






Low-Order Mechanistic Models for Volumetric and Temporal Capnography: Development, Validation, and Application

Elizabeth K. Murray , Graduate Student Member, IEEE, Carine X. You , Student Member, IEEE, George C. Verghese , Life Fellow, IEEE, Baruch S. Krauss , and Thomas Heldt , Senior Member, IEEE

Abstract—Objective: Develop low-order mechanistic models accounting *quantitatively* for, and *identifiable* from, the *capnogram* — the CO₂ concentration in exhaled breath, recorded over time (Tcap) or exhaled volume (Vcap). **Methods:** The *airflow* model's single “alveolar” compartment has compliance and inertance, and feeds a resistive unperfused airway comprising a laminar-flow region followed by a turbulent-mixing region. The *gas-mixing* model tracks mixing-region CO₂ concentration, fitted breath-by-breath to the measured capnogram, yielding estimates of model parameters that characterize the capnogram. **Results:** For the 17 examined records (310 breaths) of airflow, airway pressure and Tcap from ventilated adult patients, the models fit closely (mean rmse 1% of end-tidal CO₂ concentration on Vcap; 1.7% on Tcap). The associated parameters (4 for Vcap, 5 for Tcap) for each exhalation, and airflow parameters for the corresponding forced inhalation, are robustly estimated, and consonant with literature values. The models also allow, *using Tcap alone*, estimation of the entire exhaled airflow waveform to within a scaling. This suggests new Tcap-based tests, analogous to spirometry but with normal breathing, for discriminating chronic obstructive pulmonary disease (COPD) from congestive heart failure (CHF). A version trained on 15 exhalations from each

of 24 COPD/24 CHF Tcap records from one hospital, then tested 100 times with 15 random exhalations from each of 27 COPD/31 CHF Tcap records at another, gave mean accuracy 80.6% (stdev 2.1%). Another version, tested on 29 COPD/32 CHF, yielded AUROC 0.84. **Conclusion:** Our mechanistic models closely fit Tcap and Vcap measurements, and yield subject-specific parameter estimates. **Significance:** This can inform cardiorespiratory care.

Index Terms—Capnography, deadspace, respiratory model, spirometry, ventilators.

I. INTRODUCTION

CAPNOGRAPHY records the *partial pressure of exhaled* CO₂ (PeCO₂) as a function of time (temporal capnography, Tcap) or exhaled volume (volumetric capnography, Vcap) [1], [2]. Tcap can be easily obtained, with the subject breathing normally while a continuous sample of the exhaled breath is drawn into the instrument — the capnograph — through a nasal cannula. This is known as a *sidestream* configuration. (The inhaled breath is also sampled, but is not of clinical or physiological interest.) The CO₂ partial pressure of the sample is measured in the capnograph as a function of time, typically using an infrared sensor. (See Supplementary Material to review the relationship of partial pressure to concentration.) Tcap is routinely used in many clinical settings (operating room, intensive care unit, emergency department, ambulance service) to monitor ventilation, determine respiratory rate, and measure end-tidal CO₂ partial pressure (EtCO₂).

Vcap is uncommon in the clinical setting, as it requires measuring the entire expired airflow, with PeCO₂ typically measured close to the mouth, in the main external tubing; this is the *mainstream* configuration. PeCO₂ can instead be measured on gas sampled sidestream from the main tubing. Vcap is primarily used in the pulmonary function laboratory or in surgical or intensive-care settings with mechanically ventilated patients, and is derived from simultaneous recordings of Tcap and airflow. The pressure at the mouth (or “airway pressure”) may also be recorded. Vcap enables evaluating ventilation and gas exchange, which are reflective of deadspace (i.e., unperfused lung volumes), cardiac output, metabolic status, and related variables [3], [4], [5], [6], [7].

Manuscript received 27 April 2022; revised 7 September 2022 and 30 December 2022; accepted 12 March 2023. Date of publication 29 March 2023; date of current version 30 August 2023. (Elizabeth K. Murray and Carine X. You contributed equally to this work.) (Corresponding author: Thomas Heldt.)

Elizabeth K. Murray was with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, USA. She is now with the Department of Electrical Engineering and Computer Science, University of California, USA.

Carine X. You was with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, USA. She is now with the Department of Economics, Massachusetts Institute of Technology, USA.

George C. Verghese is with the Research Laboratory of Electronics and Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, USA.

Baruch S. Krauss is with the Division of Emergency Medicine, Boston Children's Hospital, USA, and also with the Department of Pediatrics, Harvard Medical School, USA.

Thomas Heldt is with the Institute for Medical Engineering and Science and Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: thomas@mit.edu).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TBME.2023.3262764>, provided by the authors.

Digital Object Identifier 10.1109/TBME.2023.3262764




Fig. 1. (a) Temporal waveforms of CO₂ partial pressure (black, in mmHg), airflow (red, in deciliters/s), and airway pressure (blue, in mmHg) in a ventilated adult patient during surgery (exhalation #2 of record 365 A in [8]). (b) Vcap plot of PeCO₂ versus expired volume for this exhalation, parameterized along the curve by time t from the start of exhalation at 2.1 s.

A. Tcap and Vcap From Ventilator Data

For orientation, Fig. 1(a) shows temporal waveforms of airflow, airway pressure, and PeCO₂ at the airway, for one respiratory cycle of a ventilated adult surgical patient under anesthesia, from the CapnoBase database [8], [9].

