The American Journal of Surgery xxx (2017) 1–6

Contents lists available at ScienceDirect

The American Journal of Surgery

journal homepage: www.americanjournalofsurgery.com

Seeing the forest beyond the trees: Predicting survival in burn patients with machine learning

Adrienne N. Cobb [a, b], Witawat Daungjaiboon [c], Sarah A. Brownlee [b], Anthony J. Baldea [a], Arthur P. Sanford [a], Michael M. Mosier [a], Paul C. Kuo [a, b, *]

[a] Loyola University Medical Center, Department of Surgery, 2160 S. 1st Avenue, Maywood, IL 60153, USA
[b] One:MAP Section of Surgical Analytics, Department of Surgery, Loyola University Chicago, 2160 S. 1st Avenue, Maywood, IL 60153, USA
[c] DePaul University, College of Computing and Digital Media, Department of Predictive Analytics, 243 South Wabash Avenue, Chicago, IL 60604, USA

A R T I C L E I N F O

Article history:
Received 8 July 2017
Received in revised form
4 October 2017
Accepted 5 October 2017

This study was presented as an oral presentation during the Midwest Surgical Association Annual Meeting, Chicago, Illinois July 30-August 1, 2017.

Keywords:
Burns
Survival
Machine learning
Random forest
Outcomes

A B S T R A C T

Background: This study aims to identify predictors of survival for burn patients at the patient and hospital level using machine learning techniques.
Methods: The HCUP SID for California, Florida and New York were used to identify patients admitted with a burn diagnosis and merged with hospital data from the AHA Annual Survey. Random forest and stochastic gradient boosting (SGB) were used to identify predictors of survival at the patient and hospital level from the top performing model.
Results: We analyzed 31,350 patients from 670 hospitals. SGB (AUC 0.93) and random forest (AUC 0.82) best identified patient factors such as age and absence of renal failure ($p < 0.001$) and hospital factors such as full time residents ($p < 0.001$) and nurses ($p = 0.004$) to be associated with increased survival.
Conclusions: Patient and hospital factors are predictive of survival in burn patients. It is difficult to control patient factors, but hospital factors can inform decisions about where burn patients should be treated.

1. Introduction

Burn patients require complex care involving a delicate balance among resuscitation, stabilization and rehabilitation. Their injuries can range from superficial burns only requiring local wound care to more severe burns that may require surgery and are potentially complicated by respiratory failure or sepsis. The predictors of outcome in burn patients are well established, correlating well with the Baux index of age and percent surface area burned as well as the revised Baux index which also takes into account inhalation injury.[1–4] Therefore, multiple factors must be considered when caring for these patients. Traditionally, the focus has been on preventing mortality by controlling or treating as many patient specific or disease specific factors as possible. Scoring tools such as the Baux Score or the Abbreviated Burn Severity Index (ABSI) have been designed to predict morbidity and mortality based on total body surface area (TBSA) age, sex and the presence of inhalation injury.[5] While patient characteristics are key, there are other factors that contribute to patient outcomes.

Studies suggest that system characteristics within institutions such as staffing and technology can impact the ability of hospitals to provide optimal care for patients.[6,7] Since many studies examining the prediction of survival in burn patients are completed in a single center, there is little data surrounding what systems characteristics may have contributed to the survival of burn patients. The use of a large all payer, administrative database linked to hospital level data provides a different perspective. Additionally, the use of machine learning allows us to uncover patterns or associations not otherwise elucidated with traditional linear statistical techniques. A study utilizing artificial neural networks to predict survival in burn patients found non-linear techniques better suited to address complex questions regarding prognosis due to

* Corresponding author. Department of Surgery, Loyola University Medical Center, 3rd Floor EMS Building, Health Sciences Campus, 2160 S. First Avenue, Maywood, IL 60153, USA.
E-mail addresses: adcobb@lumc.edu (A.N. Cobb), dwitawat@gmail.com (W. Daungjaiboon), sbrownlee@luc.edu (S.A. Brownlee), abaldea@lumc.edu (A.J. Baldea), asanford@lumc.edu (A.P. Sanford), mmosier@lumc.edu (M.M. Mosier), paul.kuo@luhs.org (P.C. Kuo).

https://doi.org/10.1016/j.amjsurg.2017.10.027
0002-9610/

their ability to “observe” the real events or facts then evaluate the relative influence of variables on each other and the whole case.[8]

The present study addresses two understudied topics: factors that predict the survival of burn patients beyond traditional burn specifications and the contribution of systems infrastructure by examining hospital characteristics that predict survival. We identified a heterogeneous group of burn patients, created various models to predict survival based on patient and hospital characteristics, and chose the model that performed best with the goal of informing clinical decision-making.

