



**Journal of Contemporary
Healthcare Analytics**
DLpress is a publisher of
scholarly books and
peer-reviewed scientific
research. With a dedication
to academic excellence,
DLpress publishes books and
research papers on a diverse
range of topics spanning
various disciplines, including
but not limited to, science,
technology, engineering,
mathematics, social sciences,
humanities, and arts.
Published 12, June, 2022

An End-to-End Predictive and Intervention Framework for Reducing Hospital Readmissions

Dhinakaran Rajendaran¹

¹PharmD, MSHI; Lead BI Developer, Henry Ford Health

RESEARCH ARTICLE

Abstract

Hospital readmissions within a given days of discharge are an enduring issue to the healthcare systems of the world. Early readmissions are largely viewed as indicators of poor quality of care, and they impose high economic costs and penalties to providers, valued at billions of dollars of avoidable expenditure annually. This paper presents an end-to-end solution that merges predictive analytics with accurate clinical interventions in a bid to reduce 30-day hospital readmission rates. The model consists of data importation from electronic health records, a predictive module based on machine learning to forecast individual readmission risk at discharge, and an intervention engine that accepts predicted risk and translates it into actions such as enhanced discharge planning and post-discharge follow-up. The system also consists of closed-loop monitoring to learn from patient outcomes and iteratively refine its predictions and interventions. This work validates the framework on a large clinical dataset of diabetic patient admissions, demonstrating how careful feature engineering and a variety of machine learning models (k-nearest neighbors, logistic regression, and random forests) can identify high-risk patients. The results suggest that the prediction model has good discriminative performance (with area under the ROC curve of around 0.70) and that the addition of an intervention strategy based on model predictions could decrease observed readmission rates considerably in our assessment. This end-to-end solution showcases an effective roadmap for hospitals to prospectively identify high-risk patients and intervene promptly, using evidence-based actions to improve patient outcomes and avoid unwanted readmissions.

Keywords: Clinical interventions, Electronic health records, Hospital readmissions, Machine learning, Predictive analytics, Risk prediction, System validation

1 Introduction

Thirty-day hospital readmission upon discharge is a persistent issue in healthcare due to its effect on the health of the patient and system costs [1]. Readmission typically means that the initial hospital stay or return home was not entirely successful in stabilizing the health of the patient. Evidence has shown that within a span of one year, there are millions of patients across the world who are readmitted soon after discharge and it costs a great deal (hcup-us). For example, in the US there were approximately 3.3 million readmissions for 30 days in 2011, associated with over \$40 billion in expenses. Besides costs, these multiple hospitalizations expose patients to other morbidities such as hospital-acquired infections and psychological distress. As a result, reducing avoidable 30-day readmissions has become a measure of high-priority healthcare quality, where policymakers and payers have started reimbursement penalties against high-readmission hospitals. Health care providers are thus highly encouraged to identify effective ways of forecasting and preventing unnecessary readmissions within that brief post-discharge window.

Readmission prevention is challenging, however, since return hospital admission reasons are different. A patient's likelihood of early readmission is established by various factors including the severity of illness, comorbidities, quality and discharge planning of care offered, management of medications, and availability of outpatient follow-up and social support. The majority of readmissions are potentially avoidable with improved coordination of care [2]. For instance, if

OPEN ACCESS Reproducible Model

Edited by
Associate Editor

Curated by
The Editor-in-Chief

Submitted 18, March, 2022

Accepted 11, June, 2022

Citation
Rajendaran, D. (2022)
An End-to-End Predictive and
Intervention Framework for
Reducing Hospital
Readmissions




Figure 1. Global Impact and Predictive Prevention of 30-Day Hospital Readmissions

a patient with heart failure is discharged without the necessary adjustment in their medication or without an early follow-up appointment, their risk level of decompensation and readmission is heightened. A patient with diabetes who is discharged without knowledge of how to check their blood sugar level or without any method of accessing medications to which they were instructed can deteriorate rapidly and be readmitted. These scenarios outline how breaches along the inpatient to outpatient transition period can lead to avoidable readmissions.

Historically, hospitals have employed a wide range of tools to identify patients at risk for readmission. Simple risk scoring tools such as the LACE index or HOSPITAL score utilize a few parameters (e.g., length of stay, admission severity, comorbidities, and prior emergency visits) to estimate readmission risk. While these traditional scores are easy to calculate and use, they are not very predictive since they capture only a portion of the patient information and may not generalize across different patient populations. The expanding utilization of electronic health records (EHRs) over the past several years has created enormous amounts of data containing demographic data, diagnoses, laboratory studies, medication, and history of previous utilization for individuals. This abundance of data allows the possibility for more sophisticated prediction modeling with machine learning techniques to better stratify readmission risk. Indeed, many authors have derived machine learning algorithms for predicting 30-day readmissions, usually with higher predictive performance than conventional risk scores. These models are able to capture nonlinear interactions and associations between risk factors that would be missed with more rudimentary methods [3]. But without timely interventions, availability of a predictive model is of no use to reduce readmissions. More and more, there is an appreciation that an integrated strategy combining prediction and action is needed. A high risk prediction is of value only if it results in an appropriate response to alter the course of the patient. For example, knowing that a patient has a 50% chance of being readmitted does not amount to much unless doctors do something about it by having the care plan modified or providing additional support after discharge. It is hence crucial to have a complete end-to-end system, where data-derived insights guide clinical decisions and interventions in an end-to-end, seamless flow. An effective system would screen at-risk patients during or before discharge and apply to them targeted preventive interventions—e.g., more intensive discharge counseling, medication reconciliation, early ambulatory follow-up visits, or home care—to lower the risk. Moreover, it would be dynamic, evolving in response to each patient's outcome to make the predictions and interventions better in the future. This builds

a "learning health system" loop where the process and model continually improve based on real-world results.

While the intuitive attraction of combining predictive analytics with intervention strategies is strong, making this happen within a hospital setting poses several challenges. It involves combining disparate data, stabilizing models, designing intervention protocols, and merging the workflow into clinical practice without causing alert fatigue or overloading staff. And there is the problem of clinicians coming to believe in and making sense of the predictions—therefore, the importance of model transparency and interpretability. The first predictive risk tool deployments had mixed success with regard to preventing readmissions sometimes because they did not possess a strong system for acting upon the risk stratification. Our work is motivated by such gaps: absence of an end-to-end system that predicts risk of readmission not only that, but systematically acts on such predictions and tracks the consequences.

In this paper, we propose an end-to-end predictive and intervention method of hospital readmission reduction within 30 days. The method addresses the problem from data acquisition all the way through to on-the-ground care interventions and constant progress. We describe how patient data can be utilized to anticipate readmission risk with machine learning and more importantly, how that prediction can trigger a series of pre- and post-discharge clinical interventions to avoid avoidable returns. We describe the system components, which are comprised of data ingestion and preprocessing, predictive analytics module, an intervention engine that encompasses discharge readiness assessment and post-discharge follow-up interventions, and a closed-loop monitoring process for outcome feedback. To demonstrate the performance of the prediction component of the framework, we present an experiment on a real-world dataset of hospital admissions of diabetic patients, comparing a few machine learning models and their performances. We then outline the deployment architecture needed to deploy this framework in practice [4]. With the incorporation of analytics and proactive care management, our approach aims to show a sustainable method for reducing early readmissions and improving care continuity for at-risk patients.

2 Problem Statement and Objective

Thirty-day rehospitalization in the hospital is one of the key health care measures because normally it reflects poor continuity or quality of care. A patient newly discharged from hospital and readmitted within a month means that possibly the underlying pathology was not dealt with properly or something went very wrong within the immediate postdischarge period. From a clinical point of view, multiple early readmissions can indicate potential quality of care issues — e.g., perhaps the patient was discharged too quickly, or without adequate instructions, or without proper arrangements for follow-up and home care. These readmissions also injure patients, subjecting them to additional medical risks, stress, and interruption of their recovery. Elevated 30-day readmissions are a loss of hospital resources from a systems perspective, using up beds and staff time that can be spent on other patients, and generating unnecessary healthcare cost. Payers and regulators increasingly view early readmissions as wasteful costs; for instance, insurance plans and government health programs have enacted policies that financially penalize hospitals that have greater-than-average 30-day readmissions for certain conditions. In short, the problem is that there are too many patients being readmitted to the hospital soon after discharge, and it imposes a multidimensional burden on patients, providers, and the system.