The initial portion of these waveforms corresponds to the start of an *inhalation*, with the ventilator-controlled flow going negative and holding essentially constant for most of the inhalation, till 2.1 s in the figure, while the airway pressure at the mouth increases. The inhaled gas is essentially CO₂-free, so the corresponding segment of the capnogram, with no discernible delay, begins its descent to 0 mmHg from the EtCO₂ value of around 27 mmHg that it had at the end of the previous exhalation. When inhalation ends at 2.1 s, (passive) *exhalation* begins and the flow changes to positive. Peak airflow is reached very soon after, and airflow then decays to essentially 0 before the next inhalation begins.

Measured PeCO₂ remains at 0 mmHg for 0.2 s after the start of exhalation, then begins to rise at 2.3 s, as CO₂-rich alveolar gas progressively mixes in with previously inhaled air and arrives at the CO₂ sensor. PeCO₂ finally reaches a value that reflects the mean CO₂ partial pressure in alveolar gas, typically 25–50 mmHg (27 mmHg here). Factors such as noise in the measurement, uneven emptying of the alveoli, and unsteadiness in the flow near the end of the exhalation can make the rise non-monotonic.

Integrating the exhaled airflow waveform with respect to time on any particular exhalation, and choosing the start of the exhalation as the time origin $t = 0$, yields the cumulative exhaled volume $V(t)$ up to time t in that exhalation, with the value at the end of the exhalation being the *tidal volume*, V_T . Airflow is thus the time-derivative $\dot{V}(t)$. Plotting PeCO₂(t) as a function of $V(t)$ then yields the volumetric capnogram Vcap, as shown in Fig. 1(b) for the exhalation in Fig. 1(a), with $V_T = 610$ mL, and parameterized along the curve by time t from the start of exhalation.

B. Prior Work on Modeling Tcap and Vcap

The Tcap and Vcap literatures have been fairly disjoint, but both have focused on extracting landmark features of capnogram

morphology (slopes, points of inflection or intersection, abscissas and ordinates corresponding to these, and various associated areas) with the aid of piecewise-linear fits to various segments of the capnogram or, in the case of [10] for example, functional fits to the full capnogram (not derived from any consideration of the underlying physiology). These features are then associated or correlated with particular phases, processes, parameters, and pathologies of respiration [11].

The interest in capnogram morphology is partly because particular cardiorespiratory conditions tend to be associated with characteristic capnogram shapes, e.g., Tcap from patients with chronic obstructive pulmonary disease (COPD) tends to have a “shark’s fin” appearance [12]. Thus, analysis of capnogram shape has diagnostic utility. We showed in earlier work [13] that simple machine-learning algorithms (specifically, quadratic discriminant analysis) operating on just 4 Tcap waveform features from each of 80 exhalations in a record can distinguish between patients with congestive heart failure (CHF) and those with COPD, with accuracies around 75–80%. Some recent work using machine-learning methods, [14] and related papers by the same authors, reports significantly better results for this task, but differences in data selection/labeling complicate comparison with [13].

Tcap suffices to determine respiratory rate and EtCO₂. The main clinically relevant parameters inferred from Vcap are the anatomical and physiological deadspace volumes. Anatomical deadspace corresponds to the portion of the respiratory tract with no alveoli, namely the conducting upper airway. Physiological deadspace reflects alveolar regions that are unperfused (for a variety of pathophysiological reasons). Classical modeling calculations for estimating deadspaces involve tracking CO₂ mass balance in various compartments at two or three discrete epochs of respiration [3], [4], [5], [6], [7].

Most prior work on modeling Tcap and Vcap stops well short of fitting the entire capnogram trace using a model that (i) mechanistically reflects dominant aspects of the underlying anatomical and physiological processes, and (ii) is identifiable from the capnogram alone, for a subject-specific representation of the capnogram.

The papers [15], [16], [17] point in the required direction with their *continuous* rather than discrete tracking of CO₂ mass in

anatomically and physiologically relevant compartments over a tidal breathing cycle, using coupled differential equations. Of these, only [17] aims to model the capnogram; assuming known airflow, the model can synthesize Tcap (and in principle Vcap, though this is not explicitly studied). The paper presents a particle-filtering approach to estimating model parameters, illustrated using simulated data with a sinusoidal airflow.

Our previous work [18], [19] on mechanistic modeling of the capnogram was in the spirit of [17], but aimed at modeling sidestream Tcap with no accompanying airflow measurements. A first- or second-order differential equation for the airflow dynamics associated with a single-“alveolus” model described the discharge (over time t) of alveolar gas into an unperfused *mixing deadspace* region during exhalation. We assumed constant CO₂ partial pressure \bar{p}_A for this alveolar gas over the entire exhalation, ignoring secondary effects from any spatial variation of alveolar CO₂ partial pressure across the lung or temporal variation during exhalation. A separate first-order differential equation described the progressive mixing of the injected alveolar gas with previously inhaled air. The waveform of rising CO₂ partial pressure in the mixing region constituted Tcap in this model. Model parameters, reflecting quantities of physiological or clinical significance, could then be estimated by numerically fitting the model’s capnogram to the measured one. We demonstrated good discrimination between COPD and CHF patients, and also those with normal respiration, using interpretable model parameters estimated from the capnogram, rather than using more generic morphological features.

C. Outline of This Paper

We introduce significant refinements of our models in [18], [19], to enable quantitative modeling and comparison of Vcap and Tcap in a common framework. We develop models of sufficiently low complexity to be robustly identifiable in a subject-specific way through capnography alone, yet capturing gross features of the underlying respiratory anatomy and physiology. These are necessarily highly aggregated compartmental models. A more detailed model may better represent known respiratory anatomy and physiology, but unless its parameters can be reliably estimated from the capnogram, the parameters will not be subject-specific, and the model is unlikely to be a clinically helpful representation of the capnogram, even if useful for purposes such as simulation [15].

