ANESTHESIOLOGY

Machine Learning Prediction of Postoperative Emergency Department Hospital Readmission

Velibor V. Mišić, Ph.D., Eilon Gabel, M.D., Ira Hofer, M.D., Kumar Rajaram, Ph.D., Aman Mahajan, M.D., Ph.D.

EDITOR'S PERSPECTIVE

What We Already Know about This Topic

- Unplanned hospital readmissions are a focus of quality improvement, national benchmarking, and payment incentives in the United States
- The accuracy of commonly used peer-reviewed readmission prediction algorithms at specific hospitals may be limited by hospital-specific factors
- The potential value of novel machine learning techniques capable of incorporating hundreds of patient, process, and hospital attributes is unclear

What This Manuscript Tells Us That Is New

- Hospital-specific 30-day surgical readmission models using machine learning techniques provide clinically usable predictions when applied to future patients
- A parsimonious approach limiting which data elements are considered performs as well as more comprehensive models

ABSTRACT

Background: Although prediction of hospital readmissions has been studied in medical patients, it has received relatively little attention in surgical patient populations. Published predictors require information only available at the moment of discharge. The authors hypothesized that machine learning approaches can be leveraged to accurately predict readmissions in postoperative patients from the emergency department. Further, the authors hypothesize that these approaches can accurately predict the risk of readmission much sooner than hospital discharge.

Methods: Using a cohort of surgical patients at a tertiary care academic medical center, surgical, demographic, lab, medication, care team, and current procedural terminology data were extracted from the electronic health record. The primary outcome was whether there existed a future hospital readmission originating from the emergency department within 30 days of surgery. Secondarily, the time interval from surgery to the prediction was analyzed at 0, 12, 24, 36, 48, and 60 h. Different machine learning models for predicting the primary outcome were evaluated with respect to the area under the receiver-operator characteristic curve metric using different permutations of the available features.

Results: Surgical hospital admissions (N = 34,532) from April 2013 to December 2016 were included in the analysis. Surgical and demographic features led to moderate discrimination for prediction after discharge (area under the curve: 0.74 to 0.76), whereas medication, consulting team, and current procedural terminology features did not improve the discrimination. Lab features improved discrimination, with gradient-boosted trees attaining the best performance (area under the curve: 0.866, SD 0.006). This performance was sustained during temporal validation with 2017 to 2018 data (area under the curve: 0.85 to 0.88). Lastly, the discrimination of the predictions calculated 36 h after surgery (area under the curve: 0.88 to 0.89) nearly matched those from time of discharge.

Conclusions: A machine learning approach to predicting postoperative readmission can produce hospital-specific models for accurately predicting 30-day readmissions *via* the emergency department. Moreover, these predictions can be confidently calculated at 36 h after surgery without consideration of discharge-level data.

(*Anesthesiology* 2020; XXX:00–00)

Unplanned hospital readmissions have been a pressing concern with respect to patient burden and high cost.[1] These factors are compounded when a patient is readmitted *via* the emergency department where they may consume valuable resources. This is in contrast to postoperative readmissions that occur *via* direct admission after clinic visits that are often planned or facilitated by a surgeon. Although it is not always clear whether a given readmission is preventable, it is certain that emergency department–based 30-day readmissions are sentinel events, poor markers of quality, and are typically attributable to conditions present at discharge.[2,3]

Models created to detect patients at risk for unplanned readmission have been developed, but many of them have a narrow focus on single disease states and cannot be applied to the postoperative population as a whole.[4–6] Furthermore, most of these models were exclusively created

for nonsurgical patients, making them difficult to validate in different cohorts,[7] and many of them require substantial data from the moment of hospital discharge.[8,9] Waiting until hospital discharge inhibits providers from better optimizing factors associated with readmission in parallel with the routine postoperative care.[4,10–12] This delay can also lead to ineffective transitional care coordination or possible prolongation of a hospital stay. Moreover, given that timeliness is one of the five rights of clinical decision support, waiting until the moment of discharge dramatically decrease the clinical efficacy of any model.[13]

A potential path to developing patient-risk models is machine learning with its ability to process extremely large numbers of input features and produce accurate predictive models.[14,15] More specifically, tree-based machine learning methods are able to model nonlinear relationships and interactions, typically outperforming standard logistic regression.[16] Lastly, machine learning models which are calibrated using an institution's individual data demonstrate higher accuracy than models engineered for generalized patient cohorts that ordinarily overlook unforeseen institution-specific nuances.[17]

In this study, we describe the creation and comprehensive validation of a machine learning–based methodology for predicting a patient's risk for 30-day readmission *via* the emergency department in the postoperative period. Our primary hypothesis is that machine learning methods are capable of producing hospital-specific readmission prediction models with excellent discrimination. Our secondary hypothesis is that machine learning can produce models that accurately predict postoperative readmission without dependence on data from the time of discharge.

Materials and Methods

Data Extraction

Our data extraction methodology qualified for University of California Los Angeles (Los Angeles, California) institutional review board (No. 18-000630) exception status ("waiver of consent") by virtue of having no direct patient contact and using a deidentified dataset. All study data were acquired *via* our previously published Department of Anesthesiology and Perioperative Medicine at University of California Los Angeles's Perioperative Data Warehouse.[18] The Perioperative Data Warehouse is a structured reporting data schema that contains all the relevant clinical data entered into the EPIC (EPIC Systems, USA) electronic health record system. Data were acquired *via* Clarity, the relational database created by EPIC for data analytics and reporting. Although Clarity contains raw clinical data, the Perioperative Data Warehouse was designed to organize, filter, and improve data so that they can be used reliably for creating these types of metrics. Other published articles deriving data from the Perioperative Data Warehouse can be found in the References section.[18–22]

Data extraction was restricted to the University of California Los Angeles Ronald Reagan Medical Center for developing the model in a tertiary center. This was followed by a later extract of University of California Los Angeles's Santa Monica Hospital for external validation of the methodology.

Model Endpoint Definition

We defined a readmission *via* the emergency department as any patient who enters University of California Los Angeles through any of the emergency departments within 30 days of a surgical case and is then transferred to a subsequent non–emergency department location. This definition was intended to capture patients who return to the hospital in an unplanned fashion and require inpatient or observation level of care, compared with those who were sent home after an emergency department–based evaluation.

There was a significant effort to mimic the Center for Medicare and Medicaid Services definition of postoperative readmission, but to match definitions exactly, we would need access to the proprietary algorithms created by the various third-party vendors. Our health center currently uses one such service to analyze our own data to discern which of the 30-day readmissions are exempt based on an allowable disease condition or surgical procedure. Despite not having access to the aforementioned proprietary algorithms, we did adopt the Center for Medicare and Medicaid Services exclusion of cases that were discharged and readmitted on the same calendar day.

