

Towards Non-Invasive Monitoring of Hypovolemia in Intensive Care Patients *

Alexander Roederer
Dept. of Computer &
Information Sc.
University of Pennsylvania
roederer@seas.upenn.edu

James Weimer
Dept. of Computer &
Information Sc.
University of Pennsylvania
weimerj@seas.upenn.edu

Joseph Dimartino
Penn Presbyterian Medical Center
University of Pennsylvania Health System
joseph.dimartino@uphs.upenn.edu

Jacob Gutsche
Penn Presbyterian Medical
Center
University of Pennsylvania
Health System and the
Perelman School of Medicine
jacob.gutsche@uphs.upenn.edu

Insup Lee
Dept. of Computer &
Information Sc.
University of Pennsylvania
lee@cis.upenn.edu

ABSTRACT

Hypovolemia caused by internal hemorrhage is a major cause of death in critical care patients. However, hypovolemia is difficult to diagnose in a timely fashion, as obvious symptoms do not manifest until patients are already nearing a critical state of shock. Novel non-invasive methods for detecting hypovolemia in the literature utilize the photoplethysmogram (PPG) waveform generated by the pulse-oximeter attached to a finger or ear. Until now, PPG-based alarms have been evaluated only on healthy patients under ideal testing scenarios (e.g., motionless patients); however, the PPG is sensitive to patient health and significant artifacts manifest when patients move. Since patient health varies within the intensive care unit (ICU) and ICU patients typically do not remain motionless, this work introduces a PPG-based monitor designed to be robust to waveform artifacts and health variability in the underlying patient population. To demonstrate the promise of our approach, we evaluate the proposed monitor on a small sample of intensive care patients from the Physionet database. The monitor detects hypovolemia within a twelve hour window of nurse documentation of hypovolemia when it is present, and achieves a low false alarm rate over patients without documented hypovolemia.

Categories and Subject Descriptors

G.3 [Probability and Statistics]: Time series analysis;
J.3 [Life and Medical Sciences]: Medical Information
Systems

1. INTRODUCTION

Patients who present to emergency rooms with trauma, patients undergoing surgery, and post-operative patients in intensive care units frequently experience hemorrhage. Persistent internal hemorrhage can, over time, cause a decrease in the volume of blood in the circulatory system, a condition known as *hypovolemia* [12].¹ Hypovolemia is common

*This research was supported in part by NSF CNS-1035715. Copyright retained by the authors

¹Hypovolemia can also be caused by loss of body fluids, or by

among post-operative patients, and bleeding-related complications are a major cause of prolonged length of stay and death in hospitals [29, 18].

After 15-30% of total blood volume is lost [10],² the patient is at high risk of experiencing hypovolemic shock. Shock occurs when the heart can no longer pump enough blood to the body; blood pressure plummets, and tissue perfusion is not capable of sustaining aerobic metabolism, leading to widespread ischemic tissue damage [12]. Hypovolemic shock can rapidly be fatal, and is a complication that arises often in at-risk patients [18].

Unfortunately, assessment of decreasing volume status is one of the most difficult tasks in clinical medicine [15]. At the onset of hypovolemia, a patient's body will attempt to compensate by increasing heart rate (to continue to deliver oxygen to the extremities with the diminished amount of blood). However, heart rate alone is not sufficient to determine the need for emergent interventions [4], as it is specific. Respiratory rate may rise [12], but this is often a late change [8], and can be masked by mechanical ventilation. Blood vessels in the periphery constrict, which decreases perfusion in the extremities [32], causing subtle changes in patient visage, such as paleness, long capillary refill times, or poor skin turgor. However, these changes are easy for clinicians to overlook, and no definitive studies on the predictive ability or pathogenesis of these signs seem to exist [16]. Thus, the body's hemodynamic compensation mechanisms can mask underlying deterioration the patient is experiencing.

