24

Trauma Outcome Prediction in the Era of Big Data: From Data Collection to Analytics

Shiming Yang, Peter F. Hu, and Colin F. Mackenzie

CONTENTS

24.1 Introduction

Traumatic brain injury (TBI) and hemorrhage shock are the leading causes of morbidity and mortality after injury both on the battlefield and in civilian care [1,2]. In parallel with new clinical care protocol and technique development, increasing efforts in continuous physiological monitoring and data analysis are reported [3,4]. The ultimate goal of Big Data's application in this area is to optimize the limited medical resource allocation through comprehensive analysis of a large amount of data, hence improving patients' outcomes and reducing healthcare cost.

Early recognition and mitigation of secondary injury and hemorrhage could ameliorate the effect of brain injury or prevent death caused by massive bleeding. Various statistical modeling and machine learning methods have been used to explain TBI outcome associations and to predict future conditions with massive data sets. Research and funding agencies, including National Institute of Neurological Disorders and Stroke (NINDS) and the U.S. Department of Defense (DoD), are also investing more resources in developing reliable TBI outcome and transfusion prediction models and practically usable tools based on intensive analysis of large data collections. Military medicine considers these approaches as the future way to develop combat casualty autonomous resuscitation [5,6] and enhance real-time field decision making [7].

The volume of real-time physiological patient data has proliferated with each advance in computer hardware and medical sensor technology. High-fidelity data are streamed into physiological monitors for care planning, clinical decision support, quality improvement, and remote patient monitoring. Processing and extracting useful and actionable knowledge from these data also require consideration of the techniques used to store, manage, and analyze massive data. In this chapter,

we will review continuous physiological data collection, signal processing, and data analysis methods for predicting trauma patient outcomes related to actionable therapeutic interventions and translating this into autonomous resuscitation. In Section 24.2, we describe a reliable large-scale physiological data collection system, which is the basis of any Big Data study. In Section 24.3, we review some medical signal processing and feature extraction methods for signals typically used in predicting trauma patients' short- and long-term outcomes. In Section 24.4, we provide a description of a comprehensive framework for a machine learning approach to TBI outcomes and blood transfusion prediction, with example projects for the purpose of early detection and automated resuscitation.

24.2 Automated Medical Big Data Monitoring

Modern hospitals are often equipped with bedside monitors collecting various physiological data in a real-time, continuous, and automated way. Data ranging from routine intermittent observations to high-fidelity waveforms can be recorded and streamed into monitors for care planning, clinical decision support [8], quality improvement [9], and reduced hospital mortality [10]. With massive storage capability, those data can also be stored as part of the electronic health records (EHRs) for retrospective analyses such as physiological pattern discovering [11,12] and prediction modeling [13,14]. One example is the PhysioBank, a large collection of biomedical databases, which inspires studies in cardiovascular time series dynamics, modeling intracranial pressure (ICP) for noninvasive estimation, and more [15–18]. However, in a busy resuscitation or healthcare environment, collecting more complete data is not the top priority of health providers. System failure or manual data entry errors can result in missing values, causing difficulty in application of decision-support algorithms; or such failures can cause a loss of data associated with rare events. Therefore, a reliable system that simplifies and automates the hospital-wide data collection process is necessary.

24.2.1 Data from a Real Trauma Center

In the R Adams Cowley Shock Trauma Center, a level I regional trauma center located in downtown Baltimore, Maryland, 94 GE Marquette Solar 7000/8000® (General Electric, Fairfield, CT) patient vital sign (VS) monitors are networked to provide collection of real-time patient VS data streams in 13 trauma resuscitation units (TRU) and 9 operating room (OR), 12 postanesthesia care unit (PACU), and 60 intensive care unit (ICU) individual bed/monitor units. Each patient monitor collects real-time 240 Hz waveforms and 0.5 Hz trend data, which are transferred via secure intranet to a dedicated BedMaster® server (Excel Medical Electronics, Jupiter, FL) and archived [19]. This process generates approximately 20 million data points/day/bed or roughly 30 terabits/year of data. Physiological data collected through this system, when they are displayed on the GE Marquette monitor, include electrocardiographic (ECG), photoplethysmographic (PPG), carbon dioxide (CO_2), arterial blood pressure (ABP), and ICP data, among others. Trends include heart rate (HR), respiratory rate (RR), temperature, oxygen saturation (SPO_2), end-tidal CO_2 ($EtCO_2$), and ICP, among many others. They cover the categories of brain pressure, cardiac, perfusion, and respiratory.

In addition to continuous data, other heterogeneous data with different formats and temporal resolutions are also collected and organized in databases, including ordinal or categorical data [e.g., Glasgow Coma Scale (GCS), age, sex]; radiological images; text (medical records, clinical notes); bed movement (admission, discharge and transfer time, and location); and other important data (adverse events, treatments, response, etc.).

24.2.2 Redundant VS Collection System with a Diagnosis Viewer

These heterogeneous data are available and need to be collected with different formats and temporal resolutions. In the current hospital environment, these data are in a loosely organized decentralized network. One approach to manage VS waveforms and trend data, used at a level I trauma center that admits more than 8000 severely injured patients annually, is to design a triple-redundant data collection server for high fault tolerance to maximize these data collection rates to nearly 100%.

To minimize the impact of individual server collection failure, we installed three dedicated BedMaster servers in parallel to simultaneously collect physiological patient data from the network of patient monitors described above. Figure 24.1 diagrams the data streams from multiple individual bed units to three BedMaster servers arranged in parallel. This triple modular redundancy architecture permits fast switch over time and high system availability [20]. One server is selected as a principal or "backbone" server. When it fails, values from a second sever will fill in. When two servers fail, values from the third one will be used.

The triple-redundant data collection system could increase data availability. However, a tool for fast system diagnosis is still lacking. To address the need for ongoing

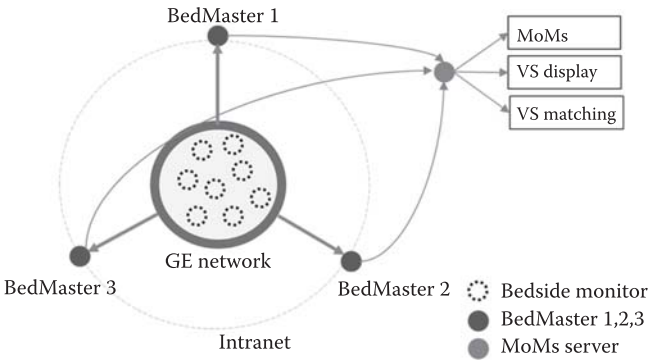

FIGURE 24.1
The MoMs system architecture with triple modular redundancy design using three BedMaster servers.

system status monitoring and real-time presentation critical clinical data, we developed the MoMs (Monitor of Monitors) information representation layer over the VS collection system. Using the current data collecting architecture and a minimum-instrument approach, we stream the most recent record from the BedMaster server from each bed to a dedicated data server, the MoMs server (Figure 24.1). A high-performance database hosts those data items labeled with data server name, bed unit, time stamp, and admission status.

The front end (MoMs viewer) is designed as a web-based application so that users can access it from any location in the hospital. IP address white list and user login modules are used for information security. Each bed collection status is summarized and pushed to the MoMs viewer through the Ajax (asynchronous JavaScript and XML) techniques every minute. All participating patient bed units are represented by individual cells in each of three spreadsheet blocks representing one of the three redundant BedMaster servers. Figure 24.2 shows a block of the web-based monitoring system corresponding to bedside collections from monitors in the TRU, ORs, and neurotrauma critical care (NTCC), and multitrauma critical care (MTCC) units. The background color of each cell represents the associated bed's data collection status. Green indicates that the data stream has been alive in the last 5 minutes. Yellow indicates that the last time stamp from data from that bed/monitor is 5 minutes to 4 hours old and that a problem may exist. Dark red indicates a time stamp gap