During the study interval, approximately 40% of the Ronald Reagan Medical Center 30-day readmissions arrived at the health center *via* the emergency department, another 40% entered though a perioperative location, and the remaining 20% of the readmissions were direct to an inpatient location. It was decided that the 40% of patients arriving *via* a perioperative location were not consistent with our primary outcome because it would be exceedingly difficult to identify which procedures were done as corrective revisions (unplanned) and those that occurred as part of staged procedure (planned). One such example is cataract surgery, because our ophthalmologists seldom operate on both eyes during a single case, but instead stagger the cases two weeks by design. As for the 20% that are direct admissions, the same dilemma applies because we would not definitively be able to distinguish which patients were admitted from clinic because of complication *versus* those needing adjuvant medical treatment (*i.e.*, postoperative chemotherapy).

Model Input Features

For each admission, we used a large collection of independent variables originating from different sources, which we summarize below.

General Data. We considered variables that describe the surgery at admission, such as the volume of blood loss and

the duration of surgery. We also considered variables that summarize overall health, such as the American Society of Anesthesiologists (Schaumburg, Illinois) score, as well as variables that describe specific aspects of patient health, such as the Acute Kidney Injury Network stage and whether the patient had received a consultation from the pain management service during a past admission.[23] We also considered nonmedical, demographic variables such as the patient's age, ethnicity, race, and primary language (Supplemental Digital Content table 1, http://links.lww.com/ALN/C173).

Lab Data. We considered a collection of commonly ordered lab tests: bilirubin, creatinine, glucose, hematocrit, hemoglobin, INR, P_{CO_2}, platelets, P_{O_2}, potassium, sodium, urea nitrogen, and leukocyte count. The process of identifying and grouping the labs was manually done by a physician informaticist on our research team with expertise in programming. The grouped labs are summarized in Supplemental Digital Content table 2 (http://links.lww.com/ALN/C174). For each type of lab, the Perioperative Data Warehouse records different subtypes of labs. These subtypes arise out of how the lab is ordered (for example, bilirubin may be measured as part of complete metabolic panel, or simply as a stand-alone lab result). For each lab subtype, we created two variables: one variable that measures the maximum deviation (delta) from the normal range observed in that lab subtype over the whole admission and one 0/1 variable to indicate whether this lab subtype was ever ordered. In addition, we created one variable to count the total number of unique lab subtypes ordered over the entire admission, and one 0/1 variable to indicate whether no labs were ever ordered during the admission. At the time of data collection and algorithm development, Logical Observation Identifiers Names and Codes codes were not available within our electronic medical record implementation. This has since changed and Logical Observation Identifiers Names and Codes are currently being implemented. Once complete, we will be able to test the substitution of Logical Observation Identifiers Names and Codes for component identifiers.

Medication Data. During each admission, patients are typically administered or prescribed many different classes of medications. For each medication, we had access to the pharmaceutical class (total of 99) and subclass (total of 537) as defined by our institution's pharmacy service. For each pharmaceutical class, we created one 0/1 indicator variable to indicate whether any medications of that class were ever prescribed during admission, one 0/1 variable to indicate whether any medications of that class were ever taken during the admission, and one 0/1 variable to indicate whether any discharge medications of that class were prescribed during the admission. We also created variables to count the number of unique classes prescribed, the number of unique classes taken by the patient, and the number of unique classes among the patient's discharge medications.

Team Data. For each admission, we knew which of the 170 surgical, medical, or consulting teams at the Ronald Reagan Medical Center was assigned to the patient during a given admission. For each team, we created a 0/1 indicator variable of whether that team was assigned to the admission. We also define a variable to count the total number of teams assigned to each admission.

Current Procedural Terminology Data. For each surgical case, we generated a list of all the current procedural terminology codes that were billed for by the surgeons. Because of computational considerations, we focused only on current procedural terminology codes that appear in at least 100 surgical admissions in the data set, resulting in 215 current procedural terminology codes. For each such code, we created a 0/1 variable to indicate whether that code is assigned to each principal surgical admission. We also created a variable to count the total number of current procedural terminology codes assigned to each admission.

Data Preprocessing

For a minority of admissions, the values of some variables were missing or null. For categorical variables, we added another category indicating that the value was missing. For some numeric variables, we were able to infer that a missing value indicates that the variable value is zero. When this was not possible, we created a new 0/1 indicator of whether the variable value is missing for each admission, and we set the original missing values to zero.

We divided the data randomly into a training set and testing set according to a 70–30 split.[15] We used the training set (70%) to estimate and the test set (30%) to evaluate each model. Each split was done to preserve the proportion of readmitted/not readmitted cases found in the whole data set. We repeated this random splitting ten times.

Model Development

We considered three different types of models: regularized logistic regression, random forest, and gradient boosted trees; the latter two are classified as tree ensemble models.[14,24,25] We provide additional technical information on the methods and parameter settings in the Supplemental Digital Content (http://links.lww.com/ALN/C175). These models differ from classical logistic regression in their ability to scale to large numbers of features without overfitting. Although the models accomplish this in different ways, the essential idea is that each model's estimation procedure has some mechanism for controlling the model complexity and how sensitive the model is to the data.[14]

Regularized Logistic Regression. Logistic regression models the probability of a readmission as a logistic function applied to a linear function of the independent variables. Ordinary logistic regression models are estimated by minimizing the negative log–likelihood. In this work, we did not use ordinary logistic regression, because it can overfit when the number of features is large relative to the training

set size and perform poorly out of sample. We instead used regularized logistic regression, which differs from ordinary logistic regression in that the estimation minimizes the negative log–likelihood plus an additional term (the regularization term). We used L1 regularization, which uses the L1 norm of the coefficient vector as the regularization term.[26] L1 regularization has two special properties: It induces shrinkage (it reduces the magnitude of the coefficients) and sparsity (it returns models where many coefficients are set to zero).[15] Both of these properties lead to good predictive performance when the number of independent variables is very large relative to the number of observations. This form of regression is also known as Least Absolute Shrinkage and Selection Operator logistic regression. Alternatives include L2 regularized logistic regression (also known as ridge regression), wherein the regularization term is the sum of the squares of the coefficients, and elastic net logistic regression, wherein the regularization term is a weighted combination of the Least Absolute Shrinkage and Selection Operator and ridge regularization terms.

Random Forest. The random forest algorithm works by estimating an ensemble of randomized classification trees. A classification tree is a predictive model where one follows a sequence of true/false queries along a tree to make a prediction.[27] Given an observation, each tree in the forest is used to make a prediction (readmission or no readmission) and the fraction of trees in the forest predicting readmission is then the predicted readmission probability. A stylized example of a random forest model being used for prediction is given in Supplemental Digital Content figure 1 (http://links.lww.com/ALN/C176).

