4-14-2015

Early Detection of Critical Pulmonary Shunts in Infants

Radoslav Ivanov

University of Pennsylvania, rivanov@seas.upenn.edu

James Weimer

University of Pennsylvania, weimerj@seas.upenn.edu

Allan F. Simpao

University of Pennsylvania, simpaoa@email.chop.edu

Mohamed A. Rehman

University of Pennsylvania, rehman@email.chop.edu

Insup Lee

University of Pennsylvania, lee@cis.upenn.edu

Follow this and additional works at: https://repository.upenn.edu/cis_papers

Part of the [Computer Engineering Commons](#), and the [Computer Sciences Commons](#)

Recommended Citation

Radoslav Ivanov, James Weimer, Allan F. Simpao, Mohamed A. Rehman, and Insup Lee, "Early Detection of Critical Pulmonary Shunts in Infants", *6th International Conference on Cyber-Physical Systems (ICCPS 2015)*, 110-119. April 2015. <http://dx.doi.org/10.1145/2735960.2735962>

6th International Conference on Cyber-Physical Systems (ICCPS 2015) <http://iccps.acm.org/2015/> in conjunction with CPS Week 2015 <http://www.cpsweek.org/2015/>, Seattle, WA, April 14-16, 2015

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/cis_papers/780
For more information, please contact repository@pobox.upenn.edu.

Early Detection of Critical Pulmonary Shunts in Infants

Abstract

This paper aims to improve the design of modern Medical Cyber Physical Systems through the addition of supplemental noninvasive monitors. Specifically, we focus on monitoring the arterial blood oxygen content (C_aO_2), one of the most closely observed vital signs in operating rooms, currently measured by a proxy - peripheral hemoglobin oxygen saturation (S_pO_2). While S_pO_2 is a good estimate of O_2 content in the finger where it is measured, it is a delayed measure of its content in the arteries. In addition, it does not incorporate system dynamics and is a poor predictor of future C_aO_2 values. Therefore, as a first step towards supplementing the usage of S_pO_2 , this work introduces a predictive monitor designed to provide early detection of critical drops in C_aO_2 caused by a pulmonary shunt in infants.

To this end, we develop a formal model of the circulation of oxygen and carbon dioxide in the body, characterized by unknown patient-unique parameters. Employing the model, we design a matched subspace detector to provide a near constant false alarm rate invariant to these parameters and modeling uncertainties. Finally, we validate our approach on real-patient data from lung lobectomy surgeries performed at the Children's Hospital of Philadelphia. Given 198 infants, the detector predicted 81% of the critical drops in C_aO_2 at an average of about 65 seconds earlier than the S_pO_2 -based monitor, while achieving a 0.9% false alarm rate (representing about 2 false alarms per hour).

Keywords

Time series analysis, Medical Information Systems

Disciplines

Computer Engineering | Computer Sciences

Comments

6th International Conference on Cyber-Physical Systems (ICCPS 2015) <http://iccps.acm.org/2015/> in conjunction with CPS Week 2015 <http://www.cpsweek.org/2015/>, Seattle, WA, April 14-16, 2015

Early Detection of Critical Pulmonary Shunts in Infants*

Radoslav Ivanov James Weimer
Dept. of Computer & Information Sc.
University of Pennsylvania
{rivanov, weimerj}@seas.upenn.edu

Allan Simpao Mohamed Rehman
Dept. of Anesthesiology and Critical Care Medicine
Children's Hospital of Philadelphia
{simpaoa, rehman}@email.chop.edu

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

ABSTRACT

This paper aims to improve the design of modern Medical Cyber Physical Systems through the addition of supplemental noninvasive monitors. Specifically, we focus on monitoring the arterial blood oxygen content (C_aO_2), one of the most closely observed vital signs in operating rooms, currently measured by a proxy – peripheral hemoglobin oxygen saturation (S_pO_2). While S_pO_2 is a good estimate of O_2 content in the finger where it is measured, it is a delayed measure of its content in the arteries. In addition, it does not incorporate system dynamics and is a poor predictor of future C_aO_2 values. Therefore, as a first step towards supplementing the usage of S_pO_2 , this work introduces a predictive monitor designed to provide early detection of critical drops in C_aO_2 caused by a pulmonary shunt in infants.

To this end, we develop a formal model of the circulation of oxygen and carbon dioxide in the body, characterized by unknown patient-unique parameters. Employing the model, we design a matched subspace detector to provide a near constant false alarm rate invariant to these parameters and modeling uncertainties. Finally, we validate our approach on real-patient data from lung lobectomy surgeries performed at the Children's Hospital of Philadelphia. Given 198 infants, the detector predicted 81% of the critical drops in C_aO_2 at an average of about 65 seconds earlier than the S_pO_2 -based monitor, while achieving a 0.9% false alarm rate (representing about 2 false alarms per hour).

Categories and Subject Descriptors

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

1. INTRODUCTION

With the proliferation of measuring devices in modern intensive care units (ICUs) and operating rooms (ORs), much effort has recently been devoted to the development of Medical Cyber Physical Systems (MCPS). By interpreting multiple vital signs quickly and efficiently, they not only provide another dimension of information to clinicians but may also be used to close the loop and perform certain tasks independently (e.g., analgesia infusion [23]). In addition, by analyzing

the trends in vital signs, these systems can be predictive in nature, thereby allowing physicians to take proactive measures to address deterioration in a patient's condition.

In this work, we explore one aspect of the design of MCPS, namely developing robust vital sign monitors and, in particular, the challenge of detecting sharp decreases in arterial blood oxygen content (C_aO_2)¹ caused by pulmonary shunts in infants. A shunt is a condition in which a lung does not participate in pulmonary exchange, i.e., it is not supplied with fresh air. Shunts may lead to a drop in the amount of O_2 and are especially hazardous in infants because of their limited O_2 reserve and higher O_2 consumption rate per kilogram of body weight when compared to adults.

A shunt may occur naturally (e.g., due to pulmonary edema), but it also occurs in surgical settings, when mechanical ventilation through an endotracheal tube is required. While it may happen inadvertently (i.e., the tube slips and does not ventilate one lung), a surgical shunt (also called an isolation) is often intentional as governed by the type of operation. For instance, during a lung lobectomy procedure for the excision of a cystic lung lesion [25],² a surgeon can perform either an open thoracotomy or a thoracoscopic approach. In the former, two-lung ventilation is common since the surgeon can directly see the lung when removing the cyst through an incision in the pleural space of the chest. The latter approach, which usually results in less blood loss but may lead to a drop in C_aO_2 , requires one-lung ventilation to provide a still operating environment and better visualization via an endoscope inserted through a small incision on the side of the chest [17]. During a lobectomy, especially in the presence of one-lung ventilation (and subsequent shunt), careful monitoring of C_aO_2 is necessary to ensure the patient's health does not deteriorate.

Currently, C_aO_2 is measured directly using blood gas analysis; however, this method requires drawing blood from the infant and cannot be utilized continuously. As a non-invasive alternative, C_aO_2 is monitored based on a measurement of the oxygen saturation (in percent) in the peripheral capillaries, denoted by S_pO_2 . Yet, when a shunt occurs, a decrease in the C_aO_2 is likely to also take place; physicians might not see it happen, however, until they receive a measurement of low S_pO_2 (typically measured at the tip of a finger),

*This research was supported in part by NSF CNS-1035715 and by the DGIST Research and Development Program of the Ministry of Science, ICT and Future Planning of Korea (CPS Global Center).