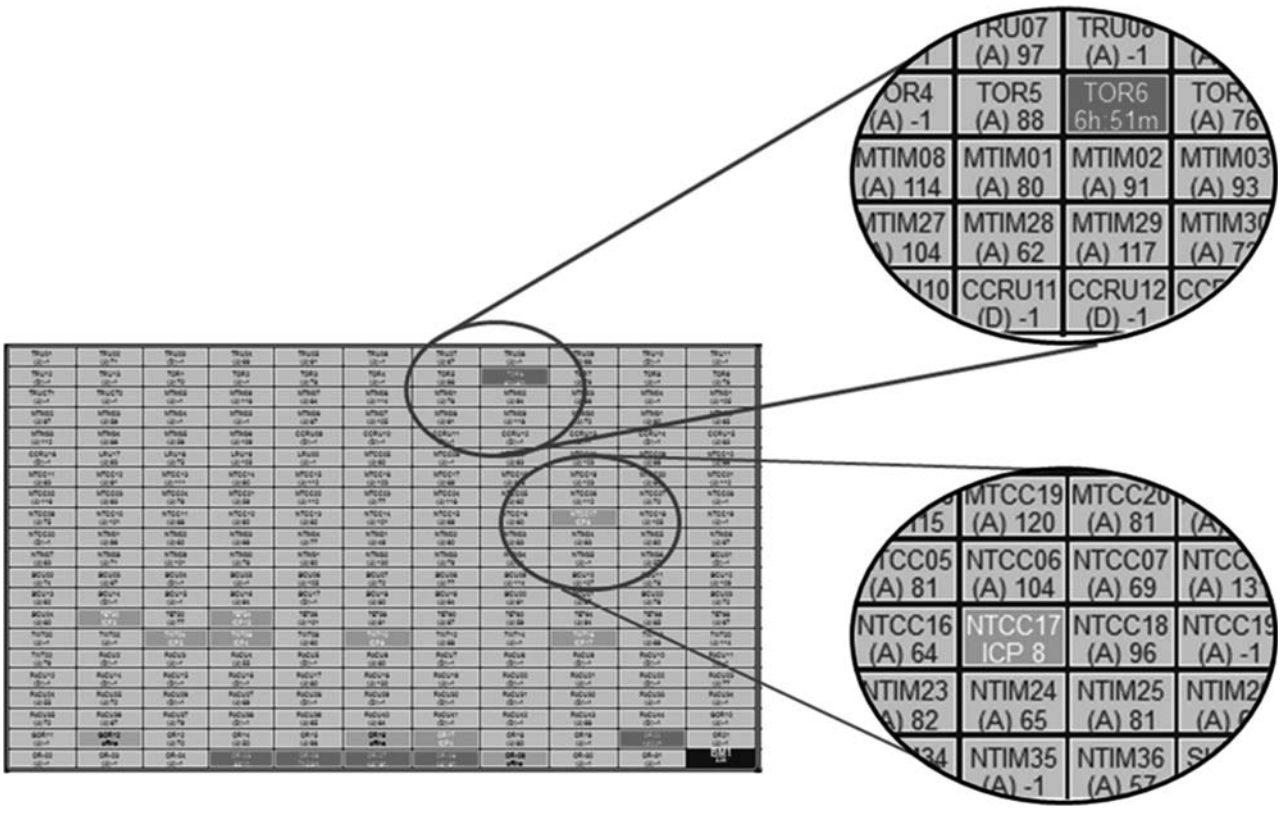

FIGURE 24.2
A portion of MoMs viewer for data collection status. Light gray (shown): collection is active (within last 5 minutes); silver (not shown): collection was active 5 minutes to 4 hours ago; dark gray (shown): no data collection has occurred in more than 4 hours; gray (shown): bedside collection is offline. In each cell, letter "A" means admitted; letter "D", discharged. The medium gray background cell indicates a patient with an intracranial monitor in place; ICP value appears in white.

greater than 4 hours and that action should be taken to remedy the problem. Report of an elapsed data collection gap includes the duration of collection failure.

There are different elements in each colored cell to indicate bed unit occupancy. Often, nurses may press a bedside button for admission (A) to or discharge (D) from this bed. This allows for a cross-check on potential causes for information gaps such as the device being temporarily inoperable or no patient being monitored. In addition, bed occupancy can be verified by real-time physiological values, such as HR. It can be used as second evidence for us to infer if a bed unit is currently occupied by a patient. If one bed unit is offline, the gap between now and its last reported time will be shown. Figure 24.2 shows one example in the unit OR-6, which is highlighted in a red cell with a time gap of 6 hours and 51 minutes.

The easy configuration of the MoMs dashboard viewer also allows it to be used for identifying and displaying clinical information of special research interest. For example, ICP monitoring is an important VS for traumatic brain-injured patients and is not often collected due to its invasive nature. To receive early notification of ICP-monitored cases, the MoMs viewer can extract ICP data from all bed/monitor units' data streams and display these data using a predefined color code. In Figure 24.2, those pink cells with white bold font text show real-time ICP values from the corresponding bed/monitor units. For example, at the time we viewed the MoMs system, the unit NTCC-17 was monitoring ICP with an instant value of 8 mmHg.

In a 12-month study period, single-server collection rates in the 3-month period before MoMs deployment averaged 81.4%. Of the 18.6% collections lost, most (18%) were brief periods, 5 minutes to 4 hours. Reasons for gaps included collection server failure, software instability, individual bed setting inconsistency, and monitor servicing. In the 6-month post-MoMs deployment period, average collection rates were 99.96%. The triple-redundant patient data collection system with real-time diagnostic information summarization and representation improved the reliability of massive clinical data collection to near 100% in a level I trauma center.

24.3 Data Analytics in Trauma Patients

In treating trauma patients, clinicians often wish to know if a patient has elevated ICP; if some lifesaving interventions are needed in the near future, such as massive transfusion; or if a patient has unstable neurologic status. Although there is a high volume of data monitored for trauma patients, they are often underused to answer those questions. Those data are often only used for bedside instant physiological status view. Due to the difficulty in storing, accessing, and analyzing, those data streams are hardly used beyond the VS readings or instant waveform morphism analysis. Clinicians often utilize those medical data in an empirical way, which may consume their extra attention while not fully unlocking the value of the data. Validated, automated medical data processing and analysis could aggregate massive amounts of data and assist clinicians to quickly recognize changes in physiological status and prioritize care.

Processing medical data/signals is an important initial step before building predictive models. In prehospital and even in hospital, data are collected in a hostile environment, which adds noise to the original signals. Numerous methods and algorithms have been studied to remove outliers and noise from the signal, to smooth or sample from high rate signals, and to explore signal characteristics in time, frequency, and joint domains [4,21–23]. In this section, we assume that data/signals have been preprocessed as "clean data" and focus on a few topics on feature design, feature selection, and model evaluation, which are less tackled in many medical data processing handbooks.

24.3.1 Features Based on Descriptive Statistics

Maintaining patients' VSs within normal ranges is a basic task for clinicians. Some treatment protocols also give guidelines on the thresholds of VSs to be watched during patient care. For example, for the management of severe TBI, it is recommended to begin initial treatment when ICP is above 20 mmHg, and CPP is suggested to remain above 70 mmHg [24]. The guidelines for field triage of injured patients recommended the use of SBP <90 mmHg or RR <10 or >29 breaths per minute as a part of the physiological criteria when considering if a patient should be transported to a facility with the highest level of care [25]. When adding the age factor, SBP <110 mmHg was recommended in the National Trauma Triage Protocol for geriatric trauma patients (age > 65 years old) [25,26]. Empirical thresholds are also used as simple predictors or decision triggers by clinicians. In field triage, shock index (SI) > 1 is used as an indicator of circulation failure and a predictor of critical bleeding, since it has high specificity and is easy to calculate [27].

Those thresholds that pass muster with the clinical experts indeed can serve as domain knowledge to design features from physiological time series. Such type of features may hold clinical meanings that are easy to interpret by humans. To quantify the cumulative effect of VS away from normal range, the "pressure-times-time dose" (PTD) is defined as the integrated area enclosed by

the VS curve and the threshold line within a given time interval. Sometimes, to compare between patients, it is also calculated as averaged PTD in unit time, which is the PTD normalized by time duration. Even thought two patients may be monitored for different time durations, their PTDs in unit time are still comparable. In predicting TBI patients' outcomes, the PTDs of ICP >20 mmHg and CPP < 60 mmHg have been shown to be good predictors of in-hospital mortality and length of ICU stay [28].

Similar to the idea of PTD, some descriptive statistics of VS within a given time period are also informative for clinician use or may contribute to outcome prediction in a model. With the assumption that the observed data are approximately normally distributed, mean and variance are often used to sketch the VS value distribution. Standard deviation (SD) is used to quantify the variability of observed VS data. The coefficient of variance, which is the SD divided by the mean, is a unit-less value that is suitable to compare between data sets with widely different means. Robust statistics, such as percentiles or quartiles, are also used to quantify the shape of the VS data distribution. The median (50th percentile or 2nd quartile) is one of the most commonly used statistics in VS feature calculation.

Those descriptive statistics, plus the PTD value of VS in a given time period, have the following merits. First, those features can easily incorporate domain knowledge and hence possess straightforward meaning. Stein et al. described a scheme of designing 588 features from 9 typically used VSs, using various thresholds and time windows [14]. Second, they could aggregate large amounts of data or simplify complex data as a few summary quantities, which can be used by many statistical prediction models. For example, to identify patients with critical bleeding risk, Mackenzie et al. summarized patients' first 15 minutes 240 Hz PPG by applying the above descriptive statistics to the peak-to-valley distances of PPG. In such a way, each patient's 216,000 waveform data points are aggregated into dozens of features as inputs for logistic regression [29]. Third, those values can be calculated efficiently even if the data are of high volume and high velocity.

Despite these listed merits, we need to understand that some of these features rely on certain assumptions. For example, normality is often assumed when we use the mean and SD to sketch the shape of VS data distribution. Besides, we may assume that the thresholds suggested by guidelines can be applied to any patient group at any time period. In reality, the thresholds may vary along time, no matter how slow. Also, some hidden factors may result in a large difference in threshold values between age groups or sex groups. Therefore, we need to be careful when incorporating clinical prior knowledge into those features, by examining the applicable conditions.

24.3.2 Features Based on Variability

Continuous noninvasive ECG and PPG sensors are ubiquitous in both prehospital emergency medical service and in-hospital healthcare for TBI patients. Waveforms measured from both sensors capture rich information on cardiovascular, circulatory, and respiratory systems. Heart rate variability (HRV), derived from ECG waveform, has long been used in studies of prehospital lifesaving interventions [30] and neurologic disorder [31,32], and the association between the autonomic nervous system (ANS) and cardiovascular mortality is well established [33]. ANS is a complex life-sustaining system, which plays a role in nearly every organ and disease [34], including regulating cardiac activity, respiration, and pupillary response. ANS dysfunction is a potential complication following severe TBI [35]. While neurological diseases can lead to many changes in cardiac function, the major changes noted are arrhythmias and repolarization changes. Goldstein et al. [36] noted that either increased ICP or decreased CPP can be associated with ANS dysfunction. Baguley et al. [37] observed that in TBI, patients with or without dysautonomia showed HRV differences compared to the control group.

