

Prediction of Critical Pulmonary Shunts in Infants

Radoslav Ivanov, *Student Member, IEEE*, James Weimer, *Member, IEEE*,
Allan F. Simpao, Mohamed A. Rehman, and Insup Lee, *Fellow, IEEE*

Abstract—As a first step toward the development of closed-loop medical cyber-physical systems, this paper presents a monitor for blood oxygen concentration that predicts critical drops in oxygen levels caused by pulmonary shunts in infants. Although blood oxygen concentration is one of the most closely monitored vital signs in modern operating rooms, it cannot be measured noninvasively and is currently monitored by a time-delayed proxy—the hemoglobin oxygen saturation. To predict sharp drops in blood oxygen concentration, we employ available noninvasive respiratory measurements and build a parameterized physiological model of the circulation of these gases through the cardiopulmonary system. Since the model parameters (e.g., metabolic rate) are unknown and vary greatly across patients, we utilize a parameter-invariant detector designed to provide a constant false alarm rate for different patients regardless of the values of the parameters and robust to missing measurements. Finally, we evaluate the performance of the detector on real patient data collected during surgeries performed at the Children’s Hospital of Philadelphia. As evaluated on 61 patients experiencing a drop in blood oxygen concentration, the detector achieves a detection rate of about 85% with a potentially life-saving early warning of 90 s on average. In addition, it achieves a false alarm rate of 0.95 false alarms per hour (about 0.5% of the tests) across 314 patients who did not experience a pulmonary shunt.

Index Terms—Medical cyber-physical systems (MCPSs), parameter-invariant detectors, time-series analysis.

I. INTRODUCTION

THE multitude of sensors and measuring devices in modern Intensive Care Units (ICUs) and operating rooms (ORs) present a great opportunity for developing medical cyber-physical systems (MCPSs) to aid clinicians. Through analyzing trends and correlations of vital signs over time, MCPSs provide critical event detections and decision

Manuscript received July 13, 2015; revised December 23, 2015; accepted January 23, 2016. Manuscript received in final form February 29, 2016. This work was supported in part by the National Science Foundation (NSF) and Intel–NSF Partnership for Cyber-Physical Systems Security and Privacy under Grant CNS-1505799, in part by the Global Research Laboratory Program through the National Research Foundation under Grant 2013K1A1A2A02078326, and in part by the Daegu Gyeongbuk Institute of Science and Technology (DGIST) Research and Development Program through the Cyber-Physical Systems Global Center within the Ministry of Science, ICT and Future Planning. Recommended by Associate Editor A. Chiuso. Preliminary versions of some of the results in this paper appeared at ICCPS’15 [1]. (Radoslav Ivanov and James Weimer contributed equally to this work.)

R. Ivanov, J. Weimer, and I. Lee are with the Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA 19104 USA (e-mail: rivanov@seas.upenn.edu; weimerj@seas.upenn.edu; leeg@seas.upenn.edu).

A. F. Simpao and M. A. Rehman are with the Department of Anesthesiology and Critical Care Medicine, The Children’s Hospital of Philadelphia, Philadelphia, PA 19104 USA (e-mail: simpaoa@email.chop.edu; rehman@email.chop.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCST.2016.2538207

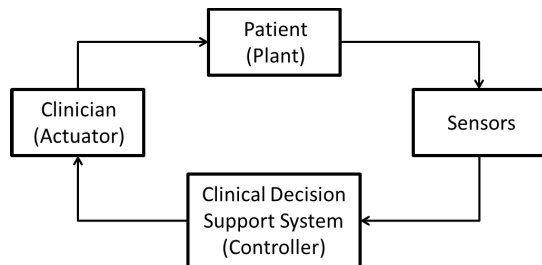


Fig. 1. MCPS control loop.

support, thus allowing clinicians to be proactive as opposed to be reactive in their treatment of patients. Ultimately, these systems can be used to close the loop in an automated and theoretically sound fashion, thereby improving quality of care [2].

Building such an MCPS, however, introduces challenges at all levels of the typical plant–controller loop (illustrated in Fig. 1). In particular, deriving accurate plant models can be difficult due to the complexity of human physiology; even when models exist, they are often highly nonlinear and parameterized by physiological variables that differ across patients, which makes the precise analysis challenging. In addition, the model state maps on physiological variables (e.g., oxygen content in the blood), which cannot be usually sensed noninvasively and in real time, thus forcing the use of available proxy measurements and approximate models. Moreover, there are stringent requirements on the clinical decision support system, for example, a detector must not only have a low missed detection rate but also a low false alarm rate, regardless of the individual (unknown) patient parameters, since alarm fatigue is a serious problem in ICUs [3]. Finally, physiological actuation is restrictive (i.e., only a few points of actuation), often time delayed, and subject to the clinician’s interpretation of the situation, which varies with personal experience and intuition.

In this paper, we focus on a first step toward a closed-loop MCPS, namely, the development of robust detectors for critical events. Specifically, we address the detection of critical drops in the oxygen (O_2) concentration in arterial blood (denoted by C_aO_2 and also referred to as O_2 content) caused by pulmonary shunts in infants. C_aO_2 is one of the most closely monitored physiological variables in modern ICUs and ORs, where too low C_aO_2 (hypoxia) can lead to organ failure (e.g., brain damage) and too high C_aO_2 (hyperoxia) can result in atelectasis (i.e., collapse of the lungs). A pulmonary shunt is a condition in which only one lung is participating in pulmonary exchange (i.e., in contact with fresh air); shunts occur frequently in ORs—they can be natural (e.g., caused by pulmonary edema) but most often occur during mechanical ventilation when an endotracheal tube is used. One-lung ventilation, especially in infants with underdeveloped lungs, may

not supply enough O_2 for the body, thus causing sharp drops in C_aO_2 ; therefore, careful monitoring of C_aO_2 is crucial in one-lung ventilation scenarios in order to ensure the patient's safety.

Currently, C_aO_2 can be directly measured using blood gas analysis but this requires drawing blood from the patient and takes at least several minutes. Instead, clinicians noninvasively monitor the hemoglobin oxygen saturation in the peripheral capillaries, denoted by S_pO_2 , which is a good approximation of the O_2 content where it is measured (usually at a finger tip). However, it takes time (at least tens of seconds) for the blood to reach the peripheral capillaries, and thus S_pO_2 is a delayed measure of the O_2 content in other locations. Therefore, a low C_aO_2 might not be observed until it is measured through S_pO_2 , at which point the patient may already be in a critical state.

Thus, in this paper, we aim to use other available vital signs in order to predict sharp decreases in C_aO_2 caused by shunts. In particular, there are several reliable real-time pulmonary measurements available, notably the partial pressures of O_2 and carbon dioxide (CO_2) in inhaled and exhaled air. However, we are unaware of any models describing the circulation of gases through the cardiovascular and pulmonary systems. Therefore, our first contribution is the development of a parameterized model of the circulation of O_2 and CO_2 in both the blood and airways using physics laws and published physiological data trends in the medical literature. This model can then be used to develop a detector of critical drops in C_aO_2 caused by shunts.

Building such a detector is challenging as the parameters in the model (e.g., metabolic rate) can greatly vary across patients and even in the same patient over time. In such cases, one can still design powerful detectors by deriving tests whose outcomes are invariant to the values of these parameters [4]. Thus, we can estimate the likelihood of each of the two hypotheses in question (i.e., shunt or no shunt) and make a decision which achieves a desired level of false alarm rate for all patients, thus avoiding outliers with multiple false alarms.

In addition to varying parameters, the detector's performance is also affected by the fact that measurements are often wrong or missing altogether. These artifacts may be due to many reasons, including the patient moving or being temporarily taken off mechanical ventilation. Therefore, rather than not making a decision at such times, we enhance the detector's power by treating such measurements as unknown parameters and again designing tests whose outcomes remain the same regardless of these additional parameters.

Finally, we evaluate the detector's performance using real patient data collected at the Children's Hospital of Philadelphia (CHOP). For 61 patients experiencing a critical shunt, the detector correctly detects about 85% of the drops in C_aO_2 , on average, about 90 s before a corresponding drop in S_pO_2 was observed. In 314 control patients not experiencing a shunt, the detector achieved a false alarm rate of about 0.95 false alarms per hour (i.e., 0.5% of the tests). A full description of the testing scenario and results is provided in Section VII.

In summary, the contributions of this paper are: 1) the development of a physiological model of the circulation of O_2 and

CO_2 in both the presence and absence of a pulmonary shunt; 2) the first application, to our knowledge, of parameter-invariant detectors to MCPSs; and 3) a case study evaluation on real patient data obtained at the CHOP. In addition, the extensions of this paper over [1] include the handling of missing and bad measurements as well as an almost doubled data set that confirms the quality of the developed detector.

The rest of this paper is organized as follows. In Section II, we describe some of the related work on the MCPS. Section III formulates the problem considered in this paper and outlines the current approach, including its deficiencies. In Section IV, we derive the physiological model, and in Section V, we apply it to the parameter-invariant detector, which is also enhanced to handle bad/missing data (Section VI). Section VII presents the case study evaluation, and Section VIII concludes this paper.

II. RELATED WORK

The purpose and performance of MCPSs for a given application are closely related to the availability of accurate models. When available, such models allow researchers to perform powerful analysis with strong theoretical guarantees; a few examples include the cardiac pacemaker [5], [6], the artificial pancreas [7]–[9], and the verification of the infusion pump [2], [10]. Since a summary of all MCPS applications is beyond the scope of this review, we leave the in-depth review of medical monitoring and control to the various surveys in [11]–[13]. In this section, we review commonly employed approaches to system modeling and monitoring in both the contexts of critical pulmonary shunt prediction and physiological monitoring in general.

Before discussing monitoring techniques that require system models, we note that the most common monitors in hospitals use threshold alarms [14]. Threshold alarms monitor a single physiologic signal as measured by a biosensor [15]. When the signal crosses a predefined threshold, the system produces an alarm that can be observed by a clinician. Threshold alarms are popular because they are simple, they are easy to implement, and easy for humans to understand. However, research has shown that single sensor threshold alarms are severely limited [16]–[18]. Various studies have documented sensor threshold false alarm rates ranging from 57% to 99% of all alarms, causing alarm fatigue in caregivers [19]–[21]. Thus, threshold alarms can ultimately fail to provide clinicians with a reliable understanding of their patients' state.

In contrast, developing physiological models, coupled with the robust monitor design, provides a systematic approach to incorporating the time history of multiple measurements into the MCPS. In general, physiological modeling belongs to the broad class of system identification. System identification is a well-studied area in control theory and other domains [22]–[24], where techniques can be broadly classified as white-box (i.e., first principles) [25], gray-box [26]–[28], and black-box models (i.e., data driven) [29]–[31].

Physiological modeling is generally performed using compartmental modeling [32], which pertains to white-box and gray-box system identification. In this setting, the model consists of compartments whose states interact with each other through differential or time-differenced dynamics



Fig. 2. Measurement devices currently available to clinicians. (a) Blood gas analyzer [54]. (b) Pulse oximeter [55]. (c) Standard anesthesia machine [56].

(i.e., a state-space model). Depending on the model, a compartment may or may not correspond to a single physiological location, e.g., a compartment may be used to represent a single lung, but it may also be used to represent the transport of blood from the heart to the tissues. Example applications of these models include the cardiac [33] and insulin-glucose [34] systems. A fundamental challenge of compartmental modeling is the balance between physiological accuracy and model identifiability. On one hand, high-fidelity models accurately capture patient physiology, but the model parameters may not be identifiable through the training data; on the other hand, parsimonious models can be identifiable through the training data, but their accuracy may be poor. In cases where a balance between physiological accuracy and model identifiability can be achieved, classical approaches to robust monitoring can be applied [35], [36].

Due to the unpredictability of human physiology and the recent explosion of patient data, data-driven (or black-box) approaches to medical modeling and monitoring have gained popularity [37]–[42] and shown improved performance over competitive techniques [43]. However, there are often practical challenges to data-driven approaches in the medical domain, namely, algorithms require rich patient data representing the entire population with accurate annotations [44]. Satisfying this constraint is especially difficult in the surgical settings considered in this work where accurate annotations are not a primary clinician concern [19], [45]. Moreover, temporal reasoning over clinical data using data-driven techniques is still an open area of research [31], [43]. Thus, it is unlikely that a purely data-driven approach will perform well as a critical pulmonary shunt predictor, since data are scarce and frequently missing, and the dynamic response is affected by a high number of potential variables (e.g., body mass and lung development) [3], [40], [46].