¹ C_aO_2 is the quantity of oxygen in arterial blood that is both bound and unbound to hemoglobin.

²A congenital cystic lung lesion is a dangerous tissue malformation that may cause infection or malignancy later in life [25], and thus warrants surgical resection in children.

at which point the child may already be in a critical state. Thus, while S_pO_2 is a good measure of oxygen content in the extremities, it is a delayed measure of C_aO_2 .

Therefore, in this work, we aim to use other available measurements to predict a critical decrease in C_aO_2 as caused by a shunt. This problem is made challenging by the lack of observability – while multiple pulmonary measurements are available (e.g., the partial pressure of CO_2 in exhaled air, denoted by $EtCO_2$), S_pO_2 is the only blood variable measured routinely in real time. Consequently, we leverage laws of physics (e.g., ideal gas law) and published physiological data trends to develop a formal parameterized model of the dynamics of the partial pressures of O_2 and CO_2 in the blood and relate them to the corresponding measured values in the airways. The challenge is that many of these parameters are unknown (e.g., baseline values of $EtCO_2$) and some cannot even be measured (e.g., lung development).

Even with uncertain models, one can still design accurate monitors by deriving sufficient statistics that are invariant to the unknown parameters [30]. We apply this technique to obtain such a statistic based on the expected behavior of the measured vital signs according to the model developed in this work. We employ the statistic to design a matched subspace detector that estimates the likelihood of the two tested scenarios (i.e., shunt vs. no shunt) and makes a decision based on observed data in favor of the more likely one. Since “alarm fatigue” is known to be a major issue in ORs, the detector is specified to have a constant false alarm rate for all patients. Additionally, such a system is a first step towards predictive closed-loop control – by identifying the cause of low C_aO_2 , the controller can take the appropriate remedial action, e.g., increase the amount of delivered O_2 , which is the typical response applied by physicians. To our knowledge, this work represents the first application of this class of detectors to MCPS monitors.

Finally, we validate our approach on data from open thoracotomy and thoracoscopy surgeries performed at the Children’s Hospital of Philadelphia (CHOP) over the last decade. In particular, we examine cases when a shunt results in a drop in S_pO_2 and record the detector’s predictive accuracy as compared with S_pO_2 measurements. Based on 26 thoracoscopic cases, the detector is able to predict 81% of the critical drops in C_aO_2 at an average of about 65 seconds before a drop in S_pO_2 occurs, and achieves an average false alarm rate of 0.9% on the 172 open thoracotomy cases (corresponding to 2.15 false alarms per hour).

In summary, this work’s contributions include: (a) the development of a dynamic model for the blood and pulmonary partial pressures in the presence of shunts; (b) the first application of parameter-invariant monitors in MCPS; (c) a case study evaluation on real-patient data.

The remainder of this paper is organized as follows. The following section describes the related work on MCPS. Section 3 specifies the problem addressed in this paper as well as the current approach. Section 4 presents a physiological model of O_2 and CO_2 circulation and describes which variables are currently measured. Section 5 introduces the monitor used in this work. A case study is provided in Section 6 and the final section provides conclusions.

2. RELATED WORK ON MEDICAL CPS

The related work on MCPS is divided into two main areas: formal verification and anomaly detection. In the former, researchers begin with a known model of the system, e.g., a

pacemaker [9, 18] or an infusion pump [23], and use formal methods to verify that it does not endanger the patient’s safety [6]. In the latter, a system is already in place that may lead to an adverse event, and the challenge is to detect or predict that event. This work fits in the second area, hence we review the related work on anomaly detection.

There are many approaches to performing anomaly detection and hypothesis testing. Model-based approaches have long been the standard approach [11, 39]. For example, one may assume each scenario follows a certain distribution so that testing can be performed using the likelihood ratio [39]. Yet, most MCPS applications are modeled using parameters that contain uncertainty. A common approach in such cases is to optimize the worst-case performance [18]. In our work, however, there exists no known model that would allow one to perform worst-case analysis.

When models are not available, data-driven machine learning (ML) approaches leverage known classifiers [8]. These have been successfully applied in many medical applications [10, 12, 21, 24, 28, 29]; however, they all require rich training data with accurate event annotations. Yet, during surgery clinicians’ priority is to serve the patient’s needs rather than provide accurate annotations, thus event labels are often delayed or missing altogether [32]. Therefore, due to the high number of potential variables (e.g., body mass, lung development, metabolic rate) and label inaccuracies, it is unlikely that unaided ML will yield accurate monitors [24]. The challenges and potential methods for improving data-driven approaches in MCPS are discussed in [5, 20].

An alternative approach that performs well with scarce data and uncertain models is the parameter-invariant detector [30]. It utilizes sufficient statistics that are invariant to unknown parameters to achieve a constant false alarm rate [30]. Although not previously applied to MCPS monitors, these detectors have been shown to work well in other CPS applications with structured dynamics and unknown parameters, specifically in detecting faults in networked systems [36, 37], building heating, ventilating and air conditioning (HVAC) systems [34], and smart grids [35].

3. PROBLEM FORMULATION

This section describes the problem addressed in this work. We consider infants undergoing surgery under general anesthesia. Depending on the type of operation, oxygen is delivered to patients either through one- or two-lung ventilation. Regardless of surgery type, a shunt may occur, either due to an intended isolation of a lung or a slip of the tube into a mainstem bronchus. During an isolation, only one lung supplies oxygen to the blood, hence C_aO_2 may decrease. In this section we describe the current approach to monitoring C_aO_2 , and present the problem statement for this work.

3.1 Current Approach to Monitoring

Currently, clinicians do not have a wide variety of tools to measure C_aO_2 ; however, it relates to other physiological variables through the oxygen content equation [31]:

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

where S_aO_2 is the arterial hemoglobin oxygen saturation in percent, Hb denotes the amount of hemoglobin in grams per deciliter, and P_aO_2 is the arterial partial pressure of dissolved oxygen measured in *mmHg* (millimeters of mercury). Note that the majority of oxygen in the blood is bound to hemoglobin, making S_aO_2 a good indicator of the C_aO_2

Figure 1: Measurement devices currently available to clinicians.

value [31]; however, S_aO_2 is not easily measured. Since the oxygen saturation (and content) is constant as the blood travels to the peripheral capillaries [38], clinicians use S_pO_2 , as measured in real time by a non-invasive pulse oximeter (see Figure 1a), to monitor C_aO_2 .

In addition to the S_pO_2 measurement, there are multiple real-time measurements of pulmonary variables available through the anesthesia machine (Figure 1b). In particular, it provides readings of the fraction of inspired O_2 (F_iO_2), i.e., the percentage of O_2 in inhaled air, as well as the partial pressure of CO_2 at the end of the breath, denoted by $EtCO_2$. The machine also provides clinicians with the ability to control certain variables. First of all, they can vary F_iO_2 through an O_2 flow control valve; in this way, they can deliver more O_2 in the case of a desaturation [3]. Moreover, there are two ways to control the patient's ventilation: (1) through the respiratory rate (RR) and tidal volume (V_t) or (2) through peak inspiratory pressure (PIP) and positive end-expiratory pressure ($PEEP$). However, clinicians interpret the machine's measurements predominately based on intuition and previous experience, hence the prediction of drops in C_aO_2 , and the resulting response, varies [14].

