

ORIGINAL RESEARCH

Open access

Explainable Machine Learning Framework for Predicting Preventable Appointment Cancellations Using Scheduling Notes, Patient Communication History, Weather Conditions, Transportation Barriers, and Prior Attendance Behavior

Anna Svensson^{1*}, Erik Lindberg¹

Abstract

Appointment cancellations undermine clinic efficiency, disrupt continuity of care, and reduce access for patients waiting for limited appointment slots. Many cancellations may be preventable when risk is recognized early enough for staff to intervene with reminders, rescheduling support, transportation assistance, or telemedicine conversion. Existing appointment-risk models often emphasize historical attendance and demographic information while underusing scheduling notes, patient communication history, weather conditions, and transportation barriers. They also frequently provide risk scores without patient-specific explanations that staff can translate into meaningful outreach. This article proposes an explainable machine learning framework for predicting preventable appointment cancellations before the appointment occurs. The framework is designed to identify not only which appointments may be at risk, but also why the cancellation risk is elevated. The proposed framework uses a gradient-boosted classification model trained on structured scheduling variables, prior attendance behavior, communication history, weather-linked features, transportation indicators, and natural language processing outputs from scheduling notes. SHAP-based explanation layers would decompose each prediction into interpretable drivers that can be reviewed by scheduling staff, clinic managers, and governance teams. Conceptually, the framework would output a cancellation risk score together with a natural-language explanation of the dominant drivers. These outputs could support targeted interventions such as reminder escalation, proactive rescheduling, transportation support, or conversion to a virtual visit when appropriate. An explainable framework for preventable appointment cancellation prediction could shift patient access management from reactive backfilling toward proactive retention. By combining heterogeneous access signals with transparent attribution, clinics could better align outreach resources with patient-specific barriers.

Keywords Machine learning, Explainable AI, SHAP, Appointment cancellations, Patient access, Scheduling notes

*Correspondence:

Anna Svensson
anna.svensson@gmail.com

¹ Department of Clinical Informatics and Health Analytics, Faculty of Medicine, Uppsala University, Uppsala, Sweden

Introduction

Preventable appointment cancellations create unused clinical capacity, delay care for other patients, and weaken continuity between patients and care teams. The distinction between unavoidable cancellations and preventable cancellations is important because the latter may be influenced by timely communication, transportation support,

rescheduling options, or care-modality changes. Prior studies of outpatient nonattendance and cancellation behavior show that appointment adherence is shaped by operational, behavioral, and access-related factors rather than by a single cause [1-3]. An explainable approach is therefore needed to separate appointments that are merely statistically high risk from those that are high risk for reasons that can be acted on before the visit date [4, 5].

Most predictive models for missed visits have relied heavily on demographic characteristics, appointment type, lead time, and prior attendance history, which are useful but incomplete signals. Studies using machine learning for no-show or cancellation prediction have demonstrated the value of historical and scheduling data, but many frameworks leave real-time access signals underdeveloped, including scheduling-note content, current reminder response, anticipated weather disruption, and transportation difficulty [6-8]. Models that depend mostly on static history may identify patients who have missed care before, yet they may fail to identify emerging barriers affecting an otherwise adherent patient [9, 10]. A more complete framework should integrate both stable behavioral patterns and dynamic pre-appointment signals that can change shortly before the scheduled encounter [11, 12].

Explainability is central to patient access prediction because staff must understand why an appointment is flagged before choosing an intervention. A risk score alone may indicate that outreach is warranted, but it does not reveal whether the appropriate response is a reminder call, a transportation resource, a rescheduling offer, or a telemedicine option. Interpretable and explainable approaches in appointment prediction and hospital-access modeling suggest that patient-facing and staff-facing decision support must make model reasoning visible enough to support trust, oversight, and workflow adoption [13-15]. SHAP-based explanations are especially relevant because they can connect patient-level predictions to specific features such as prior no-shows, cancellation timing, reminder nonresponse, weather exposure, or access barriers [16, 17].

This article proposes an explainable machine learning framework for predicting preventable appointment cancellations by fusing scheduling notes, patient communication history, weather conditions, transportation barriers, and prior attendance behavior. The framework is conceptual rather than experimental, and it is intended to describe how heterogeneous data streams could be transformed into actionable cancellation-risk explanations. Its design is grounded in prior work on outpatient no-show prediction, intervention-oriented appointment management, and emerging explainable AI approaches for patient access. By prioritizing transparent attribution, the framework aims to support targeted retention interventions while preserving the ability to audit model behavior across patient groups and clinic settings.

Background

Appointment cancellations and their preventability

Appointment cancellations and no-shows are related access failures, but preventability depends on whether the barrier can be anticipated and addressed before the visit. Predictive modeling studies show that missed appointments are often associated with scheduling patterns, prior behavior, appointment characteristics, and patient-level access constraints, which means some cancellations could be reduced through targeted outreach rather than handled only after the slot is lost [1, 3, 11]. Intervention-oriented reviews further suggest that predictive models should be connected to operational actions such as reminders, overbooking strategies, rescheduling support, or personalized outreach rather than treated as stand-alone analytics outputs [18, 19]. In an explainable framework, preventability would be operationalized as elevated cancellation risk accompanied by modifiable drivers that staff can address before the appointment [20].