The random forest model is attractive over ordinary logistic regression for two reasons. First, random forests are able to automatically learn potentially nonlinear relationships between each feature and the readmission probability, as well as interactions between the features.[14,28] Second, random forests yield good performance when the number of features is large relative to the number of observations, possibly being as large or larger than the number of observations.[24,28,29]

Gradient Boosted Trees. The gradient boosted tree model is another type of tree ensemble model and also works by growing multiple classification trees. However, rather than growing each tree randomly, the trees are grown sequentially, where each new tree is selected to most reduce the error of the current ensemble. The ensemble outputs a probability between 0 and 1 of a readmission occurring, and the ensemble is trained to minimize logistic loss. Like random forests, gradient boosted trees are well suited to prediction problems with large numbers of features, and can automatically handle nonlinearities and interactions.[14]

Predischarge *versus* Postdischarge Prediction

We performed an additional set of experiments to evaluate our models' performance using only predischarge information. We focused on the general data and the lab data. For the general data, we excluded any features that would only be available at discharge. For the lab data, we computed the features in the same way as before, but ignored any labs drawn later than T hours after the completion of surgery, where T is a parameter that we vary. We restricted T to the range $\{0, 12, 24, 36, 48, 60\}$. (For example, a value of $T = 24$ corresponds to using all data from the time of admission, to 24 h after the completion of surgery.) In the case that a patient is discharged before the cutoff point of T h after surgery completion, we did not use any additional information that may become available at or after discharge (such as length of stay). Note that the medication, team, and current procedural terminology data were not used in these additional experiments.

Model Performance

To evaluate our models, we considered two predictive metrics: area under the receiver-operator characteristic curve (also known as the c statistic or simply area under the curve) and Brier score. The area under the curve is the probability of correctly distinguishing between a randomly chosen admission from the test set that results in a readmission, and a randomly chosen admission from the test set that does not. The area under the curve is bounded between 0.5 and 1.0, with higher values being better. The Brier score is the mean squared difference between the predicted probability of readmission and the actual outcome (0 or 1 where 1 indicates a readmission). The Brier score is bounded between 0 and 1, with lower values being better. Both the area under the curve and Brier score metrics were calculated using the test data set and averaged over the ten random splits of the data. We additionally compared the models by plotting their receiver-operator characteristic curves and their precision-recall curves.

We compared our area under the curve values against the HOSPITAL (low hemoglobin at discharge, discharge from an oncology service, low sodium at discharge, [International Classification of Disease, ninth revision, clinical modification coded] procedure during hospital stay or not, index admission type is elective or not, number of admissions in previous year, and length of stay) score, a recent model for readmission risk prediction developed in a general, nonsurgical patient population.[30] We acknowledge that the HOSPITAL score was not originally formulated for surgical readmission prediction; our interest in evaluating HOSPITAL is simply to obtain a reasonable benchmark. We also compared our area under the curve values against the LACE (length of stay, acuity, comorbidity, emergency department use) score. The LACE score and the later LACE+ score (which additionally uses patient age and sex, hospital teaching status, acute diagnoses and procedures during admission, number of days on alternative level of care, and number of admissions to the hospital in the previous year) have been applied to populations consisting of both medical and surgical patients.[9,31]
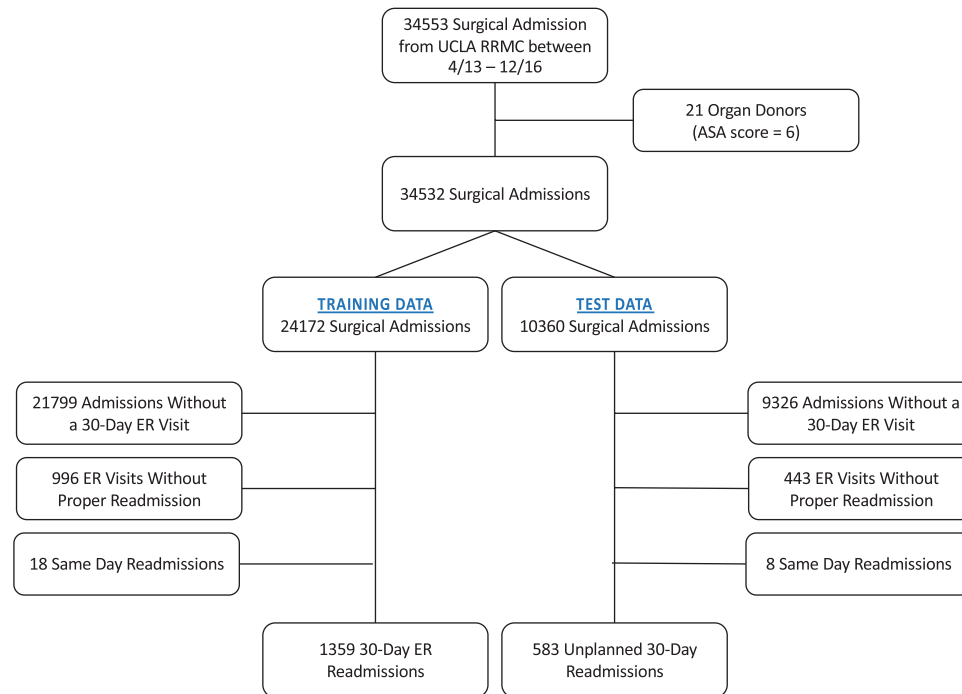
Fig. 1. Inclusion and exclusion criteria for overall data set and emergency department readmissions at the University of California Los Angeles (UCLA) Ronald Reagan Medical Center (RRMC). ER, emergency room.

External Validation

As an additional exercise in validating the overall methodologic approach, we also applied it to the University of California Los Angeles Santa Monica Hospital. The Santa Monica Hospital differs greatly from the Ronald Reagan Medical Center as it functions much more like a community hospital with primarily orthopedic cases and lower patient acuity. Typically, patients with significant comorbidities or those requiring extensive surgeries are transferred by ambulance from the Santa Monica Hospital to the Ronald Reagan Medical Center. The Santa Monica Hospital has a fraction of the intensive care capabilities and a fraction of the subspecialty resources.

To validate the overall methodologic approach, we extracted identical features using identical inclusion/exclusion criteria from the Perioperative Data Warehouse for the Santa Monica Hospital. After splitting the data in the same way as for the Ronald Reagan Medical Center, we retrained our models using the Santa Monica Hospital training sets and tested their performance on the Santa Monica Hospital test sets.