Not all readmission can be prevented — some will always need acute care again no matter optimal management — but many early readmissions are perceived as preventable with better care coordination [5]. The system's purpose in this paper is to detect and intervene with high-risk 30-day readmission patients beforehand, in the hope of breaking the cycle of recurrent hospitalization whenever possible. Using data-driven predictive modeling, the system will detect high-risk individuals at the time of hospital discharge, and then follow through by providing those individuals with personalized interventions aimed at their individual risk factors and needs. In effect, the aim is to shift the paradigm from a response strategy (curing patients once they relapse and recur) to a prevention strategy (interacting during and shortly following the first admission




Figure 2. Day Hospital Readmission Cycle and Its Implications

to prevent relapse). The system is to achieve measurable reductions in 30-day readmissions by integrating accurate risk predictions with prompt, effective interventions of care. Another important goal is to design a closed-loop learning system: the outcome for each patient (whether they got readmitted or stayed home) is fed back into the system, allowing constant improvement of the predictive algorithm and the intervention techniques. A triumph for this system is less patients unnecessarily readmitting to the hospital within 30 days, healthier and happier patients during the post-discharge phase, and finally fewer clinical and economic burdens of preventable readmissions.

3 System Overview

The end-to-end solution has been formulated, by attracting the data-based predictions towards clinical interventions regarding any feedbacks in a closed loop. The framework basically constitutes of four broad stages, which would be working in sequence. First, a data ingestion and preprocessing stage gathers all relevant patient information from hospital systems and prepares it for analysis: the information from patient demographics, medical history, diagnoses, laboratory results, and medication records, current hospitalization, cleans it, transforms and prepares the inputs for modeling. Second, at discharge or discharge-predictive moment, a predictive analytics module that processes these data generates risk probabilities of a 30-day readmission computed for every patient [6]. This engine operates on the output risk score or probability under circumstance the individual patient considering their profile and clinical factors from machine learning models trained on a historical record of patient outcomes.

Third, the intervention engine takes the model risk output and translates it into action. This engine applies a set of rules and assessments to determine what interventions, if any, are warranted for the patient. For instance, an enhanced discharge planning might be flagged to the patient if his predicted readmission risk is above certain threshold. The intervention engine envelops

components like the discharge readiness assessment into which systematically evaluates whether or not the patient is prepared to go home safely and a decision logic that couples risk levels with a specific post discharge interventions. These interventions can vary from scheduling an early follow-up appointment or arranging a home health visit to providing the patient additional education, check-in phone calls, or other available support to keep from returning to the hospital. Patients with low predicted risk may just receive regular discharge processes while the higher-risk category patients will receive escalated care coordination.

Close-loop monitoring is the fourth major stage. This is the stage in which the system monitors outcomes for patients after discharge as well as the recommended uses of interventions [7]. Such a patient's readmission enters outcome information back into the data repository, which in time is used to refine the predictive model and retrain it to better serve future patients. Even when a patient isn't readmitted, the system can still learn from successes-in other words, determine which interventions were tied to stable high-risk patients at home. The closed-loop design ensures that the framework would learn continuously in real time from its performance, adjusting itself accordingly to the emerging trend or population of patients.

4 Proposed System

4.1 Data Ingestion and Preprocessing




Figure 3. Overview of the Data Ingestion and Preprocessing Pipeline

The first step of the system is to aggregate and prepare patient data, a process known as data ingestion and preprocessing. Modern hospitals collect vast amounts of data for each patient encounter, but these raw data are often scattered across different subsystems and formats. The framework's data ingestion component pulls together all pertinent information about a patient's hospitalization and relevant history. This includes demographic details (age, gender, etc.), clinical data such as the patient's primary diagnosis and comorbid diagnoses (often coded in systems like ICD-9/10), results of laboratory tests and imaging, vital signs trends, treatments and procedures performed, medication orders (including any changes to medication regimens during the stay), and utilization history (e.g., the number of previous hospital admissions or emergency visits in the past year). Additional contextual data like the patient's insurance type or social factors noted in their record can also be included if available, as these may influence readmission risk.

Once the data is gathered, the preprocessing subcomponent cleans and transforms it into features suitable for the predictive model [8]. Cleaning involves handling missing or erroneous entries (for instance, if a lab value is missing, the system might impute a typical value or mark it as unknown). Categorical variables, such as the type of admission or discharge disposition, are converted into a machine-readable format (for example, encoding categories as dummy variables or mapping codes

to broader groupings). For instance, the discharge disposition code that indicates “left against medical advice” can be turned into a binary feature capturing whether a patient self-discharged, since that scenario often correlates with higher readmission rates. Likewise, diagnosis codes are grouped into fewer, clinically meaningful categories (cardiac, pulmonary, diabetic complications, etc.) to reduce dimensionality while preserving salient information.

Feature engineering is a critical part of preprocessing in this framework. Rather than relying solely on raw fields from the electronic health record, the system creates derived features that encapsulate key risk factors. For example, one engineered feature could be a “history of utilization” index, calculated from the number of hospital admissions, emergency room visits, and outpatient visits the patient has had in the past year. A patient with multiple prior hospitalizations would have a high utilization index, indicating chronic health issues and a higher baseline risk of readmission. We also derived a “severity of illness” score for the current admission by combining factors like length of stay, number of procedures, and number of medications (the idea being that a hospitalization involving many interventions or a prolonged stay likely signals a more severe or complicated condition, which might predispose the patient to complications after discharge) [9]. Additionally, we captured medication-related changes: the dataset includes indicators if any medication dosage was changed (“Up” or “Down”) during the stay. By including such a “medication changes count” feature, the model can factor in those adjustments, under the hypothesis that patients requiring medication titrations are at risk of instability. Similarly, lab results like an elevated HbA1c (indicating poor long-term glucose control in a diabetic patient) were converted into categorical risk flags.

After feature engineering, the preprocessing stage may apply normalization or scaling to numeric features, especially if algorithms like k-nearest neighbors or logistic regression are used (since they are sensitive to feature scales). For example, the number of lab tests and number of medications have different numeric ranges; scaling them to a standard range ensures that one feature doesn’t unduly dominate the model due to scale alone. In tree-based models like random forests, scaling is less critical, but maintaining consistent preprocessing steps is important so that the live data fed into the model is handled the same way as the training data. The end result of this data ingestion and preprocessing pipeline is a clean, rich set of features for each patient at discharge time. This feature vector is then passed to the predictive analytics module. Importantly, the entire preprocessing pipeline is configured as an automated process, so that when a new patient’s data comes in, the same transformations (mapping codes, computing indices, etc.) are applied in real-time without manual intervention. This ensures the model’s inputs are consistent and that results are available in time for discharge decisions.

4.2 Predictive Analytics Module

After preprocessing the data into a consistent feature set, the predictive analytics module determines the risk of readmission for the patient. The module is essentially a machine learning classifier trained on past hospital data to learn patterns associated with 30-day readmissions. We selected a model that struck a nice balance between predictive ability and pragmatism (i.e., a random forest, which has the ability to learn complicated interactions but provides us with measures of feature importance). Trained, the model takes as input an individual patient’s feature vector (produced by the preprocessing module) and outputs a risk prediction in the form of a probability or score between 0 and 1 that represents the patient’s risk of 30-day readmission.

Training of the model occurs on a large dataset of past patient encounters for which the 30-day outcome (readmitted or not readmitted) is already determined. The system learns from the examples of patients who were and were not readmitted, tweaking model parameters to minimize prediction errors. Along the way, it automatically figures out which features are most predictive – for instance, it may learn that patients with multiple recent admissions, or those with certain high-risk diagnoses paired with long hospital stays, have a high probability of returning. As readmission occurrences are fewer compared to non-readmissions in most datasets, class imbalance is dealt with in the training process. This may involve techniques such as placing higher importance on readmission cases when computing the model’s loss function or oversampling the cases during training, so the model isn’t biased towards always predicting “no readmission.”