An alternative approach to handling inaccurate models and sparse data is the parameter-invariant detector [4]. It uses maximally invariant statistics to develop tests whose statistical performance is invariant to the unknown model parameters. Having invariant statistical performance provides the appealing feature that these detectors can be specified to have a desired level of false alarm rate and can alleviate the alarm fatigue problem [21]. This class of detectors has been shown to work well in other applications with parameterized models such as fault detection in networked systems [47], [48] as well as the development of heating, ventilating, and air conditioning systems [49] and smart grids [50]. Beyond the application of parameter-invariant detectors in [1], more recently

parameter-invariant detectors have been shown to work well in medical monitoring applications for hypovolemia [51], [52] and meal detection in type I diabetics [53].

III. PROBLEM FORMULATION

This section formulates the problem addressed in this paper. We consider infants on mechanical ventilation in an ICU or OR setting, in which a pulmonary shunt can occur, whether inadvertently or at the request of the surgeon. In turn, this could lead to a decrease in C_aO_2 , i.e., the amount of O_2 in the arteries, which is especially dangerous in infants who may have underdeveloped immune systems or limited cardiopulmonary reserve. We outline the current approach to monitoring and controlling C_aO_2 and highlight its deficiency before stating the precise problem for this work.

A. Current Approach to Monitoring and Control

As described in Section I, clinicians do not have a way to monitor the values of C_aO_2 in real time. When suspecting a problem, they may draw blood from the patient and use it in a blood gas analyzer [Fig. 2(a)] for extraction of various properties of the blood, e.g., acidity and partial pressure of dissolved O_2 . This process is, however, invasive and time consuming (ranging from several minutes to more than 30).

As a real-time proxy for C_aO_2 , clinicians monitor the hemoglobin saturated with O_2 in the peripheral capillaries, i.e., S_pO_2 , which is the percent of hemoglobin binding sites occupied by O_2 . S_pO_2 is measured noninvasively and in real time (data are recorded every 15 s at the CHOP) by a pulse oximeter [Fig. 2(b)]. The relation between S_pO_2 and C_aO_2 is indirect and is governed by the oxygen content equation [57]

$$C_aO_2 = 1.34S_aO_2Hb + 0.003P_aO_2 \quad (1)$$

where Hb is the amount of hemoglobin in grams per deciliter of blood, P_aO_2 is the arterial partial pressure of dissolved oxygen measured in mmHg (millimeters of mercury), and S_aO_2 denotes the arterial hemoglobin oxygen saturation. Equation (1) shows that for normal values of P_aO_2 around 80–110 mmHg and Hb around 12–17 g/dL [58], the majority of O_2 is bound to hemoglobin, and hence, S_aO_2 is a good proxy for C_aO_2 . Since the hemoglobin oxygen saturation is constant as blood travels to the peripheral capillaries [58], clinicians use S_pO_2 (peripheral capillaries) as a time-delayed measure of S_aO_2 (arteries) and, consequently, C_aO_2 . On the other hand, S_aO_2 may be different from S_pO_2 during transient periods (e.g., following a shunt), and hence, S_pO_2 may be a delayed measure of C_aO_2 .

Delayed detection of low C_aO_2 causes clinician control actions to be reactive (versus proactive) and results in the degraded control of C_aO_2 . Once low C_aO_2 is detected, the corrective control action varies with the situation and clinician. For unintentional pulmonary shunts, the typical control action is to remove the shunt quickly. For induced pulmonary shunts on ventilated patients, clinicians can actuate the C_aO_2 by adjusting the inspiratory/expiratory pressures, the tidal volume, and the inspiratory oxygen concentration through the anesthesia machine [Fig. 2(c)]; however, the applied control is based predominantly on intuition and previous experience, and hence, it varies according to the clinician [59].

B. Problem Statement

As described above, the current approach to monitoring C_aO_2 is reactive and may lead to late recognition of critical events and a delayed control response. To improve the control of C_aO_2 during the shunt, we observe that in addition to providing a means to actuate C_aO_2 , the anesthesia machine provides multiple pulmonary measurements, namely, the partial pressures of O_2 and CO_2 in inhaled and exhaled air, as well as tidal volume and respiratory rate. Thus, the problem addressed in this paper is to develop, using only pulmonary measurements, a supplemental detector that predicts critical-shunt-induced drops in C_aO_2 before they are measured by the pulse oximeter as low S_pO_2 ; such a detector would allow clinicians to be proactive in their corrective response.

Formally, we assume the patient's state at sampling time k , $\mathcal{M}(k)$, has one of two discrete modes, shunt or no shunt, denoted by \mathcal{M}_S and \mathcal{M}_{NS} , respectively. Thus, for any given time K , we consider a detection window of M consecutive time steps and write the hypothesis testing problem as

$$\begin{aligned} H_0(K) : \mathcal{M}(K - M + k) &= \mathcal{M}_{NS}, \quad 1 \leq k \leq M \\ H_1(K) : \mathcal{M}(K - M + k) &= \begin{cases} \mathcal{M}_{NS}, & 1 \leq k \leq M_T \\ \mathcal{M}_S, & M_T < k \leq M \end{cases} \end{aligned} \quad (2)$$

where M_T is the shunt transition time.¹ The following sections present a solution and evaluation to the hypothesis testing problem in (2).

IV. PHYSIOLOGICAL MODELING

Solving the hypothesis testing problem specified in (2) requires discriminatory information regarding the presence or absence of a pulmonary shunt. Observing that the partial pressures of O_2 and CO_2 in the respiratory and cardiovascular systems are coupled, this section develops a model that relates the available pulmonary measurements to the shunt dynamics. To our knowledge, there does not exist a model describing the dynamics of the circulation of these gases—while nominal steady-state relationships exist for a healthy human [60], we are unaware of any model describing the dynamics of the partial pressures during a typical cycle as well as their transient response to the occurrence of a shunt. Thus, this section represents a major contribution of this paper.

¹The values of M and M_T are presented in Section VII.




Fig. 3. Typical oxygen–hemoglobin dissociation curve [57]. It shows the shape of the relationship between the partial pressure of dissolved O_2 and hemoglobin saturation.

Note that the development of high-fidelity models in the MCPS domain is difficult in large part due to varying physiology over patients and even in the same patient over time; moreover, physiological variables are often unidentifiable through available sensors. While discriminatory physiological models can be learned, this requires rich training data over the entire patient population, which are unavailable in our application. Hence, we develop a physiological model for testing (2) by capturing general trends and first-order effects as governed by laws of physics and observed physiological properties in the medical literature. The remainder of this section develops the model by first presenting physiological preliminaries and an overview of the physiological dynamics, followed by mathematical descriptions of the partial pressure dynamics and the shunt dynamics, respectively.

A. Physiological Preliminaries

We begin by noting that O_2 appears in only two forms in the blood—it is either bound to hemoglobin or dissolved in the blood, with the contribution of each to the total O_2 content governed by the oxygen content equation in (1). In addition to (1), the relationship between the partial pressure of dissolved O_2 and the hemoglobin saturation is also described by the oxygen–hemoglobin dissociation curve (Fig. 3). It shows that when the saturation is high, large decreases in the partial pressure are correlated with small decreases in saturation, but there is an inflection point after which small changes in partial pressures can result in large changes in saturation. It is important to note that the magnitude of the dissociation curve varies across patients, but it always retains the same shape [58]. Thus, by detecting sharp decreases in the partial pressure of dissolved O_2 , one can predict the decreases in saturation and hence in overall O_2 content.

Therefore, it seems natural to develop a model for the partial pressures of O_2 in the blood and airways and use the measured values in exhaled air, also provided in partial pressures. However, developing a model for O_2 is difficult due to the fact that there is no known closed-form expression for its diffusion, which is the movement of gases between the blood and air in the lungs. While diffusion has a known differential equation governed by Fick's first law, solving it is not possible due to insufficient blood measurements.




Fig. 4. Simplified schematic of CO_2 partial pressures in the respiratory and cardiovascular systems.

TABLE I
CARDIOPULMONARY PARTIAL PRESSURES OF CO_2

Variable Name	Physiological Location
$P_i\text{CO}_2$	Airways (inspiration)
$P_A^L\text{CO}_2$	Left alveoli
$P_A^R\text{CO}_2$	Right alveoli
$P_v^L\text{CO}_2$	Left pulmonary veins
$P_v^R\text{CO}_2$	Right pulmonary veins
$P_a\text{CO}_2$	Arteries
$P_p\text{CO}_2$	Peripheral Capillaries
$P_v\text{CO}_2$	Veins
$P_e\text{CO}_2$	Airways (expiration)

On the other hand, it is possible to build a complete model for the partial pressures of CO_2 , which has diffusion properties better suited to the development of such a model. In addition, the partial pressure of CO_2 is closely related to the O_2 content in the blood—as described in the following sections, in a patient with large lung capacity, a shunt may not lead to a decrease in the O_2 content and may only cause a small increase in the partial pressure of CO_2 ; however, in an infant with low lung capacity, a shunt will cause both a drop in O_2 content and a large increase in the partial pressure of CO_2 . Thus, detecting sharp increases in the partial pressure of CO_2 correlates with detecting decreases in the O_2 content.

B. Overview of Physiological Dynamics

To overview the modeled physiological processes, consider the simplified schematic model of the cardiovascular and respiratory systems, as shown in Fig. 4.² In the following overview, when a physiological location is described, the variable denoting the corresponding partial pressure of CO_2 is introduced in parentheses. For reference, all variable names and the corresponding locations are listed in Table I.

²Note that for better illustration, Figs. 4 and 5 show the pulmonary veins merging before entering the heart, whereas in healthy humans, they connect to the left atrium directly.

The cycle begins with inhalation ($P_i\text{CO}_2$) as air travels to the lungs and the alveoli; we distinguish between the partial pressures in the alveoli on the left-hand and right-hand sides ($P_A^L\text{CO}_2$ and $P_A^R\text{CO}_2$, respectively) as those may be different if there are differences between the two lungs (e.g., during a shunt, as will become apparent in Section IV-D) [58]. Gases enter the blood stream through diffusion in the alveoli and the pulmonary veins on both sides ($P_v^L\text{CO}_2$ and $P_v^R\text{CO}_2$, respectively). When the blood in the pulmonary veins reaches the heart, it is pumped in the arteries ($P_a\text{CO}_2$) [58]. The arteries take the blood to the peripheral capillaries ($P_p\text{CO}_2$), where metabolism occurs and converts O_2 into CO_2 [58]. Following metabolism, the blood enters the veins ($P_v\text{CO}_2$) and is carried back to the heart, thus finishing the cardiovascular cycle [58]. Finally, the pulmonary cycle ends with expiration ($P_e\text{CO}_2$) [58].

It is important to note that the high diffusive capacity of CO_2 ensures that diffusion reaches equilibrium and the partial pressures in the blood and lungs equalize. Thus, in a healthy human with both lungs functioning correctly, the partial pressure of CO_2 in exhaled air is equal to its arterial partial pressure, i.e., $P_e\text{CO}_2 = P_a\text{CO}_2$.³ Thus, in a healthy human monitoring, $P_e\text{CO}_2$ is nearly equivalent to monitoring $P_a\text{CO}_2$; however, in the presence of a shunt, the two are no longer equal, which justifies the need for the dynamical model developed in the next section.

C. Partial Pressure Dynamics

This section describes the dynamical model of the partial pressure of CO_2 in the cardiopulmonary system. Since we are unaware of the exact dynamics, we refer to published trends in the medical literature as well as laws of physics when developing the model. We begin by explaining the model of a healthy human breathing with both lungs and then emphasize the changes introduced by a shunt. Fig. 5 can be used as a graphical reference containing all equations used in the model.

1) *Cardiovascular Dynamics*: The first part of the model captures the partial pressure dynamics in the cardiovascular system, i.e., the cycle from the pulmonary veins to the pulmonary arteries; diffusion and the pulmonary system are discussed next. We note that the pulmonary veins are short in comparison with the rest of the cardiovascular system [58]; hence, we model $P_a\text{CO}_2$ as an instantaneous average of the partial pressures in the pulmonary veins, that is

$$P_a\text{CO}_2(t) = \frac{P_v^L\text{CO}_2(t) + P_v^R\text{CO}_2(t)}{2}. \quad (3)$$

Note that the body usually tries to optimize blood flow in the direction of highest O_2 intake; however, this process is less pronounced during mechanical ventilation [62].

While there exist fluid dynamics models that can explain the flow of blood through the blood vessels [63], they contain multiple parameters that are difficult to obtain or estimate (e.g., blood vessel wall thickness); since accurate models

³Note that the partial pressure in exhaled air might be smaller than that in the alveoli due to dead space, i.e., the volume of air in the airways that is not in contact with blood. However, dead space is about 5% of tidal volume [61], and hence, it is not considered in this work.