We remark here that this is different from how external validation is commonly done. Ordinarily, external validation entails taking a model trained on one data set (in this case, the Ronald Reagan Medical Center data) and evaluating its performance on a different data set (in this case, the Santa Monica Hospital data). We emphasize that our

goal is to validate our overall methodology—the process that transforms a data source, such as the Perioperative Data Warehouse, into a data set that can then be turned into a predictive model—and not to validate the specific models that arise from the Ronald Reagan Medical Center data set. Given the hospital-specific nature of readmissions, it is unreasonable to expect the existence of a single universal model that will achieve good predictive performance across a wide range of institutions. Furthermore, even if such a model were to exist, an institution may still wish to leverage its own data to construct a model tailored to its own patient population and surgical practices, to achieve even better performance. By validating our methodology on the Santa Monica Hospital, a hospital with a different patient population and surgical specialties from the Ronald Reagan Medical Center, we intended to verify whether our methodology is applicable to institutions outside of the Ronald Reagan Medical Center.

Temporal Validation

In addition to external validation, we also carried out a temporal validation of our methodology. We extracted admissions occurring at the Ronald Reagan Medical Center in 2017 and 2018, after the period spanned by our base dataset (2013–2016), and derived the endpoint and the features in exactly the same way as described for the base dataset. We then used the models we developed using the ten random

splits of the 2013 through 2016 Ronald Reagan Medical Center data to make predictions on these admissions. The predictions were evaluated using the area under the curve as well as the positive predictive value and negative predictive value. The positive predictive value is calculated as the number of true positives divided by the total number of positives, whereas the negative predictive value is calculated as the number of true negatives by the total number of negatives. The cutoff used for the positive predictive value and the negative predictive value was 0.20 (*i.e.*, a positive is defined as an admission with a predicted readmission probability of at least 0.20, whereas a negative is an admission with a predicted probability of less than 0.20). The purpose of this temporal validation was to evaluate the performance of the models when they are used prospectively (*i.e.*, to predict on observations arising from the same institution in the future), thereby mimicking how these models would be used in practice.

Results

Data Extraction

Patients aged 18 or older (n = 28,728) were extracted from the Perioperative Data Warehouse from April 2013 to December 2016. This resulted in 34,553 admissions. Of these, 21 were removed for being organ donors (American Society of Anesthesiologists score of VI). The resulting 34,532 admissions constituted our complete data set.

To define our endpoint, we considered the 3,407 admissions that led to a 30-day emergency department visit; 1,439 were excluded for not resulting in a readmission that led to a transfer to a unit beyond the initial emergency department encounter. Lastly, 26 were excluded because of the emergency department visit occurring on the same calendar day as discharge. This resulted in 1,942 admissions with a readmission *via* the emergency department; thus, 5.6% of the admissions in the complete data set of 34,532 admissions met the endpoint definition. Figure 2 summarizes the inclusion/exclusion criteria.

Admission Characteristics

Table 1 summarizes the general admission characteristics, and table 2 summarizes the lab data. The total number of lab subtypes is 119; therefore, there are 119 maximum lab deviation variables, 119 lab presence indicator variables, one variable counting the total number of lab subtypes, and one variable indicating no labs drawn, for a total of 240 lab-related variables. Of the 34,532 admissions, 27,071 admissions (78.4%) had at least one lab drawn. Patients were administered 92 different pharmaceutical classes; following our feature creation procedure, this resulted in a total of 279 variables. The average number of unique classes prescribed per admission was 19.7, with a minimum of 0 and a maximum of 59; an average of 16.8 classes were taken by each admission, and an average of 4.3 classes were prescribed at

discharge. Our data set spanned 158 surgical teams; following our feature creation procedure, this resulted in a total of 159 team-related variables. For each admission, the number of teams ranged from 1 to 36, with a mean of 1.9 and a median of 1.

Model Performance

Predictive Metrics. Table 3 displays the area under the curve and Brier score metrics for the different predictive models and using different sets of variables. Using only the general data resulted in area under the curve values on the order of 0.73 to 0.76. Adding the lab data appreciably increased area under the curve values to the 0.85 to 0.87 range. The medication, team, and current procedural terminology data did not lead to significant improvements in the area under the curve. The HOSPITAL score achieved an area under the curve of 0.73, which is lower than the area under the curve values of our models. This value is comparable with that in the original HOSPITAL paper, where it was tested with a general medical population.[30] The area under the curve for LACE was also 0.73, which is slightly higher than in the original LACE paper.[9] The same qualitative behavior is also observed for the Brier score.

To further compare the models, figure 2 plots the receiver-operator characteristic curves for our models using the general data and the lab data, as well as the HOSPITAL and LACE scores, for one random split of the data. Figure 3 similarly plots the precision–recall curves for our models using the general and lab data, as well as HOSPITAL and LACE, for the same random split of the data. Our three models gave similar performance, while outperforming both HOSPITAL and LACE. Lastly, Supplemental Digital Content figure 2 (http://links.lww.com/ALN/C176) plots calibration curves for our models using the general and the lab data, for the same random split of the data. Our models generally produced probability predictions that closely match the actual readmission probabilities. The random forest model appears to have overestimated the readmission risk when it is in the 0.05 to 0.30 range, whereas both L1 regularized logistic regression and the gradient boosted tree model appear to have slightly underestimated the readmission risk in the same range. Calibration results in table form for the same random split are provided as Supplemental Digital Content table 3 (http://links.lww.com/ALN/C177).

Feature Importance. Supplemental Digital Content figure 3 (http://links.lww.com/ALN/C176) shows the top 30 most important variables of the random forest model for a single random split of the data. The majority of the top 30 variables were maximum delta and indicator variables associated with a variety of lab subtypes. In addition to the lab-based features, the top 30 for the random forest model also included length of stay, the relative time to surgery, and the duration of surgery. Supplemental Digital Content figures 4 and 5 (http://links.lww.com/ALN/C176) similarly shows

Fig. 2. Regularized logistic regression, random forest, and gradient boosted tree model were trained using *general* and *labs* features, for a single random split of the data. The true positive rate (TPR) and false positive rate (FPR) are computed from the test set. For comparison, the curves for the HOSPITAL (low hemoglobin at discharge, discharge from an oncology service, low sodium at discharge, [International Classification of Disease, ninth revision, clinical modification coded] procedure during hospital stay or not, index admission type is elective or not, number of admissions in previous year, and length of stay) and LACE (length of stay, acuity, comorbidity, emergency department use) scores are also shown.

the top 30 variables of the gradient boosted tree model and the regularized logistic regression model, respectively.

Predischarge versus Postdischarge Prediction. Supplemental Digital Content figure 6 (http://links.lww.com/ALN/C176) shows how the area under the curve varies as the parameter *T*, which determines when the prediction is to be made beyond the start of surgery, varies. As *T* increased, all three models improved and plateaued after $T = 1.5$ days. After $T = 1.5$ days (*i.e.*, 36 h after the start of surgery), the area under the curve values were comparable with those achieved when making the prediction after discharge.