Fluid management is an important part of patient management in critical care, in part to mitigate the risk of hypovolemia [11]. Patient fluid inputs and outputs are closely monitored for changes [2], as common medical practice holds that these changes may reflect changes in blood volume. These changes only occur, however, after significant blood is lost [28]. Blood pressure is obfuscated by numerous bodily

external bleeding. However, both of these causes are much more apparent to clinicians when they occur, and rarely go undiscovered. Thus, we do not focus on these cases.

²The percentage varies with the patient's physiological compensatory mechanisms; healthier patients are more readily able to compensate for loss in blood volume.

compensatory mechanisms as well as anesthetics and other medication, which makes monitoring blood pressure unsuitable for ascertaining if the patient’s blood volume is changing [7]. Change in heart rate is often subtle and, depending on the size of the bleed, the change may present over a number of hours. Periodic monitoring of heart rate may not capture this trend, and there are numerous confounding etiologies for tachycardia [4, 7].

Detection is made still more challenging by inter-patient variability. Blood volume varies from patient to patient, and the way the body responds to blood loss from internal bleeding varies based on the severity and location of the bleed. Administration of IV fluids is often used as a first-line therapy under suspicion of hypovolemia, in an attempt to improve cardiac output [25, 14]. However, as previously described, once symptoms are apparent, the patient has likely lost significant amounts of blood and may have progressed to a damaging state of shock.

The proliferation of inexpensive digital sensing technology has allowed continuous monitoring of many vital signs to become the standard of care in many major ICUs [30]. Continuous monitoring has been shown to decrease length of stay and improve patient outcomes [5]. Commonly monitored physiological data, specifically the photoplethysmogram (PPG) waveform, may be used to continuously monitor patients for hypovolemia by estimating their blood volume, providing earlier and more accurate detection. Several such clinical monitoring systems have been developed over homogenous populations of test patients in controlled environments [7, 28, 14, 21]. In this work, we present a detector for hypovolemia which can be run continuously using only PPG. Over a set of five ICU patients (two with hypovolemia, three without) the detector alerts for hypovolemia within a twelve hour window of documented hypovolemia in patients who had it, and with only two false alarms over three and a half days of monitoring non-hypovolemic patients.

The remainder of this paper is organized as follows. Section 2 describes related work that motivates our approach. Section 3 provides a more detailed description of the problem scenario. Section 4 describes a basic physiological model which captures how the body compensates for blood loss over time, and an explanation of the PPG signal. In Section 5 we introduce the parameter-invariant detector, which utilizes the PPG waveform to detect changes from non-hemorrhagic to hemorrhagic states. Section 6 describes how we selected a small subset of hypovolemic and non-hypovolemic patients from the Physionet MIMIC II database [22, 9], and presents performance results of the statistic in detecting hypovolemia in these two groups of patients. We conclude in Section 7 with a summary and description of future work.

2. RELATED WORK

As prevention of hypovolemic shock is a serious and widespread ailment, much work has been done to both understand the etiology of the condition, and to attempt to detect and treat it.

The pathophysiology of hypovolemia has been described extensively in the medical literature [26, 16], though precise definitions are difficult to establish. Increase in pulse rate coupled with severe dizziness or fluid input/output mismatch are the major distinguishing characteristics proven to have predictive ability for the condition. These changes are common, in particular once large amounts of blood have been lost, but are often absent after moderate amounts of

blood loss [28].

Many changes commonly taken as signs of hypovolemia have been shown to be aspecific. Decrease in systolic blood pressure is aspecific over small quantities of lost blood, and there are insufficient studies to calculate its effectiveness at predicting higher volumes of blood loss. Non-severe dizziness is also poorly sensitive [16]. Most other signs (skin turgor, capillary refill, supine tachycardia) have no proven individual diagnostic role [28]. These papers also make the important point that patients with dehydration show similar operating characteristics to those with fixed-volume blood loss [16, 28, 12]. The optimal method of resuscitation has not been clearly established and depends on the severity of the blood loss [12].

Because conventional hemodynamic variables have proven unreliable as predictors, there have been a number of efforts to develop new predictive measures which can aid in identification of hypovolemia.