Figure 4. Architecture of the Predictive Analytics Module for 30-Day Readmission Risk Prediction. The system processes patient feature vectors through an optimized random forest classifier, generating calibrated risk scores with interpretable feature contributions. Training incorporates class balancing and cross-validation, while operational deployment enables real-time inference with closed-loop retraining.

The module is also evaluated on hold-out data and validated via cross-validation techniques to ensure that it generalizes well to new patients and doesn't overfit to quirks of the training set [10]. In our experimental setup described later, we compare a number of models and carry out hyperparameter tuning to choose the most appropriate for this task.

An important aspect of the predictive module in a healthcare application is transparency and interpretability. Clinicians are more likely to trust and utilize the tool if they can understand the reasoning behind a risk score. So, even if the underlying model is complex, the system provides interpretable output alongside the risk score. For example, the model can provide a list of the top contributing factors for why a given patient received a high risk prediction – it might state that "prior inpatient visits = 3" and "discharge to skilled nursing facility" were driving factors that raised the risk. These explanations enable the care team to validate the prediction and decide which specific risk drivers to address [12]. In our system, we prefer models that are somewhat interpretable (a random forest can yield feature importance rankings, and a logistic regression offers directly interpretable coefficients). We also calibrate the model output such that the predicted probability is well-calibrated to observed risk, as required when thresholds are used to trigger interventions.

Operationally, when the system is live in the hospital setting, the predictive analytics module runs either real-time or batch at specified intervals (e.g., every morning or discharge planning time). It can be implemented as a service that takes in a patient's feature data and returns a risk score within seconds. The hospital IT infrastructure would integrate this service into the workflow so that, for every patient who is going to be discharged, the care team can fetch the readmission risk

Algorithm 1: Readmission Risk Prediction Pipeline

Patient feature vector $\mathbf{x} \in \mathbb{R}^d$, training dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$ Risk score $p \in [0, 1]$

Step 1: Data Preparation ;

Compute class weights: $w_c = \frac{N}{CN_c}$;

// Class balancing

Perform K -fold cross-validation ;

Evaluate model performance: $\mathbb{E}_{k \sim K} [AUC_k]$;

Step 2: Model Training ;

Train a Random Forest classifier f_θ on $\mathcal{D}$ with loss $\mathcal{L}(y, f_\theta(\mathbf{x}))$;

Step 3: Prediction ;

Obtain model output: $f(\mathbf{x}) = \sum_{t=1}^T g_t(\mathbf{x})$;

Compute risk score: $p = \sigma(f(\mathbf{x}))$;

// Sigmoid normalization

Step 4: Calibration and Interpretation ;

Calibrate probability using Platt scaling: $\text{Platt}(p) = \frac{1}{1 + e^{-(a + bp)}}$;

Compute feature importance: $\frac{\partial p}{\partial x_i}$;

Step 5: Operationalization ;

Deploy model behind REST API for real-time inference ;

[11] Step 6: Feedback Loop ;

Monitor operational performance ;

 If performance drops below threshold: Retrain model using updated data ;

 Update deployment with retrained model ;

score immediately. The model can be retrained from time to time (e.g., every few months or when enough new data has been collected) so that it stays current. This retraining would occur with the accumulated data and new cases, as a component of the closed-loop learning that we outline below. Briefly, the predictive analytics module is the "brain" of the system, distilling complex patient information into one actionable metric (the risk of readmission) that will inform the next intervention decisions.

4.3 Intervention Engine




Figure 5. Discharge Readiness Assessment Workflow with Risk Monitoring

The second significant component of the framework is the intervention engine, the component of the framework that translates risk predictions into clinical action. It consists of a series of functions that together aim to minimize the patient's risk. [13] One critical component of this




Figure 6. Risk-Stratified Intervention Protocol for Readmission Prevention

engine is the discharge readiness assessment. Even excluding consideration of predictive model scores, the hospitals benefit by regularly considering whether each patient truly is ready to go home or to the next level of care. The system has an express discharge readiness checklist or assessment algorithm, which is especially important for patients who the model has designated as high-risk. The assessment considers both clinical and pragmatic aspects of readiness. On the health side, it assures that the patient's status has become stable or sufficiently improved: i.e., that their vital signs are stable, they are no longer on acute therapies, and any operative sicknesses or wounds are under control. It also assures that key inpatient milestones have been attained (e.g., completion of planned treatments or therapy goals). Practice-wise, the assessment also assures that education has been received by the patient (and caregiver, as the case may be) regarding the diagnosis and plan for post-discharge care. This involves awareness of their medications (what to take, how much, and new medications or changes administered during hospital stay), identification of warning signs or significant symptoms to be watched for, understanding on the maintenance of any dressing on wounds or medical devices, and well-communicated details regarding diet, activity, and follow-up.

Our readiness for discharge assessment is also partially automated by querying data within the electronic record – for example, a check to determine if a follow-up appointment date has been added, or if patient education documentation has been signed off by nursing staff [14]. It can point out gaps, such as a missing follow-up visit or lack of a documented education session on insulin use in a diabetic patient, prompting the care team to complete these gaps before discharge. In high-risk patients, the system can require an explicit sign-off that the readiness checklist is ready and all the needs are met before permitting the discharge order to be signed. This creates a safety net: if there's a high-risk patient who hasn't yet had a required consultation (e.g., a diabetes educator) or hasn't shown the ability to self-administer a new medication, the system notifies it for immediate review. By requiring meticulous preparation, the system removes the risk of discharging patients home unstable or uneducated – the leading causes of premature complications and readmissions. In brief, discharge readiness assessment is an upstream intervention: it seeks to catch and fix issues at the point of discharge itself, so the patient leaves the hospital as prepared as possible to continue recovery safely at home.

A second key element of the intervention engine is application of rule-based thresholds to risk-stratify and trigger varying levels of intervention. The predictive model will provide a real-valued risk score, but in practice clinical pathways will tend to need discrete categories (it makes sense to define tiers of "low", "moderate", and "high" risk with standard policies for each, for example). In our system, we create threshold values along the risk score to define these categories. For instance, we could classify patients with an estimated readmission risk above some percentage cut-off (e.g., top quintile of risk scores) as "high risk", the next segment as "moderate", and the rest as "low". The actual cut-offs are determined to the hospital's capacity and desired sensitivity [15]. Once the risk category is assigned, the engine maps each category to a pre-specified intervention plan. A high-risk patient will automatically receive a full bundle of interventions, a moderate-risk patient a limited bundle, and a low-risk patient the core standard of care. For example, the policy could be that any high-risk patients identified should have a follow-up visit set within one week of discharge, a phone call from a care coordinator within 48 hours of discharge, and referral for home health. Moderate-risk patients could have an appointment in two weeks and a one phone call at one week post-discharge. Low-risk patients may simply get standard discharge instructions and perhaps a referral to come back to their own primary care physician within the standard period. Each hospital can individually tailor the details of interventions for each group based on what proves effective and is feasible in their setting.

The rule-based system provides an open, unambiguous decision-making algorithm: for a patient's specific risk score, the appropriate action bundle is known immediately. This avoids ambiguity and ensures that high-risk cases are not left behind. It also translates clinical best practices into if-then rules that the system automatically applies [16]. Importantly, while the rules are automated, they are transparent and can be inspected or overridden by clinicians when needed. For instance, if the patient is at high risk but the doctor knows there is some context (e.g., the patient is at high risk mainly due to social factors which have recently been resolved, or the patient is being sent to a specialized rehabilitation center which would lower the risk), the treatment team can change the plan to reflect this. The system is intended to assist, not replace, clinical judgment – it is a safety net and a standardizer. Over time, the threshold mechanism imposes best practice while providing a framework that can evolve and adapt with experience. For example, if many "moderate" patients are still being readmitted, the high-risk threshold might be lowered to catch more patients; or if resources are being stretched too thin on patients who are usually fine, the threshold is raised to hit fewer, really highest-risk patients.