Another universally used sensor, pulse oximetry, generates a PPG waveform that carries rich physiological information, such as HR, SPO_2, and even RR [38]. The peaks in PPG are almost synchronized and corresponding to the R peaks from ECG in the time domain. Therefore, the peak–peak interval from PPG can be used as an alternative to the normal-to-normal (NN) interval calculated from ECG recordings. Lu et al. compared HRV and PPG variability (PPGV) in both time- and frequency-domain-based on data collected from 20 healthy subjects. They found that PPGV was highly correlated to HRV and could serve as an alternative measurement [39].

To find QRS peaks from ECG, the Pan–Tompkins method is used [40]. The ECG signal is smoothed with the Savitzky-Golay filter to increase the signal-to-noise ratio. The R peaks are then detected based on a threshold that is adaptive to the signal. To detect PPG signal peaks and valleys, the Savitzky–Golay filter is applied to smooth the signal [41]. The peaks can be found through two rounds of searching. In the first round, the peaks are roughly found through MATLAB built-in routine "findpeaks." The median distance between two consecutive peaks PD_{median} is then calculated. In the second round, any small peaks within a range of $0.6 \times PD_{median}$ from a large peak would be ignored. Figure 24.3 shows a typical PQRST segment from ECG, with identified peaks and five items that we use for calculation. The NN interval illustrated by item 1 is the time interval between two consecutive R peaks. HRV variables in the time domain and

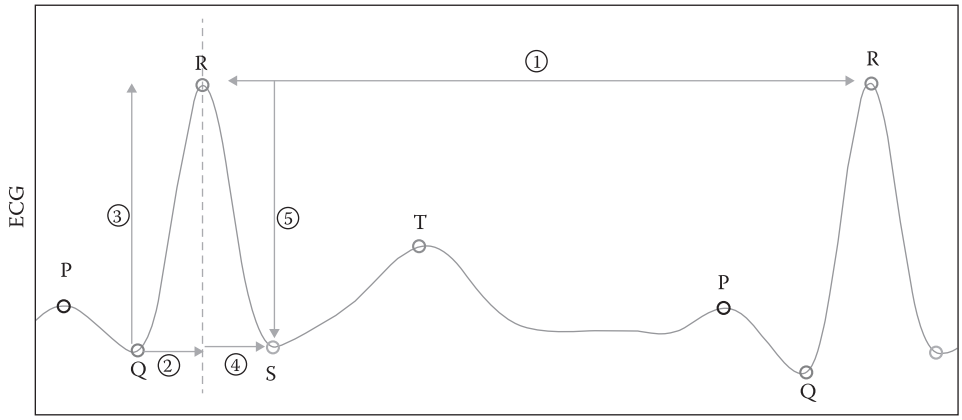

FIGURE 24.3
An exemplary ECG segment with identified P, Q, R, S, T peaks. Five items from the segment are used for ECG feature calculation. Item 1 is the NN interval. Items 2 and 4 are Q-to-R rising time and R-to-S falling time. Items 3 and 5 are Q-to-R rising amplitude and R-to-S falling amplitude.

nonlinear dynamics can be calculated based on the Task Force of the European Society of Cardiology and the North American Society of Pacing and Electrophysiology [33]. From items 2 and 4 in Figure 24.3, the rising time from Q to R and the falling time from R to S can be calculated. Similarly, from items 3 and 5, the rising and falling amplitudes from Q to R and R to S can be calculated.

Because signals may be collected when the patient goes through resuscitation, or has significant movement, artifacts exist and do not reflect the patient's real physiological status. To flag out signals with a large amount of artifacts, signal quality can be evaluated based on R peaks in ECG and the peaks in PPG. The assumption is that signals have normal distributed RR intervals. RR intervals from segments of low quality are detected as outliers using the *z*-test. Figure 24.4 illustrates two segments of signals. The left side shows ECG and PPG with low signal quality flagged with horizontal bars. The right side shows precisely identified peaks in both signals.

PPGV and waveform morphology features can also be designed similarly with expansion based on PPG unique characteristics. Figure 24.5 (top subplot) shows a normal PPG segment with identified peaks and valleys. Item 1 illustrates a peak-to-peak time interval, which is analogous to the NN interval in ECG. PPGV variables and morphology features were calculated from items 1–4 as they were for ECG. PPG waveform also has a unique dicrotic notch, and its shape has been studied and shown to be related to arterial stiffness and aging [42]. To measure the deceleration and acceleration near the dicrotic notch, the first and second derivatives of PPG were calculated through three-point central difference [43,44]. Item 6 in Figure 24.5 (shadowed area) depicts the region where PPG height changes its moving direction while reducing acceleration and speed until it reaches a local peak. The duration and amplitude change in this movement were derived as two features [45].

Similarly, blood pressure variability (BPV) has been derived from noninvasive blood pressure (NIBP) as a cardiovascular risk factor [46]. The ABP waveform also

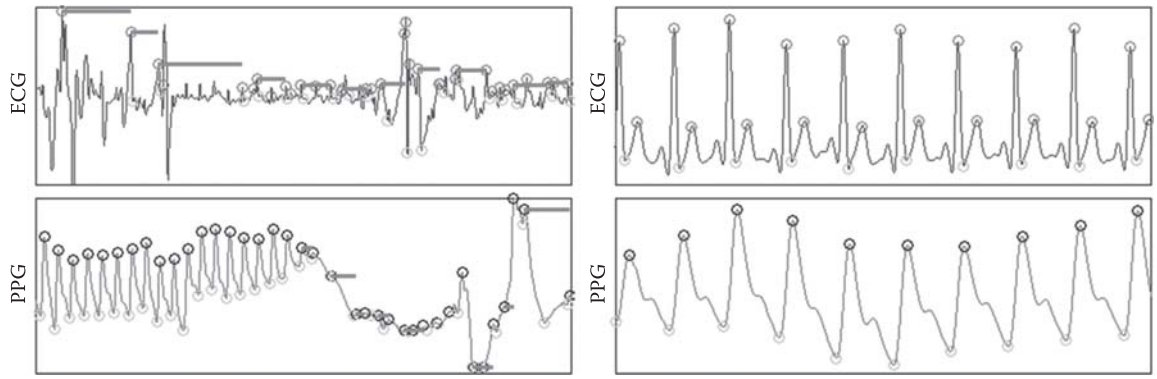

FIGURE 24.4
Illustration of ECG and PPG segments of bad and good signal quality, evaluated by *z*-test on NN intervals. On left side, horizontal bars flag ECG and PPG segments of bad signal quality.

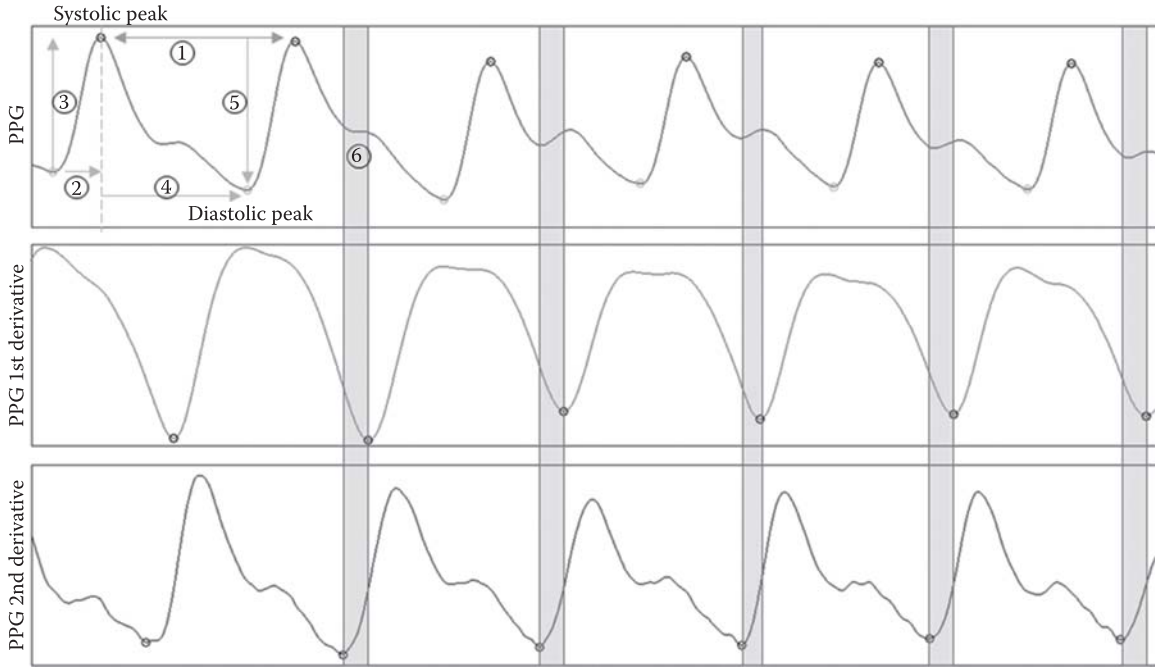

FIGURE 24.5
(Top) An exemplary PPG segment with identified peaks and valleys. (Middle) the first derivative of PPG signal. (Bottom) The second derivative of PPG signal. Item 1 is the peak–peak interval. Item 2 is the valley–peak rising time. Item 3 is the valley–peak rising amplitude. Items 4 and 5 are peak–valley falling time and amplitude. Item 6 is the notch area, where the time duration and amplitude change can be calculated.

carries a pulse wave. With identified NN intervals, NIBP or even ABP could be used to calculate BPV in the time domain and frequency domain, as well as the nonlinear dynamics features.