Scheduling notes and free-text data in patient access

Scheduling notes, cancellation reasons, and call-center documentation may contain patient-expressed barriers that structured appointment fields do not capture. Although much appointment-prediction research has focused on coded variables, recent modeling work in imaging and access settings indicates that free-text and language-model approaches could help surface clinically and operationally meaningful context when applied carefully [21, 22]. Natural language

processing could identify mentions of transportation difficulty, work conflicts, childcare constraints, uncertainty about appointment purpose, illness, or preference for rescheduling, all of which may support more precise intervention selection. In the proposed framework, free-text notes would not replace structured predictors but would enrich them by converting patient access narratives into interpretable themes available to the model [23].

Patient communication history as an engagement signal

Patient communication history provides a near-real-time signal of engagement because reminder delivery, confirmation status, portal response, and phone contact attempts can change as the appointment approaches. Text-message reminder studies and model-based intervention reviews suggest that communication channels can influence attendance and can also identify patients who may need additional support beyond standard automated reminders [18, 24]. A patient who confirms a reminder, responds through the portal, or completes a phone interaction may present a different risk profile from a patient with the same historical no-show pattern but no recent communication response. Encoding communication history as an engagement trajectory would allow the framework to distinguish persistent nonadherence from a temporary failure to connect with the clinic [25, 26].

Weather conditions and transportation barriers

Weather and transportation barriers are central to preventable cancellations because they can interfere with physical access even when a patient intends to attend. Predictive studies and equity-focused analyses of missed appointments indicate that access constraints such as travel burden, clinic location, and environmental disruption should be considered alongside patient history and appointment characteristics [20, 27, 28]. Weather-linked variables could represent forecasted severe conditions, precipitation, temperature extremes, or local travel disruption near the patient or clinic, while transportation variables could represent distance, public transit complexity, parking constraints, or documented reliance on assisted transportation. In an explainable framework, these features would be most useful when they generate concrete outreach options, such as offering a different time, transportation support, or a virtual visit [19].

Prior attendance behavior and interpretable risk models

Prior attendance behavior remains one of the most important conceptual inputs for cancellation prediction because it summarizes repeated interactions between patients and scheduling systems. Studies across outpatient, pediatric, neurology, dental, imaging, and primary care settings consistently treat past no-shows, cancellation history, appointment lead time, and rescheduling behavior as core predictors of future attendance risk [2, 6, 15-17]. However, historical behavior can be ethically and operationally problematic if it is treated as a fixed patient trait rather than as a signal of prior access barriers and system interactions. Explainable models should therefore unpack whether prior attendance history is acting alone or in combination with modifiable factors such as reminder nonresponse, transportation difficulty, or weather exposure [23, 29].

Framework Overview

High-level prediction pipeline

At a predefined interval before each scheduled appointment, the proposed framework would ingest current-state data from scheduling systems, communication platforms, weather feeds, transportation indicators, and prior attendance records. The model would then generate a cancellation-risk score and a ranked set of contributing factors that explain why the appointment may be at risk at that moment. This design follows the logic of prior machine learning appointment-prediction systems while extending them toward actionable, patient-level explanation rather than retrospective classification alone [4, 7, 12]. The output would be intended for workflow support, not autonomous decision-making, with staff retaining responsibility for interpreting explanations and selecting outreach actions [18, 21].

Core input features

The core input features would include NLP-extracted themes from scheduling notes, the patient's response status to recent reminders, weather and travel-disruption indicators, transportation difficulty measures, appointment characteristics, and prior attendance patterns. Prior modeling studies support the relevance of historical attendance, appointment metadata, and operational variables, while intervention-focused work highlights the need to incorporate communication and modifiable access signals [1, 10, 19]. Equity-oriented and interpretable modeling studies further suggest that social and transportation-related variables should be handled transparently because they may improve outreach targeting while also raising fairness concerns [20, 28]. The framework would therefore treat each feature group as both a predictive signal and a possible explanation requiring governance review [27].

Design principles

The framework would be designed around transparency, actionability, privacy protection, and compatibility with existing scheduling and outreach workflows. Transparency would require that each prediction be accompanied by a concise explanation, while actionability would require that explanations map to practical interventions such as reminder escalation, transportation support, or rescheduling assistance. Privacy protection would require limiting access to sensitive social and communication-derived features, and workflow compatibility would require that model outputs appear within the scheduling tools staff already use [18, 19]. These principles are consistent with explainable and intervention-oriented appointment prediction approaches that emphasize implementation usefulness rather than prediction alone [21, 23].

