

Estimation of Blood Oxygen Content Using Context-Aware Filtering

Radoslav Ivanov Nikolay Atanasov James Weimer
School of Engineering & Applied Sciences
University of Pennsylvania
Email: {rivanov, atanasov, weimerj}@seas.upenn.edu

Miroslav Pajic
Dept. of Electrical & Computer Eng.
Duke University
Email: miroslav.pajic@duke.edu

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

George Pappas
Dept. of Electrical &
Systems Eng.
University of Pennsylvania
Email: pappasg@seas.upenn.edu

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

Abstract—In this paper we address the problem of estimating the blood oxygen concentration in children during surgery. Currently, the oxygen content can only be measured through invasive means such as drawing blood from the patient. In this work, we attempt to perform estimation by only using other non-invasive measurements (e.g., fraction of oxygen in inspired air, volume of inspired air) collected during surgery. Although models mapping these measurements to blood oxygen content contain multiple parameters that vary widely across patients, the non-invasive measurements can be used to provide binary information about whether the oxygen concentration is rising or dropping. This information can then be incorporated in a context-aware filter that is used to combine regular continuous measurements with discrete detection events in order to improve estimation. We evaluate the filter using real-patient data collected over the last decade at the Children's Hospital of Philadelphia and show that it is a promising approach for the estimation of unobservable physiological variables.

I. INTRODUCTION

Modern operating rooms (ORs) are equipped with multiple devices that measure various vital signs and provide clinicians with ample information about the patient's state. Analyzing this data in real time, however, may be challenging when the physician is focused on the patient, especially when trends over time and correlations between variables must be observed. This presents a great opportunity for developing Medical Cyber-Physical Systems (MCPS) that aid clinicians by estimating unobservable variables or providing detections of critical events that might be unnoticed otherwise [9].

In this work, we address the problem of estimating the concentration of oxygen (O_2) in the blood, one of the most closely monitored variables in ORs. The O_2 concentration (also referred to as O_2 content) has to be maintained within

safe ranges – if it is too low it may lead to organ failure and brain damage, whereas too high values could be toxic. Therefore, one of clinicians' main concerns during surgery is to keep the patient's O_2 concentration and end-organ perfusion stable.

Currently, the O_2 content cannot be measured precisely by non-invasive means as it can only be directly measured by drawing blood from the patient. As a result, clinicians monitor the O_2 concentration using a proxy – the hemoglobin-oxygen saturation in the peripheral capillaries (S_pO_2), which is measured non-invasively and in real time. While S_pO_2 is a good measure of the concentration in the location where it is measured (e.g., a finger tip), it is not a good indication of the O_2 content in other parts of the body as there may be differences in perfusion (e.g., as caused by a tourniquet on a limb). In addition, monitoring S_pO_2 is reactive – clinicians take action when low S_pO_2 is observed, at which point the patient may already be in a critical state.

As a proactive way of monitoring the O_2 concentration, clinicians also estimate another correlated variable, namely the partial pressure of oxygen dissolved in arterial blood (P_aO_2). As opposed to S_pO_2 , P_aO_2 can be used as a predictive measure of the concentration as P_aO_2 drops significantly before major decreases in the concentration are observed. Thus, in this work we propose to estimate both the O_2 content and P_aO_2 as a convenient way of observing the trends in O_2 content. At the same time, similar to the O_2 concentration, P_aO_2 cannot be measured non-invasively and in real time.

To overcome this problem, we use other measurements available in real time in modern ORs, namely the fractions of O_2 and carbon dioxide (CO_2) in inhaled and exhaled air, respectively, the pressure and volume of inhaled air, and the respiratory rate. By correlating these variables with P_aO_2 and by modeling O_2 dynamics through the body, one can estimate P_aO_2 without multiple blood draws from the patient. However, an additional challenge with this approach is that models describing the circulation of O_2 in the blood and airways are

This work was supported in part by NSF CNS-1035715. This research was supported in part by Global Research Laboratory Program (2013K1A1A2A02078326) through NRF, and the DGIST Research and Development Program (CPS Global Center) funded by the Ministry of Science, ICT & Future Planning. This work was supported by TerraSwarm, one of six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA.

imprecise and contain multiple parameters that vary widely across patients, e.g., metabolic rate, lung membrane thickness, arterial wall thickness. While it may be possible to learn some of these parameters given enough data, most of them are not identifiable using non-invasive measurements only. Thus, standard filtering techniques are not suitable for this application.

To perform estimation in systems with imprecise models and sensors, our previous work introduced the context-aware filter [8]. The context-aware filter combines continuous and discrete measurements by modeling the probability of event detection as a function of the state. We developed recursive Kalman-like filters for two classes of detection functions and evaluated them in simulation of robotics scenarios. In this work, we employ the context-aware filter by using relative measures of the O_2 concentration. In particular, we monitor the O_2 diffusion, i.e., transfer of O_2 from lungs to the blood, and raise alarms when relative changes are observed. These alarms are then included as measurements in the context-aware filter in order to improve estimation.

Finally, we evaluate the P_aO_2 estimator on real-patient data collected at the Children’s Hospital of Philadelphia (CHOP). We compare the context-aware filter’s estimates with blood-gas data collected during some of these cases. We conclude that the filter can be used to improve upon current estimation techniques and is also a promising direction for future research in the area of estimation of unobservable physiological variables.

Related work in the MCPS domain can be broadly divided in three areas: verification, detection and estimation. When precise models are available, it is possible to use formal methods in order to ensure that the system does not endanger the patient’s safety [2], [15]. On the other hand, when models are parameterized by patient-specific parameters, one may use parameter-invariant detectors in order to achieve the same detection performance regardless of the values of these parameters [9], [22]. Finally, estimation is more challenging as most models are unidentifiable given current measurements; in some applications, it is possible to use invasive measurements to learn the parameters [12] or to use machine learning when sufficient data is available [5], [16]. Since the problem addressed in our paper does not have rich data or frequent invasive measurements, we generate context measurements based on other available data.

