Gleaning Knowledge from Data in the Intensive Care Unit

Michael R. Pinsky[1] and Artur Dubrawski[2]

[1]Department of Critical Care Medicine, University of Pittsburgh, Pittsburgh, Pennsylvania; and [2]Auton Lab, School of Computer, Science, Carnegie Mellon University, Pittsburgh, Pennsylvania

Abstract

It is often difficult to accurately predict when, why, and which patients develop shock, because signs of shock often occur late, once organ injury is already present. Three levels of aggregation of information can be used to aid the bedside clinician in this task: analysis of derived parameters of existing measured physiologic variables using simple bedside calculations (functional hemodynamic monitoring); prior physiologic data of similar subjects during periods of stability and disease to define quantitative metrics of level of severity; and libraries of responses across large and comprehensive collections of records of diverse subjects whose diagnosis, therapies, and course is already known to predict not only disease severity, but also the subsequent behavior of the subject if left untreated or treated with one of the many therapeutic options. The problem is in defining the minimal monitoring data set needed to initially identify those patients across all possible processes, and then specifically monitor their responses to targeted therapies known to improve outcome. To address these issues, multivariable models using machine learning data-driven classification techniques can be used to parsimoniously predict cardiorespiratory insufficiency. We briefly describe how these machine learning approaches are presently applied to address earlier identification of cardiorespiratory insufficiency and direct focused, patient-specific management.

Keywords: big data; machine learning; complexity modeling; functional hemodynamic monitoring

Managing acutely ill patients is often challenging. Bedside assessments of cardiorespiratory status and evolving systemic processes, such as sepsis, acute lung injury, and hemorrhage, are always disguised by the body's own defensive mechanisms, which evolved to sustain life in these stressful states without the benefit of medicine or external life support devices. Thus, the body often hides from the casual observer serious internal pathological processes until they are far advanced. To a large extent, this is very good, because it reflects marshaling of host adaptive and defense processes to sustain homeostasis and enable survival. However, within the setting of acute illness and artificial support environments common to modern medicine and hospital care, it provides a veil of varying transparency that obfuscates the accurate and timely identification of potentially reversible processes until they are well along in their course, and once end-organ and systemic failure develop. Timely detection of pathological processes, as well as rapid identification of proper treatments and verification that those treatments are working as presumed, are the major challenges faced by acute care clinicians. Although new and more powerful diagnostic tools and specific targeted therapies are being developed, there is much that can be ascertained about the nature of a specific patient's physiological status, subliminal pathological processes, and potential responses to therapy that can be gleaned from existing monitoring tools and treatments that can be applied now in the care of this high-risk/high-reward patient population. The approaches reflect using both functional hemodynamic monitoring principles and more sophisticated analyses of existing measured physiologic data streams now reported in the medical record and applied in diagnosis and management in a rudimentary fashion.

Importantly, the body's response to pathologic stress is often formulaic. For example, disease processes associated with both a low cardiac output and impaired cardiovascular reserve, such as hypovolemia (e.g., due to hemorrhage), cardiac pump failure (ischemic cardiomyopathy), and obstruction (pulmonary hypertension), are all associated with increased sympathetic tone-induced vasoconstriction and tachycardia. In contrast, vasodilatory states (e.g., sepsis, hypoglycemia, spinal shock), though also displaying increased sympathetic output causing tachycardia, have reduced peripheral vascular responsiveness. Thus, these disease states manifest themselves as decreased vasomotor tone, increased

(Received in original form April 21, 2014; accepted in final form July 28, 2014)

Supported, in part, by National Institutes of Health grants HL67181 and NR013912, and by National Science Foundation awards 0911032 and 1320347.