24.3.3 Features Based on Correlations

Components in a biological system interact with each other in a sophisticated way. Discovering patterns that sketch those interactions could help to understand and forecast the biological system's behavior. Correlation is a type of statistical quantity that can indicate predictive relationship, linear or functional, bivariate or multivariate. Hence, it is a nice tool to explore system component interactions. In this section, we review the applications of different correlations in quantifying physiological system behaviors.

Maintaining normal cerebral perfusion and oxygenation is important in managing severe TBI patients. Monitoring of the cerebral autoregulation could inform clinicians if a patient has lost the ability to maintain a constant perfusion when blood pressure changes [47]. Cerebrovascular pressure-reactivity index (PRx) was proposed as an indicator of loss of autoregulatory reserve [48,49]. PRx is calculated as a moving correlation coefficient between the mean arterial pressure (MAP) and ICP. Given a short time window, about 40 consecutive data points of MAP and ICP in 4–5 minutes are used for calculation [50]. When cerebral autoregulation is intact, cerebral blood flow (CBF) remains a normal constant speed and does not change significantly with mean blood pressure. In such a situation, PRx should be close to zero, indicating no or weak linear correlation between MAP and ICP. When cerebral autoregulation is damaged after severe head injury, CBF increases or decreases with blood pressure. The absolute value of PRx moves away from zero, indicating strong linear correlation between MAP and ICP. In this way, PRx can serve to continuously monitor the existence of cerebral autoregulation. When indicators of autoregulation are plotted against cerebral perfusion pressure (CPP), a U-shaped curve is generated, consistent with a loss of autoregulation in conditions of hypoperfusion or hyperemia. A study in 2002 by Steiner et al. took advantage of this relationship to construct curves of CPP against PRx, and hypothesized that the minima of these would indicate an ideal CPP at which pressure reactivity is maximized [51]. This group and others since have validated this model by finding a correlation between patients' deviation away from this optimal CPP and eventual neurologic outcome.

Another example of using linear correlation in brain trauma study is the calculation of pressure–volume compensatory reserve index, called RAP. It is the moving linear correlation coefficient (R) between the amplitude (A) of a frequency component corresponding to HR in the ICP waveform (Figure 24.6) and the mean ICP pressure (P). Given a short time window of length 4–5 minutes,

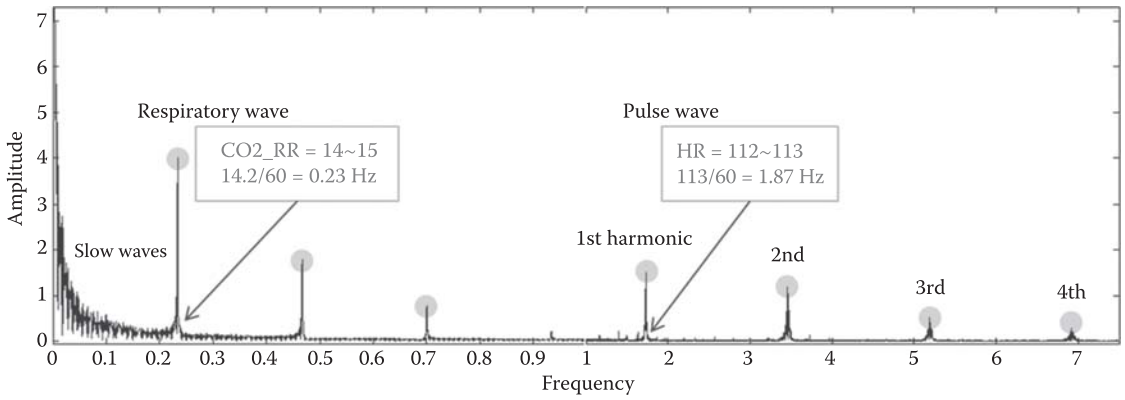

FIGURE 24.6
Fourier transform of 10 minutes ICP waveform. Pulse wave components show four harmonics, with the first one corresponding to heart rate. The other two components, respiratory wave and slow waves, are both identifiable from the frequency domain.

the ICP waveform is transformed into the frequency domain by Fourier transform. The amplitude of the frequency component relating to the HR can be found. RAP indicates the relationship between ICP and changes in volume of the intracranial space [49]. When ICP is low (e.g., ICP < 20 mmHg), near-zero RAP means that a change in cerebral blood volume has no or small impact on ICP, which indicates a good pressure–volume compensatory reserve. After severe head injury, mean pulse amplitude (AMP) and ICP may be negatively correlated, indicating the loss of cerebral autoregulation.

24.3.4 Features Based on Entropy

Physiological data are a succession of values outputted from the information source, namely, the patient. Entropy is a measure of the amount of information from the source. Given a discrete random variable X with its all-possible-element set S_X and probability function $p(x)$, the Shannon entropy $H(X)$ of the random variable X is defined as

$$H(X) = -\sum_{x \in S_X} p(x) \log p(x). \tag{24.1}$$

Given two random variables X and Y, the joint entropy is a generalization to measure the uncertainty of their joint distribution, which is defined as

$$H(X,Y) = -\sum_{x \in S_X} \sum_{y \in S_Y} p(x,y) \log p(x,y). \tag{24.2}$$

It can also be extended to the joint distribution of a set of more than two random variables. The conditional entropy of random variable X given Y measures the uncertainty of x when y is known, which is defined as

$$H(X|Y) = -\sum_{x \in S_X} \sum_{y \in S_Y} p(x,y) \log p(x|y). \tag{24.3}$$

There are more extensions derived from above definitions, such as relative entropy, mutual information, conditional mutual information, etc. [52]. They could serve as useful mathematical tools to extract features from one or jointly more medical time series to quantify the uncertainty and interactions of the systems that they represent.

Approximate entropy (ApEn) and sample entropy (SampEn) are two widely used entropies in measuring similarity in time series, especially in HRV [53,54]. For a physiological time series, even short ones, ApEn can be used to estimate the rate of new information it generates [53,55]. Given a time series $x_1, x_2, \ldots x_n$, its Takens embedding vectors with dimension m are $\boldsymbol{X}_m(i) = (x_i, x_{i+1}, \ldots, x_{i+m-1})$, where $i = 1, \ldots, n - m + 1$ [56]. The distance $d(\boldsymbol{X}_m(i), \boldsymbol{X}_m(j))$ between any two such points $\boldsymbol{X}_m(i)$ and $\boldsymbol{X}_m(j)$ is the maximum absolute element-wise difference between them. With a preset threshold r, B_i is the number of $\boldsymbol{X}_m(j)$ such that $d(\boldsymbol{X}_m(i), \boldsymbol{X}_m(j)) \le r$. Let $\Phi^m(r) = (n-m+1)^{-1} \sum_{i=1}^{n-m+1} \log(B_i/(n-m+1))$. The ApEn is defined as $ApEn = \Phi^m(i) - \Phi^{m+1}(i)$. Larger ApEn means that the time series is more irregular or a smaller chance of repeated template sequences in the time series. However, ApEn overestimates the similarity since it counts $\boldsymbol{X}_m(i)$ itself when finding B_i, which is called the self-match [53]. The SampEn was designed to reduce the bias caused by self-matching. The SampEn of HRV has been shown to be useful to detect sepsis in the neonatal ICU [54].

The interactions between physiological time series can also be sketched by the permutation entropy, which is the complexity of a symbolic sequence after embedding the time series into a high-dimensional space. First we need to define the ordinal pattern. Given a sequence of numeric elements $x_1, x_2, \ldots x_n$, its ordinal pattern is the permutation $\pi = (i_1, i_2, \ldots i_n)$ that sorts the elements in ascending order such that $x_{i_1} \le x_{i_2} \le \ldots \le x_{i_n}$. Given a time

series $x_{1\ldots N}$ and a time window of length L, the order L permutation entropy can be calculated as follows. Let π_t be the ordinal pattern (i.e., the sorting permutation) for the segment of the time series under the sliding window of length L that ends at x_t, i.e., the subsequence $x_{t-L+1}, \ldots, x_t$. Let $S_L = \{\pi_t\}$ be the set of all those unique (alphabet) ordinal patterns π_t. The time series $x_{1\ldots N}$ corresponds to the sequence $\langle \pi_t : t = L, \ldots, N \rangle$ of $N - L + 1$ ordinal patterns from the alphabet S_L. The entropy of this sequence of ordinal patterns is the permutation entropy of the time series $x_{1\ldots N}$. For example, the Shannon permutation entropy is defined in Equation 24.1,

$$H_L = -\sum_{\pi_k \in S_L} P(\pi_k) \log(P(\pi_k)), \tag{24.4}$$

where $P(\pi_k)$ is the frequency of π_k in the sequence $\langle \pi_t \rangle$. In the work presented here, we use the Rényi entropy with parameter α of the sequence $\langle \pi_t \rangle$ defined as

$$R_L^\alpha = \frac{1}{1-\alpha} \log\left(\sum_{\pi_k \in S_L} P(\pi_k)^\alpha \right). \tag{24.5}$$

The parameter α in the Rényi entropy acts as a selector of probabilities. It assigns almost equal weight to each possible probability when α is sufficiently close to zero. When α is larger, it puts more weight on higher probabilities. We can use this parameter to assign different weights on events of different probabilities.