Estimation with both continuous and discrete (context) measurements has been addressed in the target tracking community [13] where, however, the more general models lead to computationally expensive solutions. Other related problems are sensor scheduling [10], [21] and filtering with intermittent measurements [20] but these works only handle continuous measurements. In robotic applications, techniques have been developed to use both kinds of measurements through information fusion [1], [4], [24] and bounded distributions [18]; however, the former works do not incorporate system dynamics whereas the latter approach results in unwieldy distributions that require approximations. Finally, a similar problem to

the one addressed in this paper can be solved using particle filters [3]; yet, particle filters may suffer from particle-deprivation problems in high-dimensional spaces whereas the context-aware filter provides exact estimates regardless of dimensionality.

In summary, the contributions of this work are: 1) the estimation of O_2 content and partial pressure of dissolved O_2 by only using non-invasive measurements (once properly initialized); 2) the first application of the context-aware filter to the MCPS domain; 3) a case-study evaluation on real-patient data from CHOP.

This paper is organized as follows. Section II describes the current approach to monitoring the O_2 concentration and presents the P_aO_2 estimation problem. In Section III we describe a minimal parameterized physiological model of the circulation of O_2 in the cardiopulmonary system and present ways of adding additional context measurements. Section IV introduces the context-aware filter and Section V presents the final model, complete with context measurements. Section VI gives the case-study evaluation of the context-aware filter and Section VII concludes the work.

II. PROBLEM STATEMENT

In this section, we outline the current approach to monitoring O_2 concentration, note its drawbacks, and formulate the problem that is addressed in this paper.

Currently, clinicians have only one available real-time measurement of blood oxygen, namely the hemoglobin-oxygen saturation in the peripheral capillaries (S_pO_2), measured by a pulse oximeter at an extremity (usually a finger tip). The saturation is a good measure of the O_2 concentration in the location it is measured because of the oxygen content equation [23]:

$$C_pO_2 = 1.34S_pO_2Hb + 0.003P_pO_2, \quad (1)$$

where C_pO_2 is O_2 concentration in the peripheral capillaries, Hb is the amount of hemoglobin in g/dL , and P_pO_2 is the partial pressure of dissolved oxygen in the peripheral capillaries measured in $mmHg$. According to (1), O_2 appears in only two forms in the blood – it is either bound to hemoglobin or dissolved in the blood. Equation (1) shows that, for normal values of P_pO_2 around 80-200 $mmHg$ and of Hb around 12-17 g/dL [23], the majority of O_2 is bound to hemoglobin. Thus, S_pO_2 is a good measure of the O_2 concentration in the peripheral capillaries. By assuming that C_pO_2 is just a delay of C_aO_2 (the O_2 concentration in the arteries), S_pO_2 can also be used as a good proxy for C_aO_2 . At the same time, however, the saturation is usually constant at 100% in healthy people. Thus, when reduced saturation is observed, the O_2 content has already decreased; hence, monitoring the O_2 concentration through S_pO_2 is reactive in nature.

In contrast, monitoring the partial pressure of dissolved O_2 is proactive. In addition to (1), dissolved O_2 and hemoglobin-bound O_2 are related according to a well-studied hemoglobin dissociation curve [17]. Figure 1 shows an example dissociation curve. While the magnitude of the curve may vary across




Fig. 1: A typical hemoglobin dissociation curve for O_2 . It shows the shape of the relationship between the partial pressure of dissolved O_2 and hemoglobin saturation. The curve is true for any physiological location, e.g., in the peripheral capillaries the horizontal axis label would be P_pO_2 and the vertical would be S_pO_2 .

patients, the overall S-shape remains the same. Figure 1 shows that for large values of the partial pressure, the saturation is close to 100%; at the same time, any noticeable decrease in saturation (and consequently the O_2 concentration) can be observed only after a large decrease in the partial pressure. Thus, monitoring the partial pressure of dissolved O_2 in arterial blood (P_aO_2) provides clinicians with a proactive way of addressing changes in O_2 content before they are reflected in changes in S_pO_2 .

Estimating P_aO_2 , however, is challenging because it cannot be measured non-invasively and in real time. Instead, we focus on other real time measurements, available in the modern ORs, as means to infer P_aO_2 . At CHOP, the anesthesia machine provides several pulmonary measurements, namely the fractions of O_2 and CO_2 in inspired and expired air, the volume and pressure of inspired air, respiratory rate and others. While it is possible to model the relationship between these variables and P_aO_2 (e.g., Fick’s laws of diffusion), such models contain multiple parameters that vary widely across patients. Instead, in this paper we use population-average values for these parameters and aim to incorporate the context information provided by the pulmonary measurements in order to improve estimation and correct inaccuracies introduced by wrong parameter values.

Problem. *The problem considered in this work is to develop an estimator for P_aO_2 and C_aO_2 by using the noninvasive real-time inputs (fraction of O_2 in inspired air, volume and pressure of inspired air, respiratory rate) and pulmonary measurements (partial pressure of CO_2 in exhaled air) available to clinicians.*

Remark: Our solution uses one blood-gas analysis in order to initialize the estimator.

TABLE I: Summary of cardiopulmonary partial pressures and blood concentrations. Partial pressures begin with the letter “P” whereas concentrations begin with “C”.