Correspondence and requests for reprints should be addressed to Michael R. Pinsky, M.D., 606 Scaife Hall, 3550 Terrace Street, Pittsburgh, PA 15261. E-mail: pinskymr@upmc.edu

Am J Respir Crit Care Med Vol 190, Iss 6, pp 606–610, Sep 15, 2014
Copyright © 2014 by the American Thoracic Society
Originally Published in Press as DOI: 10.1164/rccm.201404-0716CP on July 28, 2014
Internet address: www.atsjournals.org

unstressed circulatory blood volume, and, in acute inflammatory states (e.g., due to sepsis), vascular endothelial injury, causing increased plasma translocation into the interstitial spaces. Classic physiology training underscores these differences, as quantified by hemodynamic monitoring and blood analysis, as central to the diagnosis and treatment of circulatory shock (1). Unfortunately, once circulatory shock has resulted in inadequate tissue perfusion manifest by hyperlactecemia and metabolic acidosis, end-organ dysfunction and injury are already occurring. Thus, waiting for obvious signs of circulatory shock to occur before starting corrective therapies exacerbates the risks of prolonging end-organ injury. Such delayed therapy is often associated with a greater need for increased intensity of therapy, with its own set of possible complications. All in all, identifying patients earlier in their disease state and reversing those pathological processes before they become resistant to therapy and induce end-organ injury must reduce morbidity, cost, and mortality.

Three levels of aggregation of information can be used to aid the bedside clinician in this task (Table 1). One is the analysis of derived parameters of existing measured variables using simple bedside calculations. The second is using prior physiologic data of similar subjects during periods of stability and disease to define quantitative metrics of level of severity. The third approach is to use libraries of responses across large and comprehensive collections of records of diverse subjects whose diagnosis, therapies, and course is already known to predict not only disease severity, but also the subsequent behavior of the subject if left untreated or treated with one of the many therapeutic options. This final approach is probably beyond the intellectual scope of all but the most experienced and astute observers. However, all three levels of integration are at our disposal, and they can be used to diagnose impending cardiorespiratory instability, determine the likely etiological process, and simultaneously determine the best therapies and monitor response to therapy.

Table 1. Levels of Real-Time Data Integration to Derive Clinically Relevant Information

1. Physiologic time series analysis based on changes in measured variables to a defined perturbation (Functional Hemodynamic Monitoring) (19)
   - Changes in CVP with spontaneous inspiration
     - Decreasing CVP: volume responsive
     - Increasing CVP: right heart failure, tamponade, pulmonary hypertension
   - Changes in coupled arterial pulse pressure and left ventricular stroke volume with breathing or arrhythmias (PPV/SVV)
     - Normal ventriculo-arterial coupling 1.2–2
     - Increased vasomotor tone > 2
     - Vasodilation < 0.9
   - Changes in cardiac output with passive leg raising
     - If $>15\%$ volume responsive
     - If $<15\%$ not volume responsive
   - PPV or SVV during positive pressure ventilation
     - If PPV $>13\%$ or SVV $>10\%$ volume responsive
   - Dynamic changes in St_{O_2} induced by a vascular occlusion test
     - St_{O_2} deoxygenation slope is decreased impaired cardiovascular regulation
       - Normal value thenar eminence: -12.8 (-16.2, -11.3) % min^{-1}
     - St_{O_2} reoxygenation slope is decreased impaired cardiovascular reserve
       - Normal value thenar eminence: 297.2 (213.7, 328.6) % min^{-1}
2. Fused parametric measures of multiple physiologic variables
   - APACHE (10), LODS (20), SOFA (21), etc.
     - All these scores fuse chronic disease, age, organ function (as assessed by laboratory values) with bedside physiologic measures to derive a severity of illness score and prediction of future need for acute care
   - MEWS (22) delayed cardiorespiratory instability prediction
     - Requires bedside clinician to hand enter data for calculation
   - Vital sign index on real time cardiorespiratory instability
     - Fused parameter of 5-min averaged values for heart rate, respiratory rate, Sp_{O_2} and intermittently measured (dyssynchronous) blood pressure
3. Machine learning–based pattern recognition
   - Heart rate variability
     - Predicts sepsis in neonates (23)
   - Not available commercially multiple variable bedside prediction systems
   - Potential machine learning–based tools for multifactor analysis, pattern discovery, and display for real-time monitoring, event detection, event forecasting, and tracking
     - Parametric and nonparametric classification and regression
       - Support vector machines
       - K nearest neighbors
       - Random forests
     - Null-space, state space and clustering models, probabilistic graphical models, and spectral methods
       - Principal component analysis
       - Gaussian process models
       - Markov random fields
       - Hierarchical clustering
       - K-means
       - Gaussian mixture models

Definition of abbreviations: APACHE = Acute Physiology and Chronic Health Evaluation; CVP = central venous pressure; LODS = Logistic Organ Dysfunction Score; MEWS = Modified Early Warning System; PPV = pulse pressure variation; SOFA = Sepsis-Related Organ Failure Assessment; Sp_{O_2} = oxygen saturation as measured by pulse oximetry; St_{O_2} = tissue O_2 saturation; SVV = stroke volume variation.