Permutation-based partitions are more robust to noise and other nonlinear distortions and artifacts than value-based fixed-size partitions of the state space, since they depend on the relative order rather than the exact values of the time series. Furthermore, in order to obtain reliable entropy estimates with fixed-size partitions, one needs long time series (in the order of 2^m in order to cover all blocks of such fixed-size partitions); permutation-based entropy estimates do not require long time series. The robustness of permutation entropy makes it particularly attractive for mining VSs collected in real clinical settings, without expensive preprocessing and cleaning of such signals.

24.4 Machine Learning Framework and Its Applications

In the medical field, statistical models are traditionally used to test casual hypotheses, with the purpose focused on explanation [57]. Consider that a new surgical procedure is being tested for its effect of improving patients' long-term neurologic outcome. Doctors may firstly form a statistical null hypothesis that the new procedure causes no statistically significant improvement of patients' long-term neurologic outcome compared with the old procedure. Sample size and confounding factors will be identified, before the randomized case enrollment. With all collected data, an appropriate statistical testing method will be used to test if the null hypothesis should be rejected at a certain significance level. There are many other statistical models to explain the association between factors and outcomes, such as multivariate regression and Bayesian models. No matter what the approaches, model interpretability is always emphasized in the explanatory studies.

Another type of purpose is to predict. In daily patient care, clinicians also want to know, for example, if a patient with blunt injury needs massive transfusion after admission, or if a TBI patient will need a decompressive craniectomy procedure after 6 hours. With all available data already observed for those patients, reliable prediction can assist clinicians in making early decisions, such as calling the blood bank to prepare a sufficient amount of blood product or scheduling early the use of the OR and notifying the operation team. As the decisions based on the predictions have immediate consequences and cost, the prediction must be of high accuracy. In such situations, interpretability could be sacrificed for high prediction accuracy.

Machine learning methods are a collection of algorithms that can discover patterns from data and use them to predict future values. They can be roughly categorized as supervised, unsupervised, and semisupervised learning. Usually, there are preprocessing steps to convert raw data into a feature set that characterizes the observed object or task. In the past, an experienced clinician distilled concise and practical rules from years of observation and clinical practice. With massive data, such a process can be accelerated with automated machine learning algorithms. However, the algorithms may generate counterintuitive models, misled by outliers, missing values, biased data, and incorrect assumptions [58,59]. Therefore, *a priori* knowledge and model validation are essential in building and selecting models. Moreover, a machine learning method could learn from a training data set and attempt to minimize its error on that set. For practical use, we hope the learned model also makes small errors on previously unseen data. A prediction model that can perform well on new data has good generalization ability.

For comprehensive learning algorithms, interested readers can find useful information from Murphy, Bishop, or Hastie's books and their references [58,60,61]. In this section, we focus more on methods that increase the model generalization and the metrics of evaluating the models. With those tools, we could better validate the models and create actionable models.

24.4.1 Feature Selection

Instead of using raw data as inputs, machine learning methods often use variables that characterize the data or learning task's properties (features). Feature selection is an important part of data preprocessing before building a simple and robust model. In many applications, including the Big Data scenario, we can find many attributes to characterize the data, without knowing which ones may contribute to the improved accuracy of prediction. Features of little or no importance not only increase the model complexity but also have a negative effect on parameter estimation for those features that are of importance to the prediction [62]. With massive data, such overfitting is more likely to happen than underfitting.

One class of feature selection methods, which are called the filter-type methods, are model-independent. They evaluate features according to a certain criterion, before adopting any learning algorithm. There are various criteria that can be used. For example, zero or near-zero variance features are often excluded. This is convenient and works in most situations. However, we should be careful, if there is evidence showing that such near-zero variance features may have strong prediction of the outcome. For example, a certain type of drug, like barbiturates, is only used as a last resort when patients are in critical condition. If one binary feature indicates if a patient was administrated barbiturates in hospital, it may have almost near-zero variance since most patients would not receive this drug. But this feature could be a strong predictor of unfavorable outcome, such as mortality.

Some filter-type feature selection methods use the known outcome in the training set. For classification problems, one feature selection approach often seen from medical literature is the use of receiver operating characteristic (ROC) analysis. Besides ROC, information gain can also be used to evaluate each feature's importance to decrease the uncertainty about the outcome. For continuous value prediction, correlation could be used to find a subset of features that each has strong correlation with the outcome while this subset of features are weakly correlated internally. Those methods have low computational complexity, which grows linearly with the number of features. Since the selection is independent of the learning algorithm, the filter-type feature selection methods are good for initial fast screening of a large amount of features to filter out possibly nonimportant ones. Also, the ignorance of interaction between features may result in redundant selection.

Another class of feature selection methods, which are called the wrapper-type methods, search for the important features within the context of models. Given a learning method and a performance metric, the wrapper methods search for an optimal subset of the features that maximizes model performance. Arguably, stepwise selection is the most frequently used wrapper-type feature selection method. In the forward selection step, each feature that is not in the current model is added to a temporary model. The new model is evaluated based on a certain criterion. For example, the most debated criterion is the use of p-value as the inclusion condition. If the hypothesis test calculates that the p-value of the newly added term is less than the inclusion threshold, then keep the feature in the next round of evaluation. Other inclusion criteria include the Akaike Information Criteria (AIC) and Bayesian Information Criteria (BIC), which are considered better than p-value-based selection. The search halts when there are no features that meet the inclusion criteria remaining outside the model. In the backward selection step, the full model that contains all candidate features, or all selected features from the forward selection, has features iteratively removed based on the exclusion criteria, which could be p-value, AIC, BIC, etc.

Permutation is another way to evaluate feature importance in a model. Intuitively, if a feature is of real importance, then randomly permuting its values will greatly reduce the model performance. Given a trained model and performance metric, the trained baseline model has its performance P_b. Each feature is randomly permuted, while other features remain unchanged. The model is then trained again with the altered ith feature and evaluated for a new performance P_{ri}. The change of performance metric is calculated as $\Delta P_i = |P_b - P_{ri}|$. To have a robust estimation of the performance change, we can repeat such random permutation many times and use the averaged difference. In this way, we can rank all features according to the change of performance metric. Then the selected features are optimal in terms of that metric.

The third class of feature selection methods, which are called the embedded methods, train the model and select features simultaneously. Many machine learning algorithms have such inherent feature selection mechanism, such as random forest, relevance vector machines, and decision trees. A more general approach is to add a penalty term for the model complexity to the cost function using a Lagrangian multiplier. Take a data set $D = \{(\boldsymbol{x}_i, y_i)\}$, where $\boldsymbol{x} \in R^d$ and $y \in R$. A learning method finds the optimal parameter set $\boldsymbol{w}$ through minimizing a loss function f by $\min_{\boldsymbol{w}} f(\boldsymbol{x}, y, \boldsymbol{w})$. Through a Lagrangian multiplier, the optimization target can be extended to be

$$\min_{\boldsymbol{w}} f(\boldsymbol{x}, y, \boldsymbol{w}) + \lambda ||\boldsymbol{w}||_1, \tag{24.6}$$

where $\lambda > 0$ is called the regularization parameter. The penalty term $||\boldsymbol{w}||_1$ is the ℓ_1 norm of the model parameter set $\boldsymbol{w}$. Equation 24.6 is well known as the least

absolute shrinkage and selection operator (LASSO) [63]. It has a nice property to suppress variable coefficients, which introduces sparsity while learning the parameter set. It has been shown to be efficient in creating parsimonious and robust models [60,64].

All three classes of feature selection methods have their own applicable scenarios. The filter types are computationally efficient. Because they can be independent of the model, during the initial stage of experiment design in medical studies, those methods can be used to quickly filter out possibly useless variables, even if there are no known outcome labels yet. This can help reduce the amount of expensive data collection. However, filter methods lack knowledge of the interaction between features and thus may include redundant features. The wrapper-type methods iteratively search for the optimal subset of features by testing each variable's "contribution" to a selected model performance metric. The search scheme takes variables' interactions into consideration. But each evaluation iteration requires the model to be trained again, which could be time consuming. The embedded methods have much more desired advantages. First, they include feature interaction for consideration. Second, with efficient optimization solvers, parameter learning and regularization can be done at the same time efficiently. Moreover, through special design of the regularization terms, not only feature importance but also the structure among feature groups could be selected [65].

24.4.2 Performance Metrics

An appropriate model performance metric allows defining the right learning objectives and can increase the chance of building generative models. Two levels of performance are crucial in creating a good predictive model. First, we need to evaluate the model's generalization capability, namely, can we have a high expectation on its performance on future unseen data? Second, we want to know the model's performance on the training data, namely, how well has this model learned from what it could observe?

In practice, we always have finite observations. One reason is that collecting data could be expensive. Another reason is that it is infeasible to sample the entire population in most situations. To know models' expected performance on unseen data, we have mainly two approaches, called external and internal validations. The former uses data from external sources, such as from other geologically different clinical facilities or from other historical time points. The latter reserves an internal small portion of the collected data and sets it aside only for testing.

