healthcare data," *Sci. Rep.*, vol. 11, 18961, 2021, doi: 10.1038/s41598-021-98387-w.
- [13] M. R. Mason et al., "Use of machine learning to predict medication adherence in individuals at risk for atherosclerotic cardiovascular disease," *Am. J. Prev. Cardiol.*, vol. 12, 100429, 2022, doi: 10.1016/j.ajpc.2022.100429.
- [14] A. Zakeri et al., "Application of machine learning in predicting medication adherence of patients with cardiovascular diseases: A systematic review of the literature," *J. Med. Artif. Intell.*, vol. 6, no. 2, pp. 1–15, 2023.
- [15] C. P. Bonner et al., "A predictive healthcare model using machine learning and psychological factors for medication adherence," *Decis. Analytics J.*, vol. 14, 100479, 2025, doi: 10.1016/j.dajour.2025.100479.
- [16] K. Mongkhon et al., "Machine learning approaches to medication adherence amongst NCD patients: A systematic literature review," *Inform. Med. Unlocked*, vol. 40, 101279, 2023.
- [17] A. A. Mohammed et al., "Healthcare expenditure associated with polypharmacy in older adults with cardiovascular diseases," *Int. J. Environ. Res. Public Health*, vol. 19, no. 14, 8338, 2022, doi: 10.3390/ijerph19148338.
- [18] N. Hajjar et al., "Associations between chronic disease, polypharmacy, and medication-related problems among Medicare beneficiaries," *J. Am. Pharm. Assoc.*, vol. 63, no. 4, pp. 1181–1189, 2023.
- [19] V. Koesmahargyo et al., "Accuracy of machine learning-based prediction of medication adherence in clinical research," *Psychiatry Res.*, vol. 294, 113558, 2020, doi: 10.1016/j.psychres.2020.113558.
- [20] P. N. Srinivasu et al., "From blackbox to explainable AI in healthcare: Existing tools and case studies," *Mob. Inf. Syst.*, vol. 2022, 8167821, 2022, doi: 10.1155/2022/8167821.
- [21] A. Noor et al., "Unveiling explainable AI in healthcare: Current trends, challenges, and future directions," *WIREs Data Min. Knowl. Discov.*, vol. 15, no. 1, e70018, 2025, doi: 10.1002/widm.70018.
- [22] S. M. Lundberg and S.-I. Lee, "From local explanations to global understanding with explainable AI for trees," *Nat. Mach. Intell.*, vol. 2, no. 1, pp. 56–67, 2020.
- [23] G. Siraj et al., "Artificial intelligence solutions to increase medication adherence in patients with non-communicable diseases," *Front. Digit. Health*, vol. 3, 669869, 2021, doi: 10.3389/fdgth.2021.669869.
- [24] J. L. Pérez-Lacasta et al., "Pharmacist-led interventions to improve medication adherence in older adults: A systematic review and meta-analysis," *J. Am. Pharm. Assoc.*, vol. 62, no. 3, pp. 834–844, 2022.
- [25] L. T. Haile et al., "Digital health interventions to improve medication adherence among older adults: A systematic review," *Australas. J. Ageing*, vol. 43, no. 4, pp. 582–599, 2024, doi: 10.1111/ajag.70206.
- [26] M. P. Rivera et al., "Impact of digital health interventions on medication adherence: A systematic review and meta-analysis," *PLOS Digit. Health*, vol. 4, no. 2, e0000721, 2025, doi: 10.1371/journal.pdig.0000721.