Data Sources and Feature Engineering

Extraction and NLP of scheduling notes

Scheduling notes would be processed using NLP methods that identify entities, phrases, and themes related to cancellation reasons, transportation needs, work or caregiving conflicts, uncertainty about appointment purpose, and prior rescheduling conversations. The framework could use controlled vocabularies for known access barriers together with topic modeling or language-model-assisted classification to discover recurring free-text themes that are not captured in structured fields. Recent work using language models and tree-based approaches for diagnostic imaging access suggests that unstructured text can be incorporated into appointment-risk prediction when outputs are constrained to operationally meaningful features [22]. The extracted note features would be stored as interpretable indicators, such as “transportation mentioned,” “patient requested callback,” or “rescheduling uncertainty,” rather than opaque text embeddings alone [21].

Communication history and reminder status encoding

Communication history would be transformed into structured engagement metrics that summarize reminder channel, delivery timing, confirmation status, portal response, phone contact completion, and unresolved outreach attempts. Prior research on targeted text reminders and predictive model-based interventions supports the idea that reminder interactions are not merely interventions but also informative signals of patient engagement and support needs [18, 24]. For example, a recent reminder confirmation could lower concern when other risk factors are present, while repeated failed contact attempts could suggest the need for a personal call or alternate communication channel. These features should be time-aware so that the model reflects the patient's latest interaction with the clinic rather than relying only on older attendance history [25].

Weather and transportation barrier variables

Weather variables would be linked to the appointment date, clinic location, and patient travel context, while transportation variables would represent travel burden, distance, likely transit dependency, or documented mobility barriers. Prior

appointment-prediction and equity-focused studies support incorporating access-related context because missed visits may reflect environmental and socioeconomic constraints rather than patient preference or disregard [20, 27, 28]. The framework would avoid treating transportation variables as blame-oriented patient attributes and would instead encode them as potential intervention signals. When severe weather and transportation difficulty appear together, the explanation layer could direct staff toward practical options such as rescheduling flexibility, reminder escalation, or telehealth conversion [19].

Table 1 summarizes the proposed input structure, temporal feature representation, and deep learning logic required to convert hospital pharmacy verification activity into short-horizon backlog forecasts.

Table 1. Input Structure, Temporal Representation, and Deep Learning Logic for Pharmacy Verification Backlog Forecasting

Model component	Manuscript-specific content	Operational meaning in pharmacy verification	Example variables or representations	Role in deep learning workflow	Practical value for pharmacy supervisors
Medication order stream	Timestamped medication orders entering the pharmacy verification queue	Captures real-time demand for pharmacist review	Order entry time, verification time, medication class, dose form, route, urgency, ordering unit	Forms the primary temporal demand sequence	Shows when incoming medication demand is beginning to exceed available review capacity
Medication order complexity	Structured estimate of verification effort required per order	Distinguishes simple queue volume from clinically demanding workload	IV status, high-alert flag, special handling, titratable medication, order-set association, number of line items	Adds complexity-weighted workload signal to each time window	Helps supervisors recognize that a small queue of complex orders may be more concerning than a larger queue of routine orders
High-alert medication flags	Identification of medications requiring heightened safety attention	Marks orders that may need faster or more careful verification	Insulin, anticoagulants, opioids, concentrated electrolytes, chemotherapy, vasoactive agents, local high-alert list indicators	Increases risk sensitivity of backlog forecast	Supports prioritization of medication-safety-relevant verification workload
Pharmacist staffing and capacity	Measures of available verification resources	Represents service capacity against incoming demand	Scheduled pharmacist count, logged-in verification users, shift phase, meal-break vulnerability, recent verification rate	Provides capacity-side features for demand-capacity balance modeling	Allows the forecast to distinguish high demand with adequate capacity from high demand

					with insufficient coverage
Patient acuity and location signals	Near-term predictors of order complexity and order bursts	Captures demand emerging from ICU, ED, perioperative, and transfer activity	ICU location, ED boarding, admission-discharge-transfer feed, perioperative status, recent transfers, service line	Helps anticipate future medication order surges before they fully appear in the queue	Gives shift leads earlier awareness of likely medication demand from high-acuity areas
Historical queue dynamics	Current and recent state of verification queue	Represents queue momentum and recovery pattern	Current unverified order count, age of pending orders, arrival intensity, clearance trend, urgent-order proportion	Enables temporal learning of backlog growth, persistence, or recovery	Helps identify whether congestion is transient, worsening, or slowly resolving
Calendar and shift context	Operational context affecting ordering and verification rhythms	Captures routine and anomalous temporal patterns	Time of day, day of week, weekend, holiday, shift change, post-rounding period	Provides temporal anchors for sequence model interpretation	Prevents normal predictable peaks from being confused with abnormal backlog risk
Sequence construction	Aggregation of all variables into repeated time bins	Converts pharmacy workflow into model-ready temporal structure	Consecutive 5-, 10-, 15-, or 30-minute windows containing demand, capacity, complexity, acuity, and queue state	Creates multivariate input sequence for LSTM, GRU, or temporal convolutional network	Allows the model to learn queue trajectory rather than isolated snapshots
Deep learning encoder	Recurrent or temporal convolutional model	Learns nonlinear relationships among demand, complexity, staffing, acuity, and queue state	LSTM, GRU, temporal convolutional network, hidden queue-state representation	Transforms sequential inputs into forecast-ready representation	Detects latent queue momentum that may not be visible from raw queue count alone
Forecast output layer	Predicted backlog depth and backlog probability	Converts hidden temporal state into operationally actionable forecast	Predicted queue length, probability of exceeding threshold, worsening/stable/improving risk direction	Produces short-horizon forecast for upcoming operational windows	Supports earlier staffing and prioritization decisions before backlog becomes disruptive