Determining Clinically Relevant Physiological Parameters from Time Series Analysis

Clinicians routinely use derived physiological variables in their assessment of patient status, although often without full understanding of their determinants. Measures of mean airway pressure, central venous pressure (CVP), and mean arterial pressure are routinely reported. However, mean airway pressure is a function of multiple interacting processes, including airway resistance, lung and chest wall compliance, tidal volume, positive end-expiratory pressure, inspiratory flow pattern, spontaneous respiratory efforts, and artificial airway characteristics; yet it is still used effectively as a primary target in pressure-limited ventilatory strategies. CVP is a function of pericardial restraint, right ventricular diastolic and systolic function, vasomotor tone, blood volume, and pulmonary artery pressure. Still, dynamic changes in CVP during spontaneous breathing can identify volume responsiveness (2) or right ventricular failure by its associated decrease or increase (Kussmaul's sign), respectively, during inspiration.

The systemic arterial pressure waveform also has many determinants, including left ventricular (LV) stroke volume (SV), heart rate, central arterial elastance, vasomotor tone, and cardiac contractility. Many physiologic parameters can be acquired from the arterial waveform analysis. Mean arterial pressure is calculated as the diastolic arterial pressure plus one-third the pulse pressure (PP; diastole to the subsequent systole). Mean arterial pressure is usually considered as the input pressure for all organs of the body, except the heart, which uses diastolic pressure as its input pressure. A low diastolic pressure can coexist with a nonhypotensive mean arterial pressure if PP is high enough, so measuring diastolic, mean, and systolic pressures are all important in assessing organ perfusion pressure. A low diastolic pressure usually connotes low vasomotor tone. PP is often used as a surrogate for LV SV, but PP is also determined by central arterial elastance and cardiac contractility. The body's response to low-flow states is to increase sympathetic tone, causing both heart rate and peripheral vasomotor tone to increase, such that cardiac output is often maintained, but LV SV decreases. Because vasomotor tone increases, PP must increase for the same SV. Traumatologists use the PP-to-SV ratio as a shock index, increasing PP/SV values reflecting increased levels of circulatory stress (3). However, even greater insight comes from examining the dynamic changes in both PP and SV, as they occur during breathing or with arrhythmias. The ratio of PP variation (PPV) to paired SV variation (SVV) defines a dynamic arterial elastance, a fundamental characteristic of the central circulation. Normal dynamic elastance varies between 2 and 1.2, whereas values under 0.9 reflect profound vasodilation (4). Thus, in a hypotensive patient who is volume responsive, if the PPV/SVV is less than 0.9, then giving fluids alone will not increase arterial pressure, although cardiac output will increase. Those subjects also need primary vasopressor therapy to sustain an effective organ perfusion pressure. Furthermore, in patients with a stable heart rate and on mechanical ventilation, a PPV greater than 13% or an SVV over 10% was shown to identify those who are volume responsive (5). Although patients with PPV and SVV values less than these threshold levels may also be volume responsive, most patients with values above these thresholds will be volume responsive. Important caveats limiting the generalizable use of PPV and SVV exist, including the presence of spontaneous breathing activity, severe cor pulmonale, markedly increased chest wall stiffness, and intra-abdominal hypertension; however, as a general guiding principle, these measures can be applied in most patients requiring resuscitation. Finally, near-infrared spectroscopy measures of tissue O_2 saturation are not useful in identifying cardiovascular stress until it is well advanced; however, when coupled to a standardized vascular occlusion test that allows objective quantification of local O_2 desaturation and reoxygenation rates, they clearly identify early circulatory stress (6), subsequent response to resuscitation efforts (7), and those critically ill patients who will fail a weaning trial (8) or develop organ failure (9).