In an early example, Chen et al. [6] created a decision tool for the identification of hypovolemia in trauma patients during helicopter transport (when reliable acquisition of vital sign data is difficult). They combined linear classifiers using ensemble methods to form a robust classifier. Though the patient cohort (trauma patients) is distinct from the post-surgical population we focus on, the work demonstrates that hypovolemia can be detected through automated analysis of vital signs.

Birkhahn et al. [3] showed that calculated shock index (a ratio of heart rate to systolic blood pressure) was significantly elevated in healthy volunteers who donated blood. Heart rate and systolic blood pressure underwent significant changes, but were within “normal” limits, which means they may easily be overlooked.

Marik et al. [15] review how dynamic changes in pressure and stroke volume in patients undergoing mechanical ventilation have emerged as useful techniques to assess volume responsiveness. In particular, dynamic changes of arterial waveform-driven variables are highly accurate in predicting fluid responsiveness in critically ill patients. However, this technique was limited to patients who were under mechanical ventilation.

Pizov et al. [21] showed that respiratory variations in arterial pressure and photoplethysmogram (PPG) waveforms correlated with minor blood loss in anesthetized patients, and that the pulse oximetry PPG waveforms accurately reflected arterial waveforms during progressive hypovolemia. Loupec et al. [14] demonstrated that a variability index based on PPG predicts fluid responsiveness in patients under mechanical ventilation with circulatory insufficiency. Zöllei et al. [32] showed that mild and moderate central hypovolemia caused clinically relevant falls in pulse rate and stroke volume. Patients with slow breathing experienced greater changes, suggesting that higher respiration rates may be compensatory mechanism for central hypovolemia.

Most recently, Stewart, Mulligan, Grudic, Convertino, and Moulton [28] proposed using the PPG to continuously estimate a compensatory reserve index (CRI), a measure of the reserve remaining to compensate for reduced central blood volume. This index was shown to detect low-volume blood loss with greater specificity than more traditional physiologic measures.

A landmark paper by Convertino et al. [7] collected data on patients undergoing simulated blood loss through lower body negative pressure devices, which have been shown to

mimic reductions in central blood volume caused by hemorrhage. The authors used machine learning techniques to create a detector able to determine whether a patient was experiencing hypovolemia with high accuracy. They used multiple vital signs, including PPG, to calculate an estimate of the actual reduction in blood volume being experienced by the patient.

The major limitation of most of these efforts was that their data analysis was performed on controlled sets of relatively healthy patients without significant complications or comorbidities, whereas patients in critical care settings often have more complications. Our proposed approach extends the discussed work, and is designed to be robust to non-uniformity in the patient population and non-uniformity in the rate of blood lost.

Significant work exists in the machine learning literature on learning patient-specific classifiers; however, these techniques usually require large amounts of labeled data from clear, discrete classes. Such a data set is difficult to come by for hypovolemia (as time of onset often isn't clear, and comorbidities are frequent). An alternative approach to medical event detection that can be designed robust to patient variations is the constant false alarm rate (CFAR) detector [23, 31]. Utilizing physiological models, CFAR detectors can generate sufficient statistics that are maximally invariant to unknown variability such that a threshold test achieves a constant false alarm rate under optimal conditions [13].

3. PROBLEM FORMULATION

This section describes the deficiencies in the current approach for detecting hypovolemia in intensive care patients, and presents a clear statement of the problem considered in this work.

3.1 Current Approach

In the ICU, the current approach to detecting hypovolemia involves closely monitoring patients' heart rate, blood pressure, respiration rate, and visual status [20, 28]. An elevated heart rate coupled with dehydration and changes in the patient's skin condition may rouse suspicion of hypovolemia. If these signs are present and a patient's blood pressure is low, fluids are typically administered [19, 27]. Patients who respond to fluids by returning to more normal blood pressures are considered likely to have been suffering from hypovolemia, but the diagnosis is often left unclear.

3.2 Problem Statement

Given common non-invasive physiological signals monitored from a patient, we aim to develop an algorithm for early detection of hypovolemia that is robust to movement artifacts and differing patient physiology. Specifically, we wish to (a) develop a statistic over the PPG waveform which is invariant to common artifacts, and (b) design a hypovolemia monitor which employs the PPG waveform statistic which is invariant to irrelevant differences in patient physiology.