Interpretable Model Architecture

Model choice and rationale

A gradient-boosted tree model would be appropriate for the proposed framework because appointment cancellation prediction requires combining heterogeneous feature types, including categorical scheduling variables, historical behavior, weather indicators, communication events, and NLP-derived flags. Prior machine learning studies of no-show and cancellation prediction have used tree-based, ensemble, and stacked approaches because they can capture nonlinear relationships among operational and patient-access variables [4, 10, 11]. The explainability layer would use SHAP-style attribution to translate model predictions into ranked feature contributions at both the population and patient levels, aligning with recent interpretable appointment-prediction work [21, 23]. This architecture would support conceptual decision support without requiring the article to report experimental performance, training details, or numerical outcomes [14].

Input feature vector and preprocessing

The input feature vector would combine appointment metadata, prior attendance behavior, communication-status encodings, weather-linked indicators, transportation-barrier variables, and NLP-derived scheduling-note themes. Categorical variables could be represented through interpretable encodings, continuous variables could be normalized when needed, and missing weather or communication fields could be represented explicitly so that absence of data is not confused with absence of risk. Prior studies show that no-show prediction often depends on feature engineering choices across appointment timing, patient history, specialty context, and access-related variables, making preprocessing an important design component rather than a purely technical step [5, 8, 13]. To preserve interpretability, each engineered feature should retain a clinically or operationally understandable label that can be displayed in explanations and audited by governance teams [17, 29].

Output: cancellation risk score and explanation

The model output would consist of a calibrated cancellation-risk score paired with a patient-level explanation that identifies the most influential drivers for the specific appointment. A SHAP waterfall-style explanation could show how factors such as prior missed visits, absent reminder confirmation, weather disruption, transportation difficulty, or scheduling-note themes contribute to the final risk estimate. Studies of predictive and explainable appointment-risk models indicate that such outputs are most useful when they connect risk attribution to possible staff actions rather than presenting risk as an isolated probability [18, 21, 23]. The explanation should therefore be rendered in plain language, for example by stating that the appointment may need outreach because recent communication is unresolved and documented transportation difficulty may interact with forecasted weather conditions [22, 28].

Figure 1 presents the proposed end-to-end workflow through which timestamped pharmacy, staffing, acuity, and queue-state data are transformed into deep learning backlog forecasts, pharmacist-reviewed decision support, and proactive verification workflow actions.




Figure 1. End-to-End Deep Learning Workflow for Forecasting Pharmacy Verification Backlogs and Supporting Proactive Staffing Decisions

Identifying Key Drivers of Preventable Cancellations

Global feature importance across the patient population

Global explanation outputs would summarize the dominant drivers of preventable cancellation risk across the patient population, allowing clinic leaders to see which barriers most often contribute to elevated risk. SHAP summary plots could conceptually identify recurring patterns such as unresolved reminder status, prior no-show behavior, longer appointment lead time, transportation difficulty, and forecasted weather disruption as major contributors when they are present in the feature set [1, 3, 7]. These summaries would be used to guide operational planning rather than to label individual patients as unreliable. For example, if communication nonresponse and transportation-linked variables repeatedly appear as high-impact drivers, the clinic could strengthen reminder workflows and transportation support before expanding appointment overbooking [18, 19].

Patient-level attribution of cancellation risk

Patient-level explanations would distinguish between appointments that appear similarly high risk but require different interventions. One patient's risk may be driven primarily by a recent pattern of late cancellations and missed visits, while another patient's risk may be driven by a current lack of reminder confirmation combined with documented transportation concerns [2, 15, 16]. Instance-level SHAP decompositions would allow schedulers to see these differences without needing to inspect the full feature vector or interpret model internals. This level of attribution is essential because a generic high-risk label may lead to inefficient outreach, whereas a patient-specific explanation can support a tailored retention action [21, 23].