Variable Names	Physiological Location
P_iO_2	Airways (inspiration)
P_AO_2	Alveoli
C_aO_2, P_aO_2	Arteries
C_pO_2, P_pO_2	Peripheral capillaries
C_vO_2, P_vO_2	Veins
C_dO_2, P_dO_2	Pulmonary capillaries
P_eO_2	Airways (expiration)

III. PHYSIOLOGICAL MODEL

In order to develop an estimator for P_aO_2 , one needs to first identify a model mapping the available measurements to P_aO_2 , as well as formalize the dynamics of the variables in the human body. While models of varying complexity exist in the literature, typically, as the model complexity increases, so does the number of unknown model parameters (e.g., lung capacity, metabolism) that vary across patients. Since these parameters are unidentifiable with current non-invasive measurements, the most popular approach is to use minimal models, i.e., models with a minimal number of parameters that still capture the first- or second-order dynamics of the system.

We develop a minimal model, building on results from the work of Kretschmer et al. [12] on estimating P_aO_2 and from our previous work [9] on detecting drops in the O_2 concentration. Our model is approximate in the sense that it captures general trends and relationships among the variables in order to reduce the number of unidentifiable parameters. We use population average values for the few remaining parameters and improve the fidelity of the model by incorporating binary context measurements (explained in Section IV).

A. Overview of Physiological Variables

This section provides an overview of the physiological variables used in this paper. For reference, all variables are summarized in Table I and shown in Figure 2. In inspired air, the partial pressure of O_2 is denoted by P_iO_2 . In the lungs, the air enters the alveoli where the partial pressure is denoted by P_AO_2 . In the alveoli, diffusion occurs, and the gas enters the blood stream at the pulmonary capillaries where the partial pressure of O_2 is denoted by P_dO_2 and the total concentration is C_dO_2 . Note that, as shown in Figure 3, some of the blood is shunted (e.g., due to blood draining directly into the cavity of the left ventricle through the thebesian veins [23]) and does not participate in diffusion. When the blood from the pulmonary veins enters the heart, it is pumped in the arteries where the partial pressure and concentration are denoted by P_aO_2 and C_aO_2 , respectively. The arteries transport the blood to the peripheral capillaries (P_pO_2 and C_pO_2), where metabolism occurs and converts O_2 into CO_2 . Finally, the veins (P_vO_2 and C_vO_2) transport the blood back to the lungs and the cardiovascular cycle repeats. The breathing cycle concludes with expiration, where the partial pressure of O_2 in expired




Fig. 2: A simplified schematic model of O_2 variables in the respiratory and cardiovascular systems.¹

air is denoted by P_eO_2 .

B. Model Dynamics

Having introduced the variables and processes at a high level, we now formalize the model dynamics and introduce the related measurements. We develop a discrete-time model for the O_2 concentration and discuss how to convert from concentrations to partial pressures in Section V.² The relationship between the variables in the airways is governed by the alveolar air equation [7]:

$$P_AO_2(k) = F_iO_2(k)(P_{ATM} - P_{H_2O}) - \frac{P_ACO_2(k)(1 - F_iO_2(k)[1 - RQ])}{RQ}, \quad (2)$$

where F_iO_2 is the fraction of O_2 in inhaled air (it can be converted to P_iO_2 using the first term on the right-hand side), P_ACO_2 is the partial pressure of CO_2 in the alveoli, P_{ATM} and P_{H_2O} are the atmospheric and water vapor pressures (in $mmHg$), respectively, and RQ is the respiratory quotient. RQ is a measure of the ratio of O_2 and CO_2 used in metabolism and varies with the type of consumed food. Note that F_iO_2 is set by clinicians, so it can be considered as input, whereas P_ACO_2 is measured by end-tidal CO_2 ($EtCO_2$), i.e., the partial pressure of CO_2 at the end of the breath.³

¹Note that, for better illustration, the figure shows the pulmonary veins merging before entering the heart, whereas in healthy humans they connect to the left atrium directly.

²Our model is discrete-time because the available sensors (at CHOP) have a discrete sampling rate. It does not model the partial pressures of dissolved O_2 directly because the required relationships are nonlinear and would unnecessarily complicate the estimation task.

³Note that $EtCO_2$ 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 [6], hence it is not considered in this work.




Fig. 3: An illustration of shunted (bottom) vs. non-shunted (top) blood dynamics in the lung. O_2 -rich non-shunted blood participates in diffusion and then mixes with CO_2 -rich shunted blood.

When diffusion occurs, O_2 usually diffuses completely so that the partial pressures are the same:

$$P_dO_2(k) = P_AO_2(k). \quad (3)$$

After diffusion, O_2 is in the blood, so its concentration needs to be computed as well. To convert from partial pressure to concentration, one uses (1)⁴ in combination with the O_2 dissociation curve (Figure 1) in order to compute the saturation corresponding to that partial pressure. Let us denote the dissociation curve by g , i.e., g is a function mapping partial pressures of dissolved O_2 to hemoglobin oxygen saturation. Thus, the O_2 concentration in the pulmonary capillaries after diffusion can be expressed as:

$$C_dO_2(k) = 1.34Hb g(P_dO_2(k)) + 0.003P_dO_2(k). \quad (4)$$

Note that g varies greatly between patients. In Section V we show how to select g based on population averages.

Continuing with the cardiovascular dynamics, the concentration in arterial blood, as shown in Figure 3, is the weighted average of the concentrations in shunted and non-shunted blood, according to the fraction f of shunted blood. Then

$$C_aO_2(k) = (1 - f)C_dO_2(k) + fC_vO_2(k), \quad (5)$$

where the shunted blood has the same O_2 concentration as venous blood.