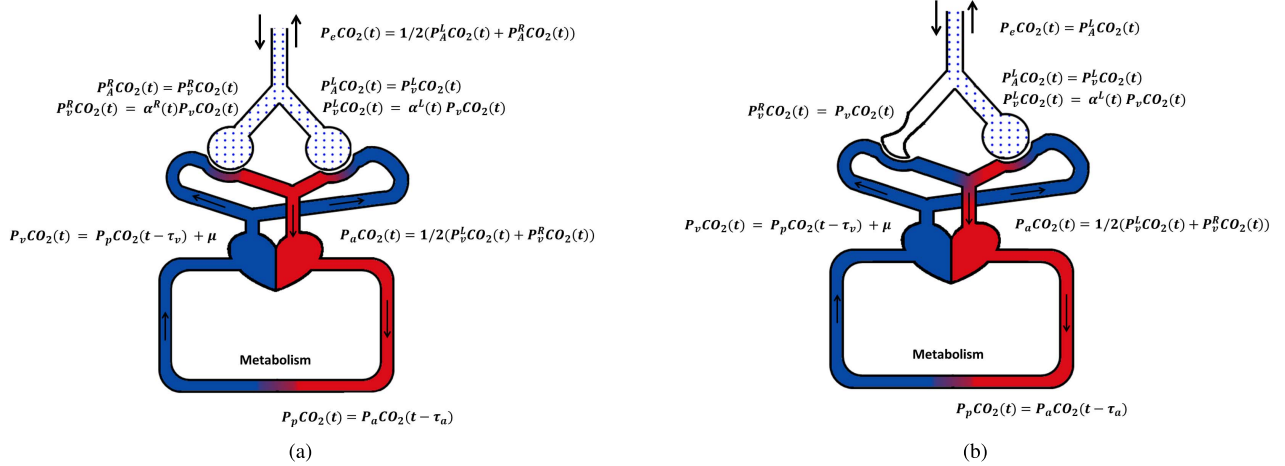


Fig. 5. Illustration of the response of the respiratory and cardiovascular partial pressures to a shunt. (a) System with two functioning lungs. (b) System with a shunt.

cannot be employed, we analyze the average blood cycle. On average, $P_p CO_2$ (i.e., the partial pressure in the peripheral capillaries) is expected to be a time delay of $P_a CO_2$, that is

$$P_p CO_2(t) = P_a CO_2(t - \tau_a) \quad (4)$$

where τ_a is the average time it takes for the blood to travel from the arteries to the peripheral capillaries.

Similarly, the partial pressure in the veins, $P_v CO_2$, can be modeled as a delay of $P_p CO_2$, τ_v , plus the effect of metabolism. Metabolism converts O_2 into CO_2 and increases the partial pressure linearly [58], that is

$$P_v CO_2(t) = P_p CO_2(t - \tau_v) + \mu. \quad (5)$$

Note that μ may vary over time but is assumed constant for a short period of time, e.g., 5 min.

Finally, we can characterize the average total circulation time of the blood

$$\tau = \tau_a + \tau_v. \quad (6)$$

This time can be approximated by dividing the blood volume (mL) by cardiac output (mL/min) [64]. Note that τ may vary over time according to heart rate; however, as will become apparent in Section VII, the anesthesia machine's sampling rate results in only one possible choice of τ , and hence, we assume it is time invariant. If faster sampling is available, one may use available filtering and downsampling techniques to obtain better estimates of the measured signals.

2) *Diffusion Dynamics*: In the body, diffusion is the movement of gases between the blood and air in the lungs from higher to lower pressures. As described above, it is not straightforward to model the diffusion of O_2 due to its low diffusive capacity and sensitivity to the percent of O_2 in inhaled air. On the other hand, CO_2 has two well-studied properties that allow us to accurately model its diffusion. The first property is that postdiffusion partial pressures are always a linear fraction of prediffusion ones [65], that is

$$P_v^R CO_2(t) = \alpha^R(t) P_v CO_2(t) \quad (7)$$

$$P_v^L CO_2(t) = \alpha^L(t) P_v CO_2(t). \quad (8)$$

As shown in [65], the fraction $\alpha^i(t)$, $i \in \{L, R\}$, varies with the tidal volume, i.e., the volume of air in the corresponding lung, denoted by V_t^i . In addition, the amount of diffused CO_2 is directly proportional to the lung area, as governed by Fick's law [58]. If the usual assumption is made that lungs are spheres, then the lung area is directly proportional to $(V_t^i)^{2/3}$, i.e., $\alpha^i(t)$ is inversely proportional to $(V_t^i)^{2/3}$.

3) *Respiratory Dynamics*: To model the respiratory dynamics of the partial pressure of CO_2 , we use its second diffusive property, namely, its high diffusive capacity. This allows CO_2 diffusion to always reach equilibrium before the end of the breath, thus equalizing the partial pressures in the alveoli and the pulmonary veins

$$P_A^R CO_2(t) = P_v^R CO_2(t) \quad (9)$$

$$P_A^L CO_2(t) = P_v^L CO_2(t). \quad (10)$$

Finally, the partial pressure of CO_2 in exhaled air is equal to the average of the alveolar partial pressures

$$P_e CO_2(t) = (1/2)(P_A^L CO_2(t) + P_A^R CO_2(t)) \quad (11)$$

i.e., the effect of dead space is ignored.

D. Shunt Dynamics

Having described the dynamics of partial pressures in a healthy human, in this section, we highlight the differences introduced by a shunt. Without loss of generality, suppose for the remainder of this paper that a shunt occurs on the right side, as shown in Fig. 5(b), the right lung does not participate in the removal of CO_2 through diffusion [58]. Therefore, the first change caused by a shunt is that (7) now becomes

$$P_v^R CO_2(t) = P_v CO_2(t). \quad (12)$$

In addition, $P_e CO_2(t)$ is now only equal to the partial pressure in the alveoli on the left side, since the right lung no longer participates in pulmonary exchange, i.e., (11) now becomes

$$P_e CO_2(t) = P_A^L CO_2(t) = P_v^L CO_2(t) \neq P_v^R CO_2(t). \quad (13)$$

This means that the blood in the pulmonary veins on the right side is rich in CO₂ and poor in O₂—when it mixes with the one on the left, it increases the overall CO₂ content and decreases the O₂ content. Note that a shunt also means $P_a\text{CO}_2 \neq P_e\text{CO}_2$, thus making it impossible to simply monitor the partial pressure of CO₂ in exhaled air in order to check for an increase in the arterial one.

This section developed the dynamics of the CO₂ partial pressures in the presence and absence of a shunt. Depending on the patient's lung capacity, a shunt may have a different effect on the level of change of the arterial partial pressure. Thus, lungs with smaller capacity can diffuse less CO₂ and O₂ (the former out of the blood and the latter in the blood), i.e., large increases in $P_a\text{CO}_2$ are correlated with decreases in the O₂ content in the blood, and provide discriminatory information for testing (2), as presented in the following section.

V. CRITICAL SHUNT DETECTOR

This section develops the critical shunt detector based on the physiological model described above. Given the hypothesis problem in (2), we develop a state-space model under each hypothesis that is then concatenated into a time-series model. The model contains the unknown parameters described in the physiological model, i.e., the metabolic rate and the diffusion ratio; therefore, we develop a detector whose output is invariant to the values of these parameters. The hypothesis test is presented for a single time step; one may develop a sequential detector by just repeating this process over time.

A. State-Space Model

Constrained by the sampling time of the anesthesia machine, we build the state-space model in a discrete-time fashion. We use the discrete variable κ to replace the continuous one τ , i.e., the average blood circulation time. Thus, the relevant equations for a healthy human from Section IV become

$$P_a\text{CO}_2(k) = (1/2)(P_v^L\text{CO}_2(k) + P_v^R\text{CO}_2(k)) \quad (14)$$

$$P_v\text{CO}_2(k) = P_a\text{CO}_2(k - \kappa) + \mu \quad (15)$$

$$P_v^i\text{CO}_2(k) = \alpha^i(k)P_v\text{CO}_2(k), \quad i \in \{L, R\} \quad (16)$$

$$P_A^i\text{CO}_2(k) = P_v^i\text{CO}_2(k), \quad i \in \{L, R\} \quad (17)$$

$$P_e\text{CO}_2(k) = (1/2)(P_A^L\text{CO}_2(k) + P_A^R\text{CO}_2(k)) \quad (18)$$

where k is the discrete time step. Note that this implies that $\tau = \kappa t_s$, where t_s is the sampling time. In addition, note that $P_e\text{CO}_2$ is the only variable in this model that is measured directly.

To develop the dynamics under \mathcal{M}_{NS} , we first estimate $\alpha^i(k)$. Note that, as argued in Section IV, it is inversely proportional to $(V_t^i)^{2/3}$, where V_t^i is the tidal volume. Furthermore, as will be apparent in Section VII, the sampling time of the bedside monitors is long enough for an infant to take several breaths. Therefore, $V_t^i(k)$ can be approximated by what may be called the cumulative tidal volume (normalized to the power 2/3 in order to capture the relationship with α^i)

$$\bar{V}^i(k) = \frac{t_s}{60} \times (\text{RR}(k) \times V_t^i(k))^{2/3}, \quad i \in \{L, R\} \quad (19)$$

where $\text{RR}(k)$ is the measured respiratory rate (in breaths per minute). Applying (19), (16) can be approximated as

$$P_v^i\text{CO}_2(k) = (\bar{\alpha}/\bar{V}^i(k))P_v\text{CO}_2(k) + \sigma n^i(k) \quad (20)$$

where $\bar{\alpha}$ is an unknown parameter representing the CO₂ diffusion ratio for fixed tidal volume (assumed the same for both lungs), σ is an unknown variance, and $n^i(k)$ denotes noise. The term $\sigma n^i(k)$ is added in order to capture the unknown error introduced by the approximation. For theoretical convenience, we assume the noise has a white Gaussian distribution, i.e., $n^i(k) \sim \mathcal{N}(0, 1)$. While this assumption may not be true in practice, we show in Section VII that the Gaussian-based detector achieves nontrivial detection rates and a *near*-constant false alarm rate across real patient data. Developing optimal parameter-invariant detectors for general distributions is part of future work. Finally, note that we cannot directly measure the tidal volume in each lung; instead, the total tidal volume $V_t(k)$ (and $\bar{V}(k)$, respectively) is measured. Thus, in a healthy human

$$\bar{V}^i(k) = (1/2)\bar{V}(k), \quad i \in \{L, R\}. \quad (21)$$

Thus, combining (20) and (21) with (14)–(18), the state-space model for \mathcal{M}_{NS} can be written as

$$\begin{aligned} \begin{bmatrix} x^L(k) \\ x^R(k) \end{bmatrix} &= \begin{bmatrix} \frac{\bar{\alpha}}{\bar{V}(k)} & \frac{\bar{\alpha}}{\bar{V}(k)} \\ \frac{\bar{\alpha}}{\bar{V}(k)} & \frac{\bar{\alpha}}{\bar{V}(k)} \end{bmatrix} \begin{bmatrix} x^L(k - \kappa) \\ x^R(k - \kappa) \end{bmatrix} \\ &+ \begin{bmatrix} \frac{2\bar{\alpha}}{\bar{V}(k)} & n^L(k) \\ \frac{2\bar{\alpha}}{\bar{V}(k)} & n^R(k) \end{bmatrix} \begin{bmatrix} \mu \\ \sigma \end{bmatrix} \\ y(k) &= [1/2 \quad 1/2] \begin{bmatrix} x^L(k) \\ x^R(k) \end{bmatrix} \end{aligned} \quad (22)$$

where $x^i(k) = P_v^i\text{CO}_2(k)$ and $y(k) = P_e\text{CO}_2(k)$.

The derivation under \mathcal{M}_S is similar and proceeds as follows, again assuming that a shunt occurred in the right lung⁴:

$$P_v^R\text{CO}_2(k) = P_v\text{CO}_2(k)$$

$$P_v^L\text{CO}_2(k) = (\bar{\alpha}/\bar{V}^L(k))P_v\text{CO}_2(k) + \sigma n^L(k)$$

$$\bar{V}^L(k) = \bar{V}(k) \quad (23)$$

where the last equation follows from the fact that only the left lung is now participating in pulmonary exchange, and similar to (20), $n^L(k) \sim \mathcal{N}(0, 1)$ and σ is the same unknown variance. Thus, the dynamics for \mathcal{M}_S are

$$\begin{aligned} \begin{bmatrix} x^L(k) \\ x^R(k) \end{bmatrix} &= \begin{bmatrix} \frac{\bar{\alpha}}{2\bar{V}(k)} & \frac{\bar{\alpha}}{2\bar{V}(k)} \\ 1/2 & 1/2 \end{bmatrix} \begin{bmatrix} x^L(k - \kappa) \\ x^R(k - \kappa) \end{bmatrix} \\ &+ \begin{bmatrix} \frac{\bar{\alpha}}{\bar{V}(k)} & n^L(k) \\ 1 & 0 \end{bmatrix} \begin{bmatrix} \mu \\ \sigma \end{bmatrix} \\ y(k) &= [1 \quad 0] \begin{bmatrix} x^L(k) \\ x^R(k) \end{bmatrix} \end{aligned} \quad (24)$$

⁴As can be observed in (22), the dynamics for x^R and x^L are the same, and hence, we assume without loss of generality that the shunt is on the right side.

where $y(k)$ and $P_e\text{CO}_2(k)$, respectively, measure only the partial pressure in the alveoli in the left lung and, hence, in the pulmonary veins on the left-hand side.