One of these interventions is the post-discharge follow-up call. For high-risk patients, the engine will also automatically place a phone call from a nurse or care coordinator two days post-discharge. Staff follows up through the phone call to track the patient's condition, confirms medication adherence, and addresses any question or early warning sign. Issues raised on the call can be referred for physician treatment (such as if the patient is complaining of concerning symptoms, they can be told to come in early or adjust their treatment with physician advice). It tends to correct misunderstandings earlier and keeps the patient on course with the plan of care rightly, which could prevent complications arising from festering. [17]

In addition to phone calls and follow-ups, the intervention engine can also provide more forceful support to those who need it. For high-risk patients, the system can recommend or automatically book home health care. This can involve sending a visiting nurse to the patient's home shortly after discharge to perform a check-up. By visiting the patient at home, the nurse is able to assess the patient's status and environment: record vital signs, inspect any surgical sites or medical devices, and check that the patient is complying with his or her therapy (e.g., taking medications properly and following dietary restrictions). Home visits provide an opportunity to identify warning signs the patient might not realize are risky, and they provide individualized education in the patient's home environment. If issues are found, the home health nurse collaborates with the patient's doctor to correct them (such as changing medication or getting an earlier office visit), quite possibly avoiding an ER visit.

The motor also coordinates solutions for social and logistical problems that otherwise would lead to readmission. For instance, it may include social services to arrange for transportation to

follow-up appointments, provide access to necessary nutrition and home care, or offer medication management support (e.g., simplifying regimens or pill box organization). Automated recall or contact (telephone or text) can be used to remind patients to adhere to their care plan. In others, a referral to a structured chronic disease management or coaching program (e.g., a diabetes care service) is made for ongoing support. Each of these interventions fills one of the identified gaps – transportation, nutrition, medication, or knowledge regarding self-care – that, if left unaddressed, can prompt a readmission to the hospital.

Worthy of mention is that the intervention engine not only schedules these interventions, but also tracks their fulfillment and their impact. It records that the home nurse visit was made and any important findings or activities from the visit. It checks off that the phone call was made and records whether the patient was or was not available by phone and whether another attempt will be required. All this information is fed back into the patient's record and into the monitoring aspect of the system. By tracking the follow-through on interventions, the system ensures accountability (no high-risk patient falls through the cracks without the promised follow-up) and information on what interventions worked. This data becomes relevant during closed-loop assessment: for example, if follow-up calls among patients who had home visits show an excessive readmission rate, the intervention could perhaps need to be increased or earlier; if the follow-up calls, however, show immense success in resolving issues, that system can be duplicated with patients with moderate risk as well. This way, the intervention engine works hand in hand with the monitoring process to improve care strategies over time.

4.4 Closed-Loop Monitoring




Figure 7. Learning System for Readmission Prevention

The final part of the system is closed-loop surveillance, which measures outcomes and feeds that back to the predictive and intervention systems [18]. After every patient discharge, the system monitors whether or not a readmission occurs in 30 days. That is done by automatically searching hospital records for a new admission for that patient within 30 days. When the patient comes back, that's an outcome to the system; if the 30th day passes with no readmission, that counts as well. The closed-loop module aggregates those outcomes for each of the patients and plots them against the initial risk predictions made and interventions undertaken.

This is fed back for a variety of purposes. First, it enables the hospital to track performance – i.e., what proportion of high-risk patients flagged actually get readmitted compared to low-risk patients. This enables checking that model stratification is effective (high-risk patients should have greater actual readmission rates than low-risk) and that interventions are actually having the intended effect. Secondly, data collected is used to retrain the predictive model itself. From time to time, the historical data set can be updated with the new entries (with each one having

the features and if it did result in readmission) and the model being retrained or fine-tuned to catch up on any trends that may have changed [19]. For example, if there is a new type of treatment or group of patients, the model will have their outcomes over time, thus keeping current. This continuous learning ensures that the validity of risk prediction is not lost with time due to population drift or changes in care process.

Besides model retraining, closed-loop monitoring can direct changes in the intervention strategies. For instance, it may be found by tracking that a subgroup of patients (e.g., more severe heart failure) remain at risk for readmission even with current treatment, suggesting that they require additional support as a subgroup (e.g., symptoms tracked daily using remote monitoring). Conversely, the data may show that some interventions are especially effective – for instance, perhaps patients who received both a home visit and a follow-up phone call had greatly decreased readmissions – so it would make sense to expand or prioritize those interventions to more patients. The hospital would be able to improve the system appropriately – move the risk thresholds, improve the intervention protocols, or add extra things to the checklist of discharges – from such evidence.

5 Experiment

To measure the performance of the predictive model and show the end-to-end pipeline, we conducted an experiment on a real-world hospital admissions dataset. The dataset used for obtaining the data is the "Diabetes 130-US hospitals" dataset, having data points on diabetic patient admissions to 130 United States hospitals between 1999 and 2008. The dataset is perfect for our application as it exactly indicates whether each of a patient's hospital stays was followed by a 30-day readmission, and it has a wide range of features on each patient's hospital stay. In total, the database holds over 100,000 inpatient encounter records [20]. Each record has the demographic information (age, race, gender) of a patient, data associated with that hospital stay (type of admission, primary and secondary diagnoses, procedure, LOS, lab results and medications prescribed, etc.), data about past utilization of health services (number of outpatient visits, emergency department visits, and hospital inpatient stay during the previous year of the current admission), and if readmitted shortly following discharge. The readmission is coded as a categorical variable taking values "<30" if the patient was readmitted within 30 days, ">30" if the patient was readmitted after more than 30 days of discharge, and "NO" if no readmission took place in the 30-day period. For our experiment, we took the target variable to be a binary indicator of 30-day readmission (coding "<30" as 1 and the other two outcomes as 0, since we want to predict readmissions within 30 days of discharge).

Before feeding this data into our predictive model, we carried out the preprocessing and feature engineering steps outlined above in the System section. We pre-processed the data by handling missing values and removing features that would not be available until after discharge time (e.g., we remove data that would be unavailable until after discharge). We then recoded categorical features: for instance, the many potential diagnosis codes were collapsed into fewer useful diagnostic categories (for instance, all cardiac-related diagnoses, all infections, etc.), and the hospital discharge disposition codes were collapsed to categories like "home", "home with healthcare services", "skilled nursing facility", "left against medical advice", etc. The motivation is to remove high-dimensional categorical noise without losing the valuable information (e.g., whether or not a patient left against medical advice or went to a nursing home can be valuable predictors for readmission risk). We also coded category variables dealing with medications; we had features indicating whether any of 24 listed diabetes drugs were prescribed and whether their dosage was altered during the visit (indicated as "No", "Steady", "Up" or "Down"). We coded them as numerical or binary features expressing whether there is a medication change or not, as a change in the medication regimen might mean uncontrolled disease. [21]

Apart from these categorical encodings, we created other features to detect patterns that would render a patient vulnerable to readmission. We calculated a prior utilization feature for each encounter, being the total of inpatient stays the patient experienced in the last year, as well as separate counts for past emergency visits and outpatient visits. Frequent past hospital use patients were given higher values for these features. We also employed the composite indices discussed in

the System section: for example, we created an "admission severity index" by combining number of procedures, number of medications, and length of stay (assuming that extensive hospitalization with a high rate of interventions likely signifies more severe illness, which would be a greater risk for post-discharge problems). We also added an indicator of whether the patient's HbA1c test result was higher than the normal range, since poor control of glycemia has been shown to be related to higher readmission risk among diabetics. After feature engineering, our dataset contained about 45 features per patient. We normalized the numeric features to place them on comparable scales for modeling.