The O_2 concentration in the peripheral capillaries is assumed to be the same as in the arteries [23], i.e., no reactions occur that change the gas concentrations:

$$C_pO_2(k) = C_aO_2(k). \quad (6)$$

Finally, the concentration in the veins is equal to that in the peripheral capillaries minus the effect of metabolism:

$$C_vO_2(k + 1) = C_pO_2(k) - \mu, \quad (7)$$

⁴Note that the O_2 content equation is true for any location in the body, i.e., one can replace C_pO_2 and P_pO_2 with C_dO_2 and P_dO_2 .

where μ captures the patient-specific metabolic rate. Note that a delay is introduced in order to model the fact that it takes time for the blood to travel from the arteries to the veins.

The whole parameterized model can now be summarized in a typical state-space equation:

$$\begin{aligned} a_{k+1} &= (1 - f)(d_{k+1}) + f(a_k - \mu) + v_{1,k} \\ e_{k+1} &= e_k + v_{2,k} \\ y_k &= e_k + w_k, \end{aligned} \quad (8)$$

where $a_k := C_a O_2(k)$, $d_k := C_d O_2(k)$, $e_k := P_A CO_2(k)$, $y_k := EtCO_2(k)$, $v_k := [v_{1,k}, v_{2,k}]^T$ is white Gaussian process noise, and w is white Gaussian measurement noise, independent of v_k . Modeling the dynamics of $P_A CO_2(k)$ more precisely is possible but introduces more parameters. A random walk model achieves two goals: 1) a linear model with few parameters is maintained and 2) $P_A CO_2$ is a system state and, hence, may vary less than the noisy $EtCO_2$ measurements.

C. Context Information

The model above captures the general trends of O_2 concentration but does not have good predictive capacity for an arbitrary patient, unless the parameters (f , μ , RQ , Hb , P_{ATM} , P_{H_2O}) are identified. However, there are other measurements and inputs that can be used to improve the estimation.

For example when $S_p O_2$ is below 100%, this information can be used to upper-bound the O_2 concentration since the majority of O_2 is hemoglobin-bound (as shown in (1)). We do not use $S_p O_2$ measurements directly in the model since mapping the saturation to the concentration of dissolved O_2 requires knowledge of the magnitude of the dissociation curve g . Instead, $S_p O_2$ can be used as a context measurement which provides alarms whenever the saturation is below 100%.

In addition, clinicians have three other ways of controlling a patient's oxygenation: tidal volume (the volume of air in the lungs), respiratory rate and the pressure of inhaled volume. Increasing either of these usually results in higher O_2 concentration; however, due to patient variability, they have different effects in different cases. Thus, clinicians adjust their combination in order to optimize O_2 uptake. While it is not straightforward to directly incorporate the inputs into the model, they can be used as context information, similar to $S_p O_2$. Specifically, alarms can be raised whenever the input variability ought to result in higher or lower O_2 concentrations. This intuition is made precise in Section V.

Context alarms can be incorporated in the estimation process by using a context-aware filter, as explained in the next section.

IV. CONTEXT-AWARE FILTER

This section presents the context-aware filter, as developed in our previous work [8]. At a high-level, the context-aware filter performs state estimation by incorporating regular continuous state measurements (e.g., $EtCO_2$) as well as discrete alarm events (e.g., a certain signal is above a given threshold). The filter was shown to work well in automotive scenarios

where detections of nearby buildings were used to improve localization.

Formally, the context-aware filter assumes the system has known linear dynamics of the form:

$$x_{k+1} = A_k x_k + B_k u_k + w_k, \quad (9)$$

where $x \in \mathbb{R}^n$ is the system state, $u \in \mathbb{R}^p$ is the applied input, $w \sim \mathcal{N}(0, Q)$ is Gaussian process noise, and A and B are matrices of appropriate dimensions.

Additionally, there are two types of available measurements: continuous state measurements as well as discrete context measurements. State measurements have the usual linear observation model:

$$y_k = C_k x_k + v_k, \quad (10)$$

where we denote state sensors' measurements by $y_k \in \mathbb{R}^m$, $v_k \sim \mathcal{N}(0, R)$ is Gaussian measurement noise, and matrix C has appropriate dimensions.

Context measurements provide information about the system's context. Context is formally defined as a finite set $\mathcal{C} = \{c_1, \dots, c_N\}$, where each c_i is a context element that can only be detected by a context sensor from certain states; example context elements include a physiological signal (e.g., $S_p O_2$) that crosses a threshold and raises an alarm or a building that is recognized using image processing. For each context element, a measurement b_k^i is received such that is $b_k^i = 1$ if c_i is detected and $b_k^i = -1$ otherwise. Each context element has a probability of being detected given the current state; we denote this probability with $p_d(c_i | x)$, i.e.,

$$b_k^i = \begin{cases} 1 & \text{w.p. } p_d(c_i | x) \\ -1 & \text{w.p. } 1 - p_d(c_i | x), \end{cases} \quad (11)$$

where p_d is a function of the system state. Thus, the set of context measurements is $b_k \in \{-1, 1\}^N$.

Given this system, the problem that the context-aware filter needs to address is the usual filtering problem, i.e., if we are given a prior probability density function on the state $p_{k|k}(x) := p(x | u_{0:k-1}, y_{0:k}, b_{0:k})$, can we compute the posterior density

$$p_{k+1|k+1}(x) := p(x | u_{0:k}, y_{0:k+1}, b_{0:k+1}),$$

describing the system's state given all available measurements and inputs.

In our previous work, we focused on two classes of probability of detection functions that model a wide variety of applications and that lead to closed-form filters. In this section, we briefly describe the one that is relevant to the problem of estimating the O_2 concentration before presenting the corresponding filter equations.