External validation is often desired in clinical studies. Data collected from one regional hospital may still have its sampling bias, mainly influenced by its local demographics. Also, data collected from a civilian hospital may not represent the military population. Hence, it is important to validate models built from single-center studies for their generalization. There are many large-scale, high-quality clinical databases maintained for public use. PhysioNet is a large physiological database archiving data contributed from worldwide [66]. The National Trauma Data Bank (NTDB) is the largest U.S. trauma registry data set, which has been used in study trauma injury epidemiology and validation of guidelines from clinical organizations [67–69).

Despite our being in the Big Data era, collecting large amounts of clinical data is still expensive, and a tremendous amount of effort is required to store and process massive data. Hence, it is luxury to keep a large part of the data outside of training. Internal validation is necessary for most single-center studies, especially when external data are not available. The internal validation randomly partitions the collected data into training and testing. A typical validation scheme is the k-fold cross-validation. A data set is randomly partitioned into roughly k equal-size nonoverlapping subsets. In the ith validation, where $i = 1 \ldots k$, the ith subset is used as a testing set, and the remaining data are used for training. Such process is iterated over all k subsets. After one round of k-fold cross-validation, there are k model evaluations. Often, we can repeat such k-fold cross-validation N times by randomly partitioning the set again. With the $k \times N$ evaluations, we can use their averaged performance or other robust statistics (e.g., SD, median, etc.) to estimate the model's future performance on unseen data. Such random subset sampling is a simulation of possibly different distributions of new data.

Commonly, mean squared error (MSE) or root mean square error (RMSE) is used in evaluating continuous outcomes. Accuracy or the confusion matrix is typically used for discrete outcomes. When prediction problems have some special issues, such as imbalance in data labels or different preference on incorrectly predicted cases, we need to carefully select the performance metrics. In this section, we discuss the ROC curve and the precision and recall curve (PRC), the two types of model performance metrics that are widely used in medical classification problems.

Accuracy as a model performance metric has some drawbacks in medical data analysis. It could be misleading when the data labels (e.g., positive and negative outcomes in binary classification) are highly imbalanced. A classifier could cheat to achieve high accuracy by just predicting all instances to be the most frequent class label. Unfortunately, in many medical problems, the outcome of interest often is a small portion. For example, in a regional level I trauma center, massive transfusion rate could be as low as 1.3%–2.2% in adult patients [70].

A classifier could predict no massive transfusion for all patients and achieve >98% accuracy. However, mispredicting for a patient who needs massive transfusion may result in severe clinical and social outcome.

To alleviate this problem, accuracy is broken into four parts for binary classifications, which is often described as the confusion matrix. The correctly predicted positive cases are true positives (TPs). The incorrectly predicted positive cases are false positives (FPs). Similarly, the true negatives (TNs) are correctly predicted negative cases; and the false negatives (FNs) are incorrectly predicted negative cases [71,72].

Given a set of instances for prediction, a classifier gives corresponding predicted values, called prediction scores. If we sort all the scores in ascending order and use each unique value as a threshold, all cases with score values smaller than the threshold are classified as negative, whereas those that are higher will be classified as positive. Then for each threshold, we can calculate the false positive rate (FPR) and the true positive rate (TPR), which constitutes a point in a 2-D coordinate. After iterating all possible thresholds, a full curve is generated, called the ROC curve. The PRC curve is created in a similar way, with positive predicted values [PPVs = TP/(TP + FP), also known as precision] and TPR (also known as recall).

The ROC and PRC curves provide a full-spectrum evaluation of a classifier by visualizing FPR, TPR, or PPV, TPR for all possible thresholds, instead of giving a single decision point. The ROC curve is widely used to evaluate prediction models in many medical studies, while the PRC curve is much less frequently seen. Davis and Goadrich [73] proved that ROC and PRC curves have one-to-one correspondence, given a data set with finite positive and negative cases. If one ROC curve has all its points above or equal to another ROC curve, it is said that the first ROC curve dominates the second one; and this is true if and only if the first PRC curve dominates the second PRC curve. However, in many situations, ROC curves from multiple models built from the same data set are twined or close to each other. When the data set is highly imbalanced, i.e., the negative cases significantly outnumber the positive cases, the PRC curve has better ability to identify classifiers that have good performance in predicting the positive cases.

From one of our projects, we took 1191 injured patients and collected continuous physiological VSs to predict the next 3-hour blood product use. Among the 1191 patients, only 7.2% received blood product. As an illustration, Figure 24.7 compares the ROC and PRC curves of three methods we used. To focus on the use of the performance curves, we skip the details of the models here. As we can see from Figure 24.7a, the darker ROC curve (Method 1) dominates the other two, as does the blue PRC curve. The other two ROC curves (Methods 2 and 3) are close to each other and intertwined. From Figure 24.7a, the two methods have very similar performance, but from Figure 24.7b, we can observe that the light gray PRC curve (Method 3) has higher precision than the gray PRC curve (Method 2). Therefore, in comparing the methods, we may consider that Methods 1 and 3 have better performance than Method 2.

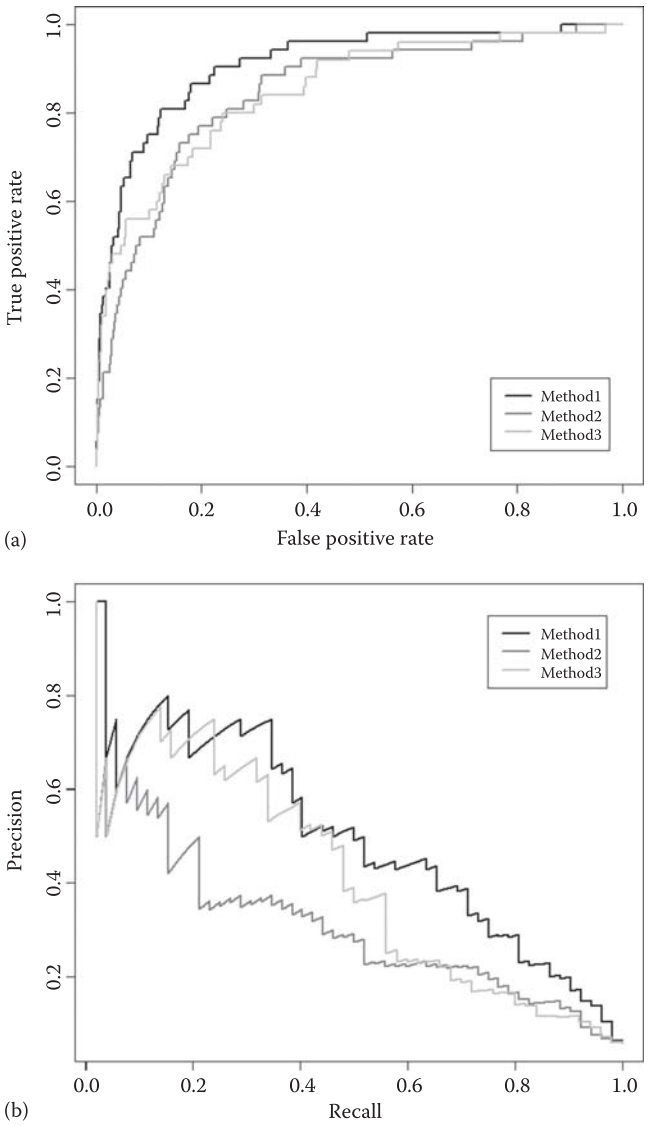

FIGURE 24.7
(a) Example ROC curves of three methods. (b) Example PRC curves of three methods.

24.5 Computing Issues

Storing, processing, and learning from massive medical data require intensive calculation. First, having a

high-performance file format to store and organize large amounts of data is critical for the input and output (I/O) of data processing. For a typical TBI patient staying 7 days in a trauma center, five 240 Hz waveforms are monitored and up to 700 million data points (equivalent to 1-gigabyte disk size, if data are stored in 12-bit format) would be collected. Traditional spreadsheet-based data management becomes less efficient within such a Big Data scenario. Hierarchical Data Format (HDF) is a high-performance data format that offers on-the-fly data compression and high I/O performance. It also supports reading from or writing to a subset of a data set, without loading the entire data file into memory.

Second, utilizing the independence among tasks allows for parallel data processing, thus making full use of multicore or many-core machines. There are mainly two levels of parallelism in typical medical prediction model training. One is between subjects, and the other is within subject. Often, the feature extraction from a patient's data is independent of others. At this level, we can distribute study cases evenly to all computing units. Within each subject, many tasks can also be done simultaneously. For example, features derived from single variables can be calculated on separate cores. Features from moving windows are also highly parallelizable. In the model learning steps, repeated cross-validation is commonly adopted, to test and validate these models' performance on new data and to prevent potential overfitting. A balanced training and testing model prediction is used to see if the model can be generalized to new previously unused data. For example, with multiple combinations of five outcomes, six feature groups, and 10-fold cross-validation repeated 10 times, about 1500–3000 multiples of model calculations and 100–300 model comparisons and statistical tests are required. Parallel training and testing can be used to speed up the learning process.

Many programming languages and scientific data analysis libraries support parallel computing, such as OpenMP (Open Multi-Processing), MPI (Message Passing Interface) for CPU parallelism, and Compute Unified Device Architecture (CUDA) for graphics processing unit (GPU) parallelism. Moreover, nowadays, cloud computing services provide scalable on-demand use of shared computing resources, including CPUs/GPUs, memory, storage, and security. The end users from hospital or medical research institutes can be freed from the burden of building and maintaining expensive high-performance data processing equipment and infrastructures. Giant vendors, such as Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP), are finding more and more applications in medical data mining and machine learning.

