

Integration of Feature Vectors from Raw Laboratory, Medication and Procedure Names Improves the Precision and Recall of Models to Predict Postoperative Mortality and Acute Kidney Injury

Ira S. Hofer (✉ ira.hofer@mountsinai.org)

University of California, Los Angeles

Marina Kupina

University of California, Los Angeles

Lori Laddaran

Frank H. Netter MD School of Medicine of Quinnipiac University

Eran Halperin

University of California, Los Angeles

Research Article

Keywords: machine learning, vectors, Featurization technique, perioperative prediction model

Posted Date: January 5th, 2022

DOI: <https://doi.org/10.21203/rs.3.rs-1193437/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Abstract

Introduction:

Manuscripts that have successfully used machine learning (ML) to predict a variety of perioperative outcomes often use only a limited number of features selected by a clinician. We hypothesized that techniques leveraging a broad set of features for patient laboratory results, medications, and the surgical procedure name would improve performance as compared to a more limited set of features chosen by clinicians.

Methods

Feature vectors for laboratory results included 702 features total derived from 39 laboratory tests, medications consisted of a binary flag for 126 commonly used medications, procedure name used the Word2Vec package for create a vector of length 100. Nine models were trained: Baseline Features, one for each of the three types of data Baseline+Each data type (, all features, and then all features with feature reduction algorithm.

Results

Across both outcomes the models that contained all features (model 8) (Mortality ROC-AUC 94.42, PR-AUC 31.0; AKI ROC-AUC 92.47, PR-AUC 76.73) was superior to models with only subsets of features

Conclusion

Featurization techniques leveraging a broad away of clinical data can improve performance of perioperative prediction models.

Introduction

The last several years has seen an explosion in the number of papers using machine learning (ML) techniques to predict a variety of perioperative outcomes. Models have been successfully developed to predict key outcomes such as hypotension^{1,2}, mortality³⁻⁶, readmission⁷, and acute kidney injury (AKI)^{4,8-11}. As a group, these papers have proven the underlying hypothesis that ML techniques can be applied to healthcare data in order to predict outcomes.

Fortunately, negative outcomes in the perioperative setting are relatively rare. For example, the overall incidence of postoperative mortality is roughly 7.7%¹², and the rate of postoperative AKI is roughly 12%¹³. As a result, most models that have been published report relatively high accuracy and areas under the receiver operating characteristic curve (ROC AUC), while having an area under the precision recall curve

(PR AUC) that are lower^{3-5, 14}. This is more than a theoretical limitation. The successful implementation of ML models into clinical practice requires the successful identification of the rare event to the exclusion of other cases. To use a concrete example if a model is to be used to change the decision of surgery, the positive predictive value (precision) must be very high, otherwise large numbers of patients will receive the incorrect treatment.

One limitation of many of the ML models that have been published, especially in the perioperative space, is that they rely on a limited number of features, hand-selected by a clinician to predict the outcome of interest. While this has proven effective, it also serves to leave a large amount of information from the electronic health record (EHR) out of the model, thus potentially limiting the model's predictive power.

In this manuscript we hypothesized that creating feature vectors from a wide variety of data types in the EHR would improve model performance. In particular we hypothesized that creating a feature vector for each of the last six months of laboratory values, patient medications and the type of procedure (as specified in the procedure name) would improve the performance of models trained to predict postoperative mortality and AKI as compared to a baseline feature set. As a primary outcome we look at the area under the Receiver Operating Characteristic (ROC) and Precision-Recall (PR) curves and as a secondary outcome we look at accuracy, precision and recall and various thresholds.

Methods

Data were extracted from the Perioperative Data Warehouse (PDW), a custom-built data warehouse containing all patients who have undergone surgery at the University of California Los Angeles (UCLA) Health since the implementation of our EHR (EPIC Systems, Madison, WI) in March 2013. The PDW has been previously described^{15, 16}. Briefly, in the first stage, data are extracted from EPIC's Clarity database into 29 tables organized to facilitate usage. In the second stage, these data are used to populate a series of measures and metrics such as procedure duration, readmissions, admission International Classification of Diseases (ICD) codes, and postoperative outcomes¹⁷⁻¹⁹. All data used for this study were obtained from this data warehouse, and institutional review board approval (18-002053) was obtained from the UCLA Office of the Human Research Protection Program, including exemption from written informed consent, for this retrospective review. All methods were performed in accordance with relevant guidelines and regulations.