When suspecting a problem, clinicians may ask for a blood gas analysis to measure C_aO_2 . In such a case, blood is drawn from the patient and is used in a blood gas analyzer (see Figure 1c) for extraction of its content and properties, including acidity, partial pressure of O_2 and CO_2 , etc. However, whereas the pulse oximeter and the anesthesia machine used at CHOP sample continuously through non-invasive means and store data every 15 seconds, a blood gas analysis is invasive and time consuming, taking up to 30 minutes.

3.2 Problem Statement

As argued above, the current approach to real-time monitoring of C_aO_2 is reactive in nature and may lead to late recognition of sharp decreases. In this work we aim to supplement the monitoring of C_aO_2 , especially during shunts. Thus, the problem addressed in this paper is to develop a non-invasive detector that predicts critical drops in C_aO_2 in infants, as caused by a shunt, before they are measured by the pulse oximeter.

4. PHYSIOLOGICAL MODEL

This section presents a dynamic model of the O_2 and CO_2 partial pressures in the respiratory and cardiovascular systems as derived from clinical data and physics; furthermore, it discusses the changes caused by the presence of a shunt. While there are detailed models of the fluid dynamics in the two systems [7, 16], they do not describe the dynamics of the gaseous content; hence, this model represents a major contribution of our work.

Figure 2: A typical hemoglobin dissociation curve for O_2 [31]. It shows the shape of the relationship between the partial pressure of dissolved O_2 and hemoglobin saturation.

Table 1: Summary of cardiopulmonary partial pressures.

Variable Names	Physiological Location
P_iO_2, P_iCO_2	Airways (inspiration)
P_AO_2, P_ACO_2	Alveoli
$P_v^LO_2, P_v^LCO_2$	Left pulmonary veins
$P_v^RO_2, P_v^RCO_2$	Right pulmonary veins
P_aO_2, P_aCO_2	Arteries
P_pO_2, P_pCO_2	Peripheral Capillaries
P_vO_2, P_vCO_2	Veins
P_eO_2, P_eCO_2	Airways (expiration)

We begin by noting that C_aO_2 varies according to Equation 1 and the hemoglobin dissociation curve for O_2 (Figure 2). Note that while the scale of the curve may vary across patients and conditions, its general shape remains constant. This suggests that if a patient is at the far right portion of the curve, a sharp decrease in P_aO_2 precedes a sharp decrease in S_aO_2 (and C_aO_2 , consequently). Therefore, one can detect a drop in C_aO_2 by monitoring the values of P_aO_2 ; as argued in Section 3.1, however, measuring P_aO_2 is invasive and infeasible in real time. Accordingly, this section provides an overview of the relevant partial pressure variables (in blood and airways) and models their dynamics in the two scenarios (i.e., shunt vs. no shunt).

4.1 Steady-State Partial Pressures

This subsection introduces the respiratory and cardiovascular partial pressures employed in the model, and describes their steady-state relationships. For easy reference, all variables discussed in this section are summarized in Table 1.

Consider the simplified schematic model of the respiratory and cardiovascular systems as shown in Figure 3.³ During inhalation, O_2 and CO_2 exert a partial pressure in the air-

³Note that, for better illustration, the figures show the pulmonary veins merging before entering the heart, whereas in healthy humans they connect to the left atrium directly.

Figure 3: A simplified schematic model of O_2 and CO_2 partial pressures in the respiratory and cardiovascular systems.

way, denoted by P_iO_2 and P_iCO_2 , respectively; in the alveoli, the partial pressures are P_AO_2 and P_ACO_2 [38]. Once in the lungs, the gases enter the blood stream through diffusion at the pulmonary capillaries [38]. Here we distinguish between two groups of variables, denoted by $P_v^L O_2$, $P_v^L CO_2$ and $P_v^R O_2$, $P_v^R CO_2$, that represent the left and right side partial pressures, respectively, immediately after diffusion but before the pulmonary veins have merged.⁴ Once the blood from the pulmonary veins enters the heart, it is pumped into the arteries, where the partial pressures are P_aO_2 and P_aCO_2 [38]. After the blood carrying the gases reaches the peripheral capillaries, metabolism converts some of the O_2 to CO_2 [38]. When in the veins, the gases' partial pressures are denoted by P_vO_2 and P_vCO_2 ; thus P_vO_2 is typically smaller than P_aO_2 , whereas P_vCO_2 is usually larger than P_aCO_2 [38]. Finally, during expiration the partial pressures are denoted by P_eO_2 and P_eCO_2 .

We note that CO_2 has a high diffusing capacity such that P_aCO_2 usually equals P_ACO_2 [38]. When $P_aCO_2 = P_ACO_2$ and the partial pressures are in equilibrium, i.e., enough O_2 is inhaled so that the body can use it in the metabolism process [38], the relationship between P_aCO_2 and P_AO_2 is captured by the alveolar gas equation [15]:

$$P_AO_2 = F_iO_2(P_{ATM} - P_{H_2O}) - \frac{P_aCO_2(1 - F_iO_2[1 - RQ])}{RQ}, \quad (2)$$

where P_{ATM} and P_{H_2O} are the atmospheric and water vapor pressures (in $mmHg$), respectively, and RQ is the respiratory quotient. RQ can be calculated as

$$RQ = \frac{CO_2 \text{ eliminated}}{O_2 \text{ consumed}}, \quad (3)$$

where "eliminated" means removed from the body. RQ is a measure of the ratio of O_2 and CO_2 used in metabolism and varies with the type of consumed food.

4.2 Partial Pressure Dynamics

This subsection presents the dynamics of the partial pressures in Figure 3. Since, to our knowledge, exact dynamics are unknown and vary from patient to patient, we model the general trends, based both on physics and published clinical

⁴Due to space limitation, only the CO_2 pulmonary-vein partial pressures are labeled in the figures.

cal data. We begin with the no-shunt model by addressing the dynamics of the cardiovascular system, diffusion, and respiratory system. We conclude the subsection with the shunt-induced partial pressure dynamics. Figure 4 is referenced throughout this section as a visual aid.

4.2.1 Cardiovascular Partial Pressure Dynamics

We begin our model development with the transport of blood from the pulmonary veins to the pulmonary arteries; the diffusion process is discussed in the next subsection. Models for fluid dynamics explain the flow of blood through the blood vessels [7], but require difficult-to-obtain parameters (e.g., arterial compliance). When accurate models cannot be employed, the expected blood circulation time, τ , can be approximated by dividing the blood volume (in mL) by cardiac output (in mL/min) [19].⁵ Thus, on average, P_pO_2 and P_pCO_2 are expected to be time delays, denoted by τ_a , of P_aO_2 and P_aCO_2 , respectively, as illustrated in Figure 4a.

Similarly, P_vO_2 and P_vCO_2 are approximated as delays, denoted as τ_v , of P_pO_2 and P_pCO_2 plus the metabolic effect.⁶ Metabolism converts O_2 into CO_2 and varies the partial pressures linearly – P_vCO_2 is larger than P_pCO_2 , whereas P_vO_2 is smaller than P_pO_2 [38]. To complete the loop, note that the pulmonary veins are short in comparison with the rest of the cardiovascular system [38], hence we model P_aO_2 as an instantaneous average of the pulmonary-vein partial pressures. Note that the body optimizes blood flow in the direction of better oxygen uptake, however this process is less pronounced during mechanical ventilation [33].

