

ORIGINAL RESEARCH

Open access

Transparent Machine Learning Model for Predicting Delayed Follow-Up After Emergency Department Visits Using Discharge Instructions, Appointment Availability, Portal Engagement, Social Risk Data, and Visit Severity

Anna Novak^{1*}, Tomas Hruby¹

Abstract

Many patients discharged from emergency departments require timely outpatient follow-up to complete diagnostic, therapeutic, or monitoring plans. When follow-up is delayed, unresolved symptoms, missed diagnoses, medication problems, and preventable return visits may occur. Current discharge workflows often rely on generic instructions and assume that patients can understand, schedule, and attend recommended care. Existing prediction approaches do not consistently combine unstructured discharge instructions, appointment access, portal engagement, social risk, and visit severity in a transparent way. This article proposes a transparent machine learning framework for predicting delayed follow-up after emergency department visits. The objective is to support patient-specific discharge planning by identifying both the likelihood of delay and the most actionable contributing barriers. The proposed model would combine structured emergency department and scheduling data with natural language processing features extracted from discharge instructions. A gradient-boosted tree model with SHAP-based explanations would provide patient-level and population-level interpretability. Conceptually, the model would identify patients at elevated risk of delayed follow-up and attribute that risk to factors such as unclear instructions, limited appointment availability, absent portal engagement, transportation barriers, or higher visit complexity. These explanations would be intended to guide targeted interventions rather than replace clinical judgment. A transparent model for delayed follow-up prediction could enable precision transitional care after emergency department discharge. By aligning predictions with actionable explanations, care teams could direct limited resources toward the specific barrier most likely to prevent timely follow-up.

Keywords Social determinants of health, Explainable artificial intelligence, Emergency department, Patient portal, Delayed follow-up, Discharge instructions

*Correspondence:

Anna Novak

anna.novak@gmail.com

¹ Department of Health Informatics and Digital Engineering, Faculty of Medicine, Czech Technical University, Prague, Czech Republic

Introduction

Delayed outpatient follow-up after an emergency department visit is a clinically important transition failure because the ED often initiates care plans that require outpatient confirmation, monitoring, or escalation. Machine learning studies of ED return visits and disposition have shown that post-discharge risk is partially predictable from clinical and utilization

data, suggesting that follow-up delay could also be anticipated before the patient leaves the ED [1, 2]. When delayed follow-up contributes to deterioration or repeat acute care use, the economic burden extends beyond the index visit to avoidable testing, crowding, and care fragmentation [3, 4]. A transparent prediction model focused on delayed follow-up would therefore address a transition point where clinical risk and operational cost intersect [5, 6].

Current practice frequently treats discharge as an information-transfer event rather than a coordinated access process. Discharge instructions may be difficult to understand, may omit concrete follow-up steps, or may not match the patient's literacy, language, or capacity to act after leaving the ED [7, 8]. Even when instructions are clinically appropriate, patients may face limited appointment availability, lack of portal engagement, unstable housing, transportation problems, or insurance barriers that make timely follow-up unrealistic [9, 10]. This mismatch between written recommendations and real-world feasibility creates a need for models that explain not only who is at risk, but why the risk exists [11, 12].

The data needed to foresee follow-up failure already exist across the electronic health record, scheduling platform, patient portal, and social risk documentation. ED discharge text can be processed to identify the presence, specificity, readability, and actionability of follow-up recommendations, while scheduling systems can indicate whether a near-term appointment is available in the appropriate clinic [7, 13]. Portal activation, test-result viewing, message access, and appointment interaction may provide behavioral signals of digital reachability after discharge [14, 15]. Social risk screens, ICD-coded social needs, and visit severity variables can further distinguish patients whose follow-up delay reflects structural barriers rather than lack of clinical urgency [10, 16].

This article proposes a transparent machine learning model that would estimate the probability of delayed follow-up after ED discharge and explain the patient-specific drivers of that prediction. The model is conceptual and model-oriented: it does not report experimental results, training outcomes, or performance statistics, but instead describes a clinically plausible architecture for integrating discharge instructions, appointment availability, portal engagement, social risk, and visit severity. By using interpretable modeling or SHAP-based explanation layers, the system could translate a risk estimate into an actionable discharge-planning recommendation. The central thesis is that prediction becomes clinically useful only when it is paired with explanations that identify modifiable barriers at the point of care.

Background

Post-ED follow-up adherence and outcomes

Post-ED follow-up adherence is shaped by the urgency of the unresolved condition, the complexity of discharge recommendations, and the patient's ability to navigate outpatient care after leaving a time-pressured acute setting. Studies predicting ED revisits, unscheduled returns, and post-discharge service use indicate that downstream outcomes are not random but reflect recognizable patterns in prior utilization, presenting condition, visit severity, and discharge context [1, 3]. Although return visits are not identical to delayed follow-up, both outcomes arise from failures or vulnerabilities in the transition from emergency care to longitudinal care [2, 4]. A transparent delayed-follow-up model would build on this evidence by shifting attention from predicting acute care reuse alone to identifying preventable gaps in planned outpatient continuity [5].