Interaction effects between environmental and behavioral factors

Preventable cancellation risk may arise from interactions among environmental, behavioral, and communication factors rather than from any single feature. For instance, forecasted severe weather may have limited practical effect for a patient who recently confirmed the appointment and has reliable transportation, but it may become more important when the patient has unresolved reminder contact and documented travel difficulty [20, 27, 28]. SHAP dependence plots could help analysts examine whether transportation barriers amplify the contribution of weather variables or whether reminder confirmation moderates the contribution of prior no-show history. Such interaction analysis would support more nuanced outreach rules and reduce the risk of assuming that all patients exposed to the same weather forecast face the same access barrier [22].

Counterfactual explanations for intervention selection

Counterfactual explanations would translate model reasoning into intervention-oriented statements that describe how risk might change if a modifiable factor were addressed. Rather than stating only that a patient is high risk, the framework could indicate that risk would be expected to decrease if the reminder were confirmed, a transportation barrier were resolved, or the visit were converted to telemedicine [18, 22]. These explanations should be framed conceptually and cautiously because they represent model-supported intervention hypotheses, not guaranteed causal effects. In practice, counterfactual reasoning would help staff choose between reminder escalation, rescheduling assistance, transportation support, and other retention strategies [19, 25].

Explainability for Stakeholders

Explanations for scheduling staff

For scheduling staff, explanations should be brief, specific, and connected to a recommended action. A dashboard could display a risk badge with a plain-language reason such as unresolved reminder contact, prior late cancellation pattern, weather-related travel concern, or transportation note, followed by an outreach suggestion [4, 6, 12]. The purpose would not be to automate staff judgment but to reduce the cognitive burden of reviewing multiple systems before deciding how to contact the patient. This approach aligns with intervention-oriented appointment prediction, where the value of the model depends on whether staff can use its output to prevent the appointment slot from being lost [18].

Explanations for clinic management

For clinic management, explainability should aggregate drivers across sites, specialties, appointment types, and time periods to identify system-level barriers. If explanations repeatedly point to transportation difficulty, delayed scheduling, reminder nonresponse, or weather-sensitive access patterns, leaders could redesign scheduling templates, adjust reminder protocols, or develop community transportation partnerships [11, 19, 27]. Management-facing explanations should therefore focus less on individual patients and more on the recurring operational conditions under which preventable cancellations arise. This aggregated view would also help determine whether the framework supports equitable access improvement or merely redistributes attention toward already well-documented patient groups [20, 28].

Audit trail and fairness monitoring

An audit trail would record each prediction, the explanation shown, the intervention selected, and the eventual appointment outcome. Such logging is necessary because features related to communication, transportation, geography, and prior attendance may correlate with socioeconomic disadvantage and could produce inequitable flagging if not monitored carefully [20, 28]. Fairness monitoring should compare explanation patterns across demographic and access-related groups, asking whether certain populations are disproportionately assigned risk for reasons that reflect structural barriers rather than modifiable appointment-level conditions. The governance process should treat explainability as a tool for accountability, not only as a way to increase user trust [21, 23].

Closed-loop feedback for model improvement

Closed-loop feedback would connect outreach actions and appointment outcomes back to model refinement. When staff apply an intervention such as a personal call, rescheduling offer, transportation support, or reminder escalation, the system should record both the action and whether the patient attended, cancelled, rescheduled, or remained unreachable [18,24]. This feedback would allow the framework to learn which explanations tend to lead to useful interventions and which explanations require revision because they are too vague, insensitive, or operationally unhelpful. Over time, the model and explanation templates could be updated to better reflect clinic workflows while remaining subject to fairness and privacy review [25, 26].

Integration Into Scheduling and Outreach Workflow

Proactive outreach triggered by high-risk predictions

The framework would trigger proactive outreach when the predicted risk and explanation indicate that a cancellation may be preventable. Depending on the explanation, the outreach pathway could involve a personal phone call, an automated rescheduling option, a reminder escalation, transportation-resource navigation, or a telemedicine conversion offer [18, 19, 24]. The threshold for action should be configured locally because clinics differ in staffing capacity, specialty urgency, patient population, and available interventions. The goal would be to direct limited outreach resources toward appointments where the model suggests that a timely, patient-specific action could preserve access [1, 11].

Embedding risk scores in scheduling software

Embedding the cancellation-risk score directly into scheduling software would reduce the need for staff to consult a separate analytics platform. A simple visual badge next to the appointment could indicate that outreach may be warranted, while a hover-over explanation could summarize the top drivers in plain language [6, 12]. This design would allow staff to review the explanation during normal scheduling, reminder, or waitlist-management workflows rather than after the appointment slot has already become difficult to recover. Integration should preserve human oversight by making the explanation visible and by allowing staff to document whether the suggested action was accepted, modified, or declined [18, 21].

Evaluation Strategy

Predictive performance metrics

Although this article does not report experiments or performance numbers, the framework should eventually be evaluated using standard predictive-performance and calibration approaches appropriate for cancellation-risk prediction. Metrics such as discrimination, precision-recall behavior for cancellation events, and calibration by clinic and demographic subgroup would help determine whether the model supports reliable operational decisions [8, 14, 29]. Evaluation should also examine whether the model remains useful across specialties, appointment types, and patient groups rather than performing well only in a narrow setting. These assessments would be necessary before deploying the framework as a decision-support tool in real scheduling workflows [5, 13].