Determination of outcome variables – mortality and AKI – from EHR data

Postoperative mortality was defined as death during the same hospitalization as the surgery as identified by either 1) a death date is noted in the EHR between hospital admission and discharge or 2) a postoperative discharge status of expired, presence of a "death note" by a treating provider, and the lack of future admissions or encounters for the patient in the EHR. Postoperative AKI was defined based upon the Acute Kidney Injury Network (AKIN) criteria creatinine criteria²⁰. The baseline creatinine was taken to be the most recent serum creatinine prior to surgery. The postoperative creatinine was the highest

creatinine value within 48 hours of surgery. If either the preoperative or postoperative creatinine was missing, the value was set to NULL. A value of 0 was used to denote no AKI while any AKIN stage 1 or above was denoted as a 1 (i.e. this was set to a binary variable).

Inclusion and Exclusion Criteria

The inclusion data set was surgeries with general anesthesia that occurred between April 1, 2013 and July 2021 that were performed at the Ronald Reagan UCLA Medical Center and Santa Monica Medical Center hospitals. Patients were filtered by their class, selecting only inpatient, same day hospitalization, emergency care and overnight recovery patients (i.e. those that spent at least one night in the hospital). Patients aged less than 18 years old or older than 89 years old were excluded from the dataset due to the institutional restrictions on data security. Cases were excluded if they had an ASA physical status score of 6 (indicating organ donors).

Model Input Features

In this manuscript, four sets of input features were defined depending on their characteristics: 1) baseline features including basic patient information and surgery specifications, 2) the most recent laboratory tests obtained before the surgery, 3) procedure description, and 4) medications taken.

Baseline Feature Vector

The features used in the baseline model were based on previous work by our group predicting postoperative mortality both before and after surgery³⁻⁵. For the purposes of this analysis, we removed those features that would be redundant with the additional feature vectors (see below). For example, lab results were removed from the feature sets used in those models because we created a separate, more comprehensive feature vector, comprised only of the labs. The list of features included in this group is in Supplementary table 1A.

Laboratory Result Feature Vector

A set of 39 commonly used laboratory results were extracted from the EHR (see Supplementary Table 2A for a complete list). These results were chosen because they are common preoperative tests (i.e. included in complete blood count, comprehensive metabolic panels or coagulation panels). Test all results were then binned in 6 months before surgery, 3 before months, and 1 month prior to surgery. Then, for each laboratory test bin the following descriptive statistics were calculated: total number of tests, cu-sum, median, variance, mean, standard score (z-score). The goal of the cu-sum was to incorporate a measure of temporal change into the descriptive statistics. The standard score (or z-score) is the number of standard deviations at which the mean of a patient's test results for each laboratory is higher or lower than the mean of that test for all patients. Thus, for each surgery we had a vector of 702 features (39 laboratory results, 3 bins, 6 descriptive statistics per bin).

Procedure Name Embedding

As part of our research, we experimented with the inclusion of clinical text data in the form of procedure names as the model inputs. Administrative codes, such as CPT codes are only available after surgery (patient discharge), thus we focused on representing the procedure name using a numerical vector available before surgery. The procedure name, as booked by the surgeon, consists of a string with a variable number of words. The number of unique words contained in all procedures names was 22,003 in the training dataset. In order to include the procedure name in the prediction model, we applied word embedding; a common method for representing words, typically in the form of a real-valued vector that encodes the meaning of the word such that the words that are closer in the vector space are expected to be similar in meaning²¹. We trained word embeddings on the clinical text data to allow the model to understand a clinical context using class Word2Vec from Gensim library²². In order to train the model, each procedure name is broken into words (tokens). The Word2Vec model takes a list of tokens for each sentence as input and returns a set of numeric vectors as output using a two-layer neural network. The size of the numeric vector is set as a parameter of the model; in this paper, it set to 100. The trained model was applied for each procedure name, returning a set of numeric vectors for each word. Since the procedure name contains multiple words, the length of the procedure name is variable. Thus, in order to get a single vector representation of the procedure name, we calculated the average vector of its words, which is the input to the main predictive model.