Importantly, we recognize that the (unperfused) anatomical deadspace region can be separated into two sequential segments: (i) a *laminar-flow* region of volume V_L in the lower conducting airways, through which laminar flow is largely maintained during exhalation, as alveolar gas initially forces out previously inhaled air in this region; and (ii) a *turbulent-mixing* region of volume V_M in the upper airways, which the laminar flow feeds into. The volume V_L is seen in Fig. 1(b) as the initial segment of the Vcap waveform (extending to a little over 150 mL) during which PeCO₂ is essentially at 0 mmHg.

A second refinement of the model here relates to the CO₂ partial pressure of alveolar gas entering the mixing volume during

exhalation: we now use an *exponential settling* to the value \bar{p}_A , rather than the step change used in our earlier work. The smoother transition is intended to represent in some form, in this single-“alveolus” model, the non-uniform emptying of the alveoli into the mixing volume. This modification substantially improves the fit of the model to measured data.

We first focus on using these refined models to represent a much richer set of data than in our previous work, namely ventilator signals such as those in Fig. 1(a). The signals during exhalations yield Tcap and Vcap, permitting detailed quantitative comparison and validation of parameters extracted from Tcap alone against those obtained from Vcap and from the airflow and airway-pressure signals. We show that the ascending part of the Vcap waveform, beyond the volume V_L and essentially all the way to V_T , is well described by the same simple one-compartment mixing dynamics used in our earlier work [18], [19]. Furthermore, the flow and pressure data during the corresponding *inhalations* allow independent estimates of parameters of the airflow model, which we compare against those determined from Tcap and in other ways.

We also extract deeper information from the model than in our previous work, recognizing that: (i) the small and large time-constants associated with the second-order airflow subsystem respectively govern the fast initial rise and subsequent slower decay of airflow during exhalation (and are associated with inertance and compliance effects respectively); and (ii) the coupled gas-mixing and airflow models allow estimation of the entire exhaled airflow waveform normalized by mixing volume, namely $\dot{V}(t)/V_M$, from Tcap alone.

The latter result motivates novel Tcap-based analogs of spirometry tests widely used for COPD–CHF discrimination and severity assessment. While spirometry is based on airflow measurements during *forced* exhalation, Tcap can be recorded during normal breathing instead, which is important for measurements on patients with respiratory difficulties. We introduce and evaluate such a test in the second part of this article.

II. MODEL DEVELOPMENT AND ANALYSIS

A. Anatomical and Physiological Basis

In setting up our model, we keep in mind Weibel’s [20], [21] classical anatomical model of the lungs, idealized to have binary branching at each of 23 successive generations, as summarized in Fig. 2. Generations 0 through 16 constitute the unperfused conducting zone, whose aggregate volume makes up the anatomical deadspace. Beyond this, the lungs transition to the perfused alveolar or respiratory region.

Detailed modeling and experimental studies, e.g., [22] [23], suggest that flow in the lungs is turbulent in the upper part of the conducting zone, generations 0 to 5, where there are larger and fewer vessels, with relatively high airflow velocities, hence higher Reynolds numbers. This is the region we consider to be the turbulent-mixing deadspace, with volume V_M ; to keep notation and terminology simple, we shall use V_M to denote both the region and its actual volume. We shall denote the partial pressure of CO₂ in V_M at time t by $p_M(t)$, and *this waveform constitutes the synthesized Tcap in our model.*

Generation	Structure	Nominal volume (mL)	
0	Trachea	30	} Mixing region, volume V_M
1-3	Bronchi	17	
4	Bronchioles	3	
5	Terminal bronchioles	3	
6-16	Terminal bronchioles	115	} Laminar-flow region, volume V_L
17-19	Respiratory bronchioles	}	} Alveolar region
20-22	Alveolar ducts		
23	Alveolar sacs		

Fig. 2. Idealized Weibel lung structure [20], [21] with successive bifurcations over 23 generations. At the start of exhalation, the perfused respiratory zone at the bottom (generations 17-23, shown in red) contains CO₂-rich alveolar gas, while the non-perfused conducting zone above is essentially CO₂-free. We model the upward flow of exhaled alveolar gas as laminar through most of the terminal bronchioles (generations 6-16, shown in blue), largely preserving the CO₂-rich/CO₂-free boundary till this arrives in the turbulent-mixing region (shown in black).

In contrast, flow tends to be laminar in the remainder of the conducting zone, generations 6 through 16, where there are smaller and much more numerous vessels, with correspondingly lower airflow velocities. This is what we take to be the laminar-flow deadspace region; we shall use V_L to denote both the region and its actual volume.

At the transition from inhalation to exhalation, we expect that the boundary between (i) air at CO₂ partial pressure 0 mmHg that has just been inhaled, and (ii) alveolar gas at average CO₂ partial pressure \bar{p}_A that is about to be exhaled, lies approximately in the transition region between generations 16 and 17, not as a sharp interface, but rather with some “blurring” due to local diffusion (see Supplementary Material).

As exhalation begins, what initially enters V_M is the previously inhaled gas in V_L , so there is initially no change in the CO₂ partial pressure in V_M , and $p_M(t)$ remains at 0 mmHg. As already mentioned, this initial interval of zero CO₂ partial pressure is seen in Fig. 1(b), extending to around 150 mL.