4. PHYSIOLOGY AND SENSING

In order to motivate choices in the monitor design described in Section 5, in this section we present a description of a high-level physiological model for hypovolemia (on which the monitor's statistic is based). We also present a



Figure 1: A pulse oximeter displaying a PPG waveform.

description of how the PPG signal (over which the statistic is run) is measured and what it captures.

4.1 Physiological Model

The mechanism behind hypovolemia involves a reaction to a drop in oxygenation caused by reduced blood volume. As the volume of blood circulating in the body decreases, the amount of oxygen reaching the control centers of the brain goes down. In an attempt to maintain perfusion to the vital organs, the brain engages hemodynamic compensation mechanisms whereby the heart rate increases (to pump more blood more quickly through the system), respiration rate increases (to deliver more oxygen to the blood), and peripheral vasoconstriction occurs, along with increased sympathetic nerve activity [17].

4.2 PPG Sensor

Photoplethysmography (PPG) is an optical measurement technique most commonly used within a pulse oximeter (as seen in Figure 1) to detect blood volume changes in the microvascular bed of tissue. The PPG has seen widespread clinical application, as it is noninvasive and can be used to measure a number of different aspects of cardiovascular function, most notably pulse rate and tissue oxygenation [1]. However, the signal also contains information about vascular distensibility, cardiac arrhythmia, systolic blood pressure, and respiratory variability [24], and, as described previously, is sensitive to changes in blood volume. Thus, PPG can provide information well suited for use in detection of hypovolemia [17].

At their simplest, devices (such as the pulse oximeter) which measure PPG contain a light emitting diode and an optical sensor. Light is emitted into flesh and either reflected off bone and back to the sensor, or is transmitted directly through flesh and into the sensor. The amount of light reabsorbed by the sensor is impacted by the scattering, absorption, reflection, transmission and fluorescence of biological tissue. In particular, recorded pulses bear a direct relationship with perfusion, as larger blood volumes produce larger attenuation in the light source [1]. The signal can be split into an *AC* component (the pulsatile component, linked to heart rate) superimposed onto a large *DC* component. While PPGs provide large amounts of information about a patient, they can be difficult to use because they often contain large amounts of artifact. The PPG sensor is very sensitive to movement and orientation against the skin as small shifts can significantly impact the measured intensity of light. These variations can lead to inaccurate data.

5. MONITOR DESIGN

In order to obtain a robust detector, this work employs

parameter invariant statistics to build CFAR detectors as described in [23, 31]. The following two sections describe the two steps for building the detector: processing the waveform to remove sections dominated by artifact, and building the CFAR detector.

5.1 PPG Waveform Processing

PPG signals often contain segments where any information in the signal is dominated by artifact introduced into the signal by a mis-aligned sensor or patient motion. Figure 2 shows a PPG signal. Consistent with previous work [7, 28, 17], the detector described in the subsequent discussion utilizes a sampled mean of the PPG signal over a 60 second window. In data that is not corrupted by artifact, the dominant non-DC frequencies of the PPG waveform correspond to the fundamental frequency of the heart rate and its harmonic frequencies, as illustrated in Figure 2a. However, when corrupted by artifact, there is significant energy in the frequencies not related to heart rate, and the heart rate’s harmonic frequencies are undetectable, as illustrated in Figure 2b. Thus, to test whether a specific range of data should be included or not invariant to the heart-beat, we test the energy ratio between the harmonic frequencies with the two greatest energies. If the largest frequency dominates the second most dominant, we include the corresponding DC value in the average DC value for the window; otherwise it is discarded.