24.6 Discussion

An unprecedented volume of data is generated daily in trauma patient care. However, one cruel fact is that healthcare resources are still very limited in both the field and the hospital. Matched blood product, operation rooms, and experienced healthcare providers such as surgeons, anesthesiologists, and nurses are always scarce. The ultimate goal of Big Data in trauma patient care is to intelligently optimize the allocation of limited healthcare resources, by reliable prediction of needs for lifesaving interventions and early decision on therapeutic plans. With automated data processing and informative data aggregation, useful knowledge from massive data can be used by clinicians in a simple way for decision making or prioritizing care in the busy hospital environment.

In this chapter, we reviewed a few critical components in large-scale medical data analysis, including reliable data collection, feature extraction, feature selection, and model evaluation. To distill reliable predictive models from massive medical data, generalizable good performance on unseen data and interpretability are the top two most important factors to consider. Clinical expert knowledge could assist in designing meaningful and useful features that may be associated with outcomes of special interest, such as transfusion, use of operation room, length of stay, and other actionable lifesaving interventions. Moreover, thorough testing with internal and external data provides some evidence of how the models may perform, before the learned models could be deployed.

References

1. Holcomb, John B. "Optimal use of blood products in severely injured trauma patients." *ASH education program book*, no. 1 (2010): 465–469.
2. Perkins, Jeremy G., Alec C. Beekley. "Damage control resuscitation." In: Savitsky E, Eastridge B, eds. *Combat Casualty Care: Lessons Learned from OEF and OIF*. Department of the Army, Office of the Surgeon General, Borden Institute, Washington, DC, 2012: 121–164.
3. Koht, Antoun, Tod B. Sloan, and J. Richard Toleikis. *Monitoring the Nervous System for Anesthesiologists and Other Health Care Professionals*. Springer, New York, NY, USA, 2012.
4. Reddy, Chandan K., and Charu C. Aggarwal, eds. *Healthcare Data Analytics*, vol. 36. CRC Press, Boca Raton, FL, 2015.
5. Palmer, Ronald W. *Integrated diagnostic and treatment devices for enroute critical care of patients within theater*. No. RTO-MP-HFM-182. Army Medical Research And Materiel Command Fort Detrick MD, 2010.

6. DuBose, Joseph J., Gallinos Barmparas, Kenji Inaba, Deborah M. Stein, Tom Scalea, Leopoldo C. Cancio, John Cole, Brian Eastridge, and Lorne Blackbourne. "Isolated severe traumatic brain injuries sustained during combat operations: demographics, mortality outcomes, and lessons to be learned from contrasts to civilian counterparts." *Journal of trauma and acute care surgery 70*, no. 1 (2011): 11–18.
7. Provost, Foster, and Tom Fawcett. "Data science and its relationship to big data and data-driven decision making." *Big data 1*, no. 1 (2013): 51–59.
8. Garg, Amit X., Neill KJ Adhikari, Heather McDonald, M. Patricia Rosas-Arellano, P. J. Devereaux, Joseph Beyene, Justina Sam, and R. Brian Haynes. "Effects of computerized clinical decision support systems on practitioner performance and patient outcomes: A systematic review." *Jama 293*, no. 10 (2005): 1223–1238.
9. Kipnis, Eric, Davinder Ramsingh, Maneesh Bhargava, Erhan Dincer, Maxime Cannesson, Alain Broccard, Benoit Vallet, Karim Bendjelid, and Ronan Thibault. "Monitoring in the intensive care." *Critical care research and practice* 2012 (2012).
10. Schmidt, Paul E., Paul Meredith, David R. Prytherch, Duncan Watson, Valerie Watson, Roger M. Killen, Peter Greengross, Mohammed A. Mohammed, and Gary B. Smith. "Impact of introducing an electronic physiological surveillance system on hospital mortality." *BMJ quality & safety* (2014): bmjqs-2014.
11. Stein, Deborah M., Megan Brenner, Peter F. Hu, Shiming Yang, Erin C. Hall, Lynn G. Stansbury, Jay Menaker, and Thomas M. Scalea. "Timing of intracranial hypertension following severe traumatic brain injury." *Neurocritical care 18*, no. 3 (2013): 332–340.
12. Kahraman, Sibel, Richard P. Dutton, Peter Hu, Lynn Stansbury, Yan Xiao, Deborah M. Stein, and Thomas M. Scalea. "Heart rate and pulse pressure variability are associated with intractable intracranial hypertension after severe traumatic brain injury." *Journal of neurosurgical anesthesiology 22*, no. 4 (2010b): 296–302.
13. Stein, Deborah M., Peter F. Hu, Megan Brenner, Kevin N. Sheth, Keng-Hao Liu, Wei Xiong, Bizhan Aarabi, and Thomas M. Scalea. "Brief episodes of intracranial hypertension and cerebral hypoperfusion are associated with poor functional outcome after severe traumatic brain injury." *Journal of trauma and acute care surgery 71*, no. 2 (2011): 364–374.
14. Stein, Deborah M., Peter F. Hu, Hegang H. Chen, Shiming Yang, Lynn G. Stansbury, and Thomas M. Scalea. "Computational gene mapping to analyze continuous automated physiological monitoring data in neuro-trauma intensive care." *Journal of trauma and acute care surgery 73*, no. 2 (2012): 419–425.
15. Goldberger, Ary L., Luis AN Amaral, Leon Glass, Jeffrey M. Hausdorff, Plamen Ch Ivanov, Roger G. Mark, Joseph E. Mietus, George B. Moody, Chung-Kang Peng, and H. Eugene Stanley. "PhysioBank, PhysioToolkit, and PhysioNet components of a new research resource for complex physiological signals." *Circulation 101*, no. 23 (2000): e215–e220.
16. Kashif, Faisal M., George C. Verghese, Vera Novak, Marek Czosnyka, and Thomas Heldt. "Model-based noninvasive estimation of intracranial pressure from cerebral blood flow velocity and arterial pressure." *Science translational medicine 4*, no. 129 (2012): 129ra44.
17. Li-wei, H. Lehman, Ryan P. Adams, Louis Mayaud, George B. Moody, Atul Malhotra, Roger G. Mark, and Shamim Nemati. "A physiological time series dynamics-based approach to patient monitoring and outcome prediction." *IEEE journal of biomedical and health informatics 19*, no. 3 (2015): 1068–1076.
18. Saeed, Mohammed, Mauricio Villarroel, Andrew T. Reisner, Gari Clifford, Li-Wei Lehman, George Moody, Thomas Heldt, Tin H. Kyaw, Benjamin Moody, and Roger G. Mark. "Multiparameter Intelligent Monitoring in Intensive Care II (MIMIC-II): A public-access intensive care unit database." *Critical care medicine 39*, no. 5 (2011): 952.
19. Excel Medical Electronics, LLC, BedMasterEx Operator's Manual, Jupiter, FL, 2013.
20. Shooman, Martin L. *Reliability of Computer Systems and Networks: Fault Tolerance, Analysis, and Design*. John Wiley & Sons, New York, 2003.
21. Khawaja, Antoun. *Automatic ECG Analysis Using Principal Component Analysis and Wavelet Transformation*. Univ.-Verlag Karlsruhe, 2006.
22. Kamath, Markad V., Mari Watanabe, and Adrian Upton, eds. *Heart Rate Variability (HRV) Signal Analysis: Clinical Applications*. CRC Press, Boca Raton, FL, 2012.
23. Van Drongelen, Wim. *Signal Processing for Neuroscientists: An Introduction to the Analysis of Physiological Signals*. Academic Press, 2006.
24. Brain Trauma Foundation, American Association of Neurological Surgeons; Congress of Neurological Surgeons. Guidelines for the management of severe head injury. *Journal of neurotrauma 24*, Suppl 1 (2007): S1–S106.
25. Sasser, Scott M., Richard C. Hunt, Mark Faul, David Sugerman, William S. Pearson, Theresa Dulski, Marlena M. Wald et al. Guidelines for field triage of injured patients: Recommendations of the National Expert Panel on Field Triage, 2011. *MMWR. Recommendations and reports: Morbidity and mortality weekly report. Recommendations and reports/Centers for Disease Control, 61(RR-1)*, 1–20, 2012.
26. Brown, Joshua B., Gestring, Mark L., Forsythe, Raquel M., Stassen, Nicole A., Billiar, Timothy R., Peitzman, Andrew B., & Sperry, Jason L. "Systolic blood pressure criteria in the National Trauma Triage Protocol for geriatric trauma: 110 is the new 90." *The journal of trauma and acute care surgery 78*, no. 2 (2015): 352.
27. Olaussen, Alexander, Todd Blackburn, Biswadev Mitra, and Mark Fitzgerald. "Shock Index for prediction of critical bleeding post-trauma: A systematic review." *Emergency medicine Australasia 26*, no. 3 (June 2014): 223–228.
28. Kahraman, S., Dutton, R. P., Hu, P., Xiao, Y., Aarabi, B., Stein, D. M., & Scalea, T. M. "Automated measurement of 'pressure times time dose' of intracranial hypertension best predicts outcome after severe traumatic brain