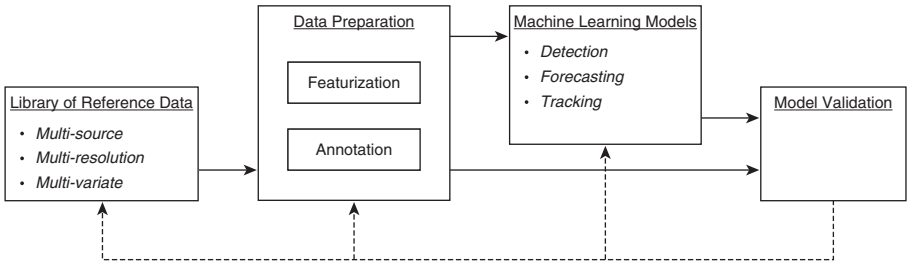

Figure 1. Flow diagram of a typical operation of the model learning phase of a machine learning–based pattern-recognition system. The available library of clinical data is first prepared for training. Streams of potentially diverse data types received from multiple sources—bedside, clinical records, medical history, patient demographics, etc., represented at varied temporal resolutions and recencies—are featurized using any combination of expert rules, statistical characterization, data mining, and anomaly detection techniques to extract and characterize potentially informative patterns in data. Occurrences of the events of interest (e.g., episodes of instability) are adjudicated by expert clinicians and annotated to serve as training examples. Annotated featurized data are then processed by machine learning algorithms to produce reliable models for particular tasks of clinical relevance (e.g., adverse event detection, forecasting instability, or tracking response to treatment). Resulting models are often empirically validated using set-aside test data sets, and further validated by expert clinicians. Results of validation can be used as feedback (*dashed lines* in the flow diagram) to tune structures of the models themselves, as well as to improve feature extraction and annotation processes (self-diagnostic algorithms, known as active learning, are often used to identify particularly informative yet unlabeled incidents for expert annotation), and they can also be used to inform improvements in the source data acquisition procedures, and help address data quality issues.

Fused Parametric Measures to Define Physiologic State

The use of fused parameters to assess patient disease severity is not new. For example, the Acute Physiology and Chronic Health Evaluation (APACHE) scoring system has evolved into an accurate means to predict in-hospital mortality (10). Multiple other scoring systems have been developed and validated. However, can these approaches be used proactively to identify patients at risk for instability or with lower levels of instability before overt organ system failure? Numerous

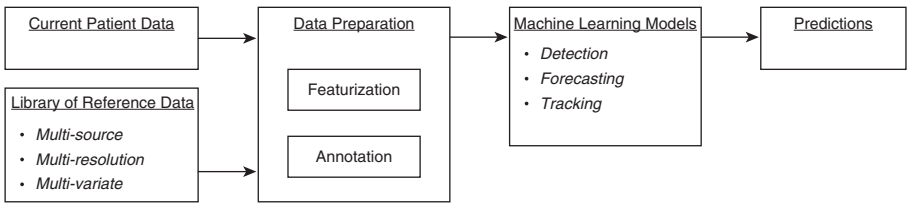

Figure 2. Flow diagram of a typical operation of the performance phase of a machine learning–based pattern-recognition system. Current observations of the monitored patient are featurized in the same way as the data from the reference library (access to which, depending on the types of machine learning models in use, may be required), and processed by the previously trained models to produce predictions.

studies have documented that critical physiologic changes are seen in 51–86% of patients who suffered a subsequent cardiopulmonary arrest on general wards, often several hours before the arrest event (11). Clearly, improved detection is critical in impacting the potential to rescue of unstable patients, and earlier detection can improve outcomes. Using a data fusion approach based on a review of 8 weeks of naive, noninvasive vital sign data collection in a step-down unit, Hravnak and colleagues (12) proposed a biosign index to score instability based on algorithms developed by Tarassenko and colleagues (13). When coupled to a nurse-driven alert and management protocol, this real-time bedside fused parameter alert system reduced significant physiologic insufficiency by 400% over the subsequent 8-week interval, and markedly reduced nursing time needed to treat these patients. Such approaches can also be easily crafted for other acute care environments.