5.2 CFAR Detector

Over a series of values extracted from the DC signal, we now build a parameter-invariant statistic to distinguish between signal patterns indicative of hypovolemia and those which are not indicative of hypovolemia. We aim to construct models to test for hypovolemia by observing that the PPG signal of a non-hypovolemic patient drifts naturally over time, while a hypovolemic patient experiences a decrease in PPG [7]. To capture the effects of these hypotheses we establish a null hypothesis H_0 representing a non-hypovolemic patient, in which the patient’s DC signal experiences simple Brownian motion, and an alternative hypothesis H_1 , in which a patient’s DC signal experiences first-order decay:

$$\begin{aligned} \mathcal{H}_0 : y(k+1) &= y(k) + \sigma_0 n(k) \\ \mathcal{H}_1 : y(k+1) &= \alpha y(k) + \beta + \sigma_1 n(k) \end{aligned} \quad (1)$$

with $0 < \alpha < 1$ representing the rate of decay in the DC signal, $\beta < 0$, and σ_0 and σ_1 representing the unknown variances of the noise under each hypothesis.

In words, the model for \mathcal{H}_0 assumes that changes in the PPG signal are caused by a random drift, consistent with observed PPG signals [7, 28], where σ_0 is an unknown constant which defines the energy of the drift. The variance, σ_0 , is different for each patient based on physiology (i.e. weight, general health, stress, etc.), thus, we do not assume any knowledge about its unique value. Alternatively, the model for \mathcal{H}_1 captures noisy first-order decays in the PPG signal associated with hypovolemia. The parameters α and β have qualitative relationships with physiology, where β is proportional to the total amount of potential blood loss while α is proportional to the blood loss rate. Lastly, and similar to σ_0 in \mathcal{H}_0 , the parameter σ_1 denotes the variance of the noise and, consistent with \mathcal{H}_0 , α , β , and σ_1 are all unknown to the detector.

To design a detector capable of handling the unknown pa-

rameters α , β , σ_0 , and σ_1 in (1), a one-sided test statistic using the techniques described in [31] is produced resulting in a t-statistic test. A threshold cutoff for this statistic can then be defined; if the statistic exceeds this threshold, the null hypothesis is rejected and the patient is considered to be in a hypovolemic state. This threshold is tuned to achieve a desired false alarm rate, at the expense of sensitivity. A parameter invariant statistic provides invariance to the unknown parameters in the above model which provide no information about which hypothesis is in effect. Results of the application of this technique on a small set of retrospective patient data are discussed in Section 6.

6. CASE STUDY

We aimed to select several representative hypovolemic and non-hypovolemic patients to serve as a case study for the proposed technique. Data was drawn from the matched subset of the Physionet MIMIC II Waveform Database, a freely available database of comprehensive clinical data collected from tens of thousands of intensive care unit patients between 2001 and 2008 at a single tertiary teaching hospital. The matched subset contains patients who have both physiologic waveforms and clinical notes [22, 9].

To select patients with a high likelihood of hypovolemia, we searched for patients with documented ICD9 codes for hypovolemia on discharge (ICD9 276.52) with accompanying notes documenting approximate time of suspected hypovolemia. From these patients, we then selected patients who had available PPG waveforms (documented as *PLETH* signals in the database). Time of hypovolemia was annotated to be the timestamp of the note describing suspected hypovolemia in the patient’s file.³ We ran the algorithm only on data from the ICU stay which contained documented hypovolemia.

To select healthy patients, we chose patients from the database who had *PLETH* waveforms, did not die, and had four or fewer ICD9 codes (indicating they suffered from few comorbidities, leading to simpler hospital stays).

6.1 Results

We ran the techniques described in Section 5 over the selected patients. A summary of the results for each patient can be found in Table 1, and a graph of the statistic and alarms generated for each patient can be found in the appendix. In both of the hypovolemic patients, our detector presented a higher-than-average number of alarms within a 24 hour envelope of hypovolemia was first documented. Patient 00618 had alarms one hour and 18 hours after hypovolemia documentation. Patient 12351 had alarms five hours before and five hours after hypovolemia documentation. These patients also experienced numerous alarms long after hypovolemia detection. We consider these other alarms inconclusive, as no clear indication of when hypovolemia ended for these patients. Total alarm duration for these two patients was 161 minutes (2.7 hours) out of a total time in the ICU of 229.3 hours (9.6 days).