Discharge instructions as a communication tool

Discharge instructions are central to ED follow-up because they translate the clinician's plan into patient action, yet their effectiveness depends on clarity, readability, completeness, and the presence of concrete steps. Research on emergency discharge communication has shown that patients may misunderstand or fail to recall instructions, particularly when the information is complex or not tailored to their needs [8, 13]. Recent informatics work using large language models and multilingual instruction generation underscores the growing interest in improving the completeness and comprehensibility of discharge documents [7, 17]. NLP-derived features from discharge instructions could therefore serve as predictors and

explanations, identifying whether a delayed-follow-up risk is driven by missing follow-up tasks, unclear timing, or inaccessible wording [18].

Appointment availability and portal engagement

Timely follow-up depends not only on what the discharge instructions recommend, but also on whether outpatient appointments are actually available and whether the patient can engage with digital tools used to coordinate care. Patient portal studies show that portal access, test-result viewing, and appointment-related interactions can reflect important differences in communication access and care navigation after ED or outpatient encounters [14, 15]. At the same time, web-based portals may reduce missed appointments for some patients while excluding others who lack activation, stable internet access, or trust in digital systems [11, 12]. A delayed-follow-up model should therefore treat portal engagement and appointment availability as contextual feasibility indicators rather than simple measures of patient motivation [9, 19].

Social risk data and health disparities

Social risk factors such as housing instability, food insecurity, transportation barriers, lack of insurance, social isolation, and unstable contact information may directly influence whether a patient can complete follow-up. Studies of unmet social needs and missed appointments suggest that nonattendance is often a structural problem rather than an individual preference, especially when barriers involve transportation, competing basic needs, or financial instability [10, 20]. ED-based social determinant screening further indicates that missed or incomplete social risk documentation can obscure the very factors most relevant to transitional care [16]. Including social risk data in a transparent model would allow care teams to distinguish clinical complexity from access barriers and to route patients toward social work, transportation support, or outreach [21, 22].

Transparent machine learning in clinical care

Transparent machine learning is especially important in emergency care because predictions are generated in fast-moving environments where clinicians must understand why a recommendation is being made. ED machine learning models for disposition, return visits, and disease-specific outcomes demonstrate that predictive algorithms can support clinical risk stratification, but opaque scores may be difficult to trust or act on without interpretable feature attribution [1, 6]. Explainable approaches such as SHAP can help identify global patterns and patient-level drivers, making the model's output more usable for care coordination [2]. For delayed follow-up, interpretability is not an optional technical feature; it is the mechanism by which risk prediction becomes a targeted intervention plan [23, 24].

Model Development Overview

High-level predictive pipeline

At the time of ED discharge, the proposed pipeline would assemble structured and unstructured data from the discharge workflow, scheduling system, patient portal, social risk fields, and visit record. The model would output a conceptual probability of 30-day follow-up delay and an explanation that attributes the prediction to specific contributing factors, such as missing follow-up instructions, long scheduling lead time, lack of portal activation, transportation barriers, or high visit complexity [2, 5]. The pipeline would be designed to operate before the patient leaves the ED, because interventions are most feasible while clinicians and care coordinators can still clarify instructions or schedule appointments [8, 14]. Rather than producing a stand-alone risk score, the system would function as an explanation-oriented discharge support tool [6].

Figure 1 illustrates the proposed transparent machine learning workflow from ED discharge data capture through SHAP-based explanation, barrier profiling, targeted intervention, and governance monitoring.




Figure 1. Transparent Machine Learning Workflow for Predicting Delayed Follow-Up After Emergency Department Discharge.

Core input features

The core feature set would include NLP-derived discharge instruction measures, appointment availability metrics, portal engagement signals, social risk flags, and visit severity indicators. Discharge instruction features would represent whether follow-up is mentioned, whether the timing is explicit, whether contact information is included, and whether the language appears readable and specific [7, 18]. Appointment and portal features would indicate whether a relevant clinic slot exists soon after discharge, whether the patient has an active portal account, and whether prior portal behaviors suggest that electronic communication is likely to reach the patient [11, 14]. Social risk and severity features would capture barriers such as transportation or housing instability alongside triage acuity, comorbidity burden, and ED disposition context [10, 16].