Medications

The last type of feature we analyzed was the medications taken by a patient before surgery. A set of high use and likely clinically predictive medications was made based on clinical judgement. Medications that were taken 24 hours before surgery were excluded. Medications were broken up into two different categories: given as an inpatient and taken at home. Combined medications (e.g. HYDROCODONE-ACETAMINOPHEN 5-325 MG PO TABS) were separated into single medications, and each prescription dose was calculated with normalized units (e.g. Hydrocodone 5 mg, Acetaminophen 325 mg). For each patient we created a binary vector, where each element of the vector indicated if the specific medication was taken by a patient. Then for each medication the Fisher's Exact Test was used to determine a significant association between medication and target variable. Medications with a p-value of less than 0.05 were included in the model. The final vector contained a set of 126 medications (see Supplementary Table 3A for a complete list).

Data preprocessing

Data were split training and testing datasets with the split ratio 75:25. To avoid information leakage, all patients that appeared in the test set were removed from the training set.

Categorical variables such as ethnicity, gender, etc in clinical data were converted to a numerical representation by applying the One-Hot-Encoding algorithm that decoded each category in a binary vector. To optimize the memory usage, a memory reduction function was implemented that validated the feasibility of the data type modification. On the training dataset, this technique decreased the dataset

used memory by almost 68% from 740 MB to 237MB, which significantly accelerated the model training process.

Model creation, training, and testing

Nine models were trained. Model 1 is the baseline model that includes only the basic features in Supplementary Table 1a. Models 2, 3 and 4 were trained only on laboratory results, procedure name embeddings, and medication features respectively. The main goal of training these models was to measure the individual contributions to the prediction model. Models 5, 6, and 7 were extensions of Models 2,3, and 4 with added features from the baseline model to the features of the respective models (i.e. baseline + laboratory, baseline + medications, etc.). Finally, Model 8 was obtained by combining features from the previous models - 1088 features in total. To simplify Model 8 we created Model 9 that included only features that were selected by the feature selection algorithm Boruta^[1].

All models were gradient boosted trees as previous work by our group has shown this technique to consistently perform well^{3, 5}. The main advantage of gradient boosted trees, as opposed to other tree based models, is that trees are created sequentially which reduces the residual error of previous trees.

The gradient boosted tree classifiers were implemented using the XGBoost package. Model hyperparameters were selected using five-fold cross-validation on the training dataset, where patients undergoing multiple surgeries appeared only in the training or testing set, but not simultaneously in both. In five-fold cross-validation, the dataset is divided into five partitions; four-fifths of the data is used to train the models and the remaining one-fifth is used as the testing set. This process is repeated so that each partition is used as a testing set only once and a training set four times. Cross-validation provides a better assessment of model performance by averaging metrics across multiple tests. The models were trained using 2000 estimators and a maximum tree depth of 7 and minimum number of 7 children.

Model performance

Prediction of both mortality and AKI were treated as binary classification problems with highly imbalanced classes. The issue of class imbalance has significant implications for metrics of model performance.

Receiver operating characteristic (ROC) curves are widely used for the estimation of predictive model performance with a binary outcome. ROC curves characterize the trade-off between true positive and false positive rates for the binary classification model by varying the discriminative threshold. However, the false positive rate is affected by the underlying rate of the event and can be deceptive for data with a large skew in the class distribution, thus making ROC curves overly optimistic.

Thus, in addition to ROC curves we considered precision-recall (PR) curves, which summarize the trade-off between true positive rate and positive predictive value by changing the prediction threshold. Positive predictive value (or precision) penalizes a model for a large number of false positives relative to the

number of true positives that makes PR curves robust even under imbalanced data. Simultaneously recall, instead of focusing on the number of predicted false positives, penalizes a model for a large number of false negative. The penalties in precision and recall are opposites, making this curve a better metric for model performance with imbalanced data.

Lastly, the F-beta score is a useful metric that calculates the weighted harmonic mean of precision and recall, reaching its optimal value at 1 and its worst value at 0. The beta parameter determines the weight of recall in the combined score^[ii] and allows data scientists to pick a threshold that optimizes the implementation tradeoffs between precision and recall. Numbers greater than 1 will give increased weight to the recall, while those less than one will give increased weight to the precision.