False alarm rate depends on the chosen threshold cutoff, which in turn was based on the number of available data

³As previously described, hypovolemia caused by internal hemorrhage is difficult to accurately diagnose, and is usually only detected once blood loss is substantial, which may be many hours after the patient has begun to lose blood. For this reason, our evaluation metrics take into account an envelope of possibility around each of the annotations.




Figure 2: Figure showing PPG signal with artifact normal PPG signal. On top, the raw voltages are displayed. Below, the N-point discrete fourier transform of the signal is shown. Note the harmonic notches that appear at 100, 200, and 300 in the fourier transform of the normal signal that are absent in the signal with artifact.

Patient ID	ICU Visit Number	Number of Alarms	Length of ICU Stay (Hours)	Hypovolemic?	Alarms within 24 (Hours from Onset)
00618	1	4	68.5	Yes	-17, 1, 18
12351	1	6	160.8	Yes	-5.6, 5.4, 21.3
03617	1	0	34.6	No	N/A
03640	1	0	21.9	No	N/A
05223	1	1	13.8	No	N/A
05223	2	0	18.1	No	N/A
05223	3	1	17.7	No	N/A

Table 1: A summary of the set of patients selected from the Physionet MIMIC II database for use in evaluation, and the results of the application of the hypovolemia detector on these patients. Patient IDs are from the matched subset. Length of stay includes relevant single ICU stay. Alarms listed are those alarms within 24 hours of documented time of suspected hypovolemia.

points in the sliding window. The three patients without hypovolemia had five total ICU stays. (Patient 05223 had three ICU stays with substantial PPG data, so we included each of them. In total, these three patients had 84.3 total hours of ICU time, and experienced only two alarms. These alarms had a duration of 57 minutes (0.95 hours).

While this analysis is promising, it is ultimately too small to make any conclusive statements about definitive sensitivity/specificity results.

7. CONCLUSION

This paper describes a method for creating a robust detector of hypovolemia in critically ill patients by creating a parameter invariant statistic over a patient’s photoplethysmogram waveform after filtering artifact from the data. Preliminary tests on a small set of retrospective patient data show that the proposed detector produces alarms near and before time of diagnosed hypovolemia while producing few false alarms in healthy patients.

Future work will include incorporating more features, either through extracting new dimensions from the PPG waveform or by using other waveforms, to improve predictive power. Future work will also involve expanding our test set of patients by identifying more cases of clear hypovolemia to improve our confidence in the technique and provide clearer estimates of specificity/sensitivity.

8. REFERENCES

- [1] J. Allen. Photoplethysmography and its application in clinical physiological measurement. *Physiological measurement*, 28(3):R1–39, Mar. 2007.
- [2] K. Bartels, R. H. Thiele, and T. J. Gan. Rational fluid management in today’s ICU practice. *Critical Care*, 17(Suppl 1):2–7, 2013.
- [3] R. H. Birkhahn, T. J. Gaeta, D. Terry, J. J. Bove, and J. Tloczkowski. Shock index in diagnosing early acute hypovolemia. *The American Journal of Emergency Medicine*, 23(3):323–326, May 2005.




Figure 3: Section of PPG waveform (bottom) and accompanying statistic value (top) for patient 00618 in the MIMICII database over time. The blue highlighted section on the upper graph shows time of suspected hypovolemia taken from the nursing notes. The upper graph shows the value of the parameter-invariant statistic, and red arrows indicate times of proposed hypovolemia alarms. The dotted line on the top graph indicates the threshold cutoff for the detector, which varies with amount of available data in the sliding window.




Figure 4: Section of PPG waveform for patient 03617 in the MIMICII database. This patient did not have hypovolemia. The upper graph shows the value of the parameter-invariant statistic. No alarms would have been produced on this patient. The lower graph shows the PPG waveform.