Machine Learning–based Pattern Recognition

The next level of the physiologic data stream analysis combines the two above approaches into a single process (14). Machine learning methodology scales up correlational analyses to potentially very highly multivariate, high-frequency, and perhaps multisource data that could help empirically discover leading indicators of instability (Figure 1). To generate potentially informative change points in monitored hemodynamic waveforms, Lonkar and colleagues (15) propose to segment each high-frequency measurement channel into sequences of nonoverlapping time intervals, and to independently characterize waveform frequency spectrum in each segment. Principal component analysis is then applied to the spectra observed during a specific physiologic state (e.g., periods of stability) to envelope the range of variability of waveform patterns that can be expected in the particular state. The resulting null hypothesis model serves as a reference for extracted spectra of hemodynamic data observed in a monitored patient. Any substantial departure of waveform spectra from the range of expected variability is then treated as a potential leading indicator of the ongoing, emerging, or ensuing episodes of hemodynamic instability. Saria and colleagues (16) detect motifs (i.e., patterns sharing a specific shape [with some variation]) in hemodynamic waveform data. They show that some of the detected signal templates can be predictive of instability in neonatal intensive care unit infants (prediction of morbidity, 86% sensitive at 96% specificity), and better performance for particular complications (infection, 90% at 100%; cardiopulmonary, 96% at 100%). Various data-mining algorithms can be used to identify which of the potentially large number (literally thousands) of patterns detected in data systematically coincide or precede particular events of interest (Figure 2). In their preliminary study, Lonkar and colleagues (15) applied probabilistic co-occurrence analysis to isolate such promising indicators that significantly correlate with the near-future tachycardia episodes that occur between tens of minutes to a few hours in the future, and report a recall rate of 85% at less than 5% false discovery rate, and an area under the receiver operating characteristic curve score of 0.857. Large numbers of anomalous patterns and events that can be gleaned from waveform data, which could potentially carry useful information, may make their clinical interpretation a daunting task. Guillame-Bert and colleagues (17) propose applying temporal rule learning methodology to extract human-interpretable logical statements expressing the most predictive combinations of patterns and states that can asynchronously appear in multivariate clinical data, irrespective of temporal resolution of their observation (high-frequency waveforms, heartbeat to heartbeat, breath to breath, clinical records, demographics, etc.). The authors report multifold improvements of recall of instability episodes minutes ahead of their onset, even though the training data used to build their model was very sparse (only 130 annotated periods of instability in an almost 1.5 patient-year stream of vital sign data) when compared with using single vital sign–based or fused metrics. The machine learning framework can also be used to probabilistically track evolution of critically ill patients. In their preliminary work, Pimentel and colleagues (18) used Gaussian process regression to represent vital sign trajectories estimated from unevenly sampled vital sign data. This representation enables clustering and classification of the specific patient trajectories to help assess current status and expected outcomes for the monitored individuals.

Advanced analytic solutions are readily available and can be applied to process existing clinical data streams in real time to derive meaningful, specific, and clinically relevant information, which, in turn, can be operationalized for improved patient care and outcomes. The basic tools to accomplish these tasks extend beyond simple data review and routine clinical care as is presently done, but they are within our grasp and are limited only by our imagination. ■

Author disclosures are available with the text of this article at www.atsjournals.org.

References

1. Vincent JL, De Backer D. Circulatory shock. *N Engl J Med* 2013;369: 1726–1734.
2. Magder S. Central venous pressure: a useful but not so simple measurement. *Crit Care Med* 2006;34:2224–2227.
3. Rady MY, Smithline HA, Blake H, Nowak R, Rivers E. A comparison of the shock index and conventional vital signs to identify acute, critical illness in the emergency department. *Ann Emerg Med* 1994;24: 685–690.
4. Monge García MI, Gil Cano A, Gracia Romero M. Dynamic arterial elastance to predict arterial pressure response to volume loading in preload-dependent patients. *Crit Care* 2011; 15:R15.
5. Michard F, Boussat S, Chemla D, Anguel N, Mercat A, Lecarpentier Y, Richard C, Pinsky MR, Teboul J-L. Relation between respiratory changes in arterial pulse pressure and fluid responsiveness in septic patients with acute circulatory failure. *Am J Respir Crit Care Med* 2000;162:134–138.
6. Creteur J, Carollo T, Soldati G, Buchele G, De Backer D, Vincent JL. The prognostic value of muscle StO_2 in septic patients. *Intensive Care Med* 2007;33:1549–1556.
7. Gómez H, Torres A, Polanco P, Kim HK, Zenker S, Puyana JC, Pinsky MR. Use of non-invasive NIRS during a vascular occlusion test to assess dynamic tissue O(2) saturation response. *Intensive Care Med* 2008;34:1600–1607.
8. Gruartmoner G, Mesquida J, Masip J, Martínez ML, Villagra A, Baigorri F, Pinsky MR, Artigas A. Thenar oxygen saturation during weaning from mechanical ventilation: an observational study. *Eur Respir J* 2014; 43:213–220.