Assumption. *Suppose the probability of context detection functions are sigmoid functions that are defined as the probit logistic function [14]:*

$$p_d(c_i | x_k) = \Phi(v_i^T x_k + a_i), \quad (12)$$

where Φ is the cumulative distribution function of the standard Normal distribution, $v_i \in \mathbb{R}^n$ is a vector of known parameters, and $a_i \in \mathbb{R}$ is a known parameter offset.

Note that (in the one-dimensional case) this function has a similar shape to the classical sigmoid function: $f(x) = 1/(1 + e^{-x})$. It is suitable for modeling threshold detection events such as a physiological signal crossing a certain threshold because the function is close to 0 for states that are far from the threshold but rises quickly when the threshold is reached and approaches 1 when it is crossed. The intuition is that once a state has crossed a certain threshold by some margin, even very imprecise sensors will detect the fact that it has crossed and will raise an alarm with high probability.

Given the probability of detection defined in (12), we can now present the resulting context-aware filter. For simplicity, only the final formulas are presented here; the entire derivation can be found in our previous work [8]. Similar to a Kalman filter, the context-aware filter has an update and a predict stage; the former, however, is split into continuous update and discrete update, depending on which measurement is received. The following three subsections present the stages of the filter, assuming that the prior $p_{k|k}$ has a **Gaussian distribution with mean μ and covariance matrix Σ** .

A. Predict

The predict stage is the classical Kalman filter prediction:

$$\begin{aligned} p_{k+1|k}(x) &= \int \phi(x; A_k z + B_k u_k, Q) \phi(z; \mu, \Sigma) dz \\ &= \phi(x; A_k \mu + B_k u_k, A_k \Sigma A_k^T + Q) \\ &= \phi(x; \mu^p, \Sigma^p), \end{aligned}$$

where $\phi(x; \mu, \Sigma)$ denotes the Gaussian density with mean μ and covariance matrix Σ .

B. Continuous Update

The continuous update is also the same as in the Kalman filter:

$$\begin{aligned} p_{k+1|k+1}^c(x) &= \frac{p(y_{k+1} | x) p_{k+1|k}(x)}{\int p(y_{k+1} | z) p_{k+1|k}(z) dz} \\ &= \frac{\phi(y_{k+1}; C_k x, R) \phi(x; \mu^p, \Sigma^p)}{\int \phi(y_{k+1}; C_k z, R) \phi(z; \mu^p, \Sigma^p) dz} \\ &= \phi(x; \mu^c, \Sigma^c), \end{aligned}$$

where

$$\begin{aligned} \mu^c &:= \mu^p + K^c(y_{k+1} - C_k \mu^p) \\ \Sigma^c &:= (I - K^c C_k) \Sigma^p \\ K^c &:= \Sigma^p C_k^T (C_k \Sigma^p C_k^T + R)^{-1}. \end{aligned}$$

C. Discrete Update

To simplify notation, we assume a single context measurement, b_k , with corresponding parameters v and a ; the approach can be straightforwardly extended to multiple measurements through repeated updates.

Proposition 1 ([9]). *The discrete update of the context-aware filter is as follows:*

$$p_{k+1|k+1}(x) = \frac{\Phi(b_k(v^T x + a)) \phi(x; \mu^c, \Sigma^c)}{Z}, \quad (13)$$

where

$$Z = \Phi\left(\frac{b_k(v^T \mu^c + a)}{\sqrt{v^T \Sigma^c v + 1}}\right).$$

As discussed in our previous work, the posterior in (13) is no longer Gaussian, yet a Gaussian with the same first two moments is a reasonable approximation.

Approximation. *We approximate the posterior distribution in (13) with a Gaussian distribution with the same mean and covariance matrix.*

The next proposition describes the first two moments, which completes the filter equations.

Proposition 2 ([9]). *The mean of the distribution in (13) is:*

$$\mu^d = \mu^c + b_k \Sigma^c v \frac{\phi\left(\frac{b_k(v^T \mu^c + a)}{\sqrt{v^T \Sigma^c v + 1}}; 0, 1\right)}{Z \sqrt{v^T \Sigma^c v + 1}}. \quad (14)$$

The covariance matrix of the distribution in (13) is:

$$\begin{aligned} \Sigma^d &= \Sigma^c + \mu_1 (\mu^c)^T + \mu^c \mu_1^T - \mu^c (\mu^c)^T \\ &\quad - b_k \Sigma^c v v^T \Sigma^c \frac{\phi\left(\frac{b_k(v^T \mu^c + a)}{\sqrt{v^T \Sigma^c v + 1}}; 0, 1\right) (v^T \mu^c + a)}{Z (v^T \Sigma^c v + 1)^{3/2}} - \mu_1 \mu_1^T. \end{aligned}$$

In summary, this section presents the context-aware filter that incorporates continuous state measurements and discrete context measurements. The following section presents the context-aware model for estimating the blood oxygen concentration.

V. CONTEXT-AWARE MODEL

Having described the parameterized model and the context-aware filter in the previous two sections, in this section we develop the full model, including both state and context measurements.