2. Methods

2.1. Data source and patient selection

The Healthcare Cost and Utilization Project (HCUP) State Inpatient Database (SID) for California 2006–2011, Florida 2009–2013 and New York 2009–2013 was used to identify adult patients admitted to the hospital with a burn diagnosis. Burns of varying severity and location were included. The SID is an administrative, all-payer data set aggregated by the Agency for Healthcare Research and Quality (AHRQ) to inform health related decisions.[9] The diagnosis codes were identified by *International Classification of Disease, 9th Revision, Clinical Modification*, (ICD9) codes for burn injury (941.20–59, 942.20–35, 942.39–45, 942.49–59, 943.20–26, 943.29–36, 943.39–46, 943.49–56, 943.59, 944.20–28, 944.30–38, 944.40–48, 944.50–58, 945.20–26, 945.29–36, 945.39–46, 945.49–56, 945.59, 946.2–5, 947.0–4, 947.8–9). These data were then merged with the 2011 American Hospital Association (AHA) Annual Survey to provide hospital level data associated with the selected burn patients. This nationwide database contains information categorizing an institution's organizational structure, facility and service lines, operation expenses, and staffing.[10] The Institutional Review Board at our institution deemed the study exempt from review as the data are de-identified, protected and publically available.

2.2. Data pre-processing

Patient level data were pre-processed to provide uniform variable formats across states using the dplyr package in R. It provides a flexible grammar of data manipulation and focuses on tools for working with data frames. Variables selected for analysis included various comorbidities, age, mortality, hospital state, insurance type, procedure codes, race, admission type, and median income quartile. Several variables were generated in the hospital level AHA dataset including Joint Commission designation, Commission on Cancer, Council of Teaching Hospitals, Level 1 Trauma Center, Nurse to Bed Ratio, surgical volume, GI intensity, radiology intensity, and ICU beds.

Missing values for both groups were replaced with column means for numerical variables and Random forest algorithm was used to impute categorical variables using the caret package.

Data were split into training (66%) and test (34%) sets. Our target variable DIED (mortality = 1, survival = 0) was extremely imbalanced in both datasets. To avoid not detecting the minority class, we used Synthetic Minority Over-sampling Technique (SMOTE) to balance both groups by up-sampling the minorities (DIED = 1) and down sampling majorities (DIED = 0).

2.3. Statistical analysis

Descriptive statistics of the study population were calculated using arithmetic means with standard deviations or median with interquartile range for continuous variables and proportions for categorical variables. Population unadjusted mortality was obtained using a simple proportion of number of inpatient mortalities by the total population. Age categories by 7-year intervals were created and plotted against mortality rates as seen in Fig. 1. Additionally, mortality rates were calculated by age range and burn type (Table 2) as seen in Fig. 2.

Multiple models were built to determine the model best suited to predict variables that impact survival in burn patients. Accuracy, sensitivity and specificity were used to evaluate the models for completeness. We also used receiver operator curve (ROC) for model comparison. The ROC demonstrates how well models separate both classes while the area under the curve (AUC) can be interpreted as the accuracy of the models. The AUC ranges from poor class separation at 0.51 to perfect class separation at 1.

We used tree-based ensemble models, such as stochastic gradient boosting (SGB) and random forest (RF), as we are able to use a variable importance measure to determine those factors that affect patient survival. The variable importance measure was used to indicate how well each variable split our target class. The stochastic gradient boosting algorithm was run with 650 trees, an interaction depth of 9, shrinkage of 0.1 and a minimum of ten observations per node. Random forest models were run with a weight class of 1:3. All analysis were completed using various packages in RStudio including randomForest, party, caret, and pROC and as mentioned above.

Gradient Boosting Machine (GBM) uses a boosting method to build ensemble trees by iteratively adding a weak classifier one at a time (in this case tree stump). In each iteration, a new tree tries to correct errors in the model from the previous iteration. New trees are added until we reach the goal of prediction. Unlike the GBM, the random forest algorithm uses bootstrap aggregation, also known as bagging, to construct a model by creating trees from sampling data from a training set with replacement and subsequently combines the trees together.

3. Results

We analyzed 31,350 patients from 670 hospitals across the three states included. The mean patient age was 40.5 years. The study population was largely male, Caucasian, and had Medicaid insurance. Hypertension was the most common comorbidity with 24.6% of the study population affected. Baseline patient and hospital characteristics are summarized in Table 1. The age distribution varied in the study population. The overall mortality rate was 2.86%, and the rate of mortality increased with age (Fig. 1). Patients with burns from multiple specified sites (ICD-9946) had the highest mortality rate among burn patients, while unspecified sites and burns to the eyes carried no mortality (Fig. 2).

At the patient level, we found that the stochastic gradient boosting model performed slightly better than the random forest algorithm with an area under the ROC curve of 0.93 and 0.90 respectively. Although, the weighted random forest model tended to pick up the minority class better than the SGB as evidenced by its superior specificity value (0.74 v. 0.71). The top five patient characteristics, as evidenced by their variable importance score, predicting survival in burn patients according to the SBG model were younger age, absence of electrolyte imbalance or coagulopathy, admission on a weekend, and absence of renal failure (Fig. 3). The top five patient characteristics that predicted survival in the random forest model were absence of electrolyte imbalance or coagulopathy, younger age, absence of congestive heart failure, and presence of weight loss. All were predicted with $p < 0.001$.

At the hospital level, the random forest algorithm far outperformed SGB with an AUC of 0.82 compared to AUC 0.62. With a specificity of 0.61, the random forest model is able to predict

Please cite this article in press as: Cobb AN, et al., Seeing the forest beyond the trees: Predicting survival in burn patients with machine learning, The American Journal of Surgery (2017), https://doi.org/10.1016/j.amjsurg.2017.10.027