Once this initial volume V_L has emptied into the mixing region V_M , the boundary layer that was previously in the transition region (between generations 16 and 17) arrives in V_M , followed by alveolar gas. During the remainder of the exhalation, this alveolar gas mixes rapidly with the gas already in V_M , causing the CO₂ partial pressure $p_M(t)$ to rise steadily.

B. Modeling Vcap

To analyze Vcap, we only require the gas-mixing subsystem model, since airflow $\dot{V}(t)$ is measured and available. We desire an expression for how $p_M(\cdot)$ varies with exhaled volume $V(t)$ rather than with t . We could write $p_M(t) = \tilde{p}_M(V(t))$ to denote this function of volume, since $V(t)$ is a monotonically increasing function of t during exhalation. However, to avoid proliferating notation, we shall simply denote it by $p_M(V)$, relying on context to resolve ambiguity. We shall similarly write $p_A(V)$ for the

partial pressure of the alveolar gas entering the region V_M after a volume V has been exhaled.

Gas-Mixing Subsystem Model for Vcap: The incremental change in CO₂ mass in the mixing volume, for an incremental exhaled volume dV , is proportional to $d(V_M p_M(V))$, by the definition of partial pressure. This must equal the incremental mass entering the mixing volume minus that leaving, so

$$d(V_M p_M(V)) = V_M dp_M(V) = (p_A(V) - p_M(V)) dV$$

or equivalently

$$\frac{d}{dV} p_M(V) = -\frac{1}{V_M} p_M(V) + \frac{p_A(V)}{V_M}. \quad (1)$$

This is a *linear constant-coefficient first-order* differential equation for the evolution of $p_M(V)$ with V during exhalation — and hence governing Vcap. Explicit solutions can be obtained under a variety of scenarios.

For the simplest case [18], [19], assume the CO₂ partial pressure in V_M is 0 mmHg till a volume V_L has been exhaled, and that $p_A(V)$ then steps instantaneously to some value \bar{p}_A for $V > V_L$, which may be thought of as the (spatial) average partial pressure of alveolar gas at the end of the preceding inhalation. The solution of (1) is then simply

$$p_M(V) = \bar{p}_A \left(1 - \exp\{-(V - V_L)/V_M\} \right), \quad V > V_L, \quad (2)$$

settling exponentially from 0 mmHg to \bar{p}_A , with settling constant V_M . This was shown in [18] for the special case of $V_L = 0$, but predicated on a particular model for airflow. We now see the solution is more general.

Our experiments have shown that (2) does not yield good fits to Vcap data, so we re-examine the assumptions of $p_M(V_L) = 0$ and a simple step increase of $p_A(V)$. Depending on how V_L is estimated (we present one choice in Section III), the actual $p_M(V_L)$ may be nonzero. Also, a smoother settling of $p_A(V)$ to a value \bar{p}_A may better (though still coarsely) capture inhomogeneity and varying alveolar ventilation and perfusion that cannot be directly represented in our single-“alveolus” model. We shall thus use a slightly more elaborate model for alveolar discharge into V_M , specifically

$$p_A(V) = p_M(V_L) + K \left(1 - \exp\{-(V - V_L)/\nu\} \right) \quad (3)$$

for $V > V_L$. Now $p_A(V)$ settles exponentially, with settling constant ν , to the asymptotic value $p_M(V_L) + K$. Our estimate of average alveolar partial pressure \bar{p}_A is thus

$$\bar{p}_A = p_M(V_L) + K. \quad (4)$$

The fundamental equation (1) has a simple explicit solution in this case too (see Supplementary Material), essentially the step response of two single-pole, linear, constant-coefficient subsystems in series, where one pole is at $-1/V_M$ and comes from (1), and the other at $-1/\nu$ comes from (3).

C. Modeling Tcap

An instrument that generates Vcap can also generate Tcap by ignoring the airflow information in Vcap and keeping only the timing information, and this is the case with our ventilator data.

Typically, however, Tcap is obtained for subjects or patients breathing normally, using sidestream capnography, and with no airflow measurement. In this setting there is insufficient information to estimate (or even recognize) V_L , because Tcap stays at 0 mmHg for essentially the entire time that the volume V_L is being exhaled, and for some time prior. Thus, the deadspace of consequence for Tcap is just V_M .

1) Gas-Mixing Subsystem Model for Tcap

For Tcap we convert the model (1) to evolve with time t rather than exhaled volume V . Simply multiplying both sides of (1) by $\dot{V}(t)(= dV(t)/dt)$ gives

$$\frac{d}{dt}p_M(t) = \dot{p}_M(t) = -\frac{\dot{V}(t)}{V_M}p_M(t) + \frac{\dot{V}(t)p_A(t)}{V_M}. \quad (5)$$

Taking $t = 0$ to be when the exhaled alveolar gas first enters V_M , i.e., when $V = V_L$, this equation holds for $t > 0$. This equation in related forms has appeared in other settings, see for instance [15], [16], [17] and literature cited in these, though only [17] uses it specifically for modeling Tcap, assuming airflow $\dot{V}(t)$ is known (V_{cap} is mentioned but not directly modeled). Using $w(t)$ to denote the V_M -normalized exhaled volume $V(t)/V_M$, we can rewrite (5) as

$$\dot{p}_M(t) = -\dot{w}(t)p_M(t) + \dot{w}(t)p_A(t), \quad t > 0. \quad (6)$$

For a given $\dot{w}(t) = \dot{V}(t)/V_M$, (6) is a *linear, time-variant, first-order* state-space equation [24] in the variable $p_M(t)$, with input $p_A(t)$ (explicit solution in Supplementary Material).