Design principles

The model would be designed around three principles: transparency, actionability, and clinical support. Transparency means that every prediction should be accompanied by understandable feature contributions rather than a black-box risk label, consistent with the use of explainable machine learning in ED prediction studies [2, 6]. Actionability means that the explanation should point to interventions that can be performed at discharge, such as clarifying instructions, booking an appointment, sending a portal message, arranging transportation, or triggering a phone call [11, 13]. Clinical support means that the model should augment, not replace, the discharge conversation by helping clinicians and care coordinators recognize barriers that might otherwise remain hidden [6, 8].

Data Sources and Feature Engineering

Extracting discharge instruction features

Discharge instruction text would be processed using NLP methods to identify whether a specific follow-up plan is stated, whether the recommended timeframe is explicit, and whether the instruction includes a phone number, clinic name, portal link, or scheduling directive. Readability and clarity features would be especially important because emergency discharge studies show that patients may leave with incomplete understanding even when instructions are provided [8, 13]. Recent work on AI-generated and multilingual discharge instructions further supports the idea that instruction completeness, language accessibility, and patient comprehension can be evaluated as measurable documentation characteristics [7, 17]. These text-derived features would allow the model to explain delayed-follow-up risk as a communication problem when the discharge plan is vague, jargon-heavy, or operationally incomplete [18].

Appointment availability and portal engagement features

Appointment availability features would be derived from scheduling systems by identifying whether the recommended service has near-term capacity, whether the lead time aligns with the discharge plan, and whether the appointment type requires additional patient action. Portal engagement features would include activation status, recent login behavior, message viewing, appointment interaction, and evidence that the patient has used the portal to access ED-related results or instructions [14, 15]. Prior studies of portal use, missed appointments, and reminder interventions suggest that digital engagement can shape follow-up behavior, but also that portal-based workflows may not benefit all patients equally [11, 12]. The model would therefore encode portal disengagement as a possible outreach need rather than as a judgment about patient responsibility [9, 19].

Social risk and visit severity features

Social risk features would combine structured screening results, documented needs, and relevant diagnostic codes that indicate barriers such as housing instability, transportation difficulty, food insecurity, insurance gaps, or social isolation. Evidence linking unmet social needs to no-show visits and missed imaging appointments supports the inclusion of social risk as a central predictor of follow-up delay [10, 20]. Visit severity features would include triage acuity, comorbidities, ED observation or hold status, diagnostic complexity, and the urgency of recommended follow-up, because higher-acuity visits may create stronger clinical need while also reflecting patients with more complex care navigation demands [6, 23]. Combining social risk and severity would help the model distinguish patients who need routine reminders from those who need intensive transition support [16, 21].

Table 1 presents the proposed feature architecture by linking each predictor domain to its interpretive meaning and to a discharge-planning action that could be taken before the patient leaves the emergency department.

Table 1. Barrier-Specific Feature Architecture for Transparent Prediction of Delayed Post-ED Follow-Up

Follow-up delay barrier domain	Manuscript-specific feature constructs	Data source	Interpretive meaning of elevated risk	Actionable discharge-planning response
Discharge instruction ambiguity	Missing follow-up timeframe; vague specialty referral wording; absence of clinic name, phone number, or scheduling directive; high reading complexity; lack of plain-language action steps	ED discharge instruction text processed through NLP	The patient may leave with an unclear plan or insufficient operational information to complete follow-up [7, 8, 13, 18].	Revise instructions before discharge; add explicit timing, clinic contact, next step, and plain-language explanation.
Appointment access limitation	No available appointment within recommended window; long specialty lead time; appointment type requiring patient self-scheduling; mismatch between recommended urgency and visible capacity	Scheduling platform and outpatient appointment system	Delay may be caused by system capacity rather than patient behavior [9, 19, 25, 26].	Schedule before discharge; escalate to transition clinic; route to access coordinator; document capacity constraint.
Limited digital reachability	No portal activation; no recent login; no message viewing; no appointment interaction; no evidence of ED-related result review	Patient portal logs	Portal-based follow-up reminders may not reach the patient reliably and should not	Use phone call, printed plan, caregiver contact, mailed instructions, or nonportal outreach.

			be assumed sufficient [11, 12, 14, 15].	
Social risk and access instability	Transportation difficulty; housing instability; insurance barrier; unstable contact information; food insecurity; social isolation; missing social risk screening	Social risk screening, ICD-coded needs, EHR documentation	Follow-up delay may reflect structural barriers or incomplete documentation rather than low motivation [10, 16, 20, 21].	Refer to social work, transportation support, community health worker outreach, benefits navigation, or care coordination.
Visit severity and complexity	Higher triage acuity; ED observation or hold status; complex diagnostic testing; comorbidity burden; urgent outpatient monitoring need	ED visit record, problem list, diagnosis and procedure data	The clinical need for follow-up may be high, while the care pathway may be difficult to navigate [6, 23].	Prioritize transitional care review; confirm clinical urgency; provide direct scheduling and closed-loop follow-up.
Prior utilization and care navigation history	Prior missed appointments; repeat ED use; fragmented outpatient care; incomplete prior follow-up; multiple specialty referrals	EHR utilization history and scheduling history	The current discharge may occur in a broader pattern of unresolved transition failures [1-4].	Trigger care coordinator outreach; consolidate referrals; verify contact method; consider higher-intensity transition support.
Missingness as signal	Missing portal data; absent social risk screen; incomplete scheduling fields; templated or copied discharge instructions	Cross-source EHR data quality review	Missing data may reflect documentation gaps, digital exclusion, or unmeasured social barriers [9, 10, 16, 21].	Flag uncertainty; avoid punitive interpretation; prompt staff to verify barriers directly with the patient.