4.2.2 Diffusion Dynamics

In the body, diffusion is the movement of O_2 between the blood and airways, from higher to lower concentrations; it occurs in the lungs. Note that both the O_2 and CO_2 ratios between the pre- and post-diffusion partial pressures, called *diffusion ratios* in this work, vary with many physiological variables (e.g., F_iO_2 , tidal volume, respiratory rate), some of which are difficult to obtain non-invasively (e.g., lung capacity, thickness of alveolar membrane) [31]. However, we are unaware of any closed-form models for the O_2 diffusion ratio; furthermore, it is challenging to develop such a model due to O_2 's relatively low diffusive capacity [31] and high sensitivity to F_iO_2 . In contrast, CO_2 has high diffusive capacity and low sensitivity to F_iO_2 ; thus, as shown in prior work [26], its diffusion ratio equals a parameter, $\alpha(t)$, that is predominantly affected by the volume of inhaled air participating in diffusion, V . V affects $\alpha(t)$ through the ideal gas law: $PV = nRT$ [38], where for constant temperature (T) and gas constant (R), an increase in V results in a lower pressure (P) and/or more diffused moles (n). Thus, the diffusion ratio, $\alpha(t)$, is inversely proportional to V .

4.2.3 Respiratory Partial Pressure Dynamics

To relate the blood and air partial pressures, first note that a model could only be developed for CO_2 , for the reasons discussed in the previous paragraph. In particular, the high diffusing capacity of CO_2 ensures that the alveolar and lung-specific pressures are the same under normal conditions [38], i.e., $P_ACO_2(t) = P_v^R CO_2(t) = P_v^L CO_2(t)$. Fi-

⁵While τ varies with cardiac output [38], we assume it is time-invariant as governed by the granularity of the anesthesia machine's sampling. This is explained in Section 6.

⁶Metabolism is constant for a small period of time (e.g., 5 minutes) [27].

Figure 4: An illustration of the response of the respiratory and cardiovascular partial pressures to a shunt.

nally, note that $P_eCO_2(t) = P_ACO_2(t)$.⁷

4.2.4 Shunt Partial Pressure Dynamics

Now consider a shunt in which one of the lungs is not ventilated. The effect on the partial pressures of CO_2 is illustrated in Figure 4b. Whereas under normal conditions both lungs participate in the removal of CO_2 , during a shunt the CO_2 -rich venous blood that passes through the unventilated lung does not receive oxygen [38] (i.e., $P_v^L CO_2 \neq P_v^R CO_2$). Since P_aCO_2 is the average of the pulmonary-vein partial pressures, it is larger than in the no-shunt case; however, the magnitude of this difference varies from patient to patient and from shunt to shunt. This implies $P_aCO_2 \neq P_ACO_2$, thus invalidating the alveolar gas equation in (2).

As discussed above, monitoring P_aO_2 (consequently, C_aO_2) in real time is not possible due to the lack of dynamic models for O_2 partial pressures. Yet, this section provides such a model for CO_2 in which a shunt results in elevated partial pressures of CO_2 . Furthermore, P_aCO_2 and C_aO_2 are related through the functional lung's capacity; for a lung with greater capacity, a shunt may not lower C_aO_2 and only cause a small rise in P_aCO_2 ; on the other hand, if the capacity is insufficient, then C_aO_2 may drop and P_aCO_2 may increase significantly. Thus, the model suggests the magnitude of increase in P_aCO_2 is related to a decrease in C_aO_2 .

5. CRITICAL SHUNT DETECTOR

In this section we present a detector for critical shunts in infants based on the relations and models developed in Section 4. A high-level introduction to the detector design problem is provided in the following subsection. The final two subsections formalize the model introduced in Section 4 and present the corresponding matched subspace detector.

5.1 Approach to Detector Design

Before formalizing the detector using the model from Section 4, we summarize the parameter invariant detection approach, closely following [30]. Given a vector of M measurements, $\mathbf{y} \in \mathbb{R}^M$, the problem is to determine the state of the system at the point of time $\mathbf{y}$ was obtained. In other

⁷Note that P_eCO_2 might be smaller than P_ACO_2 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 [13], hence it is not considered in this work.

words, suppose there exist two hypotheses describing the system's operation: the null hypothesis H_0 (e.g., no shunt) and the alternative H_1 (e.g., shunt). Under each hypothesis, $\mathbf{y}$ is drawn from a different probability distribution, F_{H_0} and F_{H_1} , respectively. The problem is to conclude which distribution $\mathbf{y}$ is more likely to have been drawn from.

As argued in Section 2, in this work we utilize the matched subspace detector due to its robustness to scarce data and unknown model parameters. Formally, the detector assumes that $\mathbf{y} = \nu \mathbf{F}\boldsymbol{\theta} + \sigma_i \mathbf{n}$, where $\mathbf{F} \in \mathbb{R}^{M \times p}$ defines the subspace of the signal under H_1 , $\boldsymbol{\theta} \in \mathbb{R}^p$ denotes the p unknown parameters (which determine the coordinates of $\mathbf{y}$ in the subspace), σ_i is an unknown gain (possibly different under each hypothesis), and $\mathbf{n} \sim N(\mathbf{0}, \mathbf{I})$ is white noise. The null hypothesis H_0 assumes $\nu = 0$, whereas $\nu > 0$ under H_1 . Thus, $\mathbf{y} \sim N(\mathbf{0}, \sigma_0^2 \mathbf{I})$ under H_0 and $\mathbf{y} \sim N(\nu \mathbf{F}\boldsymbol{\theta}, \sigma_1^2 \mathbf{I})$ under H_1 . Based on this information, one may derive a sufficient and invariant (to $\boldsymbol{\theta}$) statistic r .

Consequently, if r exceeds a certain threshold, an alarm is raised and H_0 is rejected based on the data. If the null hypothesis is rejected in error, the alarm is considered a false alarm. For this problem, there exists an r and threshold such that the Neyman-Pearson lemma ensures a specified constant false alarm rate [30].

5.2 Detector Model

Having described the matched subspace detector at a high level, we formalize the model from Section 4 and rearrange it in a form suitable for the detector. We observe that the system in Figure 4 has two discrete modes, shunt and no shunt, which we denote by $\mathcal{M}_S$ and $\mathcal{M}_{NS}$, respectively. We denote the system mode at time k by $\mathcal{M}(k)$ and define the hypothesis testing problem as

$$\begin{aligned} H_0 : \mathcal{M}(k) &= \mathcal{M}_{NS}, \quad 1 \leq k \leq M \\ H_1 : \mathcal{M}(k) &= \begin{cases} \mathcal{M}_{NS} & 1 \leq k \leq T^{NS} \\ \mathcal{M}_S & T^{NS} + 1 \leq k \leq M, \end{cases} \end{aligned} \quad (4)$$

where T^{NS} is the shunt transition time. The remainder of this section formalizes the model(s) under each hypothesis.

Recalling that $EtCO_2$ measures P_eCO_2 , in this subsection we design a detector to monitor changes in $EtCO_2$ as governed by the model of CO_2 dynamics. Recall also that τ is the expected blood circulation time; yet, as constrained by the sampling rate of the bedside monitors, denoted by t_s ,

we build the detector in a discrete-time fashion. We begin with the relevant equations from Section 4:

$$\begin{aligned} P_a CO_2(k) &= 1/2(P_v^L CO_2(k) + P_v^R CO_2(k)), \\ P_v CO_2(k) &= P_a CO_2(k - \kappa) + \mu_2, \\ Et CO_2(k) &= P_a CO_2(k) = P_v^{NS} CO_2(k), \end{aligned} \quad (5)$$

where k is the time step, κ denotes the number of sampling periods (on average) it takes the blood to circulate (i.e., $\tau = \kappa t_s$), and $P_v^{NS} CO_2$ is the pulmonary-vein partial pressure of the no-shunt lung (also equal to $P_a CO_2$ if both lungs are receiving oxygen).