To examine the performance of the different models we examined model performance with beta values of 1, 2, and 3 were chosen. The higher beta values gave more weight to the recall and penalized the number of false negatives. The F-beta score is a threshold metric; thus, scores were calculated for each model based on the selected discriminative threshold that maximizes the F-score. One of the limitations of these metrics is that they assume the distribution of the classes observed in the training dataset will match the distribution in the test set and in real data when the model is used to make predictions.

[i] https://github.com/scikit-learn-contrib/boruta_py

[ii] https://scikitlearn.org/stable/modules/generated/sklearn.metrics.fbeta_score.html

Results

Overall the dataset identified 101,070 surgeries across 93,335 admission and 79,662 patients. Patient ages ranged between 18 and 89 years with a mean age of almost 56. Overall, the rate of mortality in the dataset was 2.29% and the rate of AKI was 15.8%. ASA Physical status of 3 was the most common physical status and the most common surgical specialty was General Surgery. Detailed information of patient demographics for both the mortality and AKI models is show in Table 1.

Receiver Operating Characteristic and Precisions-Recall Performance of the Models

Table 2 and Figure 1 demonstrate the performance of the models. Overall, across both outcomes the models that contained all features (model 8) (Mortality ROC-AUC 94.42, PR-AUC 31.0; AKI ROC-AUC 92.47, PR-AUC 76.73) was superior to models with only subsets of features and model 9, which contained all features with feature selection (35% of features from model 8), performed nearly as well as model 8.

For the models with the individual feature sets (models 1-4), the baseline feature set, those features selected by clinicians, performed the best with a ROC-AUC of 92.08 for mortality and 90.51 for AKI, and a PR-AUC of 22.89 for mortality and 71.76 for AKI. The next best models for both morality and AKI was laboratory results, followed by medications and lastly procedure name. In general ROC-AUC tended to be

higher for mortality and the PR-AUC was higher for AKI. Figures 1 shows the ROC and PR curves for these models.

In looking at the combined feature sets there were fewer clear themes. With regard to mortality the addition of the laboratory results improved both ROC and PR AUC while the addition of procedure names actually lowered both ROC and PR AUCs. With regard to AKI, the addition of all feature types increased both ROC and PR AUC.

Secondary Performance Metrics – Accuracy, Precision and Recall and various Thresholds

Tables 3A and 3B show the overall F1, F2 and F3 scores as well as the accuracy, precision and recall at the various thresholds. There were fewer clear trends across the different models with much greater variability in the performance from one model to the next. Of note, Model 3 (procedure name) achieved perfect recall for both outcomes though with very low precision. For the AKI outcome some of the models achieved very high recall (above 90%) with overall accuracies above 75%. Further, Models 1&3 for mortality and Model 3 for AKI showed identical performance for the F2 and F3 scores indicating plateau effects.

Discussion

In this manuscript we successfully created feature vectors from a variety of clinical data types (medications, laboratory results, and surgical procedure name) and were able to leverage these additional features to improve the performance of models to predict postoperative AKI and mortality as compared to a baseline feature set chosen by clinician judgment. Of particular note, some features (such as procedure name) which did not necessarily perform well on their own, still enhanced the performance of the models when combined with other features. Additionally, while there were some trends in model performance, the additional benefit of a particular feature often differed between the two outcomes.

We believe that these results lead to some important conclusions. Firstly, these results demonstrate that the inclusion of more information from the electronic health record has the ability to improve model performance. Crucially, for the outcomes in this paper, the improvement in performance was not linear – i.e. more features were not always better, and features that performed poorly in one model performed better when combined with different features in another model. We believe these results point to the need for more research into a variety of modeling techniques. While this manuscript used gradient boosted trees, the performance of other techniques, such as neural networks, might be better or at least different.

Of note, there is far more information in the EHR than what we attempted to include in this manuscript. For example, we focused on a set of commonly used laboratory tests and medications as opposed to all medications and tests. Further, the medications were turned into a binary vector (thus ignoring the dose) and the laboratory results were summarized with basic descriptive statistics. Additional work examining

things like the use of medication dose, time series techniques, and/or ratios of different laboratory results may yield even more powerful results. Additionally, use of natural language processing techniques on clinician notes, image processing of radiographic results and other types of data may further improve model performance.