Transparent Model Architecture

Model choice and interpretability approach

The proposed architecture could use either an inherently interpretable model, such as regularized logistic regression or a constrained decision tree, or a gradient-boosted tree model paired with SHAP explanations. Gradient-boosted models are attractive because they can represent nonlinear relationships among discharge text features, appointment access, portal engagement, social risk, and visit severity, while SHAP values can translate those relationships into feature-level contributions [2, 6]. For clinical implementation, the model choice should prioritize clear feature effect reporting, calibration, subgroup review, and ease of explanation over marginal technical complexity [3]. In this setting, interpretability would be judged by whether the output helps care teams decide what to do before discharge, not merely by whether the algorithm is mathematically transparent [5, 24].

Input feature vector and preprocessing

The input feature vector would combine continuous scores, categorical indicators, binary flags, and text-derived measures into a standardized representation available at discharge. NLP features such as readability, follow-up specificity, presence of contact information, and clarity of timing would be aligned with structured features such as specialty type, appointment lead time, portal activation, social risk flags, and triage severity [7, 8]. Categorical variables would be encoded in ways that preserve interpretability, and missing portal data would be treated carefully because

absence of portal activity may reflect lack of access, nonactivation, or data unavailability rather than true disengagement [9, 14]. Missing social risk fields would likewise require explicit handling because incomplete screening can obscure disparities and distort explanation profiles [10, 16].

Output: follow-up delay probability and explanation

The model output would consist of a calibrated probability of delayed follow-up accompanied by a patient-level explanation, such as a SHAP waterfall plot or concise textual summary. The probability would indicate relative risk, while the explanation would identify the strongest contributors, such as no explicit follow-up timeframe, no appointment availability, no portal engagement, transportation risk, or high visit complexity [2, 6]. For usability, the explanation should be translated into discharge workflow language, so that “long scheduling lead time” prompts scheduling assistance and “unclear discharge plan” prompts instruction revision [8, 12]. This architecture would make the model both predictive and operationally meaningful, aligning the risk estimate with modifiable transition barriers [5, 19].

Identifying Key Drivers and Producing Explanations

Global SHAP analysis of delay drivers

Global SHAP analysis would summarize which features most consistently increase the predicted likelihood of delayed follow-up across the modeled ED discharge population. In a conceptually expected pattern, missing or vague follow-up instructions, absence of near-term appointment availability, limited portal engagement, transportation barriers, and greater visit complexity could emerge as prominent drivers because each represents a distinct failure point in the transition from ED care to outpatient care [2, 8, 14]. Outpatient no-show prediction studies further suggest that scheduling context, prior attendance behavior, and access-related constraints can be important predictors of missed visits, supporting the inclusion of appointment-system features in a delayed-follow-up model [25, 26]. The value of global explanations would be to guide system redesign by identifying whether follow-up delay is more strongly associated with communication, capacity, digital access, social risk, or severity-related complexity [10, 27].

Patient-level explanations for discharge planners

Patient-level explanations would convert a risk estimate into a practical discharge-planning narrative for a specific patient. For example, a SHAP waterfall could indicate that elevated predicted delay is primarily attributable to no documented follow-up timeframe, absence of portal activity, and lack of a near-term appointment, which would prompt the discharge planner to clarify instructions, schedule before discharge, and arrange nonportal outreach [8, 11, 14]. Patient portal studies show that access and engagement vary across ED and outpatient populations, so explanation logic should treat nonuse as a potential communication barrier rather than as evidence of disinterest [9, 15]. By making the reason for risk visible, the model could support tailored intervention without requiring clinicians to infer barriers from a single opaque score [2, 19].

Interaction effects between social risk and system factors

Interaction explanations would be important because social risk may amplify the effect of system-level barriers rather than operating independently. A long scheduling lead time may be manageable for a patient with stable transportation, flexible work, and portal access, but much less manageable for a patient with housing instability, food insecurity, or unreliable transportation [10, 20]. SHAP dependence plots could conceptually reveal that appointment wait time becomes more harmful when paired with social risk flags, indicating that the same operational delay may create unequal consequences across patients [2, 16]. These interaction patterns would help care teams avoid interpreting follow-up delay as a purely behavioral outcome and instead recognize how health-system constraints and social needs combine to produce inequitable transitions [21, 22].