Next, we separate the data into a test set and a training set in order to allow us to properly evaluate our models. We trained on 80% of the encounters (and within that, again used a subset for validation during model tuning) and set aside 20% of the encounters as an independent test set in order to estimate performance. The distribution of the class for all data was also highly imbalanced: fewer than 11% of encounters had a readmission follow-up 30-day readmission (the remaining 89% did not) [22]. To treat this class imbalance in training, we applied an oversampling strategy to the minority class. Specifically, in the training set, we oversampled at random encounters which had a readmission, so that classifier training would not be skewed too strongly towards predicting "no readmission" for everything. (We also experimented with SMOTE – synthetic minority oversampling – to generate synthetic minority instances, but in our case simple oversampling of real instances produced comparable results and was easier to implement.) We then fitted a range of classification models to this preprocessed training data: k-Nearest Neighbors (KNN), Logistic Regression, a linear classifier trained by Stochastic Gradient Descent (SGD), and a Random Forest ensemble. These four model types were chosen to cover from interpretable and straightforward (logistic) to intricate and potentially highly effective (random forest), with a lazy learning algorithm (KNN) and online learning algorithm (SGD).

Table 1. Performance of Classifiers on Test Set

Model	AUROC	Recall	Precision	F1-Score
k-Nearest Neighbors (k=10)	0.60	0.20	0.12	0.15
Logistic Regression	0.65	0.35	0.23	0.28
SGD Classifier	0.65	0.36	0.22	0.29
Random Forest	0.70	0.50	0.25	0.33

Hyperparameter adjustment on the training set with cross-validation was used for all models. For KNN, the most critical hyperparameter is the number of neighbors k . We tried k values of 5, 10, and 20 and found that $k = 10$ provided the best balance in performance. (Smaller k was a bit more sensitive but noisy, and larger k started to blur important distinctions.) We also used Euclidean distance for KNN after feature normalization. For logistic regression, we experimented with different regularization strengths (inverse regularization parameter C) and discovered that a moderate level of regularization (C at around 0.1) with an L2 penalty was optimal, preventing overfitting but still closely following the data. The SGD classifier was essentially training the same linear model as logistic regression but with stochastic optimization; we also tuned parameters like the learning rate schedule and number of epochs [23]. After tuning, the SGD model's performance was very close to the logistic regression (as expected since they both minimize the same objective). For Random Forest, we tried varying the number of trees (50 to 200) and the maximum tree depth. We observed that with about 100 trees, the performance gains saturate, so we picked 100 trees to a depth of 10 in our final model (to not overfit). While training the forest, we also set up the class balance parameter so each tree is trained on an imbalanced subset of the data in such a way that it effectively also resolves the imbalance by randomly undersampling the majority class for each tree.

After training, we evaluated all models on the hold-out test set (which had not been seen by the models during training or tuning). We were concerned with a few evaluation metrics: area under the Receiver Operating Characteristic curve (AUROC) as a general measure of discrimination ability, and precision, recall, and F1-score for the positive class (30-day readmissions) in order to

Table 2. Top 10 Predictive Features from Random Forest

Feature	Description
Prior Inpatient Visits	Number of hospital stays in prior year
Discharge Disposition	Type of discharge (e.g., AMA, SNF, Home)
ER Visits (Past Year)	Number of emergency department visits
Severity Index	Composite of LOS, medications, procedures
HbA1c High	Indicator of poor glycemic control
Medication Change	Change in diabetic medications during stay
Length of Stay	Duration of the current hospitalization
Number of Procedures	Total procedures performed during stay
Insulin Dose Up	Indicator if insulin dose increased
Outpatient Visits	Number of outpatient visits in past year

have a feel for how well the models were indicating those high-risk patients. Overall accuracy isn't as interesting here because always guessing "no readmission" would be 89% accurate due to class imbalance, so we emphasize the other statistics. The experiments showed a clear ordering of the models. The worst model was KNN: with $k=10$ it had an AUROC of around 0.60 and low recall (it caught well less than a quarter of the readmitted patients) [24]. Its accuracy was quite low too, i.e., most of the patients it marked as being at-risk were not readmitted. This is probably due to the fact that KNN, in a high-dimensional feature space, finds it difficult to identify pure neighborhoods – there are still a number of non-readmitted "neighbors" to dilute the prediction among many readmitted patients.

Logistic regression performed better than KNN, with an AUROC of around 0.65. On the test set, the logistic model's predictions could be tuned using the decision threshold: if we used the default threshold of 0.5 for classifying a person as high risk, the model was very specific but not sensitive (it would identify many fewer actual readmissions, but around 20% of them). We set the threshold more sensitive (predict "readmission" if $\text{prob} > 0.3$, e.g.). In that setting, logistic regression could find about 35–40% of the readmissions (recall 0.35–0.40) at a precision of about 20–25%. Practically, that translates to of all patients flagged by the model, about a quarter actually ended up being readmitted (still an improvement on random guessing, considering the base rate of 11%, and many of the "false alarms" are actually patients who may be helped by preventative treatment regardless). The SGD classifier performed essentially the same as logistic regression would have, naturally. Its AUROC was likewise in the mid-0.60s. We also found that, with sufficient training, the SGD model basically came to the same conclusion as the logistic model, supporting the fact that features are largely capturing linear relationships logistic regression can take advantage of. [25]

Table 3. Training and Inference Summary

Model	Train Time	Infer Time	Notes
KNN	None	Slow	Not scalable
Logistic Reg.	Seconds	Fast	Interpretable, calibrated
SGD	Seconds	Fast	Similar to logistic
Random Forest	Minutes	Very Fast	Best performance

Our Random Forest best model worked fairly well. On our test set, it produced an AUROC value of about 0.70, demonstrating discrimination between non-readmissions and readmissions. Most importantly, it improved the balance between sensitivity and precision. With an optimized threshold to strive for greater recall, the random forest did predict approximately 50% of the patients who would be readmitted (recall 0.50), losing precision of approximately 25%. That is, approximately one in four of its "high-risk" predictions was indeed a good selection of a patient who indeed came back within 30 days. This was a noticeable improvement in recall relative to

the logistic model with equal precision. If we compare F1-scores (balanced recall and precision), the positive class F1 of the random forest was roughly 0.33, whereas logistic and SGD were roughly 0.28–0.30, and KNN was below 0.20. All models were extremely specific (over 90% of the non-readmitted patients were correctly labeled low-risk), as we should have with class imbalance and how we calibrated our thresholds.

For a sense of what the random forest learned, we looked at its feature importance scores [26]. The characteristics that were projected from patient history of use were among the most predictive: e.g., number of prior inpatient admissions was extremely significant – patients with multiple prior hospitalizations were far more likely to come back soon again. Similarly, numerous prior emergency department visits was a risk factor. Discharge disposition was the other key feature; that is, patients discharged against medical advice were much more likely to be readmitted (basically because they get discharged before they are really cured), and the model took advantage of that. In contrast, patients admitted to a different facility (e.g., nursing home or rehabilitation center) also had increased readmission compared with discharge home, likely as a result of sicker underlying illness. Of the clinical admission characteristics, the total length of stay and the severity index (procedures and medications) were predictive – a very short stay may indicate an uncomplicated condition or planned procedure (lower risk), while a long complicated stay might indicate the patient is severely ill (higher risk). Some laboratory and medication characteristics were also predictive. For example, whether the diabetic medication of the patient was changed ("change" feature) and whether insulin was given (and especially if insulin dose was raised) were relevant to some extent, suggesting patients whose treatment had to be changed were at risk of instability at discharge. The HbA1c value result factor also had moderate significance: those with extremely high HbA1c (poor long-term glucose regulation) were more likely to have increased readmission risk, most likely because of continued difficulty controlling their condition. Specifically, age, while a variable, was not exactly in the highest category of variables after those others were adjusted for – middle-aged and elderly diabetic patients both had elevated readmission risk precipitated mainly by comorbidities and prior utilization more so than age itself.