Equation (6) already suggests what Tcap can in principle accomplish even without airflow measurements: for known $p_M(t)$ and thus $\dot{p}_M(t)$, (6) is a (linear) constraint on $\dot{w}(t)$ and $p_A(t)$ at each t , so given simple parametric models for $\dot{w}(t)$ and $p_A(t)$, we could hope to estimate the associated parameters from measurements of $p_M(t)$. That expectation is developed and confirmed in what follows.

For a parametric model of $p_A(t)$, we assumed in our prior work [18], [19] that $p_A(t)$ is a step function, rising immediately at the start of exhalation from 0 mmHg to the value \bar{p}_A , so only this single parameter was needed to describe $p_A(t)$. In the results reported in this article, we use a 2-parameter alternative to the step function, to represent an exponentially settling discharge of alveolar gas into V_M , with time-constant ϵ , so

$$p_A(t) = \bar{p}_A \left(1 - \exp\{-t/\epsilon\} \right). \quad (7)$$

This is distinct from the assumption we used for $p_A(V)$ in (3), where the partial pressure settled exponentially as a function of volume, but (7) is better adapted to modeling Tcap.

We now turn to obtaining a parametric model for $\dot{w}(t)$.

2) Airflow Subsystem Model for Tcap

We represent the lungs as a single ‘‘alveolar’’ compartment with some inertance (or mass) and compliance, connected to a resistive airway that opens to atmospheric air or, for ventilated patients, to some other airway pressure imposed at the mouth, e.g., positive end-expiratory pressure (PEEP) [25], [26].




Fig. 3. Circuit analogs for the dynamics of airflow (left) and CO₂ exchange (right) during exhalation. The inverse of the current in the airflow circuit determines the resistance of the gas-mixing circuit.

This very simple airflow model, most often without even the inertance element, is widely used to model both exhalation and inhalation [27], [28], but [17], [18], [19] appear to be unusual in applying such a model to quantitatively analyzing and interpreting Tcap. More complex models of airflow have their uses, but we require robust identification of model parameters despite the limited richness of Tcap waveforms collected — without airflow — in clinical settings. We retain the inertance element to capture the fact that airflow does not change discontinuously, as is evident from Fig. 1(a).

Airflow during exhalation is primarily driven by the passive relaxation of the compliant lung compartment from some initial inflated volume V_I attained at the end of the preceding inhalation. We use a linear and time-invariant (LTI) model with inertance, resistance, and compliance taking constant values L , R , C respectively. The exhaled volume $V(t)$ then satisfies

$$L\ddot{V}(t) + R\dot{V}(t) + \frac{1}{C}V(t) = \frac{1}{C}V_I - P(t) \quad (8)$$

for $t > 0$, with $V(0) = V_L$ (given our new choice of time origin). Here $P(t)$ denotes the pressure gradient from airway to pleural cavity, $P_{\text{aw}}(t) - P_{\text{pl}}(t)$, during exhalation. For ventilated patients, the airway pressure may include a small steady PEEP applied during the later part of the exhalation, to guarantee a sufficient lung inflation at the end of exhalation. We might model this by setting $P_{\text{aw}}(t) = P_{\text{PEEP}}(t) + r\dot{V}(t)$, where $P_{\text{PEEP}}(t)$ rises from 0 to the desired PEEP level by the end of exhalation, and r denotes the flow resistance distal to the airway pressure measurement.

a) Electrical circuit analogs: It is helpful to represent our airflow and gas-mixing subsystems (8) and (5) via the natural electrical circuit analogs shown in Fig. 3, see also [19]. These are obtained by mapping pressure or partial pressure to electrical potential, and airflow or CO₂ flow to electrical current, and accordingly defining the other component elements, variables, and parameters. The airflow circuit is LTI but has a nonlinear coupling to the gas-mixing circuit, because the resistance in the gas-mixing circuit is the reciprocal of the current in the airflow circuit (higher airflow leads to more rapid mixing).

b) Normalized airflow model: Assuming $P(t)$ attains a constant value \bar{P} towards the end of the exhalation, and that both the airflow and its derivative correspondingly settle essentially to 0, the first two terms in (8) are negligible at the end of exhalation.

The total expelled volume, which is the tidal volume, is accordingly seen from (8) to be $V_T = V_I - C\bar{P}$ as expected (because it is the initial minus final lung volume on the exhalation). To make progress, we shall approximate $P(t)$ by the constant \bar{P} for all $t > 0$ during the exhalation (i.e., for $V(t) > V_L$). With this, we can rewrite (8) as

$$\delta\tau\ddot{w}(t) + \tau\dot{w}(t) + w(t) = \alpha, \quad t > 0, \quad (9)$$

$$\delta = L/R, \quad \tau = RC, \quad \alpha = V_T/V_M. \quad (10)$$

The parameter α is the *normalized tidal volume* (or *inverse mixing-deadspace ratio*). The canonical 2nd-order LTI model in (9) has an analytical solution for given initial conditions, which should be $w(0) = V_L/V_M$, with $\dot{w}(0)$ being the normalized airflow after the volume V_L has been expelled. For reasons explained in Section III (Methods), in modeling Tcap we shall only use the solution of (9) for *zero* initial conditions.