Counterfactual explanations for resource allocation

Counterfactual explanations would translate model outputs into resource-allocation questions by asking which modifiable change would be expected to reduce follow-up delay risk for a given patient or subgroup. Instead of reporting experimental effect sizes, the model could conceptually show that earlier appointment access, clearer instructions, portal-independent outreach, or transportation support would be expected to move the prediction in a favorable direction [12, 19, 27]. This approach would be consistent with predictive no-show literature, where model-informed interventions are most useful when they guide a feasible operational response rather than merely ranking patients by risk [25, 28]. Counterfactual explanation should therefore be framed as decision support for prioritizing scarce transitional care resources, not as proof that a single intervention will work without prospective evaluation [5, 27].

Explaining Predictions to Stakeholders

Explanations for ED clinicians

For ED clinicians, the explanation should be concise enough to fit into the discharge workflow without creating alert fatigue. A brief note could state that the patient is at elevated risk of delayed follow-up because the discharge instructions lack a specific appointment timeframe, no appointment is currently visible, and prior portal engagement is absent [8, 14]. This format would align with the need for transparent clinical decision support in fast-paced ED environments, where clinicians must quickly determine whether the model's concern is plausible and actionable [2, 6]. The explanation should support conversation and documentation, allowing the clinician to revise the discharge plan before the patient leaves the ED [13, 18].

Explanations for transitional care coordinators

For transitional care coordinators, the model explanation should provide a more detailed operational view than the clinician-facing summary. A coordinator could see that the risk profile is driven by no available specialty appointment, a transportation barrier, and lack of portal activation, which would support direct scheduling, ride coordination, or post-discharge telephone outreach [9, 10, 22]. Evidence from portal reminder and appointment adherence studies suggests that follow-up interventions are more likely to be useful when they are matched to the reason a patient may miss care [19, 29]. The coordinator-facing explanation would therefore act as a task-routing tool that links model attribution to concrete follow-up support [5, 27].

Audit and fairness dashboard

An audit and fairness dashboard would continuously compare predicted risk, explanation patterns, intervention use, and observed follow-up outcomes across demographic, insurance, language, disability, and social risk groups. Because social risk documentation may be incomplete and missed follow-up may reflect structural inequities, the dashboard should assess whether the model systematically overpredicts, underpredicts, or recommends different interventions for particular groups [10, 16]. Transparent modeling can improve trust, but explainability alone does not guarantee fairness if the source data encode unequal access to appointments, portals, or screening [9, 11]. A governance process should therefore review both predictive behavior and explanation behavior, especially when social risk or portal disengagement contributes strongly to the model output [2, 21].

Patient-facing explanations (optional)

Patient-facing explanations should be optional, plain-language, and focused on support rather than risk labeling. A patient message could explain that follow-up is important because the ED visit identified a condition needing outpatient review, and that the care team can help schedule the visit, clarify the instructions, or arrange a nonportal contact method if needed [8, 12]. Patient-facing wording should avoid suggesting that the patient is personally responsible for anticipated delay, particularly when the model identifies transportation, housing, appointment capacity, or digital access barriers [9,

10]. If used, these explanations should be evaluated for comprehension, trust, stigma, and whether they improve engagement with follow-up planning [7, 17].

Integration into ED Discharge and Transitional Care Workflow

Embedding in the electronic health record discharge process

The model would run when the discharge order is initiated or when discharge instructions are finalized, using the most current information available from the ED record, scheduling system, portal logs, and social risk fields. Its output would appear within the discharge workflow as a risk estimate paired with the leading explanation factors, while high-concern explanation profiles could automatically alert the care coordination team [5, 6]. Embedding the model at this point is important because instruction revision, appointment scheduling, portal-independent contact planning, and transportation support are most actionable before the patient physically leaves the ED [8, 19]. The workflow should minimize extra clicks and present the explanation in language that maps directly to discharge tasks [2, 18].

Triggered interventions based on explanation profiles

Triggered interventions would be matched to the explanation profile rather than applied uniformly to all high-risk patients. If the main driver is unclear discharge text, the clinician could add specific follow-up timing, clinic contact information, and plain-language instructions; if the driver is appointment access, staff could schedule before discharge or escalate to a transition clinic [7, 13]. If the driver is social risk, the workflow could refer to social work, transportation support, or community health worker outreach; if the driver is portal disengagement, the system could prompt a phone call, printed plan, or alternative communication pathway [9, 10]. This explanation-based approach would make the model a routing mechanism for targeted transitional care rather than a general warning system [5, 27].