Note that the dynamics for both modes contain three unknown parameters, μ , \bar{a} , and σ , that vary across patients. Since the model is sensitive to their values, in the next sections, we design a detector whose performance is invariant to the values of these parameters.

B. Time-Series Model

Given the state-space models in (22) and (24), we construct a detector invariant to the unknown parameters. To do this, we note that the parameters appear linearly in the models under each mode; therefore, we rearrange the state-space models under each hypothesis in order to isolate the parameters. Finally, we concatenate the equations to form a time-series model that can be directly used in the detector. All proofs are given in the Appendix.

1) *Null Hypothesis*: Note that under H_0 the system is always in \mathcal{M}_{NS} .

Proposition 1: Under H_0 , one can derive the following recursive relation between the obtained measurements and the unknown parameters:

$$y(k) = \frac{2y(k-\kappa)}{\bar{V}(k)}\bar{a} + \frac{2}{\bar{V}(k)}\bar{a}\mu + \sigma_0 n_0(k) \quad (25)$$

where $\sigma_0 = \sigma/\sqrt{2}$ and $n_0(k) \sim \mathcal{N}(0, 1)$ is white noise.

Thus, with M time-concatenated measurements, we arrive at the expression

$$\mathbf{y} = \mathbf{F}_0\boldsymbol{\theta} + \sigma_0\mathbf{n}_0 \quad (26)$$

where $\boldsymbol{\theta} = [\bar{a}, \bar{a}\mu]^T$ captures the unknown parameters, \mathbf{n}_0 is white Gaussian noise, and

$$\mathbf{y} = \begin{bmatrix} y(\kappa+1) \\ y(\kappa+2) \\ \vdots \\ y(M) \end{bmatrix}, \quad \text{and } \mathbf{F}_0 = \begin{bmatrix} \frac{2y(1)}{\bar{V}(\kappa+1)} & \frac{2}{\bar{V}(\kappa+1)} \\ \frac{2y(2)}{\bar{V}(\kappa+2)} & \frac{2}{\bar{V}(\kappa+2)} \\ \vdots & \vdots \\ \frac{2y(M-\kappa)}{\bar{V}(M)} & \frac{2}{\bar{V}(M)} \end{bmatrix}.$$

Note that the first κ measurements are not included in \mathbf{y} since the initial partial pressures are unknown, thus requiring the blood to fully circulate before testing can commence.

2) *Alternative Hypothesis*: Under H_1 , the system begins in \mathcal{M}_{NS} and switches to \mathcal{M}_S at time $M_T + 1$.

Proposition 2: Under H_1 , one can derive the following recursive relations between the obtained measurements and the unknown parameters.

Case 1: $\kappa \leq k \leq M_T$

$$y(k) = \frac{2y(k-\kappa)}{\bar{V}(k)}\bar{a} + \frac{2}{\bar{V}(k)}\bar{a}\mu + \sigma_0 n_1(k).$$

Case 2: $M_T + 1 \leq k \leq M_T + \kappa$

$$y(k) = \frac{y(k-\kappa)}{\bar{V}(k)}\bar{a} + \frac{1}{\bar{V}(k)}\bar{a}\mu + \sqrt{2}\sigma_0 n_1(k).$$

Case 3: $k > M_T + \kappa$

$$y(k) = \frac{y(k-\kappa) + f(k-\kappa)}{2\bar{V}(k)}\bar{a} + \frac{2 + g(k-\kappa)}{2\bar{V}(k)}\bar{a}\mu + \sqrt{2}\sigma_0 n_1(k)$$

where $n_1(k) \sim \mathcal{N}(0, 1)$ is white noise, and $f(k)$ and $g(k)$ are defined as follows:

$$f(k) = \begin{cases} y(k-\kappa), & \text{if } M_T < k \leq M_T + \kappa \\ \frac{y(k - (L_k + 1)\kappa)}{2^{L_k}} + \sum_{n=1}^{L_k} \frac{y(k - n\kappa)}{2^n}, & \text{if } M_T + \kappa < k \leq M \end{cases}$$

$$g(k) = \begin{cases} 1, & \text{if } M_T < k \leq M_T + \kappa \\ \frac{1}{2^{L_k-1}} + \sum_{n=1}^{L_k} \frac{1}{2^{n-1}}, & \text{if } M_T + \kappa < k \leq M \end{cases}$$

where $L_k = \lfloor (k - M_T)/\kappa \rfloor \geq 1$.

Therefore, we arrive at a similar expression for \mathbf{y} under H_1

$$\mathbf{y} = \mathbf{F}_1\boldsymbol{\theta} + \sigma_0\mathbf{Rn}_1 \quad (27)$$

where \mathbf{y} is the same as in (26), $\mathbf{R} \in \mathbb{R}^{M-\kappa \times M-\kappa}$ is a diagonal matrix such that $R_{ii} = 1$ if $i \leq M_T$, and $R_{ii} = \sqrt{2}$ otherwise, \mathbf{n}_1 is white Gaussian noise, and

$$\mathbf{F}_1 = \begin{bmatrix} \mathbf{F}_0(1 : K - \kappa - 1, :) & \\ \frac{y(K-\kappa)}{\bar{V}(K)} & \frac{1}{\bar{V}(K)} \\ \vdots & \vdots \\ \frac{y(K-1)}{\bar{V}(K+\kappa-1)} & \frac{1}{\bar{V}(K+\kappa-1)} \\ \frac{y(K)}{2\bar{V}(K+\kappa)} + \frac{f(K)}{2\bar{V}(K+\kappa)} & \frac{1}{\bar{V}(K+\kappa)} + \frac{g(K)}{2\bar{V}(K+\kappa)} \\ \vdots & \vdots \\ \frac{y(M-\kappa)}{2\bar{V}(M)} + \frac{f(M-\kappa)}{2\bar{V}(M)} & \frac{1}{\bar{V}(M)} + \frac{g(M-\kappa)}{2\bar{V}(M)} \end{bmatrix}$$

where $K = M_T + 1$ and $\mathbf{F}_0(1 : K - \kappa - 1, :)$ denotes the first $K - \kappa - 1$ rows of \mathbf{F}_0 .

Thus, we have derived the model $\mathbf{y} = \mathbf{F}_0\boldsymbol{\theta} + \sigma_0\mathbf{n}_0$ under H_0 and $\mathbf{y} = \mathbf{F}_1\boldsymbol{\theta} + \sigma_0\mathbf{Rn}_1$ under H_1 . We now convert this to a form that can be used by the parameter-invariant detector.

C. Parameter-Invariant Test Statistics

In the previous section [(26) and (27)], we reformulated the hypothesis test problem stated in (2)⁵

$$\begin{aligned} H_0 : \mathbf{y} &= \mathbf{F}_0\boldsymbol{\theta} + \sigma_0\mathbf{n}_0 \\ H_1 : \mathbf{y} &= \mathbf{F}_1\boldsymbol{\theta} + \sigma_0\mathbf{Rn}_1. \end{aligned} \quad (28)$$

We now wish to design statistics for testing the hypotheses. Recalling from Section IV, the parameters $\boldsymbol{\theta}$ and σ_0 represent patient-specific unknown physiological variables corresponding to metabolism, diffusion ratio, and model uncertainty;

⁵Note that the index K from (2) is dropped in Sections V and VI in order to simplify the notation. Designing a sequential detector can be done by simply repeating over time the test described in Sections V and VI.

as argued above, none of these parameters can be accurately estimated *a priori*. The parameters induce a group of transformations, $G \subseteq \{g \mid g(\mathbf{y}) : \mathbb{R}^{M-\kappa} \mapsto \mathbb{R}^{M-\kappa}\}$, on the measurements \mathbf{y} ; thus, the statistical properties of the hypotheses in (28) vary according to the values of the parameters. In such a scenario, we ideally aim to design test statistics that are *maximally invariant* to the parameter-induced transformations.

Definition 1 (Maximally Invariant Statistic): $t(\mathbf{y})$ is maximally invariant to a group G if [4]

$$1) t(\mathbf{y}) = t(g(\mathbf{y})) \quad \forall g \in G \quad (\text{invariant})$$

$$2) t(\mathbf{y}) = t(\mathbf{y}') \longrightarrow \exists g \in G, \quad \mathbf{y} = g(\mathbf{y}') \quad (\text{maximal}).$$

Maximally invariant statistics have two attractive properties in medical applications. First, the invariance property ensures that the resulting test statistic has a similar performance for all patients, as opposed to performing well for some patients and poorly for others (due to varying physiology). Second, the maximal property ensures that only the information corrupted by the unknown parameters is removed from the hypothesis testing problem (i.e., a minimal loss of discriminatory information). This combination of invariance and maximality has been shown to work well in other domains with parameterized models [47]–[50] and in other medical applications with scarce data and unknown parameters [51]–[53].

However, for the hypothesis testing problem in (28), a maximally invariant statistic does not exist since the columns of \mathbf{F}_0 are linearly independent of the columns of \mathbf{F}_1 [4]. This implies that $\boldsymbol{\theta}$ is an unknown test parameter (i.e., provides discriminatory information) under each hypothesis rather than a nuisance parameter (i.e., provides no discriminatory information) like σ . Since no maximally invariant statistic exists, we design near-maximally invariant test statistics in two steps.

- 1) *Step 1:* Design a statistic, $t_0(\mathbf{y})$, maximally invariant to the test-parameter-induced transformations under H_0 .
- 2) *Step 2:* Design a statistic, $r_0(t(\mathbf{y}))$, maximally invariant to the nuisance-parameter-induced transformations for (28).

The remainder of this section presents how to obtain such a near-maximally invariant statistic. In the following, for an arbitrary matrix \mathbf{X} , we denote the projection matrices corresponding to the column space and null space of \mathbf{X} as $\mathbf{P}_\mathbf{X}$ and $\mathbf{P}_{\mathbf{X}^\perp}$, respectively, where $\mathbf{P}_\mathbf{X} = \mathbf{I} - \mathbf{P}_{\mathbf{X}^\perp}$.

For Step 1, we ask for maximal invariance to the parameters affecting the mean under H_0 , namely, $\boldsymbol{\theta}$, which induce the group of transformations

$$G_1 = \{g \mid g(\mathbf{y}) = \mathbf{y} + \mathbf{F}_0\boldsymbol{\theta}, \boldsymbol{\theta} \in \mathbb{R}^2\}. \quad (29)$$

Proposition 3: The statistic $t(\mathbf{y}) = \mathbf{P}_{\mathbf{F}_0^\perp}\mathbf{y}$ is a maximally invariant statistic with respect to the group G_1 defined in (29).

Remark 1: In Proposition 3, the maximally invariant statistic for (29) is proved to be the projection onto the null space of \mathbf{F}_0 . Intuitively, this makes sense since the mean of \mathbf{y} under H_0 lies in the column space of \mathbf{F}_0 , with $\boldsymbol{\theta}$ affecting only the mean's coordinates in this space. Thus, by projecting \mathbf{y} onto the null space of \mathbf{F}_0 (orthogonal to its column space), we make sure that the mean of \mathbf{y} is invariant to the value of $\boldsymbol{\theta}$.

TABLE II

DECISION SPACE FOR THE DETECTOR DEVELOPED IN THIS PAPER. THE DETECTOR'S DECISION IS GIVEN IN PARENTHESES

	$r_0 > t_0^*$	$r_0 \leq t_0^*$
$r_1 > t_1^*$	Inaccurate Model (warning)	Accept H_0 (no alarm)
$r_1 \leq t_1^*$	Accept H_1 (alarm)	Insufficient Power (warning)

Applying $t(\mathbf{y})$, we rewrite the hypothesis testing problem

$$\begin{aligned} H_0 : \mathbf{z}_0 &= \sigma_0 \mathbf{P}_{\mathbf{F}_0^\perp} \mathbf{n}_0 \\ H_1 : \mathbf{z}_0 &= \mathbf{G}_0 \boldsymbol{\theta} + \sigma_0 \mathbf{N}_0 \mathbf{n}_1 \end{aligned} \quad (30)$$

where $\mathbf{z}_0 = t(\mathbf{y}) = \mathbf{P}_{\mathbf{F}_0^\perp} \mathbf{y}$, $\mathbf{G}_0 = \mathbf{P}_{\mathbf{F}_0^\perp} \mathbf{F}_1$, and $\mathbf{N}_0 = \mathbf{P}_{\mathbf{F}_0^\perp} \mathbf{R}$. Thus, $\mathbf{z}_0 \sim \mathcal{N}(\mathbf{0}, \sigma_0 \mathbf{P}_{\mathbf{F}_0^\perp})$ under H_0 and $\mathbf{z}_0 \sim \mathcal{N}(\mathbf{G}_0 \boldsymbol{\theta}, \sigma_0 \mathbf{N}_0)$ under H_1 .