Thus, we develop the dynamics for each mode beginning with $\mathcal{M}_{NS}$. The first conclusion to be drawn is as follows:

$$P_v^i CO_2(k) = \alpha(k) P_v CO_2(k), i \in \{L, R\}. \quad (6)$$

To estimate $\alpha(k)$, we note that it is inversely proportional to the volume of air participating in gas exchange, $V(k)$, as argued in Section 4. As will be apparent in Section 6, the measuring machines' sampling time is longer than the time it takes an infant to take a few breaths. Thus, the approximate volume of air participating in gas exchange, $\bar{V}$, during the sampling period is

$$\bar{V}(k) = \frac{t_s}{60} \times RR(k) \times V_t(k), \quad (7)$$

where RR is the measured respiratory rate (in breaths per minute) and V_t is the measured tidal volume (in volume per breath). Applying this observation, Equation 6 can be approximated, for $i \in \{L, R\}$, as:

$$P_v^i CO_2(k) = (\bar{\alpha}/\bar{V}(k)) P_v CO_2(k) + \sigma n^i(k), \quad (8)$$

where $\bar{\alpha}$ is an unknown parameter representing the CO_2 diffusion ratio for fixed V , σ is an unknown gain, and $n^i(k) \sim N[0, 1]$ is a white Gaussian noise. In (8), $\sigma n^i(k)$ denotes a zero-mean noise with unknown variance to capture the unknown error in our approximation. Combining (8) with (5), we write the dynamics for $\mathcal{M}_{NS}$ as:

$$\begin{aligned} \begin{bmatrix} x^L(k) \\ x^R(k) \end{bmatrix} &= \begin{bmatrix} a(k) & a(k) \\ a(k) & a(k) \end{bmatrix} \begin{bmatrix} x^L(k - \kappa) \\ x^R(k - \kappa) \end{bmatrix} \\ &+ \begin{bmatrix} 2a(k) & n^L(k) \\ 2a(k) & n^R(k) \end{bmatrix} \begin{bmatrix} \mu_2 \\ \sigma \end{bmatrix} \\ y(k) &= \begin{bmatrix} 1/2 & 1/2 \end{bmatrix} \begin{bmatrix} x^L(k) \\ x^R(k) \end{bmatrix}, \end{aligned} \quad (9)$$

where $x^i(k) = P_v^i CO_2(k)$, $y(k) = Et CO_2(k)$ and $a(k) = \bar{\alpha}/(2\bar{V}(k))$.

A similar derivation exists for $\mathcal{M}_S$. We first observe

$$\begin{aligned} P_v^S CO_2(k) &= P_v CO_2(k) \\ P_v^{NS} CO_2(k) &= \alpha(k) P_v CO_2(k). \end{aligned} \quad (10)$$

where, $P_v^S CO_2$ and $P_v^{NS} CO_2$ denote the shunted and non-shunted pulmonary-vein partial pressures, respectively. Substituting the approximation for $\alpha(k)$, we obtain

$$\begin{aligned} P_v^S CO_2(k) &= P_v CO_2(k) \\ P_v^{NS} CO_2(k) &= (\bar{\alpha}/\bar{V}(k)) P_v CO_2(k) + \sigma n^{NS}(k), \end{aligned} \quad (11)$$

where, similar to (8), $n^{NS}(k) \sim N[0, 1]$ and σ is the same

unknown gain. Thus, the dynamics for $\mathcal{M}_S$ are

$$\begin{aligned} \begin{bmatrix} x^{NS}(k) \\ x^S(k) \end{bmatrix} &= \begin{bmatrix} \frac{a(k)}{1/2} & \frac{a(k)}{1/2} \\ 1/2 & 1/2 \end{bmatrix} \begin{bmatrix} x^{NS}(k - \kappa) \\ x^S(k - \kappa) \end{bmatrix} \\ &+ \begin{bmatrix} a(k) & n^{NS}(k) \\ 1 & 0 \end{bmatrix} \begin{bmatrix} \mu_2 \\ \sigma \end{bmatrix} \\ y(k) &= \begin{bmatrix} 1 & 0 \end{bmatrix} \begin{bmatrix} x^{NS}(k) \\ x^S(k) \end{bmatrix}, \end{aligned} \quad (12)$$

where, $x^{NS}(k) = P_v^{NS} CO_2(k)$ and $x^S(k) = P_v^S CO_2(k)$.

Without loss of generality, suppose a right-lung shunt occurs and the governing dynamics switch from $\mathcal{M}_{NS}$ to $\mathcal{M}_S$. Thus, the $\mathcal{M}_S$ pulmonary-vein partial pressures reset such that $x^{NS}(k') = x^L(k')$ and $x^S(k') = x^R(k')$, $T^{NS} - \kappa + 1 \leq k' \leq T^{NS}$. This results from the fact that the blood takes κ steps to circulate (as explained in Section 4.2.1).

The dynamics for both $\mathcal{M}_{NS}$ and $\mathcal{M}_S$ contain the unknown parameters $\bar{\alpha}$, μ_2 , and σ . Since both models are sensitive to their values, we design a detector that is invariant to them in the following subsection.

5.3 Detector Design

Given the models in (9) and (12), respectively, we construct a detector invariant to the unknown parameters. First, we convert the models into a form that can be used by the detector, i.e., $\mathbf{y} = \nu \mathbf{F} \boldsymbol{\theta} + \sigma_i \mathbf{n}$, where $\nu = 0$ under H_0 and $\nu > 0$ under H_1 , $\boldsymbol{\theta}$ contains the model parameters, σ_i is the unknown gain, and $\mathbf{n} \sim N(\mathbf{0}, \mathbf{I})$ is white noise. Under H_0 , the system is always in $\mathcal{M}_{NS}$. Thus, one can verify that

$$y(k) = \frac{\bar{\alpha}}{\bar{V}(k)} y(k - \kappa) + \frac{\bar{\alpha}}{\bar{V}(k)} \mu_2 + \sigma_0 n_0(k). \quad (13)$$

where $\sigma_0 n_0(k) = (\sigma/2) (n^L(k) + n^R(k)) \sim N[0, \sigma^2/2]$, i.e., $\sigma_0 = \sigma/\sqrt{2}$. Hence, $\boldsymbol{\theta} = [\bar{\alpha}, \bar{\alpha}\mu_2]^T$; with M measurements, $\mathbf{y} = \mathbf{F}_0 \boldsymbol{\theta} + \sigma_0 \mathbf{n}_0$, where

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

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.