Table 4. Logistic Regression Threshold Adjustment Effects

Threshold	Recall	Precision
0.5 (default)	0.20	0.30
0.3 (adjusted)	0.35–0.40	0.20–0.25

For the purpose of training performance but also practicality interests, all the models were reasonably manageable on this dataset (100k records) [27]. The logistic and SGD models trained in a matter of seconds. The KNN model actually does not have a training step but in deployment would be slow to compute distances to all training points for each prediction; this would be impractical in a real-time hospital setting with a large dataset, which is part of the reason we likely would not deploy KNN even though it is so simple. The random forest model trained in a few minutes using 100 trees but, once trained, predicted very fast (on the order of milliseconds per patient), which is appropriate for a live hospital setting. We also checked the calibration of the models – the random forest probability estimates were reasonably well-calibrated (when it estimates 30% risk, roughly that proportion of similar patients did readmit), especially after we had applied methods like Platt scaling on validation data, whereas raw outputs of some models can be less interpretable as probabilities. For our uses, a probability that corresponds reasonably well to actual risk is helpful both in threshold setting and in communicating risk to clinicians (e.g., telling a clinician "this patient has approximately one in three chance of readmission" makes intuitive sense).

The test did validate that the predictive analytics module is able to determine those patients with significantly high readmission risk, and that using a more advanced model (random forest) with carefully designed features does provide a noticeable performance boost over simpler approaches. Although a 0.70 AUROC signifies that the task is challenging (perfect prediction isn't possible since there are lots of unpredictable variables in a patient's life post-release), even this level of precision

allows the model to accurately differentiate more-risky patients from less-risky patients. When incorporated into the complete architecture, as we suggest, these forecasts allow the hospital to aim interventions – in this case, our model might indicate around 20% of discharged patients for increased follow-up, and that group would contain approximately 50% of patients who will in fact be readmitted. If the hospital takes action on those predictions (through the intervention engine's protocols), the hospital might be able to avert many of those forecasted events [28]. The experiment also illustrated the necessity to address class imbalance and optimize for positive recall of the positive class: a model that optimizes solely on total accuracy or AUC regardless of the actual need in the real world to catch most readmissions may not be as useful in the real world. By tuning, we controlled the model and threshold selection so as to be well-matched for the system objective – to maximize detection of high numbers of risk-threatened patients within reasonable bounds on false alarm. The results from the feature importance analysis also provide us with confidence that the model is making rational decisions (like prior use and severity of disease), which enhances confidence in its application as part of clinical decision support.

We also performed a rapid error analysis to understand the limitations of our predictive approach. A few readmissions that the model failed to detect (false negatives) were instances of patients who seemed pretty low risk on paper but subsequently had some surprise complications at home – like a patient with an otherwise uneventful brief hospitalization who stopped taking a critical medication after discharge and relapsed. Those conditions (such as non-adherence with medications or inadequate family support) are not reflected fully in our dataset, and therefore, it cannot be claimed that the model is going to capture all readmissions perfectly without detailed data. Some of those identified by the model as being high-risk were not subsequently readmitted (false positives); numerous individuals had risk indicators (such as frequent previous admissions) but were able to avoid return to the hospital through proper post-discharge management or other advantageous situations. This only makes the argument that the output of the model should be used to guide further care and not be taken to represent an absolute fate – the implication of high risk is "this patient needs to be monitored" rather than "this patient will definitely return." Machine-learning-ically, the model suffers similarly from data quality and diversity. For instance, our dataset does not contain socio-economic parameters or thorough outpatient follow-up, which in reality have large effects on readmissions. Such types of parameters might be added to further enhance the model's precision in the future [29]. Finally, we noted that some features, which one would expect to highly predict, like gender or race, did not significantly influence the predictions even after controlling for clinical variables. This suggests that risk of readmission in this patient population was more driven by history and health status than by demographics. The analysis of error illustrates that as much of a utility as the model is in risk stratification, it should be present as part of an integrated strategy that involves clinical judgment and patient engagement. As an example, we tried to take an ensemble of averaging random forest and logistic regression predictions to try to exploit complementary strengths. The performance of the ensemble was basically identical to the random forest alone (no significant increase in AUROC or recall), meaning that the random forest was already catching a large majority of the signal present in the logistic model. In another setting with more heterogeneous model types (e.g., with the inclusion of a neural network or gradient-boosted trees), an ensemble might be beneficial, but from our current models, we opted for maintaining the sole best model (random forest) as the backbone of the predictive module.

Model's statistics are on par with other research that used the same data. Past research in this diabetes readmission dataset has indicated trouble attaining an AUROC higher than in the low 0.70s using basic machine learning algorithms, due to some readmissions' inherent randomness. Thus, our results (AUROC 0.70) are a state-of-the-art level of accuracy for this task [30]. The relatively low precision and recall figures underscore the fact that while the model performs much better than chance, it does not eliminate uncertainty – to be expected in clinical prediction problems. This makes our framework's approach of combining the predictive model with broad safety-net interventions reasonable: even those patients who were classified as high-risk but perhaps were not readmitted still get extra care, and the cost of that is weighed against the benefit of not having to deal with those cases the model correctly identifies. In general, the experiment

confirms that a well-tuned machine learning model can stratify readmission risk to a significant degree, and it forms a good basis for the proactive intervention strategies that we've integrated into the overall framework.

5.1 Deployment Architecture

The deployment of the framework to a real-life hospital setting involves the integration of data systems, analytical services, and clinical workflow tools in a combined and secure manner. A likely deployment architecture for our end-to-end readmission reduction system would consist of multiple layers: the hospital's electronic health record (EHR) system serving as both data source and point-of-care interface, a back-end analytics server (or cloud hosting) housing the predictive model and rules engine, and a set of care coordination apps or modules that execute the interventions and tracking.

In fact, how this could actually work in practice is as follows. Whenever a patient is ready to be discharged, their data (demographics, diagnoses, lab results, etc.) are automatically transferred from the EHR into the predictive model. This can be done in real time via an integration interface (e.g., by using an API that the EHR calls with patient data). The model would be hosted either on the local server within the hospital data center or a HIPAA-compliant cloud infrastructure that aligns with healthcare data privacy laws. Since computational complexity is low, even a typical hospital server would be adequate for scoring the patients in real time [31]. The model service outputs a risk score to the EHR.

The front-end would involve clinicians seeing the outcome as part of their existing workflow. For instance, a discharge planning screen within the EHR could display a field labeled "30-day Readmission Risk" and the degree of risk and perhaps the percentage likelihood to give some context. Apart from the risk, the system can provide recommended interventions depending on the logic of the intervention engine: i.e., "High risk – ensure follow-up in 7 days, home health visit arranged, and pharmacist consultation on discharge." These recommendations can be inserted into the discharge checklist for the patient. The care team (physicians, nurses, case managers) is informed – they view the recommendations, and because the rules are hospital-approved guidelines, in most cases they would agree and implement them. The system may offer override or modification of interventions (with a recorded reason) if the clinician decides otherwise for a specific case.

In the background, the intervention engine would be a rules service that communicates with other hospital systems. For example, when the engine suggests that there is a need for a follow-up appointment within 7 days, the system creates an appointment request in the outpatient scheduling system. If there is a need for home health care, it creates a referral order [32]. For a callback phone call, it creates an assignment for the transitional care nursing staff. The engine would supply the hospital's care management application with these assignments and alerts. Alerts can also be sent: a high-risk discharge can prompt the system to send a notification to the case management team to draw special attention to that patient.

Data storage is also an important feature. All data required for prediction and outcomes need to be kept confidential. The hospital can have a readmission risk repository or database of characteristics for each discharge, risk score, and planned intervention stored therein. This can be within the hospital's enterprise data warehouse. This database allows for closed-loop analysis (model retraining and performance monitoring) and allows auditing if interventions were done (e.g., ensuring that all of the high-risk patients had their follow-up calls). For model retraining, new data and results can be batched off from this database and used to retrain the model (e.g., through monthly retrain jobs).

Security and confidentiality are the most critical when applying such a system [33]. Assuming a cloud service for the predictive model, the hospital must make patient data encrypted and the cloud environment healthcare-compliant. On-premises deployment might be the choice for most hospitals to store data locally, which is feasible because the compute load is moderate. The model and engine can be run on local servers by the hospital's IT department.