For Step 2, we obtain the maximally invariant statistic for (30) as follows [4]:

$$r_0 = c_0 \frac{\mathbf{z}_0^\top \mathbf{P}_{\mathbf{G}_0} \mathbf{z}_0}{\mathbf{z}_0^\top (\mathbf{I} - \mathbf{P}_{\mathbf{G}_0}) \mathbf{z}_0} \quad (31)$$

where $c_0 = \text{rank}(\mathbf{G}_0) / (M - \kappa - \text{rank}(\mathbf{F}_0) - \text{rank}(\mathbf{G}_0))$ is a scaling constant. Thus, under H_0 , $\mathbf{z}_0 \sim \mathcal{F}(\text{rank}(\mathbf{G}_0), M - \kappa - \text{rank}(\mathbf{F}_0) - \text{rank}(\mathbf{G}_0))$. Here, we stress that regardless of the parameter values, r_0 has the same distribution under H_0 . Therefore, a detector can be designed that utilizes r_0 to provide a consistent performance regardless of an individual patient's physiology.

Before introducing the detector in the following section, we note that a second near-maximally invariant statistic can be generated by switching the hypotheses in Steps 1 and 2, i.e., performing Step 1 using H_1 and performing Step 2 using H_0 . By first multiplying the measurements by \mathbf{R}^{-1} and then following the same procedure as above, we can write a second near-maximally invariant statistic as:

$$r_1 = c_1 \frac{\mathbf{z}_1^\top \mathbf{P}_{\mathbf{G}_1} \mathbf{z}_1}{\mathbf{z}_1^\top (\mathbf{I} - \mathbf{P}_{\mathbf{G}_1}) \mathbf{z}_1} \quad (32)$$

where $\mathbf{z}_1 = \mathbf{P}_{(\mathbf{R}^{-1}\mathbf{F}_1)^\perp} \mathbf{R}^{-1} \mathbf{y}$, $\mathbf{G}_1 = \mathbf{P}_{(\mathbf{R}^{-1}\mathbf{F}_1)^\perp} \mathbf{R}^{-1} \mathbf{F}_0$, $c_1 = \text{rank}(\mathbf{G}_1) / (M - \kappa - \text{rank}(\mathbf{F}_1) - \text{rank}(\mathbf{G}_1))$, and $\mathbf{z}_1 \sim \mathcal{F}(\text{rank}(\mathbf{G}_1), M - \kappa - \text{rank}(\mathbf{F}_1) - \text{rank}(\mathbf{G}_1))$ under H_1 .

D. Decision Space

Using the near-maximally invariant statistics, r_0 and r_1 , this section introduces a test for evaluating (28). We begin by recalling that r_0 has a central \mathcal{F} distribution under H_0 and r_1 has a central \mathcal{F} distribution under H_1 , regardless of the unknown parameters $\boldsymbol{\theta}$ and σ_0 . Under H_0 , by choosing a threshold t_0^* at the tail of the \mathcal{F} distribution and raising an alarm only when $r_0 > t_0^*$, one can guarantee a false alarm rate for the test across all patients. Similarly, under H_1 , by choosing a threshold t_1^* and raising an alarm only when $r_1 > t_1^*$, one can guarantee a missed alarm rate. Based on these two statistics, the detector's decision space consists of four cases as shown in Table II.

In particular, if the two tests agree, then the corresponding hypothesis is accepted. If they disagree, then one of two warnings is output. When both tests reject their individual hypotheses (i.e., $r_0 > t_0^*$ and $r_1 > t_1^*$), we conclude that the model is likely inaccurate. Model inaccuracies are natural in the physiological model developed in Section IV since the model is somewhat crude and may (partly) capture other conditions that can manifest in a patient, e.g., hypovolemia. When both tests fail to reject their individual hypotheses (i.e., $r_0 \leq t_0^*$ and $r_1 \leq t_1^*$), we conclude that the test cannot make a decision to the chosen levels of false alarm and missed alarm.

This section developed a critical shunt detector for the problem in (2). The detector utilizes the physiological model in Section IV to construct time-series models for the measurements under each hypothesis. A pair of near-maximally invariant statistics is employed, one providing patient invariant performance under H_0 and the other under H_1 . The statistics are used to design a test that can either provide a decision which achieves a specified performance or provide feedback regarding why the performance specification is not satisfied. The detector is evaluated in Section VII, after being augmented to handle missing and bad data in the following section.

VI. ROBUSTNESS TO MISSING AND BAD DATA

For various reasons, vital sign measurements are often wrong or missing altogether. In particular, some sensors are greatly affected by the patient's movement (e.g., pulse oximeter), while others provide irrelevant measurements because patients are sometimes taken off mechanical ventilation so that clinicians can check for leaks. In this section, we provide a framework for handling missing measurements with regard to the model developed above. In addition, we describe how to reduce the number of alarms caused by wrong measurements by observing correlations between bad and missing data.

A. Handling Missing Measurements

One way of dealing with missing measurements is to not make a decision when a vital sign is missing. However, this may be too conservative as measurements are often missing in medical settings—for example, in [1], we use only 26 out of 91 cases with shunts due to being unable to make a decision while the event occurs. Thus, in this paper, we aim to include some of these unused cases in an effort to make the detector applicable in a real-world setting.

In order to handle missing measurements in a robust way, we treat them as unknown parameters and ask for invariance to them in the same way that physiological parameters, e.g., μ , the metabolic rate, are treated. This step is performed before the hypothesis test, i.e., when the model is in the form $\mathbf{y} = \mathbf{F}_0\boldsymbol{\theta} + \sigma_0\mathbf{n}_0$ under H_0 and $\mathbf{y} = \mathbf{F}_1\boldsymbol{\theta} + \sigma_0\mathbf{R}\mathbf{n}_1$ under H_1 .

The three measurements that are used in the model are $P_e\text{CO}_2$, V_t , and RR; they affect the model in different places, and hence, different techniques are used. We describe each one in turn, beginning with missing V_t or RR, which are treated in the same way. The following proposition makes

it clear which parts of the model are affected by missing measurements. The proof is given in the Appendix.

Proposition 4: The matrices \mathbf{F}_0 and \mathbf{F}_1 can be represented in the following way, for some matrices $\mathbf{K}_0, \mathbf{K}_1 \in \mathbb{R}^{M-\kappa \times M}$ and vectors $\mathbf{k}_0, \mathbf{k}_1 \in \mathbb{R}^{M-\kappa \times 1}$ that do not depend on the measurements:

$$\begin{aligned}\mathbf{F}_0 &= [\bar{\mathbf{V}} \circ \mathbf{K}_0 \bar{\mathbf{y}} \quad \bar{\mathbf{V}} \circ \mathbf{k}_0] \\ \mathbf{F}_1 &= [\bar{\mathbf{V}} \circ \mathbf{K}_1 \bar{\mathbf{y}} \quad \bar{\mathbf{V}} \circ \mathbf{k}_1]\end{aligned}$$

where \circ denotes the element-wise product and $\bar{\mathbf{y}} = [y(1) \cdots y(M)]^\top \in \mathbb{R}^M$ and $\bar{\mathbf{V}} = [(1/\bar{V}(\kappa+1)), \dots, (1/\bar{V}(M))]^\top \in \mathbb{R}^{M-\kappa}$.

1) *Missing Tidal Volume or Respiratory Rate:* It is now clear how a missing V_t or RR measurement affects the detector, namely, it affects the corresponding rows in \mathbf{F}_0 and \mathbf{F}_1 . In particular, suppose $\bar{V}(p)$ is missing, that is, either a tidal volume or respiratory rate measurement is missing at step p .⁶ Therefore, we ask for invariance to row p .

To achieve a maximally invariant statistic, note that $\bar{\mathbf{V}}$ can also be represented as $\mathbf{I}_{M-\kappa}\bar{\mathbf{V}}$, i.e., row p in $\bar{\mathbf{V}}$ is multiplied by column p in $\mathbf{I}_{M-\kappa}$. Thus, similar to Section V-C, one needs to find the projection to the null space of matrix

$$\mathbf{Q}_V = [\mathbf{e}_p] \quad (33)$$

where \mathbf{e}_p is the unit vector with a 1 in position p and 0 otherwise.

2) *Missing Partial Pressure of Exhaled Carbon Dioxide:* The difference when a $P_e\text{CO}_2$ measurement is missing is that the \mathbf{y} vector is affected. Note that $\bar{\mathbf{y}}$ and \mathbf{y} have different indices and different dimensions. Suppose that $y(q)$ is missing. Then the vector \mathbf{y} is affected in position $q-\kappa$, whereas $\bar{\mathbf{y}}$ is affected in position q . Therefore, the corresponding columns in the premultiplying matrices, $\mathbf{I}_{M-\kappa}$, \mathbf{K}_0 , and \mathbf{K}_1 , respectively, must be eliminated. In other words, we look for the projection to the null space of the following matrix:

$$\mathbf{Q}_E = [\mathbf{e}_{q-\kappa} \quad \bar{\mathbf{V}} \circ \mathbf{K}_0^q \quad \bar{\mathbf{V}} \circ \mathbf{K}_1^q]$$

where \mathbf{K}_0^q and \mathbf{K}_1^q denote column q in the respective matrix.

Consequently, one may build the matrix $\mathbf{Q} = [\mathbf{Q}_V \quad \mathbf{Q}_E]$ and obtain the projection to its null space, $\mathbf{P}_{\mathbf{Q}^\perp}$, similar to Section V-C. By premultiplying both models by $\mathbf{P}_{\mathbf{Q}^\perp}$, we arrive at the following models under each hypothesis:

$$\begin{aligned}H_0 : \mathbf{y}' &= \mathbf{F}'_0\boldsymbol{\theta} + \sigma_0\mathbf{n}_0 \\ H_1 : \mathbf{y}' &= \mathbf{F}'_1\boldsymbol{\theta} + \sigma_0\mathbf{R}'\mathbf{n}_1\end{aligned}$$

where $\mathbf{y}' = \mathbf{P}_{\mathbf{Q}^\perp}\mathbf{y}$, $\mathbf{F}'_0 = \mathbf{P}_{\mathbf{Q}^\perp}\mathbf{F}_0$, $\mathbf{F}'_1 = \mathbf{P}_{\mathbf{Q}^\perp}\mathbf{F}_1$, and $\mathbf{R}' = \mathbf{P}_{\mathbf{Q}^\perp}\mathbf{R}$. Finally, note that when testing H_1 against H_0 (see Section V-D), one needs to premultiply both models by \mathbf{R}^{-1} before handling missing measurements.

B. Reducing Technical Alarms

In addition to missing measurements, we have observed multiple wrong measurements caused by similar factors. Furthermore, we have also observed correlations between

⁶To keep notation simple, we consider only the case of a single missing measurement, though the approach can be straightforwardly extended to multiple measurements. In fact, windows with multiple missing measurements frequently occur in the case study presented in Section VII; there, we use the extended framework to handle multiple missing measurements.

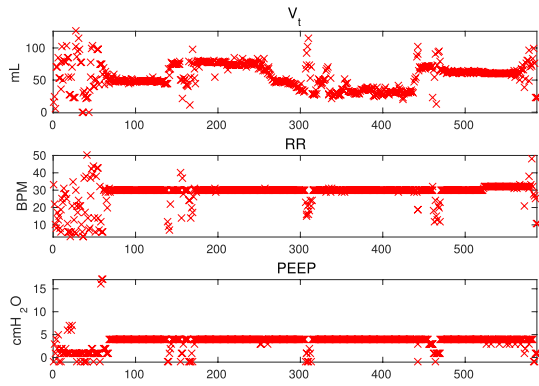


Fig. 6. Typical missing/bad data patterns in a surgery case at the CHOP. Missing measurements are set to -1 .

certain missing measurements and other bad measurements (since they are all provided by the same machine). These correlations have been incorporated in the detector in order to reduce the number of technical alarms, i.e., false alarms caused by bad measurements. As will be demonstrated in Section VII, technical alarms constitute a large portion of all false alarms.

The first observation was that positive end-expiratory pressure (PEEP) measurements are often missing when other measurements are bad, as can be seen in Fig. 6, which presents the V_t , RR and PEEP measurements in a typical case at the CHOP. Thus, while PEEP is not directly used in the model, it is used to reduce the number of false alarms. Specifically, as will be apparent in Section VII, the detector is run in a sliding window fashion, and any window with more than half missing PEEP measurements is discarded.