For $\delta \ll \tau$, which is typically the case in our data, δ turns out to be approximately the time constant that governs the rapid exponential settling of *airflow* to its peak near the beginning of exhalation, and also of *airway pressure* from its peak (see Supplementary Material). This initial phase is evident in both these waveforms in the example of Fig. 1(a), during the interval 2.1–2.2 s. The slower subsequent exponential decay of both airflow and airway pressure to the end of exhalation is, on the other hand, approximately governed by the time constant τ . (In airflow models with no inertance, so $L = 0$, $\delta = 0$, the modeled airflow rises instantly to its peak at the start of exhalation, then decays exponentially with time-constant τ .) Thus, fitting a sum of two exponentials to the airflow or airway-pressure waveform during exhalation allows estimation of δ and τ (see Supplementary Material related to Section IV-B) — but given only Tcap, one has to proceed differently.

If we consider Tcap as providing a measurement of $p_M(t)$, then (6) shows we are in a position to attempt identification of our parameterized models (9) for $\dot{w}(t)$ and (7) for $p_A(t)$. The combination of the gas-mixing model (6) and the airflow model (9) comprises a 3rd-order 3-parameter nonlinear time-invariant state-space model for Tcap, driven by the 2-parameter input in (7), with the state at time t determined by the state variables $p_M(t)$, $w(t)$ and $\dot{w}(t)$, [24]. From measurements of $p_M(t)$ over the course of an exhalation, we can hope to estimate the 5 parameters involved in the model (6), (7), and (9), namely δ , τ , α , \bar{p}_A , and ϵ . Importantly, having estimates of δ , τ , α would allow us to solve (9) for $w(t)$, and thereby estimate the normalized airflow. In Sections III and IV we show how these expectations are met.

D. New Tcap Tests for COPD–CHF Discrimination

Our earlier work [13], [18], [19] explored the use of capnography for COPD–CHF and COPD–Normal classification by using simple machine-learning approaches to discriminate between regions in a 4- or 5-dimensional parameter space that were characteristic of COPD and those characteristic of CHF. However, our observation above that a scaled version of the airflow profile

can be estimated from Tcap alone suggests a quite different approach to COPD–CHF discrimination.

Both COPD and CHF patients can present with dyspnea, and spirometry is widely used in pulmonary function testing to distinguish between them, and further to assess COPD severity [29]. It measures, among other variables, the *forced expiratory volume in one second* (FEV1) and *forced vital capacity* (FVC), which summarize key characteristics of the airflow waveform during forced exhalation. Examining whether the FEV1/FVC ratio is below a population-adjusted threshold serves as an initial confirmation of COPD and its severity. Thresholds around 0.7 are typically used, though this is adjusted based on demographic parameters and the severity level that is being screened for [30].

A natural *unforced analog* of FEV1/FVC would be $V(1)/V_T$ in our notation, i.e., the ratio u of the *unforced expiratory volume in one second* (UEV1) to the tidal volume (TV or V_T), which we are now able to estimate from Tcap alone:

$$u = w(1)/\alpha = V(1)/V_T = \text{UEV1}/\text{TV}. \quad (11)$$

Section IV-C presents exploratory results using this ratio for COPD–CHF discrimination. However, a fuller evaluation, adhering more closely to the STARD 2015 template [31] for evaluation of diagnostic accuracy, must be left to future work.

E. Modeling Forced Inhalations

In the case of ventilator data with measurements of both airflow and airway pressure, an additional window into airflow subsystem parameters is provided by the forced *inhalation* portion of the waveforms. Ventilators can operate in different modes, but the records of ventilated patients in the CapnoBase database [8], [9] dominantly involve *volume-controlled ventilation*, which maintains approximately *constant airflow* during the inhalation; this is clearly seen in the example in Fig. 1(a). All the records we study are of this type.

Our circuit model on the left of Fig. 3 and the corresponding equations require modification for the inhalation phase. As seen in Fig. 1(a), the initial airway pressure (when the airflow and its derivative are essentially 0) is the value \bar{P} across the compliance/capacitor element of our model at the end of exhalation. There is then a short transient before the airflow settles at an essentially constant negative value, $-F (< 0)$.

During intervals of constant airflow, the pressure gradient across the inertance is zero, and across the resistance is FR . The gradient across the compliance ramps up linearly, with slope F/C . If the pleural pressure does not vary significantly during this interval, we expect the linear pressure rise on the compliance to be directly reflected in a linear ramping up of airway pressure, and this is seen in the example in Fig. 1(a) during the interval of constant airflow. These features allow a direct and decoupled estimation of R and C from the available waveforms, as detailed in Section III (Methods). Our interest in estimating R and C in this article is primarily to compare the product RC estimated from inhalation data, with the value of τ estimated from Tcap alone — and separately from airflow or airway pressure — during exhalation.

III. METHODS

A. Data Sources

1) **Ventilator Data:** The ventilator data for this article are taken from the publicly accessible CapnoBase database [8], [9]. We used 17 recordings of 2 minutes each, of the sort shown in Fig. 1(a), recorded from 11 adult patients aged 24–74, under anesthesia (in volume-controlled mode, with essentially constant airflow during inhalation) during a variety of elective surgeries. Patients were on a GE SmartVent 7900 ventilator; a Datex-Ohmeda CAiOVX measurement module sampled the flow sidestream to measure CO_2 partial pressure, and corrected for delay and deformation in the sidestream tubing; the associated S/5 Collect software resampled all signals to 300 Hz. Amplitude resolution for pressure is $0.1 \text{ cmH}_2\text{O}$ (0.074 mmHg), for CO_2 partial pressure is 0.01 kPa (0.075 mmHg), for airflow is 0.1 L/min (1.67 mL/s). Exhalations preceded by unusually short ventilator-imposed inhalations, or with other unusual artifacts, were omitted, but 310 exhalations were retained.