Another consideration is uptime and fail-safes. The system should be designed such that if the predictive service is temporarily unavailable, it fails gracefully – for example, the EHR could simply indicate that risk could not be assessed, rather than delaying the discharge. Over time it is important to monitor for any alert fatigue if it flags too many patients as high risk. Thus, ongoing calibration is required. The quality improvement team would monitor the flag count and patient status on a regular basis and tweak thresholds or processes as needed to maintain alerts as meaningful.

The deployment structure consists of an analytics layer (rules engine and predictive model) closely combined with the clinical workflow layer (care coordination tools and EHR) [34]. With interfaces, pipelines, and security in place, the framework provides real-time risk scores and automatically triggers the corresponding intervention workflows for each discharge. In effect, clinicians now have a decision support tool which not only identifies patients who are at high risk but also activates the steps for lowering that risk, neatly integrated into the hospital IT infrastructure.

6 Conclusion

Readmissions, particularly 30 days after discharge, remain a stubborn problem for healthcare, but this paper demonstrates that systematic and data-driven strategy can make significant inroads against it. We demonstrated an end-to-end system not only to predict which patients were at highest readmission risk through the application of machine learning, but to incorporate those predictions into a loop of intervention and continuous improvement. The intuition is that predictive analytics and clinical action must go hand-in-hand: a model's risk score is valuable only if it triggers a response with the potential to alter the patient's course. Our system was designed with this in mind, linking a predictive module to an intervention engine and a feedback mechanism.

Through our experimentation on a large real-world data set of diabetic patient admissions, we showed that current machine learning techniques (enhanced by thoughtful feature engineering) can indeed stratify patients by readmission risk with a satisfactory degree of accuracy. The best-performing model (a random forest) identified a subset of patients who are responsible for a significant proportion of follow-up readmissions. This is consistent with the ability to risk stratify based on data: while few readmissions are perhaps inevitable, most can be forecasted based on trends within a patient's illness history and clinical profile. Perhaps equally important, we showed that it is feasible to achieve good results using less complex models (i.e., logistic regression) and that their results are consistent with clinical judgment – e.g., history of previous use and markers for illness severity were reliable predictors [35]. This harmony between model results and medical understanding is reassuring, as it suggests that the device can be more readily accepted and used by clinicians. The test also helped reinforce the limits of predictability (with an AUC of 0.70, there is still significant uncertainty regarding outcomes in a single patient), which is why a predictive model must be used to guide supportive treatments but not as an oracle.

By embedding the predictive model in a system, the framework uses the resulting risk estimates to steer quality improvement. Those patients predicted to be at high risk, based on the model, are pre-emptively given more resources – be it extra education, more frequent follow-up, or support services – than they might otherwise have received under an otherwise standardized discharge process. In principle, that means that even when the model "misses" (e.g., catching someone who in reality would not have been readmitted), the outcome is that the patient merely got more care and maybe some benefit from it; and when the model does get it right (catching someone who would have been readmitted), then the intervention will avoid that. So the cost-benefit scale tips in favor of acting on the predictions. If our model were actually implemented, a hospital would find that its 30-day readmission rate declines sharply over time – not necessarily because the model is perfect, but because it does what transitional care effort should do: allocate effort more effectively to those who need it most.

One of the strengths of our model is that it is closed-loop. Rather than a static tool, it is a learning system that will evolve with more data that comes to hand and as hospital practices evolve [36]. The more it is applied, the more outcome data it has to learn from – and this can then be used to retrain the model to enhance or maintain accuracy, and to refine intervention strategies. For

example, if a particular intervention is consistently preventing readmissions among high-risk patients, that can be scaled up or formalized; if there remain some high-risk patients slipping through, new approaches can be tried for them. So, the system promotes a culture for continual improvement, consistent with the ideals of a learning health system. There are also far-reaching implications for health care management.

Preventable readmissions reduction benefits not only patient outcomes – keeping patients safer at home rather than bringing them back to the hospital – but also has cost implications in value-based care. Hospitals now are incentivized to improve quality metrics such as readmission ratios (with programs such as national readmission reduction programs using penalties for bad performance). A system such as ours provides a concrete means to make those changes happen, by converting predictive analytics into standard care decisions. It can also increase patient satisfaction, as patients who receive improved discharge planning and follow-up are more likely to feel well cared for and experience fewer unexpected health tragedies shortly after discharge. We must acknowledge that effective implementation will require more than technology.

Change management is important: staff and clinicians need to be trained on the system, how to use risk scores, and how to implement interventions recommended [37]. Frontline providers and leadership need to be sold on the system by demonstrating its evidence base and dependability. In addition, our experiment focused on diabetic patients as an example of use; in fact, the framework may be applied to all adult medical/surgical patients or specific conditions (e.g., heart failure, where interventions can differ but the overall approach is the same). Further trials and pilot programs would be beneficial to determine the extent to which readmission prevention is achieved when the model is applied and further tailor its components in a practical clinical environment. Lastly, this research provides an entire master plan to reduce hospital 30-day readmission by an integration of predictive analysis and coordinated care interventions.

Through using hospital data to predict risk and actively managing the risk through directed post-discharge intervention, health care providers can intervene earlier in the illness path. Combining analytics with real-world workflow integration ensures that the system is not just a conceptual model but a deployable solution that aligns with clinical workflows. If scaled up and put into practice, this type of system would improve patient outcomes (fewer return hospital visits, better recovery experiences), take pressure off of hospital resources, and be part of the larger trend towards proactive, value-based care. It is one example of how data-driven wisdom can be wedded with compassionate, preventive medicine to solve one of healthcare's persistent problems. In the future, aspects of this model might be expanded out with additional sources of information (e.g., wearable sensor measurements or social determinants of health data) to further increase precision, or combined with telehealth programs to provide even more continuous post-discharge monitoring. However, regardless of the existing hospital data, this study finds that substantial declines in readmissions are possible where predictive analytics are applied systematically to care processes.

References

- [1] Junqueira ARB, Mirza F, Baig MM. A machine learning model for predicting ICU readmissions and key risk factors: analysis from a longitudinal health records. *Health and Technology*. 2019 4;9(3):297-309.
- [2] Teo K, Yong CW, Muhamad F, Mohafez H, Hasikin K, Xia K, et al. The Promise for Reducing Healthcare Cost with Predictive Model: An Analysis with Quantized Evaluation Metric on Readmission. *Journal of healthcare engineering*. 2021 11;2021:9208138-10.
- [3] Biswas S, Dinh D, Lucas M, Duffy SJ, Brennan A, Liew D, et al. Incidence and Predictors of Unplanned Hospital Readmission after Percutaneous Coronary Intervention. *Journal of clinical medicine*. 2020 10;9(10):3242.
- [4] Kirk PS, Liu X, Borza T, Li BY, Sessine M, Zhu K, et al. Dynamic readmission prediction using routine postoperative laboratory results after radical cystectomy. *Urologic oncology*. 2020 1;38(4):255-61.