Lastly, the accuracy, precision and recall at the various thresholds studies (F1, F2 and F3 score) demonstrate that there is no one perfect model. For example, the use of procedure name had relatively poor AUC results but achieved perfect recall for both outcomes, indicating its potential use in workflows where that is a key performance metrics. Thus, while summary performance metrics, such as AUC, are useful for global purposes the actual workflow and clinical tradeoffs must be evaluated when picking a specific model for clinical care. Critically, this may require the evaluation of different types of features, or modeling techniques in addition to determining the proper threshold.

This study does have some limitations. Most significantly, this is a single center trial examining a single type of model (gradient boosted trees) and two clinical outcomes – postoperative mortality and AKI. While the conclusions we draw from these results are likely applicable to other hospitals and may generalize to other modeling techniques or outcomes, this cannot be known for certain. Additionally, as noted, this study only used certain medications and laboratory results. While we believe that we have identified those most common and most likely to influence outcomes, it is possible that the results would have been different with a different medications/tests.

Overall, we believe this study adds to the body of work that demonstrates the need for more research into techniques to improve healthcare model performance. There is likely no “magic bullet” of a perfect model that always performs the best. Rather what is needed is a variety of techniques (featurization, modeling, etc.) that can be called upon for a specific clinical task to find the optimal model for that workflow. What does seem to be clear from this manuscript, is that having ways to access more data is probably better and simply relying on a small set of features thought to be clinically relevant is unlikely to create the best performing model.

Declarations

Data Availability

Due to institutional restriction regarding protected health information and patient privacy the data used in this manuscript are not publicly available. Researchers interested in obtaining the data are encouraged to contact the corresponding author who will do his best to facilitate access subject to institutional guidelines.

References

1. Hatib F, Jian Z, Buddi S, et al. Machine-learning Algorithm to Predict Hypotension Based on High-fidelity Arterial Pressure Waveform Analysis. *Anesthesiology*. Oct 2018;129(4):663-674.

doi:10.1097/ALN.0000000000002300

2. Kendale S, Kulkarni P, Rosenberg AD, Wang J. Supervised Machine-learning Predictive Analytics for Prediction of Postinduction Hypotension. *Anesthesiology*. Oct 2018;129(4):675-688. doi:10.1097/ALN.0000000000002374
3. Hill BL, Brown R, Gabel E, et al. An automated machine learning-based model predicts postoperative mortality using readily-extractable preoperative electronic health record data. *Br J Anaesth*. Dec 2019;123(6):877-886. doi:10.1016/j.bja.2019.07.030
4. Hofer IS, Lee C, Gabel E, Baldi P, Cannesson M. Development and validation of a deep neural network model to predict postoperative mortality, acute kidney injury, and reintubation using a single feature set. *NPJ Digit Med*. 2020;3:58. doi:10.1038/s41746-020-0248-0
5. Lee CK, Hofer I, Gabel E, Baldi P, Cannesson M. Development and Validation of a Deep Neural Network Model for Prediction of Postoperative In-hospital Mortality. *Anesthesiology*. Oct 2018;129(4):649-662. doi:10.1097/ALN.0000000000002186
6. Lee CK, Samad M, Hofer I, Cannesson M, Baldi P. Development and validation of an interpretable neural network for prediction of postoperative in-hospital mortality. *NPJ Digit Med*. Jan 8 2021;4(1):8. doi:10.1038/s41746-020-00377-1
7. Misic VV, Gabel E, Hofer I, Rajaram K, Mahajan A. Machine Learning Prediction of Postoperative Emergency Department Hospital Readmission. *Anesthesiology*. May 2020;132(5):968-980. doi:10.1097/ALN.0000000000003140
8. Lee HC, Yoon HK, Nam K, et al. Derivation and Validation of Machine Learning Approaches to Predict Acute Kidney Injury after Cardiac Surgery. *J Clin Med*. Oct 3 2018;7(10)doi:10.3390/jcm7100322
9. Lee HC, Yoon SB, Yang SM, et al. Prediction of Acute Kidney Injury after Liver Transplantation: Machine Learning Approaches vs. Logistic Regression Model. *J Clin Med*. Nov 8 2018;7(11)doi:10.3390/jcm7110428
10. Xue B, Li D, Lu C, et al. Use of Machine Learning to Develop and Evaluate Models Using Preoperative and Intraoperative Data to Identify Risks of Postoperative Complications. *JAMA Netw Open*. Mar 1 2021;4(3):e212240. doi:10.1001/jamanetworkopen.2021.2240
11. Zhang Z, Ho KM, Hong Y. Machine learning for the prediction of volume responsiveness in patients with oliguric acute kidney injury in critical care. *Crit Care*. Apr 8 2019;23(1):112. doi:10.1186/s13054-019-2411-z
12. Nepogodiev D, Martin J, Biccard B, Makupe A, Bhangu A, National Institute for Health Research Global Health Research Unit on Global S. Global burden of postoperative death. *Lancet*. Feb 2 2019;393(10170):401. doi:10.1016/S0140-6736(18)33139-8
13. Grams ME, Sang Y, Coresh J, et al. Acute Kidney Injury After Major Surgery: A Retrospective Analysis of Veterans Health Administration Data. *Am J Kidney Dis*. Jun 2016;67(6):872-80. doi:10.1053/j.ajkd.2015.07.022
14. Lei VJ, Luong T, Shan E, et al. Risk Stratification for Postoperative Acute Kidney Injury in Major Noncardiac Surgery Using Preoperative and Intraoperative Data. *JAMA Netw Open*. Dec 2