Under H_1 , the system begins in $\mathcal{M}_{NS}$ and switches to $\mathcal{M}_S$ at time T^{NS} such that (13) holds for $k \leq T^{NS}$, and for $k > T^{NS}$ one can derive the following relation:

$$y(k) = \frac{\bar{\alpha}}{4\bar{V}(k)} y(k - \kappa) + \frac{\bar{\alpha}}{4\bar{V}(k)} x^S(k - \kappa) + \frac{\bar{\alpha}}{2\bar{V}(k)} \mu_2 + \sigma n^{NS}(k). \quad (14)$$

Note that once a shunt has occurred,

$\mathbf{x}^S = [x^S(T^{NS} + 1) \dots x^S(M)]^T$ is unknown since $\mathbf{y}$ measures the no-shunt partial pressure. Therefore, it is only possible to estimate $\mathbf{x}^S$ by following the dynamics since the last sample obtained when no shunt was in place. To estimate $\mathbf{x}^S$ after T^{NS} , note that from (12),

$$x^S(k) = (1/2)y(k - \kappa) + (1/2)x^S(k - \kappa) + \mu_2. \quad (15)$$

Finally, we obtain the following model, starting at $T^{NS} + 1$,

$$(\mathbf{I} + \mathbf{B})\mathbf{x}^S = \begin{pmatrix} y(T^{NS} + 1 - \kappa) \\ \vdots \\ y(T^{NS}) \\ (1/2)y(T^{NS} + 1) \\ \vdots \\ (1/2)y(M - \kappa) \end{pmatrix} + \mu_2 \mathbf{1}, \quad (16)$$

where $\mathbf{I} \in \mathbb{R}^{M-T^{NS} \times M-T^{NS}}$ denotes the identity matrix, and $\mathbf{B}$ is a matrix such that $B_{ij} = -1/2$ if $i = j + \kappa$, and $B_{ij} = 0$ otherwise. Thus, by multiplying both sides by $(\mathbf{I} + \mathbf{B})^{-1}$ and expanding the parentheses, one can represent $\mathbf{x}^S$ as follows:

$$x^S(k) = f(k) + g(k)\mu_2, \quad k > T^{NS}. \quad (17)$$

Therefore, under H_1 , $\mathbf{y} = \mathbf{F}_1\boldsymbol{\theta} + \sigma_1\mathbf{n}_1$, where $\sigma_1 = \sigma$ and the first T^{NS} rows of $\mathbf{F}_1$ are equal to the corresponding rows of $\mathbf{F}_0$, and the remaining are as follows (the matrix $\mathbf{F}'_1$ is obtained by taking all rows of $\mathbf{F}_1$ starting at $K = T^{NS} + 1$):

$$\mathbf{F}'_1 = \begin{bmatrix} \frac{y(K-\tau)}{4V_t(K)} + \frac{f(K-\tau)}{4V_t(K)} & \frac{1}{2V_t(K)} + \frac{g(K-\tau)}{4V_t(K)} \\ \frac{y(K+1-\tau)}{4V_t(K+1)} + \frac{f(K+1-\tau)}{4V_t(K+1)} & \frac{1}{2V_t(K+1)} + \frac{g(K+1-\tau)}{4V_t(K+1)} \\ \vdots & \vdots \\ \frac{y(M-\tau)}{4V_t(M)} + \frac{f(M-\tau)}{4V_t(M)} & \frac{1}{2V_t(M)} + \frac{g(M-\tau)}{4V_t(M)} \end{bmatrix}.$$

Thus, we have arrived at the model $\mathbf{y} = \mathbf{F}_0\boldsymbol{\theta} + \sigma_0\mathbf{n}_0$ under H_0 and $\mathbf{y} = \mathbf{F}_1\boldsymbol{\theta} + \sigma_1\mathbf{n}_1$ under H_1 . To convert it to the form $\mathbf{y} = \nu\mathbf{F}\boldsymbol{\theta} + \sigma_i\mathbf{n}$, we premultiply both equations by $\mathbf{P}_{\mathbf{F}_0^\perp}$, i.e., the projection matrix to the null space of $\mathbf{F}_0$. Thus, we obtain the model $\mathbf{z}_0 = \nu\mathbf{G}_0\boldsymbol{\theta} + \sigma_i\mathbf{n}_0$, where $\mathbf{z}_0 = \mathbf{P}_{\mathbf{F}_0^\perp}\mathbf{y}$, $\mathbf{G}_0 = \mathbf{P}_{\mathbf{F}_0^\perp}\mathbf{F}_1$ and $\mathbf{n}_0 \sim N(\mathbf{0}, \mathbf{I})$. Thus, as shown in [30], the following is both a sufficient (i.e., independent of σ_i) and invariant (to $\boldsymbol{\theta}$) statistic

$$r_0 = \frac{\mathbf{z}_0^T \mathbf{P}_{\mathbf{G}_0} \mathbf{z}_0 / \text{rank}(\mathbf{P}_{\mathbf{G}_0})}{\mathbf{z}_0^T (\mathbf{I} - \mathbf{P}_{\mathbf{G}_0}) \mathbf{z}_0 / \text{null}(\mathbf{P}_{\mathbf{G}_0})}, \quad (18)$$

where $\mathbf{P}_{\mathbf{G}_0}$ is the projection on the range of $\mathbf{G}_0$, and null and rank are the dimensions of its null space and range, respectively. r_0 is distributed according to an F-distribution with $\text{rank}(\mathbf{P}_{\mathbf{G}_0})$ numerator and $\text{null}(\mathbf{P}_{\mathbf{G}_0})$ denominator degrees of freedom. Thus, by choosing a threshold t_0^* at the tail of the distribution (e.g., larger than 99% of the values) and raising an alarm only when $r_0 > t_0^*$, one can guarantee a false alarm rate for the detector (e.g., 1%).

At the same time, we note that the O_2 - CO_2 circulation model during a shunt developed in this section may be also (partly) matched by other conditions, e.g., hypovolemia. Therefore, to reduce the expected high number of false alarms, we compute the reverse statistic of r_0 , i.e., we assume $\nu > 0$ under H_0 and verify whether we can reject this hypothesis based on the data. The statistic is computed similarly:

$$r_1 = \frac{\mathbf{z}_1^T \mathbf{P}_{\mathbf{G}_1} \mathbf{z}_1 / \text{rank}(\mathbf{P}_{\mathbf{G}_1})}{\mathbf{z}_1^T (\mathbf{I} - \mathbf{P}_{\mathbf{G}_1}) \mathbf{z}_1 / \text{null}(\mathbf{P}_{\mathbf{G}_1})}, \quad (19)$$

where $\mathbf{z}_1 = \mathbf{P}_{\mathbf{F}_1^\perp}\mathbf{y}$ and $\mathbf{G}_1 = \mathbf{P}_{\mathbf{F}_1^\perp}\mathbf{F}_0$. One may choose a threshold, t_1^* , such that H_1 is rejected if $r_1 > t_1^*$.

Based on these two statistics, the detector's decision space consists of four cases, as indicated in Table 2. If the two tests agree, then the detector accepts the hypothesis that they agree on. If both statistics are above their thresholds and

Table 2: 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)

both hypotheses are rejected, then the data does not match either model, thus we only output a warning indicating the model is inaccurate. Finally, if neither statistic is above the threshold, then the data matches both models, and a warning is issued indicating the test lacks power.

6. CASE STUDY

To evaluate the performance of the detector, we used real patient data from lobectomy surgeries on infants performed at CHOP over the last decade. The initial data set consisted of 292 patients varying in age between 0 and 323 days, 195 of which underwent the open procedure while the remaining underwent thoracoscopy.

The relevant vital signs (S_pO_2 , V_t , RR , $EtCO_2$) were sampled every 15 seconds. Note that the time for blood to circulate the body (i.e., the parameter τ) is about a minute in adults but is much shorter in infants [22]. Thus, given the discrete model, the best choice matching observed data [22] is $\tau \approx 30$ seconds, i.e., $\kappa = 2$ time steps. Additionally, note that the detector was run in a sliding window fashion; for evaluation purposes, we chose the window size (M from Section 5) to be $M = 18$, i.e., 4.5 minutes, and the detection window size to be $M - T^{NS} = 8$, i.e., 2 minutes. For both thresholds, t_0^* and t_1^* , we specify a corresponding error rate of 1% to be the false alarm rate and miss rate, respectively.