Table 2 translates model explanations into stakeholder-specific interventions, governance checks, and evaluation indicators for operational deployment in ED discharge and transitional care workflows.

Table 2. Explanation-to-Intervention Logic for Operational Use of the Delayed Follow-Up Model

Model explanation profile	Example patient-level explanation	Primary stakeholder	Recommended operational response	Equity and governance consideration	Evaluation indicator
Communication-driven risk	“Predicted delay is mainly driven by unclear follow-up timing and missing clinic contact information.”	ED clinician or discharge nurse	Rewrite discharge instructions with specific timing, clinic name, phone number, warning signs, and next action.	Ensure instruction clarity is reviewed across language, literacy, disability, and interpreter-use groups [7, 8, 13, 17].	Proportion of high-risk discharges with revised follow-up instructions before departure.
Capacity-driven risk	“Risk is elevated because the recommended specialty clinic has no available appointment within the	Scheduling staff or transition clinic coordinator	Book appointment before discharge, escalate to urgent access pathway, or document unavailable capacity.	Avoid blaming patients for delays caused by outpatient access shortages [9, 19, 27].	Time from ED discharge to scheduled appointment; proportion scheduled before discharge.

	clinically appropriate window.”				
Digital-access-driven risk	“Risk is elevated because portal activation and prior portal engagement are absent.”	Care coordinator or nurse navigator	Use phone outreach, printed plan, caregiver-supported communication, or nonportal reminder pathway.	Treat portal nonuse as possible access limitation rather than patient disengagement [9, 11, 14, 15].	Completion of nonportal outreach among patients without portal engagement.
Social-barrier-driven risk	“Risk is elevated because transportation difficulty and unstable contact information are documented.”	Social worker, community health worker, or transitional care team	Arrange transportation support, verify phone/address, connect to community resources, and schedule follow-up with barrier support.	Monitor whether social risk features lead to supportive intervention rather than stigmatizing labels [10, 16, 20, 22].	Referral completion for transportation/social support; follow-up completion by social risk subgroup.
Clinical-complexity-driven risk	“Risk is elevated because visit severity, diagnostic complexity, and urgent monitoring need are high.”	ED clinician and transitional care coordinator	Prioritize closed-loop follow-up, confirm urgency, communicate with receiving clinic, and create escalation plan.	Ensure high-complexity patients are not lost when multiple services share responsibility [6, 23].	Completed follow-up within recommended clinical window.
Interaction-driven inequity risk	“Long appointment lead time contributes more strongly to delay when combined with transportation or housing instability.”	Quality, equity, and operations leadership	Review access policies, transportation support availability, and specialty scheduling pathways for structurally vulnerable groups.	Evaluate whether the same operational barrier creates unequal follow-up consequences across patient groups [10, 16, 21].	Subgroup calibration; disparity in delayed follow-up by transportation, housing, language, insurance, or disability status.
Uncertainty or missingness-driven risk	“Prediction relies on incomplete portal, social risk, or scheduling data.”	Governance committee and frontline team	Prompt direct barrier verification; suppress overconfident recommendation; improve documentation workflow.	Missingness may represent unmeasured need or digital exclusion and should be audited explicitly [9, 10, 16].	Missing-data rate by subgroup; frequency of uncertainty flags; staff override documentation.
System-learning signal	“Global SHAP patterns show delay is most	ED operations, ambulatory	Redesign discharge templates, expand transition clinic	Explainability should support system redesign,	Monthly change in delay rate, intervention completion, workload

	frequently driven by appointment capacity and unclear instructions.”	access leadership, and quality improvement team	capacity, adjust scheduling rules, and monitor intervention uptake.	not merely produce another risk score [2, 5, 27].	burden, and equity of benefit.
--	--	---	---	---	--------------------------------

Evaluation Strategy

Predictive performance metrics

Evaluation should examine discrimination, calibration, precision-recall behavior, and subgroup performance, but without treating numerical model performance as sufficient evidence of clinical value. Because delayed follow-up is an access-sensitive outcome, subgroup evaluation should assess whether the model performs differently by race, insurance, language, age, disability, portal access, and documented social risk [9, 10, 16]. Prior ED prediction and outpatient no-show modeling studies show that technically plausible models can be developed for transition-related outcomes, but implementation decisions require attention to calibration, generalizability, and operational relevance [1, 3, 25]. The model should therefore be evaluated as a clinical workflow tool rather than as a stand-alone statistical artifact [6, 27].