A second method to reduce technical alarms is through RR. RR is controlled by clinicians and is usually steady during the case (also shown in Fig. 6). Therefore, a measurement must be wrong if it is too far from the mode; specifically, if 90% of RR measurements are the same in a window and if another measurement is at least 2 breaths/min away from the mode, then it is treated as bad. When a bad RR measurement is detected, the threshold for raising an alarm is doubled.

VII. CASE STUDY

To evaluate the performance of the detector, we use real patient data from lobectomy surgeries performed on infants at the CHOP during the period 2005–14. A lobectomy is the incision of a cystic lung lesion and requires mechanical ventilation; sometimes, at the request of the surgeon, one-lung ventilation is performed, which leads to a shunt and a potentially critical drop in the blood O_2 content. The dataset consists of 484 cases, of which 167 had shunts. This is a significant improvement upon [1] where we had 292 patients total and 97 with shunts.

The vital signs were sampled every 15 s, and thus each step k in the discrete-time model corresponded to 15 s. Note that the average time for the blood to circulate the body (i.e., the parameter τ in the continuous-time model) was about a minute in adults but much shorter in infants [66]. Thus, in the discrete-time model, the best choice that matches observed data [66] was $\tau \approx 30$ s, i.e., $\kappa = 2$ time steps. The detector

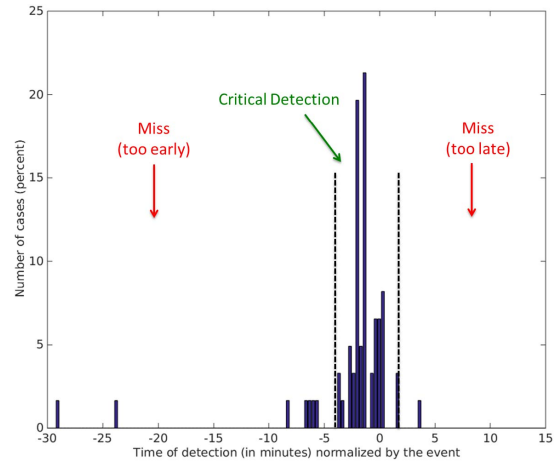


Fig. 7. Histogram showing how many minutes in advance each critical event was detected. Negative times indicate early detection.

was run in a sliding window fashion; a window of size $M = 34$ (i.e., 8.5 min) was used, whereas the detection window size was $M - M_T = 10$, i.e., 2.5 min. For the thresholds, t_0^* and t_1^* , we specify 1% and 0.01% as the desired false alarm and missed detection rates, respectively.

In the next sections, we describe the empirical detection and false alarm rates of the detector developed in this paper.

A. Detection Rate Evaluation

To evaluate the detection rate, we use all cases that had a shunt during the surgery. Note that the detector is designed to capture critical drops in the O_2 content caused by shunts; thus, a good detection occurs when the detector raises an alarm slightly before or very soon after such an event occurs. Defining these events in the data, however, is challenging; as argued in Section III, shunt annotations are not always very accurate; even when they are, shunts do not always cause immediate drops in the O_2 content. Therefore, we inspected the data and identified the nearest in time decrease in S_pO_2 as the actual event that the detector needs to predict.

The three vital signs that are used directly in our model are V_t , RR, and $EtCO_2$ ($EtCO_2$ measures the partial pressure of CO_2 at the end of the breath, i.e., P_eCO_2). Note that while the detector can handle missing measurements, it still does not make a decision if more than half of the PEEP measurements are missing in a given window, as described in Section VI-B. Thus, if an event occurs when the detector does not have enough data to make a decision, that case is removed from the dataset. Furthermore, all cases that have the event in the first 15 min were removed because measurements tend to be very noisy at the beginning of cases when patients are getting intubated and suctioned. After removing these cases, we retain 61 cases with shunts that we could evaluate the detection performance. Only 26 were used in [1]; with the current detector that can handle missing measurements, 40 of the old cases could be used.

Note that detectors are usually evaluated using receiver operating characteristic (ROC) curves. We cannot provide such a curve due to the fact that we do not know exactly when

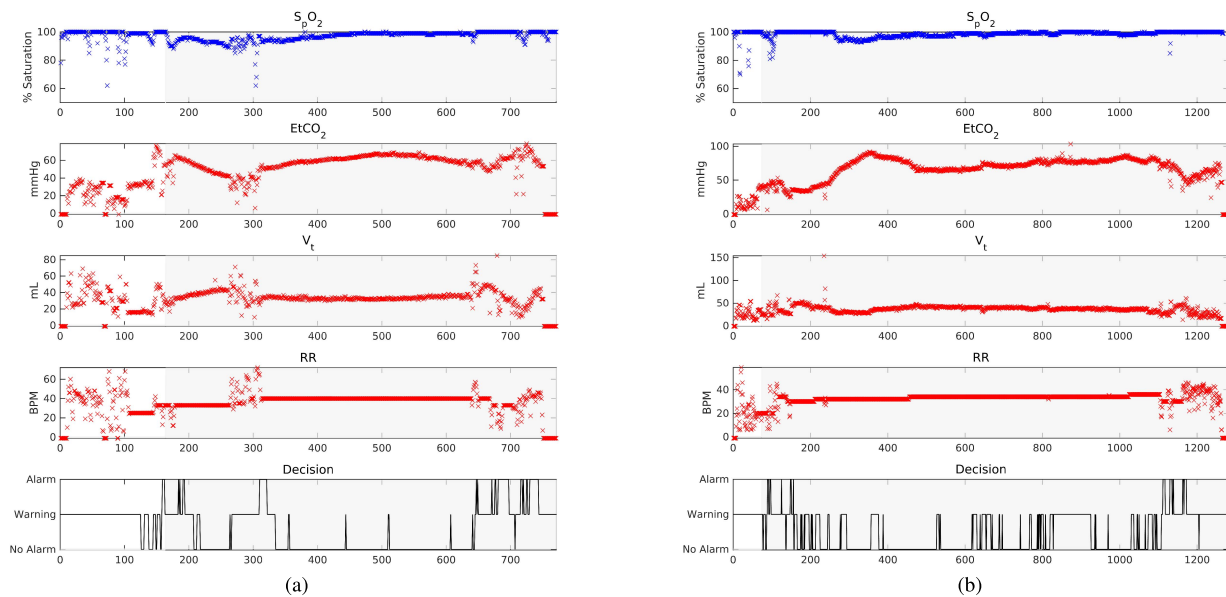


Fig. 8. Example cases for different scenarios. The shaded area denotes the occurrence of the isolation. (a) Example case with good prediction. (b) Example case with late detection.

events occur and which alarms can be treated as true or false; an ROC-like curve is shown in Section VII-C. Instead, this section provides a more practical evaluation of the detection rate—it gives the percentage of cases in which the event was predicted and it shows a histogram of how many minutes in advance the event was predicted.

To record the time of detection, note that several alarms may be raised before the event—taking the time of the last one as the time of detection may misrepresent the detector as alarming too late, whereas taking the earliest may mean it is alarming too early. Thus, we set a rule for recording the time of detection—it is the time of the alarm closest to 2 min before the event; that may provide clinicians enough time to take action and reverse (or even avoid) the effect of the shunt.

With this rule in mind, Fig. 7 provides a histogram with the detection times and how many minutes before the event the detection occurred. The vast majority of cases (87%) have alarms less than 5 min before the event and up to 2 min after. Note also that there are only two outliers that do not have alarms within 15 min of the event.

Fig. 7 shows that the detector is robust to differences across patients and to bad measurements. In addition, for the cases that have alarms less than 5 min before the event, clinicians would receive warning on average 90 s before the event. According to our clinical collaborators, this would provide them with enough time to identify the problem and apply the appropriate control response (e.g., deliver more O_2).

To better illustrate the performance of the detector, we provide example cases with good and late detection. Fig. 8(a) shows an example case with good detection, i.e., an alarm was raised that would alert clinicians to take proactive action and avoid the critical scenario. It contains the evolution of the vital signs as well as the decision made. As described above, note that the case starts with a period with noisy data, followed by the clear pattern of a rise in $EtCO_2$, which is

what eventually triggers the alarm. Note that the shunt likely occurred before the drop in S_pO_2 , but that would not be captured by the detector due to the noisy $EtCO_2$. Finally, note that there is a long period during the surgery when the detector is not alarming, which indicates that it is not prone to raise false alarms when good data are available.

Fig. 8(b) presents a case with late detection. One can observe the same bad/missing data patterns at the beginning of the case. The case is detected late because all vital signs are very noisy for a long period before the event. The detector does raise an alarm as soon as better data are available.

B. False Alarm Rate Evaluation

To evaluate the false alarm rate, we use the remaining 317 cases that did not have a shunt. Thus, most alarms in these cases are false. Yet, upon investigation of the cases, we found that many alarms are actually near periods with decreased S_pO_2 ; since these alarms would be beneficial to clinicians and since they still detect critical drops in the O_2 content, they are treated as true. Thus, we treat as true all alarms that precede by 10 min have a 2%-point (or more) drop in S_pO_2 or exceed it by 2 min. Any other alarm is considered false.

To count the false alarms in a way that would be meaningful to the medical community,⁷ the detector would be implemented as follows. If an alarm is raised, then it can be muted by a clinician; this alarm denotes the beginning of an event. If the following decisions of the detector are also alarms or warnings, then they are part of the same event, and hence, no further alarms are raised. When a no alarm decision is made, the event has ended, and the next alarm indicates a new event.

⁷The definition of false alarm rate in statistics is different from medicine. In the former, it is the number of false alarms divided by the number of tests; in the latter, it is the number of false alarms divided by the number of alarms.

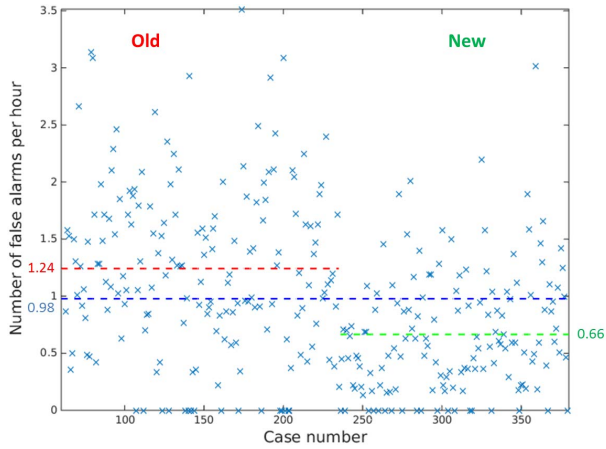


Fig. 9. Scatterplot with the average number of false alarms per hour for each of the cases with no shunts. The red dashed line denotes the average for the old dataset, the green dashed line the average for the new, and the blue dashed line the average for the both datasets combined.

For each case, we count the average number of events per hour, and all are treated as false (except the ones described above). A scatter plot with the average number of events per hour for each case is shown in Fig. 9. As can be seen in Fig. 9, the cases are reasonably uniformly distributed, with just a few cases with more than three false alarms per hour. This distribution is due to the constant false alarm rate property of the parameter-invariant detector—it achieves the same false alarm rate regardless of the unknown parameters.

For reference, the cases are divided into the ones from the old and new datasets. The overall average is 0.95 false alarms per hour, though for the old cases, it is much higher than the new ones—1.21 versus 0.64. The variance of the old dataset is visibly higher as well. This trend suggests that the improvement in sensors and measuring devices is naturally decreasing the false alarm rate of our detector by providing it with less noisy or bad data. Note also that the average false alarm rate is much lower than current threshold-based alarms [67], which caused the alarm fatigue problem.

In addition, note that the average false alarm rate in [1] with the old dataset was 2.15. In this work, through reduction of technical alarms, this is brought down to 1.21, with no cost in detection performance. We consider this rate a significant improvement and have implemented the detector in real time in an OR.

C. ROC Curves

As further evaluation of the proposed detector, we illustrate its performance using an ROC-like curve. In particular, each point on the curve represents a detection rate over the entire population (as computed in Fig. 7) corresponding to a false alarm rate over the entire population (as shown in Fig. 9). Fig. 10 shows the resulting curve. As can be seen, the detector’s power increases linearly with the false alarm rate, which once again shows the robustness of the approach.⁸

⁸Note that the curve is not monotonically increasing due to the way alarms are counted. When thresholds are varied, some alarm events are combined (or separated), thus sometimes affecting the false alarm rate differently than the detection rate.