A. System Dynamics

We begin by noting that the model in (8) is close to linear in the ranges we are interested in. To see this, note that $F_i O_2(k)$, i.e., the fraction of O_2 in inhaled air, is 21% in breathing air and usually much higher during mechanical ventilation. This means that $P_A O_2(k)$, as computed in (2), is also very high (in the extreme case when $F_i O_2(k) = 100\%$, $P_A O_2(k) = 713 \text{ mmHg}$, when we plug in normal values for $P_{ATM} = 760 \text{ mmHg}$, $P_{H_2O} = 43 \text{ mmHg}$). This in turn means that $P_d O_2(k)$ is also high, i.e., in the top right corner of the dissociation curve in Figure 1. Therefore, $g(P_d O_2(k)) \approx 1$, i.e., (4) simplifies to:

$$C_d O_2(k) = 1.34 H b + 0.003 P_d O_2(k). \quad (15)$$

Using (15) in (8), the new model becomes

$$\begin{aligned} a_{k+1} &= (1-f)(1.34Hb + 0.003(c_1u_k + c_{2,k}e_k)) \\ &\quad + f(a_k - \mu) + v_{1,k} \\ e_{k+1} &= e_k + v_{2,k} \\ y_k &= e_k + w_k, \end{aligned} \quad (16)$$

where $c_1 = (P_{ATM} - P_{H_2O})$, $u_k = F_iO_2(k)$ and $c_{2,k} = (1 - u_k[1 - RQ])/RQ$.

Thus, the above model is a linear time-varying system (note that the input u_k , which also appears in $c_{2,k}$, is multiplied by one of the states, e_k , but this does not introduce non-linearities because in this paper we are only considering the estimation problem and not the control problem). There are several parameters in the model; as argued above, these cannot be learned due to unobservability. Thus, we select population average values for the parameters (except for f) and argue that context measurements will correct model inaccuracies. More specifically, based on medical literature [23], these values were selected as: $Hb = 12$ g/dL, $P_{ATM} = 760$ mmHg, $P_{H_2O} = 47$ mmHg, $\mu = 5$ mL/dL, $RQ = 0.8$.

The parameter f , which represents the fraction of shunted blood, does not have typical ranges and can vary widely depending on the patient's condition (e.g., a pulmonary shunt leads to 50% shunted blood). Thus, we adopt an approach used in prior work [12] for the estimation of f . This requires an initializing measurement of P_aO_2 through blood-gas analysis. By obtaining this measurement, one can estimate C_aO_2 through (4), where a functional form for g is also assumed, as developed in [11]. Then, using (8) and assuming that $a_{k+1} = a_k = a$, one obtains the equation:

$$a = (1-f)d + f(a - \mu), \quad (17)$$

where d is computed from (4). This equation can now be solved for f in order to obtain the fraction of shunted blood.

B. State Measurements

Note that we have only one state measurement, namely y_k , which directly measures state e_k , i.e., we have a linear state observation model.

C. Context Measurements

For the context observation model, we use other available measurements as described in Section III-C. We develop two classes of context measurement – one class based on the hemoglobin-oxygen saturation (S_pO_2) and one class based on the other inputs.

The alarm related to S_pO_2 measurements is raised when $S_pO_2(k)$ drops below 99%. Since S_pO_2 represents the vast majority of O_2 content, according to (1), one can reasonably conclude that if $S_pO_2 < 99\%$, then $C_aO_2 < (1.34 * 0.99)Hb$. This naturally leads to a threshold alarm based on S_pO_2 ; the parameters of (12) can be set to $v_i = [1 \ 0]^T$, $a_i = -(1.34 * 0.99)Hb$.

The second class of alarms consists of several alarms due to the more complicated nature of the signal. This class of

alarms aims to use the three other inputs available to clinicians: tidal volume (V_t), respiratory rate (RR) and peak inspiratory pressure (PIP). Each of these inputs affects diffusion through Fick's law of diffusion, which can be stated as follows, adapted to this application [23]:

$$\dot{d}_k \propto cA(P_aO_2(k) - P_dO_2(k)), \quad (18)$$

where $d_k = C_dO_2(k)$ as before, c is a constant that captures the O_2 diffusive capacity and lung thickness, and A is the lung surface area. Equation (18) states that the number of diffused moles is directly proportional to the surface area and to the difference between the pressures in the lung and in the blood. Note that (18) cannot be solved because of the unknown initial condition and unknown parameters. However, one can compute the signal on the right hand side at each point in time; since it is proportional to O_2 diffusion, when the signal is higher, one would expect the O_2 concentration to increase as well.

To construct this signal, note that if we make the usual assumption that a lung is a sphere, then $A \propto V_t^{2/3}$. In addition, since a patient can take several breaths in between two measurements, the respiratory rate can be used as well in order to compute a ‘‘cumulative tidal volume’’ since the last measurement, i.e.,

$$\bar{V}(k) = \frac{t_S}{60} RR(k)V_t(k), \quad (19)$$

where t_S is the sampling time in seconds. Thus $A \propto \bar{V}^{2/3}$.

Furthermore, note that PIP is directly proportional to P_aO_2 . Thus, one can adapt (2) to include PIP as effectively increasing atmospheric pressure:

$$\begin{aligned} P_aO_2(k) &= F_iO_2(k)(P_{ATM} - P_{H_2O} + PIP(k)) \\ &\quad - \frac{P_aCO_2(k)(1 - F_iO_2(k)[1 - RQ])}{RQ}. \end{aligned} \quad (20)$$

The final piece of the ‘‘diffusion’’ signal is the initial value of $P_dO_2(k)$. Since the initial value is equal to the venous $P_vO_2(k)$, $P_dO_2(k)$ is directly proportional to $P_aO_2(k-1)$; therefore, we use the expected value of a_{k-1} to obtain an ‘‘expected’’ $P_aO_2(k)$. To obtain a rough estimate of the partial pressure, one needs to invert (1) and solve the following nonlinear equation (e.g., by using simplex methods):

$$\begin{aligned} \mathbb{E}[a_{k-1}] &= 1.34Hb g(\mathbb{E}[P_aO_2(k-1)]) \\ &\quad + 0.003\mathbb{E}[P_aO_2(k-1)], \end{aligned} \quad (21)$$

where \mathbb{E} denotes the expectation operator; note that a functional form of g must be assumed, e.g., as in [11]. Thus, the final constructed signal is:

$$s_k = \bar{V}(k)^{2/3} * (P_aO_2(k) - \mathbb{E}[P_aO_2(k-1)]). \quad (22)$$

In order to use s as a context measurement, one needs to identify changes in its baseline and raise alarms. To do this, an initial baseline of the signal is selected, and alarms are raised if the signal is too high or too low with respect to that baseline. In particular, suppose the first blood-gas measurement of P_aO_2 is received at time step q ; then the value of s_q is selected as a