Explanation quality and clinical acceptability

Explanation quality should be evaluated through structured feedback from ED clinicians, nurses, discharge staff, care coordinators, and social workers. These stakeholders should assess whether the explanation is understandable, whether the attributed drivers are clinically plausible, and whether the suggested intervention pathway matches what can realistically be done at discharge [2, 8]. Acceptability should also include whether explanations reduce cognitive burden or create new work that is not supported by staffing, scheduling capacity, or community resources [5, 22]. Because patient portal and discharge-instruction interventions can affect different patients differently, explanation review should include equity-focused assessment of whether the model frames barriers respectfully and avoids stigmatizing patients with social risk or digital-access limitations [9, 29].

Prospective impact on follow-up rates

Prospective evaluation should compare model-informed targeted discharge planning with standard discharge processes in a pragmatic design that measures completed follow-up, care-team workload, patient experience, and equity of benefit. The purpose would not be to prove that prediction alone improves care, but to test whether explanation-guided interventions help patients overcome the specific barriers identified at discharge [19, 27]. Prior work on predictive no-show interventions and outpatient appointment modeling supports the need to move beyond retrospective prediction toward implementation studies that examine whether model outputs change operations and outcomes [26, 28]. A prospective design should also monitor unintended effects, such as overreliance on the model, unequal allocation of care coordination, or failure to address structural appointment shortages [20, 25].

Limitations

Data quality and missingness

The proposed model would depend on data elements that may be incomplete, inconsistent, or only partially reflective of patient experience. Discharge instruction text may be templated, copied forward, or too generic to reveal the true quality of the discharge conversation, while social risk fields may be missing because screening was not performed rather than because no need exists [8, 10]. Appointment availability may also change after discharge, and portal engagement data may underestimate communication access for patients who rely on caregivers, telephone calls, printed instructions, or

community support [9, 14]. These limitations mean that missingness should be modeled and explained carefully, especially when data absence could reflect structural exclusion from documentation or digital systems [16, 21].

Generalizability and equity

Generalizability would be limited by differences in ED workflows, scheduling infrastructure, portal adoption, social risk screening practices, specialty access, and transitional care resources across health systems. A model developed in a system with robust portal use and transition clinics may not transfer cleanly to settings with limited outpatient capacity or different documentation norms [6, 11]. Equity concerns would remain even with transparent explanations because an accurate model can still reproduce inequitable patterns if it recommends interventions that are unavailable to under-resourced patients or clinics [10, 20]. Multi-site validation and governance would therefore be necessary to ensure that prediction and explanation improve follow-up support rather than merely documenting barriers that the health system fails to address [2, 27].

Conclusion

A transparent machine learning model for delayed follow-up after emergency department discharge could support a more precise approach to transitional care. By combining discharge instructions, appointment availability, portal engagement, social risk data, and visit severity, the model would treat follow-up completion as a system-sensitive outcome rather than a simple patient behavior.

The major strength of this framework is that it links prediction to explanation and explanation to action. Instead of identifying a patient as generically high risk, the model would indicate whether the likely barrier involves unclear instructions, lack of appointment capacity, limited digital engagement, social needs, or clinical complexity.

Important challenges would remain before implementation. Data completeness, social risk documentation quality, clinician acceptance, workflow burden, and the availability of real interventions would all determine whether the model improves care rather than merely adding another risk score.

Future multi-site implementation studies should evaluate not only predictive accuracy, but also whether explanation-guided discharge support reduces delayed follow-up and narrows health disparities. The central goal should be a learning transition system that helps patients complete the care that the emergency department recommends.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 04 Apr 2026