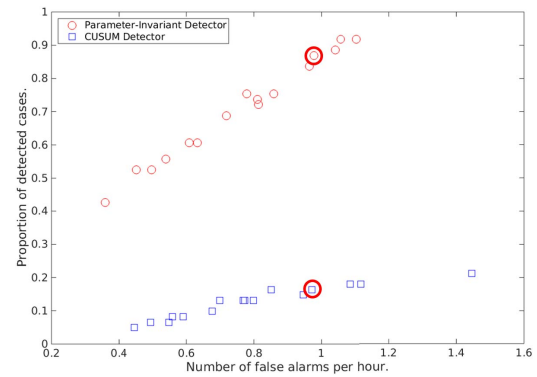


Fig. 10. ROC-like curve showing the performance of the parameter-invariant detector compared with that of a CUSUM detector. Operating points used in Figs. 7, 9, and 11 are circled in red.

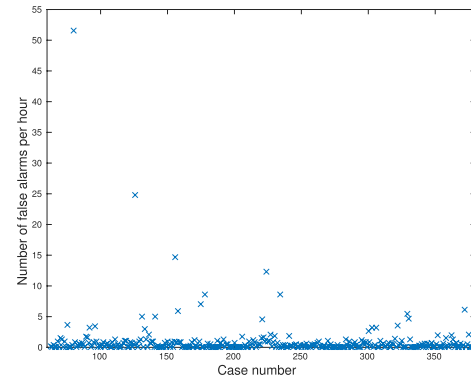


Fig. 11. Scatterplot with the average number of false alarms per hour generated by the CUSUM detector.

Finally, the detector is also compared with a standard change detection technique, namely, the cumulative sum control chart (CUSUM) detector, in order to show the merit of the parameter-invariant approach. In order to develop the CUSUM detector, the model parameters for each patient were estimated using the expectation maximization technique described in [68]; the detector algorithm was then borrowed from [69, Ch. 8.10]. As can be seen in Fig. 10, the CUSUM detector is greatly outperformed by the parameter-invariant detector. There are two main reasons for this difference.

- 1) The model developed in this paper captures general trends but is a poor predictor of the future that makes it unsuitable for model-predictive techniques such as the Kalman filter.
- 2) It is difficult to obtain good parameter estimates in the presence of noisy and missing measurements.

Finally, the CUSUM detector’s false alarm rate distribution is shown in Fig. 11 (with roughly the same false alarm rate as the parameter-invariant detector operating point shown in Fig. 9); as can be seen in Fig. 11, multiple patients receive more than 5 alarms per hour, which might create a serious burden on clinicians and might result in critical alarms being ignored.

VIII. CONCLUSION

In this paper, we addressed the problem of predicting critical drops in the blood O₂ content as caused by shunts in infants. We developed a physiological model of the cardiopulmonary system that is parameterized by several variables that vary across patients, e.g., metabolic rate. In order to guarantee a low false alarm rate regardless of the values of the parameters, we utilized a parameter-invariant detector—the first application of this class of detectors to the MCPS. To evaluate its performance, we used real patient data from surgeries performed at the CHOP and confirmed a near-constant false alarm rate across all patients. To further evaluate the quality of the detector, we are currently implementing it in real time in an OR.

APPENDIX

Proof of Proposition 1: From (22)

$$\begin{aligned} y(k) &= \frac{x^L(k) + x^R(k)}{2} \\ &= \left(\frac{2\bar{\alpha}}{\bar{V}(k)} \right) \frac{x^L(k - \kappa) + x^R(k - \kappa)}{2} + \frac{2\bar{\alpha}\mu}{\bar{V}(k)} + \sigma_0 n_0(k) \\ &= \frac{2y(k - \kappa)}{\bar{V}(k)} \bar{\alpha} + \frac{2}{\bar{V}(k)} \bar{\alpha}\mu + \sigma_0 n_0(k) \end{aligned}$$

where $\sigma_0 n_0(k) = (\sigma/2)(n^L(k) + n^R(k)) \sim \mathcal{N}(0, \sigma^2/2)$. \square

Proof of Proposition 2:

Case 1: By definition, for $\kappa \leq k \leq M_T$, $y(k)$ evolves according to \mathcal{M}_{NS} , and hence the recursive relation is the same

$$y(k) = \frac{2y(k - \kappa)}{\bar{V}(k)} \bar{\alpha} + \frac{2}{\bar{V}(k)} \bar{\alpha}\mu + \sigma_0 n_1(k).$$

Case 2: For $k > M_T$, the observation (with a right shunt) is

$$\begin{aligned} y(k) &= x^L(k) \quad (34) \\ &= \frac{\bar{\alpha}(x^L(k - \kappa) + x^R(k - \kappa))}{2\bar{V}(k)} + \frac{\bar{\alpha}\mu}{\bar{V}(k)} + \sigma n^L(k). \quad (35) \end{aligned}$$

Note that depending on the value of k , $x^L(k - \kappa)$ and $x^R(k - \kappa)$ may evolve according to either \mathcal{M}_{NS} or \mathcal{M}_S . In the case $M_T + 1 \leq k \leq M_T + \kappa$, (35) reduces to

$$y(k) = \frac{y(k - \kappa)}{\bar{V}(k)} \bar{\alpha} + \frac{1}{\bar{V}(k)} \bar{\alpha}\mu + \sqrt{2}\sigma_0 n_1(k) \quad (36)$$

since $x^L(k - \kappa)$ and $x^R(k - \kappa)$ are both observed through $y(k - \kappa)$ in (22).

Case 3: For $k > M_T + \kappa$, $y(k - \kappa) = x^L(k - \kappa)$, and therefore, (35) becomes

$$y(k) = \frac{y(k - \kappa)}{2\bar{V}(k)} \bar{\alpha} + \frac{x^R(k - \kappa)}{2\bar{V}(k)} \bar{\alpha} + \frac{1}{\bar{V}(k)} \bar{\alpha}\mu + \sigma n^L(k). \quad (37)$$

Note that, for $k > M_T + \kappa$, $x^R(k - \kappa)$ is unknown since $y(k - \kappa)$ measures the no-shunt partial pressure. Therefore, one can compute $x^R(k - \kappa)$ by following the

dynamics under \mathcal{M}_S . Thus, from (24), one can derive the following for $k > M_T + \kappa$:

$$x^R(k - \kappa) = \frac{x^L(k - 2\kappa) + x^R(k - 2\kappa)}{2} + \mu. \quad (38)$$

Using the same observation as in (35), we split the analysis in two cases, depending on whether $x^L(k - 2\kappa)$ and $x^R(k - 2\kappa)$ evolve according to \mathcal{M}_{NS} or \mathcal{M}_S .

Case 3a: $M_T + \kappa + 1 \leq k \leq M_T + 2\kappa$. This reduces similarly to (36)

$$x^R(k - \kappa) = y(k - 2\kappa) + \mu. \quad (39)$$

Case 3b: $k > M_T + 2\kappa$. Again, in this case $y(k - 2\kappa) = x^L(k - 2\kappa)$, and hence (38) now becomes

$$x^R(k - \kappa) = (1/2)y(k - 2\kappa) + (1/2)x^R(k - 2\kappa) + \mu. \quad (40)$$

Therefore, one may compute $x^R(k - \kappa)$ in the above recursion by recursively substituting the formula for $x^R(k - n\kappa)$, until a time step $p \leq M_T + 2\kappa$ is reached such that $x^R(p)$ can be computed according to Case 3a

$$\begin{aligned} x^R(k - \kappa) &= \frac{y(k - (L_k + 2)\kappa)}{2^{L_k}} + \frac{\mu}{2^{L_k - 1}} \\ &\quad + \sum_{n=1}^{L_k} \frac{y(k - (n+1)\kappa)}{2^n} + \frac{\mu}{2^{n-1}} \quad (41) \end{aligned}$$

where $L_k = \lfloor (k - M_T - \kappa)/\kappa \rfloor \geq 1$.

Thus, by combining Case 3a and 3b and by grouping the terms according to whether they contain μ , $x^R(k - \kappa)$ can be represented as [f and g are derived from (39) and (41)]

$$x^R(k - \kappa) = f(k - \kappa) + g(k - \kappa)\mu, \quad k > M_T + \kappa. \quad (42)$$

Finally, one can substitute (42) for $x^R(k - \kappa)$ in (37), thus obtaining the desired form of $y(k)$. \square

Proof of Proposition 3: To show invariance,

$$t(g(\mathbf{y})) = \mathbf{P}_{\mathbf{F}_0^\perp}(\mathbf{y} + \mathbf{F}_0\boldsymbol{\theta}) = \mathbf{P}_{\mathbf{F}_0^\perp}\mathbf{y} = t(\mathbf{y}).$$

To show optimality,

$$\begin{aligned} t(\mathbf{y}) &= t(\mathbf{y}'), \text{ i.e., } \mathbf{P}_{\mathbf{F}_0^\perp}\mathbf{y} = \mathbf{P}_{\mathbf{F}_0^\perp}\mathbf{y}' \\ \text{i.e., } \mathbf{y} &= \mathbf{y}' + \mathbf{P}_{\mathbf{F}_0}\mathbf{y} - \mathbf{y}' \\ \text{i.e., } \exists \boldsymbol{\theta} &\in \mathbb{R}^2, \mathbf{y} = \mathbf{y}' + \mathbf{F}_0\boldsymbol{\theta} \\ \text{i.e., } \exists g &\in G, \mathbf{y} = g(\mathbf{y}'). \end{aligned}$$

\square

Proof of Proposition 4: Let \mathbf{F}_0 be written as

$$\mathbf{F}_0 = [\mathbf{f}_{0,0} \quad \mathbf{f}_{0,1}].$$

Then $\mathbf{f}_{0,0}$ can be written as $\mathbf{f}_{0,0} = \bar{\mathbf{V}} \circ 2\mathbf{S}_0\bar{\mathbf{y}}$, where

$$\mathbf{S}_0 = [\mathbf{I}_{M-\kappa} \mid \mathbf{0}_{M-\kappa \times \kappa}]$$

where $\mathbf{I}_{M-\kappa} \in \mathbb{R}^{M-\kappa}$ is the identity matrix and $\mathbf{0}_{M-\kappa \times \kappa} \in \mathbb{R}^{M-\kappa \times \kappa}$ is a matrix of 0s. Thus

$$\mathbf{K}_0 = 2\mathbf{S}_0.$$

Similarly, $\mathbf{f}_{0,1} = \bar{\mathbf{V}} \circ (2 \times \mathbf{1}_{M-\kappa \times 1})$, where $\mathbf{1}_{M-\kappa \times 1} \in \mathbb{R}^{M-\kappa \times 1}$ is the vector of all 1s. Thus