Explanation quality and actionability

Explanation quality should be evaluated through user studies with schedulers, access-center staff, clinic managers, and governance stakeholders. The key question is whether explanations improve understanding, support appropriate intervention selection, and reduce unnecessary manual review, rather than whether users merely find the interface attractive [18, 21]. Staff should be asked whether explanation phrases such as unresolved reminder contact, transportation concern, or weather-sensitive access barrier are specific enough to guide action. Evaluation should also examine whether explanations are interpreted consistently and whether they avoid stigmatizing patients whose risk reflects social or structural barriers [20, 28].

Prospective impact on cancellation rates

A prospective evaluation could compare model-guided outreach with usual reminder workflows to assess whether explainable predictions improve appointment retention and slot recovery. Such a design should evaluate operational outcomes, patient experience, staff burden, and equity effects, while avoiding assumptions that every flagged appointment can or should be prevented [18, 19, 24]. The study should also record which explanation-driven interventions were used so that future refinement can distinguish between accurate prediction and effective action. In this way, the evaluation would test the complete access-management pathway: prediction, explanation, intervention, and appointment outcome [25, 26].

Table 2 outlines the evaluation strategy, governance safeguards, implementation risks, and practical operational actions needed to responsibly deploy the proposed pharmacy backlog forecasting model.

Table 2. Evaluation, Governance, Implementation Safeguards, Failure Modes, and Practical Actions for the Proposed Pharmacy Backlog Forecasting Model

Domain	Manuscript-specific evaluation or	Why it matters operationally	Recommended metric, safeguard, or process	Potential failure mode	Practical action enabled or protected

	governance issue				
Forecast accuracy	The model must estimate future verification queue depth and backlog probability	Inaccurate forecasts may create unnecessary staffing shifts or missed congestion	Mean absolute error for queue length, threshold sensitivity, specificity, calibration of backlog probability, prediction interval coverage	Forecast underestimates queue growth during sudden order surges	Earlier recognition of rising queue pressure and more timely staffing response
Temporal validation	Model testing must preserve chronological order	Random splitting can leak future workflow patterns into training	Walk-forward validation, time-based train-test split, silent prospective testing	Overstated performance due to temporal leakage	More realistic estimate of performance under live pharmacy rhythms
Surge-event performance	Model must remain useful during ICU, ED, perioperative, or high-alert medication surges	Backlog risk is most safety-relevant during abnormal demand spikes	Stratified evaluation during surge periods, STAT-order clusters, weekend shifts, holiday staffing, and reduced-capacity periods	Model performs well on routine days but poorly during operational stress	Targeted surge staffing, urgent-order prioritization, and shift-lead situational awareness
Calibration and uncertainty	Forecast should communicate confidence and instability	Pharmacists need to know when the model is uncertain	Calibration curves, prediction interval coverage, uncertainty flag on dashboard	Overconfident alert during sparse or unusual data conditions	Encourages cautious human review rather than blind reliance on model output
Explainability	Supervisors need to understand why backlog risk is rising	A risk score alone may not support practical action	Feature attribution by time window, driver summary, trend explanation, driver categories such as staffing, high-alert orders, ICU demand, slow clearance	Alert fatigue if warnings lack understandable causes	Directs supervisors toward the most relevant operational response
Human oversight	The model should support but not replace pharmacist judgment	Verification and staffing decisions remain professional responsibilities	Pharmacist-in-charge review, shift-lead confirmation, no automatic staffing change, no autonomous medication verification	Automation bias or inappropriate operational action	Preserves pharmacist accountability and contextual decision-making
Data quality	Forecast depends on reliable order,	Missing or delayed data may	Real-time data completeness checks, source-system	Staffing feed lags behind actual	Prevents misleading forecasts from

	staffing, acuity, and queue data	distort backlog estimates	reconciliation, stale-data warnings	pharmacist availability	incomplete operational data
Local customization	Hospitals differ in pharmacy workflow, staffing model, and high-alert medication definitions	A model trained in one setting may not transfer directly	Local feature mapping, site-specific high-alert list alignment, recalibration, workflow validation	Poor generalizability across centralized and decentralized pharmacy models	Supports safe adaptation to local pharmacy practice
Dashboard integration	Forecast must appear where supervisors already manage the queue	Separate systems may reduce use and increase cognitive burden	Embed forecast into pharmacy queue dashboard with trajectory, threshold, driver, and action context	Useful model ignored because it is outside routine workflow	Makes forecast available at the moment staffing and prioritization decisions are made
Alert governance	Alerts should be actionable and not excessive	Too many low-value alerts can erode trust	Tiered alert thresholds, escalation rules, alert suppression during known artifacts, review of false-positive burden	Alert fatigue or desensitization	Maintains attention to clinically and operationally meaningful backlog warnings
Implementation monitoring	Deployment should be evaluated beyond technical performance	Operational value depends on whether the forecast improves workflow	Verification turnaround time, urgent-order review timeliness, staff response rate, pharmacist trust, override reasons, alert usefulness	Technically accurate model fails to change operations	Links model success to pharmacy safety, efficiency, and staff situational awareness
Medication safety boundary	Forecasting backlog is not equivalent to approving medication orders	The model should not imply autonomous medication safety review	Clear labeling as operational decision support, no medication approval function, pharmacist final authority	Users misinterpret backlog tool as clinical verification tool	Protects the clinical safety role of pharmacist verification