2019;2(12):e1916921. doi:10.1001/jamanetworkopen.2019.16921

15. Epstein RH, Hofer IS, Salari V, Gabel E. Successful Implementation of a Perioperative Data Warehouse Using Another Hospital's Published Specification From Epic's Electronic Health Record System. *Anesth Analg*. Apr 22 2020;doi:10.1213/ANE.0000000000004806
16. Hofer IS, Gabel E, Pfeiffer M, Mahbouba M, Mahajan A. A Systematic Approach to Creation of a Perioperative Data Warehouse. *Anesth Analg*. Jun 2016;122(6):1880-4. doi:10.1213/ANE.0000000000001201
17. Gabel E, Hofer IS, Satou N, et al. Creation and Validation of an Automated Algorithm to Determine Postoperative Ventilator Requirements After Cardiac Surgery. Comparative Study. *Anesth Analg*. May 2017;124(5):1423-1430. doi:10.1213/ANE.0000000000001997
18. Hofer IS, Cheng D, Grogan T. A Retrospective Analysis Demonstrates That a Failure to Document Key Comorbid Diseases in the Anesthesia Preoperative Evaluation Associates With Increased Length of Stay and Mortality. *Anesth Analg*. Jan 15 2021;doi:10.1213/ANE.0000000000005393
19. Hofer IS, Cheng D, Grogan T, et al. Automated Assessment of Existing Patient's Revised Cardiac Risk Index Using Algorithmic Software. *Anesth Analg*. May 2019;128(5):909-916. doi:10.1213/ANE.0000000000003440
20. Mehta RL, Kellum JA, Shah SV, et al. Acute Kidney Injury Network: report of an initiative to improve outcomes in acute kidney injury. OriginalPaper. *Crit Care*. Apr 01 2007;11(2):R31. doi:10.1186/cc5713
21. Jurafsky D, James M. *Speech and language processing : an introduction to natural language processing, computational linguistics, and speech recognition*. Prentice Hall; 2000.
22. Srinivasa-Desikan B. *Natural Language Processing and Computational Linguistics: A practical guide to text analysis with Python, Gensim, spaCy, and Keras*. Packt Publishing; 2018.

Tables

Due to technical limitations, tables are only available as a download in the Supplemental Files section.

Figures

Figure 1

Demonstrate the performance of the models. Overall, across both outcomes the models that contained all features (model 8) (Mortality ROC-AUC 94.42, PR-AUC 31.0; AKI ROC-AUC 92.47, PR-AUC 76.73) was superior to models with only subsets of features and model 9, which contained all features with feature selection (35% of features from model 8), performed nearly as well as model 8.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [STable1.pdf](#)
- [STable2.pdf](#)
- [STable3.pdf](#)
- [Table1.pdf](#)
- [Table2.pdf](#)
- [Table3A.pdf](#)
- [Table3B.pdf](#)