External Validation Using the Santa Monica Hospital Data. We applied the same methodology as the Ronald Reagan Medical Center to the Santa Monica Hospital. We extracted data from 19,650 surgical admissions at the Santa Monica Hospital from the same time span. After removing organ

donors (18 admissions), the resulting 19,632 admissions constituted our complete Santa Monica Hospital data set. Of these, 820 (4.2%) admissions had an emergency department readmission and met our endpoint definition.

We divided the data into a training and testing set according to a 70-30 split. After training on 70% of the data, the test set contained 5,890 admissions resulting in 246 readmissions *via* the emergency department (see the Santa Monica Hospital Consolidated Standards of Reporting Trials diagram in Supplemental Digital Content figure 7, http://links.lww.com/ALN/C176). We trained the same three machine learning models using only the general data, and using both the general and lab data. The model results with the general and lab data were nearly identical to those for the Ronald Reagan Medical Center, with area under the curve values in the 0.86 to 0.88 range (table 4). In addition,

Table 1. Summary of Admission Characteristics in Complete University of California Los Angeles Ronald Reagan Medical Center Data Set

Admissions	34,532
Emergency department readmissions	1,942 (5.6%)
Admissions with more than one emergency department readmission	20 (0.0%)
Maximum number of emergency department readmissions per admission	2
Age	55 ± 17
American Society of Anesthesiologists score	
I	2,328 (6.7%)
II	11,668 (33.8%)
III	16,964 (49.1%)
IV	3,395 (9.8%)
V	171 (0.5%)
NA	6 (0.0%)
Admissions with primary diagnosis-related group code	25,730 (74.5%)
Number of surgery common procedural terminology codes	3 ± 2
Number of diagnosis-related group codes	4 ± 3
Number of admissions in previous year	1 ± 1
Surgical information*	
Total blood loss, ml	252 ± 1007
Total fluids transfused, ml	1,565 ± 1445
Total colloid transfused, ml	920 ± 1102
Total blood transfused, ml	137 ± 1012
Duration of anesthesia, min	251 ± 168
Hemodialysis, n admissions	600 (1.7%)
Tracheostomy	
Yes	615 (2.4%)
No	3,910 (15.24%)
NA	29,976 (86.8%)
Mechanical ventilation, n admissions	4,523 (17.6%)
Duration of mechanical ventilation,* min	6,388 ± 16099
Length of stay, h	150 ± 324
Disposition	
Home or self care	25,789 (74.7%)
Home health service	5,773 (16.7%)
Skilled nursing facility	1,566 (4.5%)
Other	1,404 (4.1%)

Values are reported as either number of admissions (% of total data set) or as mean ± SD. For variables or groups of variables marked with *, the mean and SD are taken over observations with nonmissing values for those variables.

we also used the data to train and evaluate models for pre-discharge prediction, analogously with the Ronald Reagan Medical Center. The area under the curve as a function of the cutoff parameter T is shown in Supplemental Digital Content figure 8 (http://links.lww.com/ALN/C176). This figure is consistent with Supplemental Digital Content figure 6 (http://links.lww.com/ALN/C176) for the Ronald Reagan Medical Center, showing that high area under the curve values can be obtained as soon as 36 h after surgery.

Temporal Validation at the Ronald Reagan Medical Center Using 2017 to 2018 Data. We extracted 19,343 admissions from the Ronald Reagan Medical Center in the period 2017 to 2018. Of these, 12 were removed because of organ donor status, leaving 19,331 admissions eligible for analysis. Within this set of admissions, 969 (5.0%) met our endpoint definition of an emergency department readmission. We considered all of our machine learning models built using the general and lab-based features, as this combination of features led to the best predictive performance when tested on 2013 to 2016 data. We evaluated the area under the curve of each such model when used to predict on admissions in 2017 to 2018. We also evaluated the positive predictive value and negative predictive value for a probability cutoff of 0.20. The results are shown in tables 5 and 6. These results were consistent with our earlier evaluation using 2013 to 2016 data (table 3), with all of our models achieving area under the curve values in the 0.85 to 0.88 range. In general, one would expect a model to perform worse in temporal validation, because observations in the future may behave differently from those in the training and testing data (for example, due to changes in operations or patient mix enacted since 2016). The consistency between our temporal validation results and our earlier results for 2013 to 2016 suggests that admissions in 2017 to 2018 behave similarly to those in 2013 to 2016. All three models achieved positive predictive values in the range of 0.20 to 0.40, with negative predictive values greater than 0.96. We note that the low positive predictive values are to be expected, due to the low emergency department readmission rate in the data. L1 logistic regression and gradient boosted trees achieve higher positive predictive values than random forest, at the cost of a slightly reduced negative predictive value. This is consistent with our calibration results (Supplemental Digital Content figure 2, http://links.lww.com/ALN/C176), which suggest that random forest often overestimates the readmission probability. As in our previous results, gradient boosted trees and random forest did not exhibit an edge over L1 logistic regression, despite their ability to automatically model interactions and nonlinearities.

Discussion

Our analysis of this data furnishes us with two key insights about prediction of emergency department readmission in postoperative patients. First, all of the machine learning models achieve high discrimination. Irrespective of which specific model is chosen, all of the models achieved out-of-sample area under the curve values in the 0.85 to 0.87 range using general demographic data, surgical data, and basic lab data. Most prior work in predicting readmissions in general medical patients achieved area under the curve values in the 0.6 to 0.7 range,[4] with HOSPITAL and LACE achieving area under the curve values of 0.72 and 0.68 in previous work, respectively.[9,30] Second, we find that there is virtually no loss in performance if our models are restricted to using data available within 36 h after the completion of surgery. This suggests that our models could be used to identify patients that are at high risk of readmission while still in the hospital and soon after the surgical procedure.

It is important to emphasize the difference in the definition of unplanned readmission between this manuscript and the Center for Medicare and Medicaid Services. Our prediction models are able to identify patients that are

Table 2. Summary of Lab Data for Complete University of California Los Angeles Ronald Reagan Medical Center Data Set

Lab Type	Number of Subtypes	Number of Admissions	Abnormal Definition	Average Delta	Maximum Delta
Bilirubin, mg/dl	7	2,819	> 1.2	7.2	104.1
Creatinine, mg/dl	5	7,128	> 1.3	1.8	25.1
Glucose, mM	19	24,613	< 65 or > 100	55	1,172
Hematocrit, %	5	18,228	< 35	7.6	34.5
Hemoglobin, g/dl	14	16,199	< 11	2.2	9.6
INR	9	4,718	> 1.3	0.9	148.7
P_{CO_2}, mm Hg	12	6,788	> 42	9	2,686
Platelet count, $10^3/\mu l$	4	8,222	< 143	59	142
P_{O_2}, mm Hg	12	7,632	< 98	36	88
Potassium, mM	10	11,837	< 3.6 or > 5.3	0.4	401.7
Sodium, mM	10	11,953	< 135 or > 145	3.7	131.2
Urea nitrogen, mg/dl	9	9,612	> 22	21	186
White blood cells, $10^3/\mu l$	3	19,037	< 4 or > 10	4.3	242.1

The second column indicates the number of subtypes of each lab. The third column indicates the number of admissions in the overall data set for which a lab of the given type was drawn. The last three columns indicate how abnormal values of the lab are defined, what the average deviation from normal is, and what the maximum deviation from the normal range is. The average and maximum of the deviations are taken only over abnormal lab values. The subtype refers to a unique identifier assigned to labs that changes based on ordering methodology (*i.e.*, order as a standalone element or as part of a panel) and laboratory location. Clinicians using the EPIC medical record system are blinded to the subtype being ordered/reviewed.