Limitations

Data gaps and external factors

The proposed framework depends on data streams that may be incomplete, inconsistent, or difficult to interpret across clinics. Weather data may not capture micro-climate conditions, last-minute travel disruption, or patient-specific exposure, while transportation variables may be only rough proxies unless patients directly report their access needs [20, 27]. Scheduling notes may be unevenly documented, and communication histories may omit interactions that occur outside integrated systems. These limitations mean that model explanations should be treated as decision-support signals requiring staff judgment, not as definitive accounts of why a patient may cancel [22].

Equity and privacy considerations

Transportation barriers, communication patterns, address-derived variables, and prior attendance behavior may correlate with socioeconomic disadvantage, disability, language access, or other sensitive dimensions of care access. If used without governance, these features could lead to disproportionate flagging of underserved patients or to explanations that inadvertently frame structural barriers as individual deficits [20, 28]. Privacy controls should restrict access to sensitive note-derived and communication-derived features, and fairness monitoring should assess both predictions and explanation patterns across patient groups. The ethical goal of the framework should be supportive outreach and access preservation, not surveillance or punitive scheduling decisions [21, 23].

Conclusion

An explainable machine learning framework for predicting preventable appointment cancellations could help clinics identify at-risk appointments before access is lost. By focusing on preventability, the framework would separate cancellations that may be influenced by timely support from those that are unavoidable or outside the clinic's reasonable control. Its value would lie in linking prediction to practical outreach decisions. The model would therefore function as a staff-facing access tool rather than an autonomous scheduling authority.

A central strength of the proposed framework is its integration of multiple access signals that are usually analyzed separately. Scheduling notes can reveal patient-expressed concerns, communication history can indicate current engagement, weather and transportation variables can identify near-term access barriers, and prior attendance behavior can summarize longitudinal patterns. Explainability connects these signals by showing which factors matter for each appointment. This allows the framework to support targeted intervention rather than generic high-risk labeling.

Important challenges remain before such a framework could be implemented responsibly. Data completeness, documentation variation, missing communication records, uncertain transportation measures, and changing weather conditions could all limit prediction and explanation quality. Sensitive social information must also be handled carefully so that the model supports patients rather than stigmatizing them. Field testing would be necessary to determine whether explanations are understandable, actionable, fair, and compatible with real scheduling workflows.

Implementation pilots should be conducted in diverse clinic settings, including primary care, specialty care, imaging, and safety-net environments. These pilots should evaluate not only whether predictions are technically useful, but also whether staff can act on explanations in ways that improve access and patient experience. The strongest use case is a workflow in which the system identifies a preventable barrier early and helps the clinic offer the right support before the appointment is cancelled. With careful governance, explainable cancellation prediction could become a practical tool for proactive patient access management.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 07 Dec 2024

Revised: 25 Jan 2025

Accepted: 24 Mar 2025

Published online: 20 July 2025

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

- Mohammadi I, Wu H, Turkcan A, Toscos T, Doebbeling BN. Data analytics and modeling for appointment no-show in community health centers. *J Prim Care Community Health*. 2018;9:2150132718811692.
- Dantas LF, Hamacher S, Cyrino Oliveira FL, Barbosa SD, Viegas F. Predicting patient no-show behavior: a study in a bariatric clinic. *Obes Surg*. 2019;29(1):40-7.
- Nelson A, Herron D, Rees G, Nachev P. Predicting scheduled hospital attendance with artificial intelligence. *NPJ Digit Med*. 2019;2(1):26.
- Ahmadi E, Garcia-Arce A, Masel DT, Reich E, Puckey J, Maff R, et al. A metaheuristic-based stacking model for predicting the risk of patient no-show and late cancellation for neurology appointments. *IJSE Trans Healthc Syst Eng*. 2019;9(3):272-91.
- AlMuhaideb S, Alsweilem O, Alsubaie N, Ferwana I, Alnajem A. Prediction of hospital no-show appointments through artificial intelligence algorithms. *Ann Saudi Med*. 2019;39(6):373-81.
- Chong LR, Tsai KT, Lee LL, Foo SG, Chang PC. Artificial intelligence predictive analytics in the management of outpatient MRI appointment no-shows. *AJR Am J Roentgenol*. 2020;215(5):1155-62.
- Daghistani T, AlGhamdi H, Alshammari R, AlHazme RH. Predictors of outpatients' no-show: big data analytics using Apache Spark. *J Big Data*. 2020;7(1):108.
- Carreras-García D, Delgado-Gómez D, Llorente-Fernández F, Arribas-Gil A. Patient no-show prediction: a systematic literature review. *Entropy (Basel)*. 2020;22(6):675.
- Alshammari R, Daghistani T, Alshammari A. The prediction of outpatient no-show visits by using deep neural network from large data. *Int J Adv Comput Sci Appl*. 2020;11(10).
- Aladeemy M, Adwan L, Booth A, Khasawneh MT, Poranki S. New feature selection methods based on opposition-based learning and self-adaptive cohort intelligence for predicting patient no-shows. *Appl Soft Comput*. 2020;86:105866.
- Ferro DB, Brailsford S, Bravo C, Smith H. Improving healthcare access management by predicting patient no-show behaviour. *Decis Support Syst*. 2020;138:113398.