- [5] Mehta SJ, Hume E, Troxel AB, Reitz C, Norton L, Lacko H, et al. Effect of Remote Monitoring on Discharge to Home, Return to Activity, and Rehospitalization After Hip and Knee Arthroplasty: A Randomized Clinical Trial. *JAMA network open*. 2020 12;3(12):e2028328.
- [6] Bonetti AF, Reis WCT, Mendes AM, Rotta I, Tonin FS, Fernandez-Llimos F, et al. Impact of Pharmacist-led Discharge Counseling on Hospital Readmission and Emergency Department Visits: A Systematic Review and Meta-analysis. *Journal of hospital medicine*. 2019 3;15(1):52-9.
- [7] Ko M, Chen E, Agrawal A, Rajpurkar P, Avati A, Ng AY, et al. Improving hospital readmission prediction using individualized utility analysis. *Journal of biomedical informatics*. 2021 6;119:103826-6.
- [8] Marafino BJ, Schuler A, Liu VX, Escobar GJ, Baiocchi M. Predicting preventable hospital readmissions with causal machine learning. *Health services research*. 2020 10;55(6):993-1002.
- [9] da Silva Loreto M, Lisboa T, Moreira VP. Early prediction of ICU readmissions using classification algorithms. *Computers in biology and medicine*. 2020 2;118:103636.
- [10] Hickson LTJ, Rule AD, Thorsteinsdottir B, Shields RC, Porter IE, Fleming MD, et al. Predictors of early mortality and readmissions among dialysis patients undergoing lower extremity amputation. *Journal of vascular surgery*. 2018 6;68(5):1505-16.
- [11] Swami J, Donihi AC, Siminerio LM, French EK, Delisi K, Hlasnik DS, et al. Readmission and Comprehension of Diabetes Education at Discharge (ReCoDED Study). *Diabetes*. 2018 7;67(Supplement₁).
- [12] Mori M, Mahmood SUB, Zhuo H, Yousef S, Green J, Mangi AA, et al. Persistence of risk of death after hospital discharge to locations other than home after cardiac surgery. *The Journal of thoracic and cardiovascular surgery*. 2019 3;159(2):528-35.e1.
- [13] Hu Z, Du D. A new analytical framework for missing data imputation and classification with uncertainty: Missing data imputation and heart failure readmission prediction. *PloS one*. 2020 9;15(9):1-15.
- [14] Horne BD, Roberts C, Rasmusson K, Buckway J, Alharethi R, Cruz J, et al. Risk score-guided multidisciplinary team-based Care for Heart Failure Inpatients is associated with lower 30-day readmission and lower 30-day mortality. *American heart journal*. 2019 9;219:78-88.
- [15] Milliren CE, Melvin P, Ozonoff A. Pediatric Hospital Readmissions for Infants With Neonatal Opioid Withdrawal Syndrome, 2016-2019. *Hospital pediatrics*. 2021 9;11(9):979-88.
- [16] McKechnie D, Fisher MJ, Pryor J, McKechnie R. Predictors of Unplanned Readmission to Acute Care from Inpatient Brain Injury Rehabilitation. *Journal of clinical nursing*. 2019 12;29(3-4):593-601.
- [17] Savitz ST, Leong TK, Sung SH, Lee K, Rana JS, Tabada GH, et al. Novel Data Domains and Machine Learning Modestly Improved Performance of Risk Calculators for Heart Failure Readmission. *Health Services Research*. 2020 8;55(S1):85-5.
- [18] Thomson R, Fletcher N, Valencia O, Sharma V. Readmission to the Intensive Care Unit Following Cardiac Surgery: A Derived and Validated Risk Prediction Model in 4,869 Patients. *Journal of cardiothoracic and vascular anesthesia*. 2018 4;32(6):2685-91.
- [19] Wang H, Cui Z, Chen Y, Avidan MS, Abdallah AB, Kronzer A. Predicting Hospital Readmission via Cost-Sensitive Deep Learning. *IEEE/ACM transactions on computational biology and bioinformatics*. 2018 4;15(6):1968-78.
- [20] Franckowiak TM, Raub JN, Yost R. Derivation and validation of a hospital all-cause 30-day readmission index. *American journal of health-system pharmacy : AJHP : official journal of the American Society of Health-System Pharmacists*. 2019 3;76(7):436-43.

- [21] Khoury H, Ragalie WS, Sanaiha Y, Boutros H, Rudasill SE, Shemin RJ, et al. Readmission After Surgical Aortic Valve Replacement in the United States. *The Annals of thoracic surgery*. 2020 1;110(3):849-55.
- [22] Everett E, Mathioudakis N. Association of socioeconomic status and DKA readmission in adults with type 1 diabetes: analysis of the US National Readmission Database. *BMJ open diabetes research & care*. 2019 4;7(1):e000621.
- [23] Barroilhet S, Bieling AE, McCoy TH, Perlis RH. Association between DSM-5 and ICD-11 personality dimensional traits in a general medical cohort and readmission and mortality. *General hospital psychiatry*. 2020 1;64:63-7.
- [24] Rossi I, Ross SW, May AK, Reinke CE. Readmission After Emergency General Surgery: NSQIP Review of Risk, Cause and Ideal Follow-Up. *The Journal of surgical research*. 2020 12;260:359-68.
- [25] Kwok CS, Rao SV, Potts J, Kontopantelis E, Rashid M, Kinnaird T, et al. Burden of 30-Day Readmissions After Percutaneous Coronary Intervention in 833,344 Patients in the United States: Predictors, Causes, and Cost: Insights From the Nationwide Readmission Database. *JACC Cardiovascular interventions*. 2018 4;11(7):665-74.
- [26] Chen SL, Kuo IJ, Kabutey NK, Gabra F, Fujitani RM. Perioperative risk factors for hospital readmission after elective endovascular aortic aneurysm repair. *Journal of vascular surgery*. 2018 4;68(3):731-8.e1.
- [27] Lee D, Keller MS, Fridman R, Lee J, Pevnick JM. Association between operational positive depression symptom screen scores on hospital admission and 30-day readmissions. *General hospital psychiatry*. 2021 2;70:38-43.
- [28] Noel CW, Forner D, Wu V, Enepekides D, Irish JC, Husain ZA, et al. Predictors of surgical readmission, unplanned hospitalization and emergency department use in head and neck oncology: A systematic review. *Oral oncology*. 2020 10;111:105039-9.
- [29] Hadaya J, Tran Z, Sanaiha Y, Aguayo E, Dobaria V, Press MC, et al. Impact of frailty on short term outcomes, resource use, and readmissions after transcatheter mitral valve repair: A national analysis. *PloS one*. 2021 11;16(11):e0259863.
- [30] Nazir S, Ahuja KR, Ariss RW, Chandal KH, Khuder SA, Moukarbel GV. Home health care utilization trend, predictors, and association with early rehospitalization following endovascular transcatheter aortic valve replacement. *Cardiovascular revascularization medicine : including molecular interventions*. 2021 5;36:1-6.
- [31] Shaheen AAM, Nguyen HH, Congly SE, Kaplan GG, Swain MG. Nationwide estimates and risk factors of hospital readmission in patients with cirrhosis in the United States. *Liver international : official journal of the International Association for the Study of the Liver*. 2019 2;39(5):878-84.
- [32] Galet C, Zhou Y, Eyck PT, Romanowski KS. Fall injuries, associated deaths, and 30-day readmission for subsequent falls are increasing in the elderly US population: a query of the WHO mortality database and National Readmission Database from 2010 to 2014. *Clinical epidemiology*. 2018 11;10:1627-37.
- [33] Flaks-Manov N, Srulovici E, Yahalom R, Perry-Mezre H, Balicer RD, Shadmi E. Preventing Hospital Readmissions: Healthcare Providers' Perspectives on "Impactability" Beyond EHR 30-Day Readmission Risk Prediction. *Journal of general internal medicine*. 2020 3;35(5):1484-9.
- [34] Kwok CS, Shah B, Al-Suwaidi J, Fischman DL, Holmvang L, Alraies C, et al. Timing and Causes of Unplanned Readmissions After Percutaneous Coronary Intervention: Insights From the Nationwide Readmission Database. *JACC Cardiovascular interventions*. 2019 3;12(8):734-48.
- [35] Blanc AL, Fumeaux T, Stirnemann J, Lozeron ED, Ourhamoune A, Desmeules JA, et al. Development of a predictive score for potentially avoidable hospital readmissions for general internal medicine patients. *PloS one*. 2019 7;14(7):e0219348.

- [36] McConachie SM, Raub JN, Yost R, Monday L, Agrawal S, Tannous P. Evaluation of a multidisciplinary approach to reduce internal medicine readmissions using a readmission prediction index. *American journal of health-system pharmacy : AJHP : official journal of the American Society of Health-System Pharmacists*. 2020 5;77(12):950-7.
- [37] Cousins J, Nair NP, Curtain C, Bereznicki BJ, Wilson K, Adamczewski B, et al. Preventing Adverse Drug Reactions After Hospital Discharge (PADR-AD): Protocol for a randomised-controlled trial in older people. *Research in social & administrative pharmacy : RSAP*. 2021 9;18(8):3284-9.