2) **Tcap Data (Without Airflow Measurements):** To evaluate the novel Tcap-based test suggested in Section II-D for COPD–CHF discrimination, we revisit Tcap data from [13], [18], [19].

a) **Training database:** Tcap recordings (without airflow measurements) were collected, under IRB-approved protocols, from COPD and CHF patients in the pulmonary function laboratory of Brigham and Women’s Hospital, Boston, Massachusetts, during Jul–Dec 2010. Standard spirometry measurements (FEV1, FVC) confirmed each patient’s diagnosis and severity (moderate, severe). Enrolled subjects were seated, breathing normally, and connected through a standard nasal cannula to a capnograph (Capnostream 20, Oridion Medical, Needham, Massachusetts; flow rate 50 mL/min , sampling rate 50 Hz , amplitude resolution 1 mmHg on the PeCO_2 measurement) for a 3-minute recording. 24 COPD records (12 moderate, 12 severe, 28–73 years, 12 female) and 24 CHF records (19 moderate, 5 severe, 32–84 years, 9 female) constituted the training database.

b) **Test database:** Additional data were collected as above and with an identical device, again with IRB approval, from patients presenting to the Emergency Department of Einstein Medical Center, Philadelphia, Pennsylvania, with a chief complaint of moderate/severe dyspnea. Those meeting COPD or CHF eligibility criteria (detailed in [13]) were enrolled, and informed consent was obtained. The recordings of 10–30 min each yielded our *test set*, comprising 29 COPD records (collected Feb–Dec 2007) and 32 CHF records (collected May 2012–Jan 2013), not classified by severity. We also use a *trimmed* version of this test set (27 COPD/31 CHF) that retains only records with at least 30 validated exhalations (most have a few hundred, the longest has 617).

B. Fitting Vcap; Estimating V_L , V_M , ν , \bar{p}_A

To assess how well the gas-mixing model (1) with input (3) fits Vcap, we first require an estimate of the laminar-flow deadspace volume V_L . In all the results reported here, we estimate V_L using the construction illustrated by the dashed black




Fig. 4. The full black curve is Vcap from Fig. 1(b). The dashed black line is tangent at Vcap’s point of inflection o , and its intersection with the volume axis (marked $*$) yields the estimate of V_L , 176 mL here. The optimal fit of (1) with the input (3) to this Vcap, starting at $V = V_L$ and extending to $V_T - 50 \text{ mL}$, is shown by the red curve, with associated estimate $V_M = 54.3 \text{ mL}$. The dashed blue curve is the estimated $p_A(V)$ given by (3), with $\bar{p}_A = 26.1 \text{ mmHg}$.

line in Fig. 4 for the case of the particular volumetric capnogram in Fig. 1(b). Specifically, a close 10th-order polynomial least-mean-square-error fit (not shown in Fig. 4) is constructed for the entire Vcap waveform from $V = 0$ to $V = V_T$, using the `polyfit` routine in MATLAB (MathWorks, Natick, Massachusetts). The point of inflection of the polynomial (marked o in Fig. 4) is where its slope is maximum. The intersection of the tangent line at this point with the volume axis is taken as our estimate of V_L (marked $*$ in the figure, 176 mL in this case).

Having an estimate of V_L , we use MATLAB’s `lsqcurvefit` routine to determine values of the parameters V_M , ν , and K that lead to the nonlinear least-mean-square-error fit of $p_M(V)$ to the Vcap data for $V_L < V < V_T - 50 \text{ mL}$, where $p_M(V)$ is given by the analytical solution of the model equations (1), (3) (see Supplementary Material). (We do not attempt to fit the last 50 mL of the Vcap record, as the data tend to be erratic and noisy at the end of the exhalation.) The algorithm is initialized with estimates $K = 30 \text{ mmHg}$, $V_M = 0.5V_T$ and $\nu = 0$ in every case. The particular structure of this model, as the cascade of two single-pole, linear, constant-coefficient systems with poles at $-1/V_M$ and $-1/\nu$, causes the analytical solution, and therefore the fit of the model to the data, to be unchanged if the values of V_M and ν are interchanged. Since ν was introduced to relax our prior assumption of a step increase in $p_A(V)$, i.e., to relax the assumption that $\nu = 0$, we pick the smaller of the two converged values to be our estimate of ν .

The red curve in Fig. 4 shows the resulting fit, and it appears good visually; the root-mean-square-error (rmse) of the fit is 0.18 mmHg over the span of the fit. This number can be used as a benchmark when interpreting the results presented in Section IV for the performance of the model on all the other exhalations in our set of ventilatory recordings. The dashed blue curve in Fig. 4 is the estimated $p_A(V)$ given by (3), settling exponentially to the value $\bar{p}_A = 26.1 \text{ mmHg}$.

C. Fitting Tcap; Estimating $\delta, \tau, \alpha, \epsilon, \bar{p}_A$ (And V_M)

To fit Tcap waveforms, we combine the normalized gas-mixing subsystem model (6), (7) with the normalized airflow subsystem model (9). Since the combination is nonlinear and without an analytical solution in general, the process of fitting involves embedding numerical simulation of the combined 3rd-order model within a loop in MATLAB's `lsqcurvefit` to determine the values of $\delta, \tau, \alpha, \epsilon, \bar{p}_A$ that yield a least-mean-square-error fit (using the respective starting values 0.01 sec, 1.3 sec, 3.5, 0.01 sec, 30 mmHg for all exhalations).