Table 3. Test Set Area under the Curve and Brier Score for Postdischarge Prediction at the University of California Los Angeles Ronald Reagan Medical Center

Data	Method	Area under the Curve	Brier Score
All	HOSPITAL	0.735 ± 0.009	0.051 ± 0.000
	LACE	0.732 ± 0.003	0.051 ± 0.000
general	L1 Logistic Regression	0.759 ± 0.007	0.050 ± 0.000
	Random Forest	0.744 ± 0.009	0.051 ± 0.000
	Gradient Boosted Trees	0.760 ± 0.006	0.050 ± 0.000
general + labs	L1 Logistic Regression	0.865 ± 0.005	0.045 ± 0.000
	Random Forest	0.864 ± 0.006	0.045 ± 0.000
	Gradient Boosted Trees	0.866 ± 0.006	0.045 ± 0.000
general + meds	L1 Logistic Regression	0.765 ± 0.007	0.050 ± 0.000
	Random Forest	0.746 ±0.008	0.050 ± 0.000
	Gradient Boosted Trees	0.764 ± 0.005	0.050 ± 0.000
general + teams	L1 Logistic Regression	0.758 ± 0.007	0.050 ± 0.000
	Random Forest	0.748 ± 0.009	0.050 ± 0.000
	Gradient Boosted Trees	0.759 ± 0.006	0.050 ± 0.000
general + cpt	L1 Logistic Regression	0.764 ± 0.006	0.050 ± 0.000
	Random Forest	0.748 ± 0.010	0.050 ± 0.000
	Gradient Boosted Trees	0.762 ± 0.007	0.050 ± 0.000
general + labs + meds + teams + cpt	L1 Logistic Regression	0.866 ± 0.006	0.045 ± 0.000
	Random Forest	0.860 ± 0.007	0.046 ± 0.000
	Gradient Boosted Trees	0.867 ± 0.007	0.045 ± 0.000

The area under the curve and Brier score values reported are averaged over the ten random splits of the data; mean ± SD. Under Data, *general* corresponds with the general surgical/health and demographic features; *labs* corresponds with the lab-based features; *meds* corresponds with the medication features; *teams* corresponds with the consulting team features; and *cpt* corresponds with the common procedural terminology features.

readmitted into the hospital *via* the emergency department irrespective of the cause. We believe that focusing on emergency department readmissions is more appropriate because such readmissions are unplanned in nature and costly. This is in contrast to the definition of unplanned readmission by the Center for Medicare and Medicaid Services and Horwitz *et al.*[32] Their definition of unplanned readmission excludes special cases such as transplanted patients and those undergoing chemotherapy. This highlights the clinical orientation of this manuscript rather than compliance to administrative workflows.

Choosing a Model

Interestingly, we do not find a significant difference between the three machine learning models. In particular, L1 regularized logistic regression obtained comparable performance to both random forest and gradient boosted trees. As compared with logistic regression, tree ensemble models have the advantage of being able to model nonlinear relationships and automatically incorporate interactions among a large number of features to extract accurate predictive models. Our results suggest that this level of flexibility is not necessary for modeling emergency department readmission risk.

We emphasize here that L1 regularized logistic regression is different from ordinary logistic regression, as we discussed in the Materials and Methods section. We found that in ordinary logistic regression, the estimation procedure would over fit on the training set and would produce grossly inaccurate test set predictions (for example, the average test set area under the curve with the general and labs data was 0.77, with a standard error of 0.04).

The main characteristic of the models that is crucial to the success of our approach is their ability to avoid overfitting in the presence of large numbers of features. Because of limitations of classical logistic regression, traditional approaches to developing risk models require significant manual input and expertise to filter out unnecessary features, and advocate for soft limits on the number of features,

Fig. 3. L1 regularized logistic regression, random forest, and gradient boosted tree model were trained using *general* and *labs* features, for a single random split of the data. The precision and recall are computed from the test set. For comparison, the curves for the HOSPITAL (low hemoglobin at discharge, discharge from an oncology service, low sodium at discharge, [International Classification of Disease, ninth revision, clinical modification coded] procedure during hospital stay or not, index admission type is elective or not, number of admissions in previous year, and length of stay) and LACE (length of stay, acuity, comorbidity, emergency department use) scores are also shown.

Table 4. Test Set Area under the Curve and Brier Score for Postdischarge Prediction at the University of California Los Angeles Santa Monica Hospital

Data	Method	Area under the Curve	Brier Score
general	L1 Logistic Regression	0.777 ± 0.013	0.038 ± 0.000
	Random Forest	0.762 ± 0.010	0.038 ± 0.000
	Gradient Boosted Trees	0.773 ± 0.014	0.038 ± 0.000
general + labs	L1 Logistic Regression	0.865 ± 0.009	0.037 ± 0.001
	Random Forest	0.875 ± 0.009	0.036 ± 0.000
	Gradient Boosted Trees	0.877 ± 0.007	0.036 ± 0.000

The area under the curve and Brier score values reported are averaged over the ten random splits of the data; mean ± SD. Under Data, *general* corresponds with the general surgical/health and demographic features and *labs* corresponds with the lab-based features.

such as the one-in-ten rule of not using more than one feature per ten events in the data, so as to avoid overfitting.[33–35] Machine learning methods, in contrast, are able to automatically handle large number of features and extract accurate predictive models without overfitting, thereby letting the data speak. Fortunately, there are many available resources for physicians to learn more about different machine learning fundamentals and techniques.[36–38]

Study Limitations

In this study, there were a few limitations that are inherent in these types of retrospective, machine learning projects. First, we have little insight into the reason for hospital readmission. Although we have created a model to predict that a patient will come back to the emergency room in 30 days, there is no information about how to possibly prevent the

Mišić *et al.*

Table 5. Temporal Validation Results for Postdischarge Prediction at the University of California Los Angeles Ronald Reagan Medical Center Data