- [4] K. J. Brasel, C. Guse, L. M. Gentilello, and R. Nirula. Heart rate: is it truly a vital sign? *The Journal of trauma*, 62(4):812–7, Apr. 2007.
- [5] H. Brown, J. Terrence, P. Vasquez, D. W. Bates, and E. Zimlichman. Continuous monitoring in an inpatient medical-surgical unit: a controlled clinical trial. *The American journal of medicine*, 127(3):226–32, Mar. 2014.
- [6] L. Chen, T. M. McKenna, A. T. Reisner, A. Gribok, and J. Reifman. Decision tool for the early diagnosis of trauma patient hypovolemia. *Journal of biomedical informatics*, 41(3):469–78, June 2008.
- [7] V. a. Convertino, S. L. Moulton, G. Z. Grudic, C. a. Rickards, C. Hinojosa-Laborde, R. T. Gerhardt, L. H. Blackbourne, and K. L. Ryan. Use of advanced machine-learning techniques for noninvasive monitoring of hemorrhage. *The Journal of trauma*, 71(1 Suppl):S25–32, July 2011.
- [8] V. a. Convertino, C. a. Rickards, K. G. Lurie, and

- K. L. Ryan. Hyperventilation in Response to Progressive Reduction in Central Blood Volume to Near Syncope. *Aviation, Space, and Environmental Medicine*, 80(12):1012–1017, Dec. 2009.
- [9] A. L. Goldberger, L. A. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C.-K. Peng, and H. E. Stanley. Physiobank, physiotoolkit, and physionet components of a new research resource for complex physiologic signals. *Circulation*, 101(23):e215–e220, 2000.
- [10] I. Greaves, K. Porter, T. J. Hodgetts, and M. Woollard. *Emergency care: a textbook for paramedics*. Elsevier Health Sciences, 2006.
- [11] M. P. W. Grocott, M. G. Mythen, and T. J. Gan. Perioperative fluid management and clinical outcomes in adults. *Anesthesia and analgesia*, 100(4):1093–106, Apr. 2005.
- [12] G. Gutierrez, H. D. Reines, and M. E. Wulf-Gutierrez. Clinical review: Hemorrhagic Shock. *Critical care (London, England)*, 8(5):373–81, Oct. 2004.
- [13] R. Ivanov, J. Weimer, A. Simpao, M. Rehman, and I. Lee. Early detection of critical pulmonary shunts in infants. In *Proceedings of the 6th ACM/IEEE International Conference on Cyber-Physical Systems, ICCPS '15*, 2015.
- [14] T. Loupec, H. Nanadoumgar, D. Frasca, F. Petitpas, L. Laksiri, D. Baudouin, B. Debaene, C. Dahyot-Fizelier, and O. Mimoz. Pleth variability index predicts fluid responsiveness in critically ill patients. *Critical Care Medicine*, 39(2):294–299, Feb. 2011.
- [15] P. E. Marik, R. Cavallazzi, T. Vasu, and A. Hirani. Dynamic changes in arterial waveform derived variables and fluid responsiveness in mechanically ventilated patients: a systematic review of the literature. *Critical care medicine*, 37(9):2642–7, Sept. 2009.
- [16] S. McGee, W. B. I. Abernethy, and D. L. Simel. Is This Patient Hypovolemic? *Journal of the American Medical Association*, 281(11):1022, Mar. 1999.
- [17] S. P. McGrath, K. L. Ryan, S. M. Wendelken, C. a. Rickards, and V. a. Convertino. Pulse oximeter plethysmographic waveform changes in awake, spontaneously breathing, hypovolemic volunteers. *Anesthesia and analgesia*, 112(2):368–74, Feb. 2011.
- [18] M. Moscucci, K. Fox, C. P. Cannon, W. Klein, J. Lopez-Sendon, G. Montalescot, K. White, and R. Goldberg. Predictors of major bleeding in acute coronary syndromes: the Global Registry of Acute Coronary Events (GRACE). *European Heart Journal*, 24(20):1815–1823, Oct. 2003.
- [19] M. Naccarato, S. Leviner, J. Proehl, S. Barnason, C. Brim, M. Crowley, C. Lindauer, A. Storer, and J. Williams. Clinical Practice Guideline : Orthostatic Vital Signs. *Emergency Nurses Association Clinical Practice Guidelines*, (March), 2011.
- [20] National Clinical Guideline Centre. Intravenous fluid therapy in adults in hospital. *National Institute for Health and Care Excellence (NICE)*, page 36, 2013.
- [21] R. Pizov, A. Eden, D. Bystritski, E. Kalina, A. Tamir, and S. Gelman. Arterial and Plethysmographic Waveform Analysis in Anesthetized Patients with Hypovolemia. *Anesthesiology*, 113(1):83–91, 2010.