Revised: 29 May 2026

Accepted: 14 Jun 2026

Published online: 20 July 2026

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

- Hong WS, Haimovich AD, Taylor RA. Predicting 72-hour and 9-day return to the emergency department using machine learning. *JAMIA Open*. 2019;2(3):346-52.
- Chmiel FP, Burns DK, Azor M, Borca F, Boniface MJ, Zlatev ZD, et al. Using explainable machine learning to identify patients at risk of reattendance at discharge from emergency departments. *Sci Rep*. 2021;11(1):21513.
- Lee YC, Ng CJ, Hsu CC, Cheng CW, Chen SY. Machine learning models for predicting unscheduled return visits to an emergency department: a scoping review. *BMC Emerg Med*. 2024;24(1):20.
- Hsu CC, Chu CC, Lin CH, Huang CH, Ng CJ, Lin GY, et al. A machine learning model for predicting unscheduled 72 h return visits to the emergency department by patients with abdominal pain. *Diagnostics (Basel)*. 2021;12(1):62.
- Lee SY, Eagleson RM, Hearld LR, Gibson MJ, Hearld KR, Hall AG, et al. Leveraging machine learning to enhance appointment adherence at a novel post-discharge care transition clinic. *JAMIA Open*. 2024;7(4):ooae086.
- De Hond A, Raven W, Schinkelshoek L, Gaakeer M, Ter Avest E, Sir O, et al. Machine learning for developing a prediction model of hospital admission of emergency department patients: Hype or hope? *Int J Med Inform*. 2021;152:104496.
- Gimeno A, Krause K, D'Souza S, Walsh CG. Completeness and readability of GPT-4-generated multilingual discharge instructions in the pediatric emergency department. *JAMIA Open*. 2024;7(3):ooae050.
- Hoek AE, Anker SC, van Beeck EF, Burdorf A, Rood PP, Haagsma JA. Patient discharge instructions in the emergency department and their effects on comprehension and recall of discharge instructions: a systematic review and meta-analysis. *Ann Emerg Med*. 2020;75(3):435-44.
- Holmes C, Holmes K, Scarborough J, Hunt J, d'Etienne JP, Ho AF, et al. The status of patient portal use among Emergency Department patients experiencing homelessness: A large-scale single-center observational study. *Am J Emerg Med*. 2023;66:118-23.
- Fiori KP, Heller CG, Rehm CD, Parsons A, Flattau A, Braganza S, et al. Unmet social needs and no-show visits in primary care in a US northeastern urban health system, 2018-2019. *Am J Public Health*. 2020;110(S2):S242-S250.
- Zhong X, Liang M, Sanchez R, Yu M, Budd PR, Sprague JL, et al. On the effect of electronic patient portal on primary care utilization and appointment adherence. *BMC Med Inform Decis Mak*. 2018;18(1):84.
- Graham TA, Ali S, Avdagovska M, Ballermann M. Effects of a web-based patient portal on patient satisfaction and missed appointment rates: survey study. *J Med Internet Res*. 2020;22(5):e17955.

Sheikh H, Brezar A, Dzwonek A, Yau L, Calder LA. Patient understanding of discharge instructions in the emergency department: do different patients need different approaches? *Int J Emerg Med.* 2018;11(1):5.

Turer RW, McDonald SA, Lehmann CU, Thakur B, Dutta S, Taylor RA, et al. Real-time electronic patient portal use among emergency department patients. *JAMA Netw Open.* 2024;7(5):e249831.

Foster B, Krasowski MD. The use of an electronic health record patient portal to access diagnostic test results by emergency patients at an academic medical center: retrospective study. *J Med Internet Res.* 2019;21(6):e13791.

Samuels-Kalow ME, Mayes K, Cash RE, Schiavoni KH, Vogeli C, Thorndike AN, et al. Missed screening for adverse social determinants of health and emergency department utilization. *Ann Emerg Med.* 2024;83(4):416-8.

Huang T, Safranek C, Socrates V, Chartash D, Wright D, Dilip M, et al. Patient-representing population's perceptions of GPT-generated versus standard emergency department discharge instructions: randomized blind survey assessment. *J Med Internet Res.* 2024;26:e60336.

Russell S, Jacobson N, Pavlic A. Improving patient understanding of emergency department discharge instructions. *West J Emerg Med.* 2024;25(6):917.

Atta S, Brown RB, Wasser LM, Mayer N, Cassidy J, Liu PJ, et al. Effect of a patient portal reminder message after no-show on appointment reattendance in ophthalmology: a randomized clinical trial. *Am J Ophthalmol.* 2024;263:93-8.

Sotudian S, Afran A, LeBedis CA, Rives AF, Paschalidis IC, Fishman MD. Social determinants of health and the prediction of missed breast imaging appointments. *BMC Health Serv Res.* 2022;22(1):1454.

Lefchak B, Cushwa A, Kersten H, Courts K, McPeak K. Characterization of social risk factors among newborns seen at an urban pediatric primary care predictive of appointment nonattendance during the first 6 months of life. *Health Equity.* 2022;6(1):heq-2021.

Fong A, Boxley C, Schubel L, Gallagher C, AuBuchon K, Arem H. Identifying Complex Scheduling Patterns Among Patients With Cancer With Transportation and Housing Needs: Feasibility Pilot Study. *JMIR Cancer.* 2025;11(1):e57715.

Goto T, Camargo CA Jr, Faridi MK, Yun BJ, Hasegawa K. Machine learning approaches for predicting disposition of asthma and COPD exacerbations in the ED. *Am J Emerg Med.* 2018;36(9):1650-4.

Zhou Y, Swoboda TK, Ye Z, Barbaro M, Blalock J, Zheng D, et al. Using machine learning algorithms to predict patient portal use among emergency department patients with diabetes mellitus. *J Clin Med Res.* 2023;15(3):133.

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.

Mochón F, Elvira C, Ochoa A, Gonzalez JC. Machine-learning-based no show prediction in outpatient visits.

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

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.

Goldstein S, Buuck K, Thompson H, Yu R, Small N, des Bordes J, et al. Patient Portal Access and Missed Medical Appointments: A Case-Control Study. *J Patient Exp.* 2026;13:23743735261415666.