injury." *Journal of trauma and acute care surgery 69*, no. 1 (2010a): 110–118.
29. Mackenzie, Colin F., Wang, Yulei, Hu, Peter F., Chen, Shihyu Y., Chen, Hegang H., Hagegeorge, George,… & ONPOINT Study Group. "Automated prediction of early blood transfusion and mortality in trauma patients." *Journal of trauma and acute care surgery 76*, no. 6 (2014): 1379–1385.
30. Liu, Nehemiah T., John B. Holcomb, Charles E. Wade, Mark I. Darrah, & Jose Salinas. "Utility of vital signs, heart rate variability and complexity, and machine learning for identifying the need for lifesaving interventions in trauma patients." *Shock 42*, no. 2 (2014): 108–114.
31. Ernst, Gernot. *Heart Rate Variability*. Springer, London, 2014.
32. Acharya, U. Rajendra, K. Paul Joseph, Natarajan Kannathal, Choo Min Lim, and Jasjit S. Suri. "Heart rate variability: A review." *Medical and biological engineering and computing 44*, no. 12 (2006): 1031–1051.
33. Task Force of the European Society of Cardiology. "Heart rate variability standards of measurement, physiological interpretation, and clinical use." *European heart journal* 17 (1996): 354–381.
34. Aslanidia, Theodoros. Perspectives of autonomic nervous system perioperative monitoring—Focus on selected tools. *International Archives of Medicine 8*, no. 22 (2015): 1–7.
35. Goodman, Brent, Bert Vargas, and David Dodick. "Autonomic nervous system dysfunction in concussion." *Neurology 80*, no. 7 Supplement (2013): P01–265.
36. Goldstein, Brahm, Mark H. Kempski, Donna BA DeKing, Christopher Cox, David J. DeLong, Mary M. Kelly, and Paul D. Woolf. "Autonomic control of heart rate after brain injury in children." *Critical care medicine 24*, no. 2 (1996): 234–240.
37. Baguley, Ian J., Roxana E. Heriseanu, Kim L. Felmingham, and Ian D. Cameron. "Dysautonomia and heart rate variability following severe traumatic brain injury." *Brain Injury 20*, no. 4 (2006): 437–444.
38. Allen, John. "Photoplethysmography and its application in clinical physiological measurement." *Physiological measurement 28*, no. 3 (2007): R1.
39. Lu, Sheng, He Zhao, Kihwan Ju, Kunson Shin, Myoungho Lee, Kirk Shelley, and Ki H. Chon. "Can photoplethysmography variability serve as an alternative approach to obtain heart rate variability information?" *Journal of clinical monitoring and computing 22*, no. 1 (2008): 23–29.
40. Pan, Jiapu, and Willis J. Tompkins. "A real-time QRS detection algorithm." *IEEE transactions on biomedical engineering* 3 (1985): 230–236.
41. Press, William H., Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery, Chapter 14: Statistical description of data. In: *Numerical Recipes the Art of Scientific Computing*, 3rd ed., Cambridge University Press, 2007.
42. Voss, Andreas, Rico Schroeder, Andreas Heitmann, Annette Peters, and Siegfried Perz. "Short-term heart rate variability—influence of gender and age in healthy subjects." *PloS ONE 10*, no. 3 (2015): e0118308.
43. Yousef, Q., M. B. I. Reaz, and Mohd Alauddin Mohd Ali. "The analysis of PPG morphology: Investigating the effects of aging on arterial compliance." *Measurement science review 12*, no. 6 (2012): 266–271.
44. Elgendi, Mohamed, Ian Norton, Matt Brearley, Derek Abbott, and Dale Schuurmans. "Detection of a and b waves in the acceleration photoplethysmogram." *Biomedical engineering online 13*, no. 1 (2014): 1.
45. Nitzan, Meir, Anatoly Babchenko, Boris Khanokh, and David Landau. "The variability of the photoplethysmographic signal—A potential method for the evaluation of the autonomic nervous system." *Physiological measurement 19*, no. 1 (1998): 93.
46. Voss, Andreas, Matthias Goernig, Rico Schroeder, Sandra Truebner, Alexander Schirdewan, and Hans R. Figulla. "Blood pressure variability as sign of autonomic imbalance in patients with idiopathic dilated cardiomyopathy." *Pacing and clinical electrophysiology 35*, no. 4 (2012): 471–479.
47. Aaslid, Rune, Kari-Fredrik Lindegaard, Wilhelm Sorteberg, and Helge Nornes. "Cerebral autoregulation dynamics in humans." *Stroke 20*, no. 1 (1989): 45–52.
48. Czosnyka, Marek, Piotr Smielewski, Peter Kirkpatrick, David K. Menon, and John D. Pickard. "Monitoring of cerebral autoregulation in head-injured patients." *Stroke 27*, no. 10 (1996): 1829–1834.
49. Balestreri, M., M. Czosnyka, L. A. Steiner, E. Schmidt, P. Smielewski, B. Matta, and J. D. Pickard. "Intracranial hypertension: What additional information can be derived from ICP waveform after head injury?" *Acta neurochirurgica 146*, no. 2 (2004): 131–141.
50. Czosnyka, Marek, and John D. Pickard. "Monitoring and interpretation of intracranial pressure." *Journal of neurology, neurosurgery & psychiatry 75*, no. 6 (2004): 813–821.
51. Steiner, Luzius A., Marek Czosnyka, Stefan K. Piechnik, Piotr Smielewski, Doris Chatfield, David K. Menon, and John D. Pickard. "Continuous monitoring of cerebrovascular pressure reactivity allows determination of optimal cerebral perfusion pressure in patients with traumatic brain injury." *Critical care medicine 30*, no. 4 (2002): 733–738.
52. Cover, Thomas M., and Joy A. Thomas. *Elements of information theory*. John Wiley & Sons, New York, 2012.
53. Richman, Joshua S., and J. Randall Moorman. "Physiological time-series analysis using approximate entropy and sample entropy." *American journal of physiology—Heart and circulatory physiology 278*, no. 6 (2000): H2039–H2049.
54. Lake, Douglas E., Joshua S. Richman, M. Pamela Griffin, and J. Randall Moorman. "Sample entropy analysis of neonatal heart rate variability." *American journal of physiology—regulatory, integrative and comparative physiology 283*, no. 3 (2002): R789–R797.
55. Pincus, Steven M. "Approximate entropy as a measure of system complexity." *Proceedings of the national academy of sciences 88*, no. 6 (1991): 2297–2301.
56. Takens, Floris. "Detecting strange attractors in turbulence." In *Dynamical systems and turbulence, Warwick 1980*, pp. 366–381. Springer, Berlin, 1981.

57. Shmueli, Galit. "To explain or to predict?" *Statistical science* (2010): 289–310.
58. Bishop, Christopher M. *Pattern Recognition and Machine Learning*. Springer, New York, 2006.
59. Lantz, Brett. *Machine Learning with R*. Packt Publishing Ltd, Birmingham, 2013.
60. Murphy, Kevin P. *Machine Learning: A Probabilistic Perspective*. MIT Press, Massachusetts, 2012.
61. Hastie, Trevor J.. Robert John Tibshirani, and Jerome H. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, New York, 2011.
62. Kuhn, Max, and Kjell Johnson. *Applied Predictive Modeling*. Springer, New York, 2013.
63. Tibshirani, Robert. "Regression shrinkage and selection via the lasso." *Journal of the royal statistical society. series B (methodological)* (1996): 267–288.
64. Hastie, Trevor, Robert Tibshirani, and Martin Wainwright. *Statistical Learning with Sparsity: The Lasso and Generalizations*. CRC Press, Boca Raton, FL, 2015.
65. Yuan, Ming, and Yi Lin. "Model selection and estimation in regression with grouped variables." *Journal of the royal statistical society: Series b (statistical methodology) 68*, no. 1 (2006): 49–67.
66. Moody, George B. "PhysioNet: Research resource for complex physiological signals." Available at http://ecg.mit.edu/george/publications/physionet-jecg-2009.pdf, retrieved Jul. 15, 2016.
67. Webman, Rachel B., Elizabeth A. Carter, Sushil Mittal, Jichaun Wang, Chethan Sathya, Avery B. Nathens, Michael L. Nance, David Madigan, and Randall S. Burd. "Association between trauma center type and mortality among injured adolescent patients." *JAMA pediatrics* (2016).
68. Tinkoff, Glen, Thomas J. Esposito, James Reed, Patrick Kilgo, John Fildes, Michael Pasquale, and J. Wayne Meredith. "American Association for the Surgery of Trauma Organ Injury Scale I: Spleen, liver, and kidney, validation based on the National Trauma Data Bank." *Journal of the American college of surgeons 207*, no. 5 (2008): 646–655.
69. Millham, Frederick H., and Wayne W. LaMorte. "Factors associated with mortality in trauma: Re-evaluation of the TRISS method using the National Trauma Data Bank." *Journal of trauma and acute care surgery 56*, no. 5 (2004): 1090–1096.
70. Parimi, Nehu, Peter F. Hu, Colin F. Mackenzie, Shiming Yang, Stephen T. Bartlett, Thomas M. Scalea, and Deborah M. Stein. "Automated continuous vital signs predict use of uncrossed matched blood and massive transfusion following trauma." *Journal of trauma and acute care surgery 80*, no. 6 (2016): 897–906.
71. Fawcett, Tom. "ROC graphs: Notes and practical considerations for researchers." *Machine learning 31*, no. 1 (2004): 1–38.
72. Krzanowski, Wojtek J., and David J. Hand. *ROC Curves for Continuous Data*. CRC Press, Boca Raton, FL, 2009.
73. Davis, Jesse, and Mark Goadrich. "The relationship between Precision-Recall and ROC curves." In *Proceedings of the 23rd international conference on Machine learning*, pp. 233–240. ACM, 2006.