Fig. 4: Absolute errors for each of the two compared P_aO_2 estimators, the context-aware filter and the F_iO_2 -based estimator. Red dashed line shows the average error of the context-aware filter, whereas blue dashed line indicates the average error of the F_iO_2 -based estimator.

baseline and alarms are raised at a later step k if s_k is lower than $0.5s_q$ or $0.8s_q$ or if it is higher than s_q , $1.2s_q$, or $1.5s_q$.

To select the respective C_aO_2 thresholds, we note that since s_k is directly proportional to C_aO_2 , a relative change in s_k should result in a similar relative change in $C_aO_2(k)$. Thus, we identify the baseline $C_aO_2(q)$ and set the thresholds accordingly. For example, if $s_k < 0.8s_q$, then an alarm is raised and the corresponding sigmoid parameters are $v_i = [1 \ 0]^T$, $a_i = -0.8C_aO_2(q)$. The other thresholds are derived similarly.

This fully specifies the context observation model and completes the full system model. The following Section presents the case-study evaluation of this model and of the resulting context-aware filter.

VI. CASE STUDY

This section presents a case-study evaluation of the P_aO_2 estimator developed in this work. We use real-patient data collected during lung lobectomy surgeries on children performed at CHOP. A lung lobectomy is the surgical removal of a lung lobe, often due to disease such as cancer or a cystic lung lesion; lobectomies often require one-lung ventilation (i.e., the endotracheal tube is inserted down a mainstem bronchus, so the patient breathes with one lung only) in order to keep the perioperative lung still. In children, one lung is often not enough to provide sufficient O_2 to the body, hence the O_2 concentration tends to decrease.

For evaluation purposes, we use the blood-gas samples taken during these cases and compare them with our estimates. As mentioned earlier, clinicians do not usually draw blood unless they suspect a problem, hence there are at most several measurements per case, while most cases do not have any. After removing all cases with less than two measurements (recall that one is necessary for the algorithm initialization),


we retain 51 cases overall. In each case, we initialize the context-aware filter with the first blood-gas measurement and evaluate it on the remaining ones. In addition, as described in the previous section, the diffusion signal baseline (used to define context measurement thresholds) is also computed at the time of the first blood-gas measurement. Finally, note that the available blood-gas measurements only contain P_aO_2 measurements, hence **only P_aO_2 estimates are evaluated**.

Figure 4 presents the absolute errors of the context-aware filter, with all measurements from all patients stacked together. For better evaluation, we compare the filter with a P_aO_2 estimation algorithm developed in previous work that uses a similar model and also requires one blood-gas measurement for initialization [12]; this algorithm is named here “ F_iO_2 -based estimator”. As can be seen in the Figure, the context-aware filter eliminates all of the F_iO_2 -based estimator outliers except for one (discussed below). In addition, the context-aware filter achieves a lower average error overall, 51.7 mmHg, than the F_iO_2 -based estimator’s average error, 63.3 mmHg. To put the error in perspective, note that P_aO_2 measurements are usually in the 200-400 mmHg range (due to F_iO_2 being usually close to 100%), with the exception of a few cases with infants where it is in the 100-200 mmHg range. With this in mind, errors of 100 mmHg are still significant; yet, the reasonably uniform distribution of the errors suggests that the context-aware filter is not greatly affected by inter-patient variability and is thus a reasonable choice of estimator, once a more accurate model and more precise context measurements are obtained.


To further analyze the performance of the context-aware filter, we analyze two cases, one with very good performance and one with bad performance. Figure 5a presents an example case where context measurements bring a significant improvement.⁵ It shows the estimates of each of the two estimators, together with the blood-gas samples, as well as all other measurements and inputs used in the filters. Note that after the initializing blood-gas measurement, clinicians reduce F_iO_2 (around time step 800), probably content with the patient’s current condition. Yet, other inputs (V_t , RR , PIP) do not change greatly, indicating that the patient’s O_2 concentration should not decrease significantly. This is confirmed by the diffusion signal, which only decreases by about 20%; thus the $0.8s_q$ alarm is raised but the $0.5s_q$ alarm remains silent, which causes the filter to set the estimate somewhere in between. In contrast, the F_iO_2 -based estimator is greatly affected by the reduced F_iO_2 .

As an example bad-performance case, we consider the outlier in Figure 4 for the context-aware filter. Note that, once again, the context-aware filter is not greatly affected by the decreased F_iO_2 . In this case, the problem is that the diffusion signal is actually too low at the initialization stage (around step 580), so high-signal alarms are raised later. A possible explanation for the bad performance of the filter in this case

⁵Note that the estimates prior to the first blood-gas sample are not used for evaluation but are included for completeness.



(a) Example case with good estimation by the context-aware filter.



(b) Example case with bad estimation by the context-aware filter.