9. Mesquida J, Espinal C, Gruartmoner G, Masip J, Sabatier C, Baigorri F, Pinsky MR, Artigas A. Prognostic implications of tissue oxygen saturation in human septic shock. *Intensive Care Med* 2012;38: 592–597.
10. Knaus WA, Wagner DP, Draper EA, Zimmerman JE, Bergner M, Bastos PG, Sirio CA, Murphy DJ, Lotring T, Damiano A, *et al*. The APACHE III prognostic system: risk prediction of hospital mortality for critically ill hospitalized adults. *Chest* 1991;100:1619–1636.
11. Naeem N, Montenegro H. Beyond the intensive care unit: a review of interventions aimed at anticipating and preventing in-hospital cardiopulmonary arrest. *Resuscitation* 2005;67:13–23.
12. Hravnak M, Devita MA, Clontz A, Edwards L, Valenta C, Pinsky MR. Cardiorespiratory instability before and after implementing an integrated monitoring system. *Crit Care Med* 2011;39:65–72.
13. Tarassenko L, Hann A, Young D. Integrated monitoring and analysis for early warning of patient deterioration. *Br J Anaesth* 2006;97:64–68.
14. Pinsky MR. Complexity modeling: identify instability early. *Crit Care Med* 2010;38(10 suppl):S649–S655.
15. Lonkar R, Dubrawski A, Fiterau M, Garnett R. Mining intensive care vitals for leading indicators of adverse health events. *Emerg Health Threats J* 2011;4:11073.
16. Saria S, Rajani AK, Gould J, Koller D, Penn AA. Integration of early physiological responses predicts later illness severity in preterm infants. *Sci Transl Med* 2010;2:48ra65.
17. Guillame-Bert M, Dubrawski A, Chen L, Hravnak M, Pinsky MR, Clermont G. Learning temporal rules to forecast instability in intensive care patients. *Intensive Care Med* 2013;39:S470.
18. Pimentel MAF, Clifton DA, Tarassenko L. Gaussian process clustering for the functional characterisation of vital-sign trajectories: proceedings of the 2013 IEEE International Workshop on Machine Learning for Signal Processing, Southampton, UK, 2013.
19. Pinsky MR, Payen D. Functional hemodynamic monitoring. *Crit Care* 2005;9:566–572.
20. Le Gall JR, Klar J, Lemeshow S, Saulnier F, Alberti C, Artigas A, Teres D; ICU Scoring Group. The Logistic Organ Dysfunction system: a new way to assess organ dysfunction in the intensive care unit. *JAMA* 1996;276:802–810.
21. Vincent J-L, Moreno R, Takala J, Willatts S, De Mendonça A, Bruining H, Reinhart CK, Suter PM, Thijs LG. The SOFA (sepsis-related organ failure assessment) score to describe organ dysfunction/failure. *Intensive Care Med* 1996;22:707–710.
22. Subbe CP, Kruger M, Rutherford P, Gemmel L. Validation of a modified early warning score in medical admissions. *QJM* 2001;94:521–526.
23. Moorman JR, Carlo WA, Kattwinkel J, Schelonka RL, Porcelli PJ, Navarrete CT, Bancalari E, Aschner JL, Whit Walker M, Perez JA, *et al*. Mortality reduction by heart rate characteristic monitoring in very low birth weight neonates: a randomized trial. *J Pediatr* 2011; 159:900–906, e1.