Period	Method	Area under the Curve	Positive Predictive Value	Negative Predictive Value
2017–2018	L1 Logistic Regression	0.859 ± 0.003	0.354 ± 0.016	0.964 ± 0.002
2017–2018	Random Forest	0.856 ± 0.002	0.225 ± 0.008	0.975 ± 0.001
2017–2018	Gradient Boosted Trees	0.866 ± 0.003	0.311 ± 0.013	0.968 ± 0.001
2017	L1 Logistic Regression	0.858 ± 0.002	0.328 ± 0.015	0.966 ± 0.002
2017	Random Forest	0.851 ± 0.002	0.212 ± 0.006	0.976 ± 0.001
2017	Gradient Boosted Trees	0.862 ± 0.003	0.290 ± 0.015	0.970 ± 0.001
2018	L1 Logistic Regression	0.862 ± 0.004	0.387 ± 0.021	0.961 ± 0.001
2018	Random Forest	0.862 ± 0.002	0.239 ± 0.009	0.973 ± 0.001
2018	Gradient Boosted Trees	0.871 ± 0.004	0.336 ± 0.013	0.967 ± 0.001

The models used correspond to the *general* and *labs* features. The column labeled Period indicates which years of data are used in the temporal validation (2017 only, 2018 only, or both 2017 and 2018). All metrics reported are averaged over the ten models corresponding to the ten random splits of the 2013–2016 data (*i.e.*, no model recalibration was performed), mean ± SD.

Table 6. Temporal Validation Results for Predischarge Prediction at 36 h after Surgery Completion at the University of California Los Angeles Ronald Reagan Medical Center

Period	Method	Area under the Curve	Positive Predictive Value	Negative Predictive Value
2017–2018	L1 Logistic Regression	0.877 ± 0.005	0.309 ± 0.016	0.972 ± 0.002
2017–2018	Random Forest	0.877 ± 0.002	0.230 ± 0.009	0.980 ± 0.001
2017–2018	Gradient Boosted Trees	0.878 ± 0.004	0.283 ± 0.014	0.975 ± 0.001
2017	L1 Logistic Regression	0.873 ± 0.004	0.288 ± 0.017	0.973 ± 0.002
2017	Random Forest	0.875 ± 0.002	0.220 ± 0.010	0.981 ± 0.001
2017	Gradient Boosted Trees	0.875 ± 0.005	0.265 ± 0.017	0.976 ± 0.001
2018	L1 Logistic Regression	0.882 ± 0.006	0.332 ± 0.016	0.971 ± 0.002
2018	Random Forest	0.880 ± 0.002	0.241 ± 0.009	0.979 ± 0.002
2018	Gradient Boosted Trees	0.882 ± 0.004	0.304 ± 0.013	0.974 ± 0.001

The models used correspond to the *general* and *labs* features. The column labeled Period indicates which years of data are used in the temporal validation (either 2017 only, 2018 only, or both 2017 and 2018). The metrics reported are averaged over the ten models corresponding to the ten random splits of the 2013–2016 data (*i.e.*, no model recalibration was performed), mean ± SD.

admission. Although there is a significant body of research that tries to address common risk factors associated with unplanned readmission, our models are unable to offer any insight into its decision making rationale. This is further complicated by the use of a deidentified research dataset. With future integrations with clinical workflows, better identifying patient characteristics (*i.e.*, ZIP code as a surrogate for socioeconomic status) can be used to identify factors that are easier for clinicians to identify in practice.[39]

Second, our electronic health record system only captures readmission to the University of California Los Angeles Health System. Any outside readmission would not be entered into our institution's electronic health record system and thus be invisible to the model. EPIC does have CareEverywhere, a cross-institution platform, and there are local Accountable Care organization payor databases, but these options were beyond the scope of this study.

Third, as discussed in the Material and Methods section, our definition of an unplanned emergency department–based postoperative readmission strays from the actual definition used by the Center for Medicare and

Medicaid Services. The need for the deviation stems from our inability to replicate the work done by the third-party vendors that our institution uses to calculate the reported readmission rate. One would need proprietary algorithms to further evaluate whether a given patient has an illness or underwent a surgical procedure that exempts them from being an infraction. Despite this limitation, we aligned with this standard readmission definition to the extent of our ability.

Finally, our models are scalable to other institutions that have the capability to replicate the necessary general and lab data sets with computer scientists able to implement machine learning techniques. However, this would not be a simple out-of-the box type of implementation; there would be a necessary retraining/recalibration process for each individual site that requires some level of machine learning proficiency. Although this process is not as intensive as the process of gathering the perfect dataset, it demonstrates the hospital-specific nature of hospital readmissions. When implementing the models using the Santa Monica Hospital data for validation, we consumed 70% of the available data

for training. This would imply that any future institutions would need a few thousand patients already extracted to fine-tune the chosen model. After that initial process, there would be optional recalibration to improve performance, but that would not be a mandate.

In summary, this work demonstrates the ability of machine learning techniques to produce clinically useful models for predicting readmissions *via* the emergency department in surgical patients from electronic health record data. We have shown that predictions with high discrimination can occur as soon as 36 h after the completion of surgery, giving time for intervention teams to apply readmission reduction strategies.[1] Although this process has only been implemented using retrospective data, a natural next step for future research is to implement this process in real time, and quantify the benefits of guiding interventions using machine learning.

Research Support

Competing Interests

Correspondence

Address correspondence to Dr. Gabel: Anderson School of Management, University of California Los Angeles, 757 Westwood Plaza, Suite 3325 Los Angeles, California 90095. egabel@mednet.ucla.edu. Information on purchasing reprints may be found at www.anesthesiology.org or on the masthead page at the beginning of this issue. Anesthesiology's articles are made freely accessible to all readers, for personal use only, 6 months from the cover date of the issue.