For a variety of reasons (e.g., disconnection of ventilator) the patient data contains wrong or missing measurements. In particular, if any of the vital signs directly used in the detector (i.e., $EtCO_2$, V_t or RR) has a missing measurement during a window of time, the detector does not make a decision for that window. Additionally, most cases begin with artifactual data (e.g., caused by motion artifact, tourniquet on the limb). Consequently, when an isolation is performed during this time, a lot of the measurements are wrong, hence all cases with events in the first 15 minutes are discarded since the data is unreliable. After removing bad data cases, we retain 26 cases containing shunts (for evaluating the detection rate) and 172 cases without a shunt (for evaluating the false alarm rate). The detection rate and false alarm rate are evaluated in the following sections, respectively.

6.1 Detection Rate Evaluation

To evaluate the detection rate of the critical shunt detector, we run the detector on all 26 cases with annotated shunts. We first define the events that are worth detecting from a medical point of view. Ideally, these are events that led to a physician taking action, but they are rarely clear from the data. Thus, we look for events that can be observed, namely sharp decreases in S_pO_2 correlated with shunts. However, as argued in Section 2, the annotated timestamps do not always correspond to the actual time the shunt occurred. Therefore, we identified the “nearest in time” decrease in S_pO_2 , also referred to as “inferred time”, as the actual event that we would like our alarm to precede.

To evaluate the detection rate, we note that the fact that

Table 3: A table showing the detection, annotation and inferred event times (in sampling steps) for cases with shunts.

Case #	A. Time	I. Time	D. Time	I - D
1	205	208	116	92
2	50	73	70	3
3	71	84	80	4
4	158	161	51	110
5	290	245	232	13
6	14	65	60	5
7	73	83	65	18
8	126	147	131	16
9	33	69	67	2
10	33	104	96	8
11	39	60	56	4
12	152	146	148	-2
13	230	228	228	0
14	49	79	74	5
15	49	75	114	-39
16	251	207	204	3
17	67	71	57	14
18	258	261	254	7
19	46	82	152	-70
20	217	231	192	39
21	77	62	62	0
22	63	122	71	51
23	76	121	61	60
24	224	227	221	6
25	297	312	306	6
26	130	133	123	10

a prediction occurred is arguably equally important as the time of prediction. Thus, in this section, we provide two different ways of analyzing the detection rate – (1) the proportion of cases with an alarm preceding the event (raw detection rate) and (2) a summary of results for all cases showing how much time in advance each event was detected.

To record the time of detection, we record the alarm occurring closest to 2 minutes before the drop in S_pO_2 , which corresponds to the specified detection window time. To show both ways of evaluating the detection rate, first note that the raw detection rate was 88.5%. The results for the detection, annotation and inferred times for all cases are shown in Table 3.

A histogram of the last column in the table, i.e., how much time in advance each event was detected, is shown in Figure 5, with cases labeled according to the time of detection. Most events were detected less than 20 steps (5 minutes) before the start of the event. This highlights the power of the detector – the majority of the cases, 81%, received alarms at a time that would enable clinicians to take proactive measures, at an average of 65 seconds earlier than previously possible. Only two cases had late detections, thus missing the event, while five had non-critical detection (some of which potentially false alarms).

We now examine a case from each group to illustrate the detector’s performance in greater detail. Figure 6a shows a case whose event was predicted before its occurrence. It contains the evolution of the vital signs used in the detector as well as the decision made.⁸ As shown in the figure, there is a period with noisy data at beginning of the case, followed by the clear pattern of $EtCO_2$ rising. Note that the isolation itself likely occurred before the sharp decrease in S_pO_2 (as indicated by the steady rise in $EtCO_2$) but this would not be captured by the detector due to the noisy RR

⁸Note that time starts with negative values because no decision can be made until M measurements are collected.

Figure 5: A histogram showing how many time steps (one step = 15 seconds) in advance each critical event was detected (negative values represent late detection).

measurements. Finally, note that there is a long period during the surgery (beginning around time 120) without any alarms, which indicates that when good data is available, the detector is not prone to raise false alarms.

Figure 6b shows a case with late detection. The incorrect data patterns are the same as in the previous case. This case is detected late because of the few-minute period with noisy data (especially $EtCO_2$), however an alarm is raised only 30 seconds after the beginning of the event, thus still being potentially useful to a clinician.

Finally, Figure 6c shows a case in which the event was predicted too early, possibly due to a false alarm. In this case an alarm was raised more than 50 steps (i.e., more than 12 minutes) before the event. The alarm was due to the wrong measurements of V_t and RR at that time. On the other hand, the alarm was raised around the annotated time for this case, so the detector may have actually detected the shunt but not the drop in S_pO_2 caused by it. After close inspection, it was noted that S_pO_2 dropped in this case due to a sharp decrease in F_iO_2 . Adjusting the detector to incorporate changes in F_iO_2 is an avenue for future work.

6.2 False Alarm Rate Evaluation

To evaluate the false alarm rate, we used the open cases with no annotated isolation. To count false alarms in a way that would be meaningful to the medical community,⁹ the detector is implemented as follows. If H_1 is accepted, then an alarm is raised, which can be muted by a physician. This alarm denotes the beginning of an event; if the following decisions of the detector are also either an alarm or a warning, then they relate to the same event, hence no further alarms are raised. As soon as a “no alarm” decision is made, the event has ended, and the next time H_1 is accepted denotes a new event and a new alarm.

For each case we count the average number of events detected per hour, and all are treated as false alarms. At the same time, it was noted that certain events flagged by the detector may be worth considering by physicians (e.g., partial shunts due to mucus plugging), therefore the number of false alarms shown in this work may be exaggerated.

⁹The 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.

(a) An example case with good prediction (case 9).

(b) An example case with late detection (case 12).

(c) An example case with early detection (case 23).

(d) An example case with no isolation (case 50).

Figure 6: Example cases for different scenarios. The shaded area denotes the occurrence of the isolation.

gerated. A scatter plot with the result for each case is shown in Figure 7, starting from case number 27. The cases are reasonably uniformly distributed, with no outliers with more than six alarms per hour; the low variance is due to the constant-false-alarm-rate property of our detector and presents an improvement to the high-variance bedside alarms currently used [4]. The average over all cases is 2.15 per hour, which is significantly less than most current threshold-based alarms [4].

To investigate how to further lower the number of false alarms, we present a typical open case in Figure 6d. This case has eight alarms in about two hours. However, almost every alarm is caused by wrong measurements in some variables. For example, RR , which is usually held constant, has multiple spikes that most certainly are wrong measurements because patients are frequently taken off mechanical ventilation (e.g., in order to be hand-ventilated to check for a leak at the surgical site). Thus one way of reducing false alarms is to detect wrong data and either correct it or discard the window it is encountered in.

Figure 7: A scatterplot with the average number of false alarms per hour for each of the open cases.