Fan G, Deng Z, Ye Q, Wang B. Machine learning-based prediction models for patients no-show in online outpatient appointments. *Data Sci Manag.* 2021;2:45-52.

Alshammari A, Almalki R, Alshammari R. Developing a predictive model of predicting appointment no-show by using machine learning algorithms. *J Adv Inf Technol.* 2021;12(3).

Salazar LH, Leithardt VR, Parreira WD, da Rocha Fernandes AM, Barbosa JL, Correia SD. Application of machine learning techniques to predict a patient's no-show in the healthcare sector. *Future Internet.* 2022;14(1):3.

Liu D, Shin WY, Sprecher E, Conroy K, Santiago O, Wachtel G, et al. Machine learning approaches to predicting no-shows in pediatric medical appointment. *NPJ Digit Med.* 2022;5(1):50.

Alabdulkarim Y, Almukaynizi M, Alameer A, Makanati B, Althumairy R, Almaslukh A. Predicting no-shows for dental appointments. *PeerJ Comput Sci.* 2022;8:e1147.

Peng J, Patel AD, Burch M, Rossiter S, Parker W, Rust S. Predicting patient no-shows in an academic pediatric neurology clinic. *J Child Neurol.* 2022;37(7):582-8.

Oikonomidi T, Norman G, McGarrigle L, Stokes J, van der Veer SN, Dowding D, et al. Predictive model-based interventions to reduce outpatient no-shows: a rapid systematic review. *J Am Med Inform Assoc.* 2023;30(3):559-69.

Deina C, Fogliatto FS, da Silveira GJ, Anzanello MJ. Decision analysis framework for predicting no-shows to appointments using machine learning algorithms. *BMC Health Serv Res.* 2024;24(1):37.

Yang Y, Madanian S, Parry D. Enhancing health equity by predicting missed appointments in health care: machine learning study. *JMIR Med Inform.* 2024;12(1):e48273.

Toffaha KM, Simsekler MC, Alshehhi A, Omar MA. Predicting hospital no-shows: interpretable machine learning models approach. *IEEE Access.* 2024;12:166058-67.

Rafique D, Liu X, Gong B, Belsito L, McCradden MD, Mazwi ML, et al. Predicting pediatric diagnostic imaging patient no-show and extended wait-times using LLMs, regression, and tree based models. *Front Artif Intell.* 2025;8:1652397.

Hathaway QA, Yanamala N, Narumanchi T, Narumanchi J. Bringing precision to pediatric care: explainable AI in predicting no-show trends before and during the COVID-19 pandemic. *Bioengineering (Basel).* 2025;12(3):227.

Ulloa-Pérez E, Blasi PR, Westbrook EO, Lozano P, Coleman KF, Coley RY. Pragmatic randomized study of targeted text message reminders to reduce missed clinic visits. *Perm J.* 2022;26(1):64-72.

Tuan WJ, Yan Y, Abou Al Ardat B, Felix T, Chen Q. Predicting missed appointments in primary care: a personalized machine learning approach. *Ann Fam Med.* 2025;23(4):294-301.

Alawadhi A, Jenkins D, Palin V, van Staa T. Development and evaluation of prediction models to improve the hospital appointments overbooking strategy at a large tertiary care hospital in the Sultanate of Oman: a retrospective analysis. *BMJ Open.* 2025;15(4):e093562.

Borges A, Carvalho M, Maia M, Guimarães M, Carneiro D. Predicting and explaining absenteeism risk in hospital patients before and during COVID-19. *Socioecon Plann Sci.* 2023;87:101549.

Hamdan AF, Bakar AA. Machine learning predictions on outpatient no-show appointments in a Malaysia major tertiary hospital. *Malays J Med Sci.* 2023;30(5):169-81.

Salazar LH, Parreira WD, Fernandes AM, Leithardt VR. No-show in medical appointments with machine learning techniques: a systematic literature review. *Information (Basel).* 2022;13(11):507.