References

1. Axon RN, Williams MV: Hospital readmission as an accountability measure. JAMA 2011; 305:504–5
2. Marcantonio ER, McKean S, Goldfinger M, Kleefield S, Yurkofsky M, Brennan TA: Factors associated with unplanned hospital readmission among patients 65 years of age and older in a Medicare managed care plan. Am J Med 1999; 107:13–7
3. Frankl SE, Breeling JL, Goldman L: Preventability of emergent hospital readmission. Am J Med 1991; 90:667–74
4. Kansagara D, Englander H, Salanitro A, Kagen D, Theobald C, Freeman M, Kripalani S: Risk prediction models for hospital readmission: A systematic review. JAMA 2011; 306:1688–98
5. Vest JR, Gamm LD, Oxford BA, Gonzalez MI, Slawson KM: Determinants of preventable readmissions in the United States: A systematic review. Implement Sci 2010; 5:88
6. Ross JS, Mulvey GK, Stauffer B, Patolla V, Bernheim SM, Keenan PS, Krumholz HM: Statistical models and patient predictors of readmission for heart failure: A systematic review. Arch Intern Med 2008; 168:1371–86
7. Nguyen OK, Halm EA, Makam AN: Further limitations of the HOSPITAL score in US hospitals. JAMA Intern Med 2016; 176:1232–3
8. Nguyen OK, Makam AN, Clark C, Zhang S, Xie B, Velasco F, Amarasingham R, Halm EA: Predicting all-cause readmissions using electronic health record data from the entire hospitalization: Model development and comparison. J Hosp Med 2016; 11:473–80
9. Walraven C van, Dhalla IA, Bell C, Etchells E, Stiell IG, Zarnke K, Austin PC, Forster AJ: Derivation and validation of an index to predict early death or unplanned readmission after discharge from hospital to the community. CMAJ 2010; 182:551–7
10. Jack BW, Chetty VK, Anthony D, Greenwald JL, Sanchez GM, Johnson AE, Forsythe SR, O'Donnell JK, Paasche-Orlow MK, Manasseh C, Martin S, Culpepper L: A reengineered hospital discharge program to decrease rehospitalization: A randomized trial. Ann Intern Med 2009; 150:178–87
11. Coleman EA, Parry C, Chalmers S, Min SJ: The care transitions intervention: Results of a randomized controlled trial. Arch Intern Med 2006; 166:1822–8
12. Naylor MD, Brooten D, Campbell R, Jacobsen BS, Mezey MD, Pauly MV, Schwartz JS: Comprehensive discharge planning and home follow-up of hospitalized elders: A randomized clinical trial. JAMA 1999; 281:613–20
13. Campbell R: The five rights of clinical decision support: CDS tools helpful for meeting meaningful use. J AHIMA 2013; 84:42–47 (web version updated February 2016)
14. Hastie T, Tibshirani R, Friedman J: The Elements of Statistical Learning. New York, NY: Springer New York, 2009
15. James G, Witten D, Hastie T, Tibshirani R: Statistical Learning. In: Springer, 2013:15–57
16. Elith J, Leathwick JR, Hastie T: A working guide to boosted regression trees. J Anim Ecol 2008; 77:802–13
17. Oh J, Makar M, Fusco C, McCaffrey R, Rao K, Ryan EE, Washer L, West LR, Young VB, Guttag J, Hooper DC, Shenoy ES, Wiens J: A generalizable, data-driven approach to predict daily risk of clostridium difficile infection at two large academic health centers. Infect Control Hosp Epidemiol 2018; 39:425–33
18. Hofer IS, Gabel E, Pfeffer M, Mahbouba M, Mahajan A: A systematic approach to creation of a perioperative data warehouse. Anesth Analg 2016; 122:1880–4

19. Gabel E, Hofer IS, Satou N, Grogan T, Shemin R, Mahajan A, Cannesson M: Creation and validation of an automated algorithm to determine postoperative ventilator requirements after cardiac surgery. Anesth Analg 2017; 124:1423–30

20. Childers CP, Hofer IS, Cheng DS, Maggard-Gibbons M: Evaluating surgeons on intraoperative disposable supply costs: Details matter. J Gastrointest Surg 2019; 23:2054–62

21. Gabel E, Shin J, Hofer I, Grogan T, Ziv K, Hong J, Dhillon A, Moore J, Mahajan A, Cannesson M: Digital quality improvement approach reduces the need for rescue anti-emetics in high-risk patients: A comparative effectiveness study using interrupted time series and propensity score matching analysis. Anesth Analg 2019; 128:867–76

22. Hofer IS, Cheng D, Grogan T, Fujimoto Y, Yamada T, Beck L, Cannesson M, Mahajan A: Automated assessment of existing patient's revised cardiac risk index using algorithmic software. Anesth Analg 2018:1

23. Mehta RL, Kellum JA, Shah SV, Molitoris BA, Ronco C, Warnock DG, Levin A; Acute Kidney Injury Network: Acute Kidney Injury Network: Report of an initiative to improve outcomes in acute kidney injury. Crit Care 2007; 11:R31

24. Breiman L. Random Forests. Mach Learn 2001; 45:5–32

25. Friedman J, Hastie T, Tibshirani R: Additive logistic regression: A statistical view of boosting (With discussion and a rejoinder by the authors). Ann Stat 2000; 28:337–407

26. Tibshirani R: Regression shrinkage and selection via the lasso. J R Stat Soc Ser B 1996:267–88

27. Breiman L, Friedman J, Charles S, Olshen R: Classification and regression trees. Chapman & Hall, 1993

28. Díaz-Uriarte R, Alvarez de Andrés S: Gene selection and classification of microarray data using random forest. BMC Bioinformatics 2006; 7:3

29. Svetnik V, Liaw A, Tong C, Culberson JC, Sheridan RP, Feuston BP: Random forest: A classification and regression tool for compound classification and QSAR modeling. J Chem Inf Comput Sci 2003; 43:1947–58

30. Donzé J, Aujesky D, Williams D, Schnipper JL: Potentially avoidable 30-day hospital readmissions in medical patients: Derivation and validation of a prediction model. JAMA Intern Med 2013; 173:632–8

31. Walraven C van, Wong J, Forster AJ: LACE+ index: extension of a validated index to predict early death or urgent readmission after hospital discharge using administrative data. Open Med 2012; 6:e80–90

32. Horwitz LI, Grady JN, Cohen DB, Lin Z, Volpe M, Ngo CK, Masica AL, Long T, Wang J, Keenan M, Montague J, Suter LG, Ross JS, Drye EE, Krumholz HM, Bernheim SM: Development and validation of an algorithm to identify planned readmissions from claims data. J Hosp Med 2015; 10:670–7

33. Harrell F: Regression Modeling Strategies. 1st ed. New York: Springer New York, 2001

34. Harrell FE Jr, Lee KL, Mark DB: Multivariable prognostic models: Issues in developing models, evaluating assumptions and adequacy, and measuring and reducing errors. Stat Med 1996; 15:361–87

35. Hosmer DW, Lemeshow S, Sturdivant RX: Applied Logistic Regression. 3rd ed. Hoboken, New Jersey: John Wiley & Sons, Inc., 2013

36. Kononenko I: Machine learning for medical diagnosis: History, state of the art and perspective. Artif Intell Med 2001; 23:89–109

37. Waljee AK, Higgins PD: Machine learning in medicine: A primer for physicians. Am J Gastroenterol 2010; 105:1224–6

38. Maglogiannis IG: Emerging artificial intelligence applications in computer engineering: Real word AI systems with applications in eHealth, HCI, information retrieval and pervasive technologies. IOS Press 2007

39. Nagasako EM, Reidhead M, Waterman B, Dunagan WC: Adding socioeconomic data to hospital readmissions calculations may produce more useful results. Health Aff (Millwood) 2014; 33:786–91