7. CONCLUSION

In this work we addressed the early detection of sharp decreases in C_aO_2 due to pulmonary shunts. We developed a model of the O_2 and CO_2 circulation in the blood and airways and described the effect of a shunt. Based on the model, we used a matched subspace detector with a constant false alarm rate in order to account for unknown parameters and differences between patients. Finally, we validated our approach on a real-patient data case study and demonstrated that it is a promising direction for future work in MCPs. Extensions of this work include a bad data detector to further reduce the number of false alarms and a real-world implementation in an Intensive Care Unit.

8. REFERENCES

- [1] Roche Cobas B 123 POC System. <http://www.roche.com/products/product-details.htm?region=us&type=product&id=146>.
- [2] Shenzhen Bestman Instrument Co. Pulse Oximeter. <http://www.szbestman.com/contents/76/669.html>.
- [3] The Dräger Apollo Anesthesia Machine. <http://www.draeger.com/sites/enus-us/Pages/Hospital/Apollo.aspx>.
- [4] A. Aboukhalil, L. Nielsen, M. Saeed, et al. Reducing false alarm rates for critical arrhythmias using the arterial blood pressure waveform. *Journal of biomedical informatics*, 41(3):442–451, 2008.
- [5] K.-P. Adlassnig, C. Combi, A. K. Das, et al. Temporal representation and reasoning in medicine: research directions and challenges. *Artificial Intelligence in Medicine*, 38(2):101–113, 2006.
- [6] D. Arney, M. Pajic, J. M. Goldman, et al. Toward patient safety in closed-loop medical device systems. In *Proceedings of the 1st International Conference on Cyber-Physical Systems*, pages 139–148, 2010.
- [7] D. H. Bergel. *Cardiovascular fluid dynamics*, volume 1. Elsevier, 2012.
- [8] C. M. Bishop et al. *Pattern recognition and machine learning*, volume 1. Springer New York, 2006.
- [9] P. Bogdan, S. Jain, J. M. Goyal, and R. Marculescu. Implantable pacemakers control and optimization via fractional calculus approaches: A cyber-physical systems perspective. In *Proceedings of the Third International Conference on Cyber-Physical Systems*, pages 23–32, 2012.
- [10] A. Burgos, A. Goñi, A. Illarramendi, and J. Bermúdez. Real-time detection of apneas on a pda. *IEEE Transactions on Information Technology in Biomedicine*, 14(4):995–1002, 2010.
- [11] J. Chen and R. J. Patton. *Robust model-based fault diagnosis for dynamic systems*. Springer Publishing Company, Incorporated, 2012.
- [12] V. A. Convertino, S. L. Moulton, G. Z. Grudic, et al. Use of advanced machine-learning techniques for noninvasive monitoring of hemorrhage. *Journal of Trauma-Injury, Infection, and Critical Care*, 71(1):S25–S32, 2011.
- [13] S. Cruickshank and N. Hirschauer. The alveolar gas equation. *Continuing Education in Anaesthesia, Critical Care & Pain*, 4:24–27, 2004.
- [14] S. M. Donn and W. Boon. Mechanical ventilation of the neonate: should we target volume or pressure? *Respiratory care*, 54(9):1236–1243, 2009.
- [15] W. O. Fenn, H. Rahn, and A. B. Otis. A theoretical study of the composition of the alveolar air at altitude. *American Journal of Physiology*, 146:637–653, 1946.
- [16] J. B. Grotberg and O. E. Jensen. Biofluid mechanics in flexible tubes. *Annual Review of Fluid Mechanics*, 36(1):121–147, 2004.
- [17] G. B. Hammer. Single-lung ventilation in infants and children. *Pediatric Anesthesia*, 14:98–102, 2004.
- [18] Z. Jiang, M. Pajic, and R. Mangharam. Cyber-physical modeling of implantable cardiac medical devices. In *Proceedings of the IEEE*, pages 122–137, 2012.
- [19] R. Klabunde. *Cardiovascular physiology concepts*. Lippincott Williams & Wilkins, 2011.
- [20] I. Lee, O. Sokolsky, S. Chen, et al. Challenges and research directions in medical cyber-physical systems. *Proceedings of the IEEE*, 100(1):75–90, 2012.
- [21] T. M. Mitchell, R. Hutchinson, M. A. Just, et al. Classifying instantaneous cognitive states from fmri data. In *AMIA Annual Symposium Proceedings*, pages 465–469, 2003.
- [22] B. B. Oberst and F. La Roche. Circulation time in the newborn infant, using the fluorescein dye method. *Journal of Pediatrics*, 45(5):580–582, 1954.
- [23] M. Pajic, R. Mangharam, O. Sokolsky, et al. Model-driven safety analysis of closed-loop medical systems. *IEEE Transactions on Industrial Informatics*, 10(1):3–16, 2014.
- [24] A. Pantelopoulos and N. G. Bourbakis. A survey on wearable sensor-based systems for health monitoring and prognosis. *IEEE Transactions on Systems, Man, and Cybernetics*, 40(1):1–12, 2010.
- [25] D. Parikh and M. Samuel. Congenital cystic lung lesions: Is surgical resection essential? *Pediatric Pulmonology*, 40:533–537, 2005.
- [26] A. C. Powles and E. J. Campbell. How to be less invasive. *American Journal of Medicine*, 67(1):98–104, 1979.
- [27] H. B. Richardson. The respiratory quotient. *Physiological Reviews*, 9(1):61–125, 1929.
- [28] S. Saria, D. Koller, and A. Penn. Learning individual and population level traits from clinical temporal data. In *Proceedings of Neural Information Processing Systems*, pages 1–9, 2010.
- [29] S. Saria, A. K. Rajani, J. Gould, et al. Integration of early physiological responses predicts later illness severity in preterm infants. *Science translational medicine*, 2(48):48ra65, Sept. 2010.
- [30] L. L. Scharf and C. Demeure. *Statistical Signal Processing*. Addison-Wesley Publishing Company, 1991.
- [31] S. Shafer, J. P. Rathmell, and R. Stoelting. *Stoelting's Pharmacology & Physiology*. Wolters Kluwer, 2014.
- [32] A. F. Simpao, E. Y. Pruitt, S. D. Cook-Sather, et al. The reliability of manual reporting of clinical events in an anesthesia information management system (aims). *Journal of Clinical Monitoring and Computing*, 26(6):437–439, 2012.
- [33] J. L. Theissen, S. R. Fischer, L. D. Traber, and D. L. Traber. Pulmonary blood flow regulation: influence of positive pressure ventilation. *Respiratory Physiology*, 102(2-3):251–260, 1995.
- [34] J. Weimer, S. A. Ahmadi, J. Araujo, et al. Active actuator fault detection and diagnostics in hvac systems. In *Proceedings of the Fourth Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, pages 107–114, 2012.
- [35] J. Weimer, S. Kar, and K. H. Johansson. Distributed detection and isolation of topology attacks in power networks. In *International Conference on High Confidence Networked Systems*, pages 65–72, 2012.
- [36] J. Weimer, D. Varagnolo, and K. H. Johansson. Distributed model-invariant detection of unknown inputs in networked systems. In *International Conference on High Confidence Networked Systems*, pages 127–134, 2013.
- [37] J. Weimer, D. Varagnolo, M. Stankovic, and K. Johansson. Parameter-invariant detection of unknown inputs in networked systems. In *Conference on Decision and Control*, pages 4379–4384, 2013.
- [38] J. B. West. *Respiratory Physiology: The Essentials*. Lippincott Williams & Wilkins, 2012.
- [39] A. S. Willsky. A survey of design methods for failure detection in dynamic systems. *Automatica*, 12(6):601–611, 1976.