$$\mathbf{k}_0 = 2 \times \mathbf{1}_{M-\kappa \times 1}.$$

- [13] H. T. Milhorn, *Application of Control Theory to Physiological Systems*. Philadelphia, PA, USA: W. B. Saunders, 1966.
- [14] A. Konkani, B. Oakley, and T. J. Bauld, "Reducing hospital noise: A review of medical device alarm management," *Biomed. Instrum. Technol.*, vol. 46, no. 6, pp. 478–487, 2012.
- [15] B. R. Eggins, *Biosensors: An Introduction*. Chichester, U.K.: Wiley, 1996.
- [16] F. E. Block, Jr., L. Nuutinen, and B. Ballast, "Optimization of alarms: A study on alarm limits, alarm sounds, and false alarms, intended to reduce annoyance," *J. Clin. Monitor. Comput.*, vol. 15, no. 2, pp. 75–83, Feb. 1999.
- [17] C. L. Tsien and J. C. Fackler, "Poor prognosis for existing monitors in the intensive care unit," *Critical Care Med.*, vol. 25, no. 4, pp. 614–619, Apr. 1997.
- [18] M. Gorges, B. A. Markewitz, and D. R. Westenskow, "Improving alarm performance in the medical intensive care unit using delays and clinical context," *Anesthesia Analgesia*, vol. 108, no. 5, pp. 1546–1552, May 2009.
- [19] J. Edworthy and E. Hellier, "Fewer but better auditory alarms will improve patient safety," *Quality Safety Health Care*, vol. 14, no. 3, pp. 212–215, 2005.
- [20] S. Siebig, S. Kuhls, M. Imhoff, U. Gather, J. Schölmerich, and C. E. Wrede, "Intensive care unit alarms—How many do we need?" *Critical Care Med.*, vol. 38, no. 2, pp. 451–456, Feb. 2010.
- [21] M. Cvach, "Monitor alarm fatigue: An integrative review," *Biomed. Instrum. Technol.*, vol. 46, no. 4, pp. 268–277, Jul./Aug. 2012.
- [22] K. J. Åström and P. Eykhoff, "System identification—A survey," *Automatica*, vol. 7, no. 2, pp. 123–162, Mar. 1971.
- [23] L. Ljung and S. Gunnarsson, "Adaptation and tracking in system identification—A survey," *Automatica*, vol. 26, no. 1, pp. 7–21, Jan. 1990.
- [24] R. Isermann, "Process fault detection based on modeling and estimation methods—A survey," *Automatica*, vol. 20, no. 4, pp. 387–404, Jul. 1984.
- [25] O. Nelles, *Nonlinear System Identification: From Classical Approaches to Neural Networks and Fuzzy Models*. New York, NY, USA: Springer, 2001.
- [26] H. J. A. F. Tulleken, "Grey-box modelling and identification using physical knowledge and Bayesian techniques," *Automatica*, vol. 29, no. 2, pp. 285–308, Mar. 1993.
- [27] T. Bohlin and S. F. Graebe, "Issues in nonlinear stochastic grey box identification," *Int. J. Adapt. Control Signal Process.*, vol. 9, no. 6, pp. 465–490, Nov./Dec. 1995.
- [28] N. R. Kristensen, H. Madsen, and S. B. Jørgensen, "Parameter estimation in stochastic grey-box models," *Automatica*, vol. 40, no. 2, pp. 225–237, Feb. 2004.
- [29] A. Juditsky *et al.*, "Nonlinear black-box models in system identification: Mathematical foundations," *Automatica*, vol. 31, no. 12, pp. 1725–1750, Dec. 1995.
- [30] J. Sjöberg *et al.*, "Nonlinear black-box modeling in system identification: A unified overview," *Automatica*, vol. 31, no. 12, pp. 1691–1724, Dec. 1995.
- [31] J. A. K. Suykens and J. P. L. Vandewalle, Eds., *Nonlinear Modeling: Advanced Black-Box Techniques*. New York, NY, USA: Springer, 1998.
- [32] C. Cobelli and E. Carson, *Introduction to Modeling in Physiology and Medicine*. New York, NY, USA: Academic, 2008.
- [33] N. A. Trayanova, "Whole-heart modeling: Applications to cardiac electrophysiology and electromechanics," *Circulat. Res.*, vol. 108, no. 1, pp. 113–128, 2011.
- [34] L. Magni *et al.*, "Model predictive control of type 1 diabetes: An *in silico* trial," *J. Diabetes Sci. Technol.*, vol. 1, no. 6, pp. 804–812, 2007.
- [35] J. Chen and R. J. Patton, *Robust Model-Based Fault Diagnosis for Dynamic Systems*. Norwell, MA, USA: Kluwer, 2012.
- [36] A. S. Willsky, "A survey of design methods for failure detection in dynamic systems," *Automatica*, vol. 12, no. 6, pp. 601–611, Nov. 1976.
- [37] A. Burgos, A. Goñi, A. Illarramendi, and J. Bermúdez, "Real-time detection of apneas on a PDA," *IEEE Trans. Inf. Technol. Biomed.*, vol. 14, no. 4, pp. 995–1002, Jul. 2010.
- [38] V. A. Convertino *et al.*, "Use of advanced machine-learning techniques for noninvasive monitoring of hemorrhage," *J. Trauma-Injury, Infection, Critical Care*, vol. 71, no. 1, pp. S25–S32, Jul. 2011.
- [39] T. M. Mitchell, R. Hutchinson, M. A. Just, R. S. Niculescu, F. Pereira, and X. Wang, "Classifying instantaneous cognitive states from fMRI data," in *Proc. AMIA Annu. Symp.*, 2003, pp. 465–469.
- [40] A. Pantelopoulou and N. G. Bourbakis, "A survey on wearable sensor-based systems for health monitoring and prognosis," *IEEE Trans. Syst., Man, Cybern. C, Appl. Rev.*, vol. 40, no. 1, pp. 1–12, Jan. 2010.
- [41] S. Saria, D. Koller, and A. Penn, "Learning individual and population level traits from clinical temporal data," in *Proc. Neural Inf. Process. Syst.*, 2010, pp. 1–9.
- [42] S. Saria, A. K. Rajani, J. Gould, D. Koller, and A. A. Penn, "Integration of early physiological responses predicts later illness severity in preterm infants," *Sci. Transl. Med.*, vol. 2, no. 48, pp. 48–65, Sep. 2010.
- [43] M. Stacey and C. McGregor, "Temporal abstraction in intelligent clinical data analysis: A survey," *Artif. Intell. Med.*, vol. 39, no. 1, pp. 1–24, Jan. 2007.
- [44] C. Bishop, *Pattern Recognition and Machine Learning*, vol. 1. New York, NY, USA: Springer, 2006.
- [45] A. F. Simpaio, E. Y. Pruitt, S. D. Cook-Sather, H. G. Gurnaney, and M. A. Rehman, "The reliability of manual reporting of clinical events in an anesthesia information management system (AIMS)," *J. Clin. Monitor. Comput.*, vol. 26, no. 6, pp. 437–439, Dec. 2012.
- [46] K.-P. Adlassnig, C. Combi, A. K. Das, E. T. Keravnou, and G. Pozzi, "Temporal representation and reasoning in medicine: Research directions and challenges," *Artif. Intell. Med.*, vol. 38, no. 2, pp. 101–113, Oct. 2006.
- [47] J. Weimer, D. Varagnolo, and K. H. Johansson, "Distributed model-invariant detection of unknown inputs in networked systems," in *Proc. Int. Conf. High Confidence Netw. Syst.*, 2013, pp. 127–134.
- [48] J. Weimer, D. Varagnolo, M. S. Stankovic, and K. H. Johansson, "Parameter-invariant detection of unknown inputs in networked systems," in *Proc. IEEE 52nd Annu. Conf. Decision Control*, Dec. 2013, pp. 4379–4384.
- [49] J. Weimer *et al.*, "Active actuator fault detection and diagnostics in HVAC systems," in *Proc. 4th Workshop Embedded Sens. Syst. Energy-Efficiency Buildings*, 2012, pp. 107–114.
- [50] J. Weimer, S. Kar, and K. H. Johansson, "Distributed detection and isolation of topology attacks in power networks," in *Proc. Int. Conf. High Confidence Netw. Syst.*, 2012, pp. 65–72.
- [51] A. Roederer, J. Weimer, J. Dimartino, J. Gutsche, and I. Lee, "Towards non-invasive monitoring of hypovolemia in intensive care patients," in *Proc. Med. Cyber-Phys. Syst. Workshop*, 2015, pp. 23–30.
- [52] A. Roederer, J. Weimer, J. DiMartino, J. Gutsche, and I. Lee, "Robust monitoring of hypovolemia in intensive care patients using photoplethysmogram signals," in *Proc. 37th Int. Conf. IEEE Eng. Med. Biol. Soc. (EMBC)*, Aug. 2015, pp. 1504–1507.
- [53] S. Chen, J. Weimer, M. R. Rickels, A. Peleckis, and I. Lee, "Towards a model-based meal detector for type I diabetics," in *Proc. Med. Cyber-Phys. Syst. Workshop*, 2015, pp. 13–22.
- [54] Roche Cobas B 123 POC System, accessed on Jul. 13, 2015. [Online]. Available: <http://www.roche.com/prod-ucts/product-details.htm?region=us&type=product&id=146>
- [55] Shenzhen Bestman Instrument Co. Pulse Oximeter, accessed on Jul. 13, 2015. [Online]. Available: <http://www.szbest-man.com/contents/76/669.html>
- [56] The Dräger Apollo Anesthesia Machine, accessed on Jul. 13, 2015. [Online]. Available: http://www.draeger.com/sites/enus_us/Pages/Hospital/Apollo.aspx
- [57] R. K. Stoelting and S. C. Hillier, *Pharmacology & Physiology in Anesthetic Practice*. Philadelphia, PA, USA: Lippincott Williams & Wilkins, 2006.
- [58] J. B. West, *Respiratory Physiology: The Essentials*. Philadelphia, PA, USA: Lippincott Williams & Wilkins, 2012.
- [59] S. M. Donn and W. Boon, "Mechanical ventilation of the neonate: Should we target volume or pressure?" *Respirat. Care*, vol. 54, no. 9, pp. 1236–1243, 2009.
- [60] W. O. Fenn, H. Rahn, and A. B. Otis, "A theoretical study of the composition of the alveolar air at altitude," *Amer. J. Physiol.*, vol. 146, no. 5, pp. 637–653, 1946.
- [61] S. Cruickshank and N. Hirschauer, "The alveolar gas equation," *Continuing Educ. Anaesthesia, Critical Care Pain*, vol. 4, no. 1, pp. 24–27, 2004.
- [62] J. L. Theissen, S. R. Fischer, L. D. Traber, and D. L. Traber, "Pulmonary blood flow regulation: Influence of positive pressure ventilation," *Respirat. Physiol.*, vol. 102, nos. 2–3, pp. 251–260, Dec. 1995.
- [63] D. H. Bergel, *Cardiovascular Fluid Dynamics*, vol. 1. Amsterdam, The Netherlands: Elsevier, 2012.
- [64] R. E. Klabunde, *Cardiovascular Physiology Concepts*. Philadelphia, PA, USA: Lippincott Williams & Wilkins, 2011.
- [65] A. C. P. Powles and E. J. M. Campbell, "How to be less invasive," *Amer. J. Med.*, vol. 67, no. 1, pp. 98–104, Jul. 1979.

- [66] B. B. Oberst and F. La Roche, "Circulation time in the newborn infant, using the fluorescein dye method," *J. Pediatrics*, vol. 45, no. 5, pp. 580–582, Nov. 1954.
- [67] A. Aboukhalil, L. Nielsen, M. Saeed, R. G. Mark, and G. D. Clifford, "Reducing false alarm rates for critical arrhythmias using the arterial blood pressure waveform," *J. Biomed. Informat.*, vol. 41, no. 3, pp. 442–451, Jun. 2008.
- [68] R. H. Shumway and D. S. Stoffer, "An approach to time series smoothing and forecasting using the EM algorithm," *J. Time Ser. Anal.*, vol. 3, no. 4, pp. 253–264, Jul. 1982.
- [69] F. Gustafsson, *Adaptive Filtering and Change Detection*, vol. 1. New York, NY, USA: Wiley, 2000.



Radoslav Ivanov (S'15) received the B.A. degree in computer science and economics from Colgate University, Hamilton, NY, USA. He is currently pursuing the Ph.D. degree with the Computer and Information Science Department, University of Pennsylvania, Philadelphia, PA, USA.

His current research interests include the design and control of cyber-physical systems (CPS), in particular, automotive and medical CPS, and predictive and retrospective analysis of medical patient data.



James Weimer (M'06) received the B.S. degree in electrical engineering from Purdue University, West Lafayette, IN, USA, in 2005, and the M.S. and Ph.D. degrees in electrical and computer engineering from Carnegie Mellon University, Pittsburgh, PA, USA, in 2007 and 2010, respectively.

He was a Post-Doctoral Researcher with the Department of Automatic Control, KTH Royal Institute of Technology, Stockholm, Sweden. He is currently a Post-Doctoral Researcher with the Department of Information and Computer Science,

University of Pennsylvania, Philadelphia, PA, USA. His current research interests include the design and analysis of cyber-physical systems with application to medical devices/monitors, networked systems, building energy management, and smart grids.



Allan F. Simpao received the M.D. with Jefferson Medical College, Philadelphia, PA, USA and M.B.I. with Oregon Health & Science University, Portland, OR, USA degrees.

He has lectured nationally on anesthesia and informatics topics, including the use of big data in pediatrics, social media ethics and professionalism, and physicians' use of social media for networking, education, and advocacy. He is currently a Board-Certified Pediatric Anesthesiologist and Clinical Informatician. His clinical expertise is in pediatric

and fetal anesthesia. His current research interests include clinical informatics and applying visual analytics to health care databases for outcomes analysis and quality improvement.



Mohamed A. Rehman is currently one of the first Pediatric Anesthesiologists and Critical Care Physicians formally trained in biomedical informatics. He is an Innovative Leader who established the first pediatric anesthesia data warehouse. During his current tenure at The Children's Hospital of Philadelphia (CHOP), Philadelphia, PA, USA, he provides visionary leadership by directing the first Pediatric Anesthesia and Critical Care Biomedical Informatics Group in the nation. He has been an Endowed Chair in Biomedical Informatics and

Entrepreneurial Sciences with CHOP.

Dr. Rehman received the Biomedical Informatics Training from Oregon Health and Science University, Portland, OR, USA, and became board certified in the new Accreditation Council for Graduate Medical Education subspecialty of clinical informatics.



Insup Lee (S'80–M'83–SM'97–F'01) received the Ph.D. degree in computer science from the University of Wisconsin–Madison, Madison, WI, USA.

He is currently the Cecilia Fitter Moore Professor of Computer and Information Science and serves as the Director of the PRECISE Center with the University of Pennsylvania, Philadelphia, PA, USA, where he holds a secondary appointment with the Department of Electrical and Systems Engineering. His current research interests include cyber-physical systems, real-time and embedded systems, run-time

assurance and verification, trust management, and high-confidence medical systems.

Dr. Lee received the IEEE TC-RTS Outstanding Technical Achievement and Leadership Award in 2008.