- [22] M. Saeed, M. Villarroel, A. T. Reisner, G. Clifford, L.-W. Lehman, G. Moody, T. Heldt, T. H. Kyaw, B. Moody, and R. G. Mark. Multiparameter intelligent monitoring in intensive care ii (mimic-ii): a public-access intensive care unit database. *Critical care medicine*, 39(5):952, 2011.
- [23] L. L. Scharf and C. Demeure. *Statistical Signal Processing*. Addison-Wesley Publishing Company, 1991.
- [24] K. H. Shelley. Photoplethysmography: beyond the calculation of arterial oxygen saturation and heart rate. *Anesthesia and analgesia*, 105(6 Suppl):S31–6, tables of contents, Dec. 2007.
- [25] T. Shires, D. Colin, J. Carrico, and S. Lightfoot. Fluid Therapy in Hemorrhagic Shock. *JAMA Surgery*, 88:688–693, 1964.
- [26] R. Sinert and M. Spektor. Clinical assessment of hypovolemia. *Annals of emergency medicine*, 45(3):327–9, Mar. 2005.
- [27] J. R. Spaniol, A. R. Knight, J. L. Zebley, D. Anderson, and J. D. Pierce. Fluid Resuscitation Therapy for Hemorrhagic Shock. *Journal of Trauma Nursing*, 14(3):152–160, July 2007.
- [28] C. L. Stewart, J. Mulligan, G. Z. Grudic, V. a. Convertino, and S. L. Moulton. Detection of low-volume blood loss: compensatory reserve versus traditional vital signs. *The journal of trauma and acute care surgery*, 77(6):892–7; discussion 897–8, Dec. 2014.
- [29] M. E. Stokes, X. Ye, M. Shah, K. Mercaldi, M. W. Reynolds, M. F. T. Rupnow, and J. Hammond. Impact of bleeding-related complications and/or blood product transfusions on hospital costs in inpatient surgical patients. *BMC health services research*, 11(1):135, Jan. 2011.
- [30] N. S. Ward. The accuracy of clinical information systems. *Journal of critical care*, 19(4):221–5, Dec. 2004.
- [31] J. Weimer, R. Ivanov, A. Roederer, S. Chen, and I. Lee. Parameter invariant design of medical alarms. *IEEE Design & Test*, submitted.
- [32] E. Zöllei, V. Bertalan, A. Németh, P. Csábi, I. László, J. Kaszaki, and L. Rudas. Non-invasive detection of hypovolemia or fluid responsiveness in spontaneously breathing subjects. *BMC anesthesiology*, 13(1):40, Jan. 2013.

APPENDIX

A. PATIENT DATA AND STATISTIC VALUE

All data taken from Physionet MIMICII database [22, 9]. In each figure, a section of PPG waveform is shown on the bottom. The statistic produced by the parameter invariant technique run over that waveform is shown above it.




Figure 5: PPG waveform (bottom) and parameter invariant statistic (top) for patient 00618.




Figure 6: PPG waveform (bottom) and the parameter invariant statistic (top) for patient 12351.




Figure 7: PPG waveform (bottom) and the parameter invariant statistic (top) for patient 03617.




Figure 8: PPG waveform (bottom) and the parameter invariant statistic (top) for patient 03640.




Figure 9: PPG waveform (bottom) and the parameter invariant statistic (top) for patient 05233's first ICU visit.




Figure 10: PPG waveform (bottom) and the parameter invariant statistic (top) for patient 05223's second ICU visit.




Figure 11: PPG waveform (bottom) and the parameter invariant statistic (top) for patient 05223's third ICU visit.