Fig. 5: Example cases for different scenarios. Red S_pO_2 data points indicate low- S_pO_2 alarms; blue S_pO_2 data points indicate no S_pO_2 alarms. Diffusion signal: red data points indicate $0.5s_q$ alarms; yellow data points indicate $0.8s_q$ alarms; green data points indicate no alarms; blue data points indicate $1.2s_q$ alarms; magenta data points indicate $1.5s_q$ alarms (recall s_q is the diffusion signal at the initialization point, i.e., first blood-gas analysis).

is a wrong timestamp of the first blood-gas sample; these timestamps are entered manually and are prone to significant errors, as explored in prior work [19]. In particular, note that tidal volume and respiratory rate are steadily decreasing from around step 420 onwards; thus it is not unlikely that the blood-gas sample was obtained at that time as well. As is apparent from the diffusion signal, if the baseline is set around step 420, no high-signal alarms would be raised later. Finally, note that estimation is made harder by the lack of low- S_pO_2 alarms.

Based on these results, we conclude that the context-aware filter is a promising direction for future research in the MCPS area. By incorporating auxiliary information, it is able to correct some of the deficiencies of imprecise models and results in better estimation overall, even when the variables in question are unobservable.

VII. CONCLUSION

In this work we addressed the problem of estimating the blood oxygen content given only non-invasive measurements. In addition to regular continuous state measurements, we incorporated context measurements as derived from auxiliary input signals that are not used in the system model. We included these measurements in a context-aware filter and evaluated the resulting estimator on real-patient data collected at CHOP. We concluded that the context-aware filter is a promising direction for research in MCPS, especially in the

presence of imprecise models and sensors. As part of future work, we will investigate other context measurements that will further alleviate the model's uncertainty.

REFERENCES

- [1] N. R. Ahmed, E. M. Sample, and M. Campbell. Bayesian multicategorical soft data fusion for human–robot collaboration. *Robotics, IEEE Transactions on*, 29(1):189–206, 2013.
- [2] D. Arney, M. Pajic, J. Goldman, I. Lee, R. Mangharam, and O. Sokolsky. Toward patient safety in closed-loop medical device systems. In *Proceedings of the 1st International Conference on Cyber-Physical Systems*, pages 139–148, 2010.
- [3] N. Atanasov, M. Zhu, K. Daniilidis, and G. Pappas. Semantic localization via the matrix permanent. In *Robotics: Science and Systems*, 2014.
- [4] A. N. Bishop and B. Ristic. Fusion of spatially referring natural language statements with random set theoretic likelihoods. *Aerospace and Electronic Systems, IEEE Transactions on*, 49(2):932–944, 2013.
- [5] 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.
- [6] S. Cruickshank and N. Hirschauer. The alveolar gas equation. *Continuing Education in Anaesthesia, Critical Care & Pain*, 4:24–27, 2004.
- [7] 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.
- [8] R. Ivanov, N. Atanasov, M. Pajic, G. Pappas, and I. Lee. Robust estimation using context-aware filtering. In *53rd Annual Allerton Conference on Communication, Control, and Computing*, 2015.

- [9] R. Ivanov, J. Weimer, A. Simpao, M. Rehman, and I. Lee. Early detection of critical pulmonary shunts in infants. In *Proceedings of the ACM/IEEE Sixth International Conference on Cyber-Physical Systems, ICCPS '15*, pages 110–119. ACM, 2015.
- [10] S. Joshi and S. Boyd. Sensor selection via convex optimization. *Transactions on Signal Processing*, 57(2):451–462, 2009.
- [11] G. Kelman. Digital computer subroutine for the conversion of oxygen tension into saturation. *Journal of Applied Physiology*, 21(4):1375–1376, 1966.
- [12] J. Kretschmer, T. Becher, A. Riedlinger, D. Schadler, N. Weiler, and K. Moller. A simple gas exchange model predicting arterial oxygen content for various FiO₂ levels. In *Engineering in Medicine and Biology Society (EMBC), 2013 35th Annual International Conference of the IEEE*, pages 465–468, July 2013.
- [13] R. Mahler. *Statistical Multisource-Multitarget Information Fusion*. Artech House, Inc., 2007.
- [14] H. Nickisch and C. Rasmussen. Approximations for binary gaussian process classification. *Journal of Machine Learning Research (JMLR)*, 9(Oct):2035–2078, 2008.
- [15] M. Pajic, R. Mangharam, O. Sokolsky, D. Arney, J. Goldman, and I. Lee. Model-driven safety analysis of closed-loop medical systems. *IEEE Transactions on Industrial Informatics*, 10(1):3–16, 2014.
- [16] 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.
- [17] S. Shafer, J. P. Rathmell, and R. Stoelting. *Stoelting's Pharmacology & Physiology*. Wolters Kluwer, 2014.
- [18] D. Simon and D. L. Simon. Constrained kalman filtering via density function truncation for turbofan engine health estimation. *International Journal of Systems Science*, 41(2):159–171, 2010.
- [19] A. F. Simpao, E. Y. Pruitt, S. D. Cook-Sather, H. Gurnaney, and M. Rehman. 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.
- [20] B. Sinopoli, L. Schenato, M. Franceschetti, K. Poolla, M. Jordan, and S. Sastry. Kalman filtering with intermittent observations. *Automatic Control, IEEE Transactions on*, 49(9):1453–1464, 2004.
- [21] M. P. Vitus, W. Zhang, A. Abate, J. Hu, and C. J. Tomlin. On efficient sensor scheduling for linear dynamical systems. *Automatica*, 48(10):2482–2493, 2012.
- [22] J. Weimer, R. Ivanov, A. Roederer, S. Chen, and I. Lee. Parameter invariant design of medical alarms. *IEEE Design & Test*, 2015.
- [23] J. B. West. *Respiratory Physiology: The Essentials*. Lippincott Williams & Wilkins, 2012.
- [24] K. Wyffels and M. Campbell. Negative information for occlusion reasoning in dynamic extended multiobject tracking. *Robotics, IEEE Transactions on*, 31(2):425–442, 2015.