Since airflow is measured in the case of our ventilator data, we could actually determine when exhalation begins — at flow reversal from negative to positive — and also the laminar-flow volume V_L that has to be expelled before CO_2 partial pressure in the exhaled breath starts to rise and register on Tcap. However, because we are interested in comparing with what can be determined from Tcap alone, without airflow or airway-pressure measurements, we choose the time origin within an exhalation to be when PeCO_2 first crosses up through a low threshold (and subsequently rises for the rest of the exhalation). For Tcap obtained from the ventilator data, we have chosen this threshold to be 0.5 mmHg above the baseline attained at the end of the preceding inhalation; for our remaining (lower-resolution) Tcap data, we choose it to be 1 mmHg above baseline. The values of the normalized volume $w(t)$ and normalized flow $\dot{w}(t)$ are zero at the start of exhalation, but are typically *nonzero* at this threshold crossing (because by then a volume V_L has been exhaled). Though we could determine these initial conditions in the case of ventilator data, they are not available when only Tcap is known. Accordingly, the results we report in Section IV are obtained by assuming all initial conditions are *zero*, and using the solution of (9) (given in the Supplementary Material). An example of the fit of our model to Tcap is shown in the upper part of Fig. 5, for the same exhalation as in Figs. 1 and 4. The corresponding parameter estimates are $\delta = 0.10$ s $\ll \tau = 0.82$ s, $\alpha = 4.37$, $\epsilon = 0.086$ s, $\bar{p}_A = 26.8$ mmHg. The capnogram synthesized by the model for these parameter values fits the measured Tcap well, with an rmse of 0.3 mmHg.

D. Estimating the Exhaled Airflow Profile From Tcap

Having estimates of δ, τ , and α allows solution of (9) with zero initial conditions (see expressions in the Supplementary Material) to estimate the normalized airflow waveform $\dot{w}(t) = \dot{V}(t)/V_M$ during exhalation. Some discrepancy is to be expected because the actual initial conditions (after the laminar flow volume V_L has been exhaled) are nonzero.

If a measurement of the full exhaled airflow waveform is available for comparison (as with our ventilator data), then V_M can be estimated as the multiplicative factor that best aligns the normalized and measured airflow waveforms (adjusting as needed for the difference in time origins). This is illustrated by the lower set of traces in Fig. 5, which shows the estimated flow for the same exhalation we have been using as a running example, plotted with the actual measured flow (after time-shifting to adjust for the different choices of time-origin between Vcap and Tcap, and scaling for a least-mean-square-error fit over the

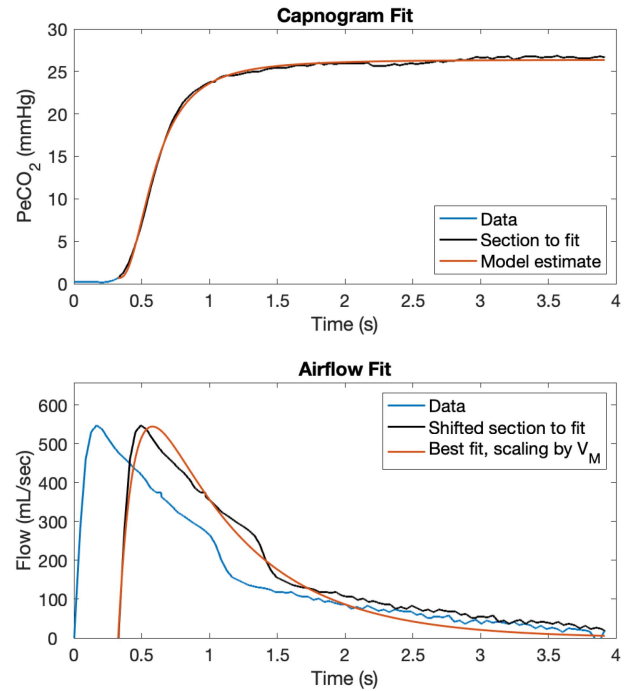


Fig. 5. The upper figure shows the fit (rmse 0.3 mmHg) of our model to Tcap for the same exhalation used in Figs. 1 and 4, $\delta = 0.10$ s, $\tau = 0.82$ s, $\alpha = 4.37$, $\epsilon = 0.086$ s, $\bar{p}_A = 26.8$ mmHg. The lower figure shows the least-mean-square-error match (rmse 30 mL/s) of the scaled normalized airflow waveform to the actual airflow waveform (adjusted for a common start time), yielding the estimate $V_M = 126$ mL as the scale factor. A further estimate of V_M is $V_T/\alpha = 139$ mL.

duration of the exhalation). The alignment of the flow profiles appears good, and the corresponding rmse is 30 mL/s. (Note also that the estimated δ and τ values we just quoted, Section III-C, are consistent with their being the approximate time constants governing the rise and fall of airflow.) Our results in Section IV establish that this performance is typical.

The estimate of V_M can be obtained more easily from the relation $V_M = V_T/\alpha$, provided the *tidal volume* V_T is known; V_T is easier to measure than the full airflow profile. This estimate of V_M should closely match that obtained by scaling the estimated normalized flow to best fit the actual flow, because these two flows are ideally related by the same scale factor as the one that relates the *integrals* of these flows, namely the normalized and actual *tidal volumes*.

E. Fitting Forced Inhalations; Estimating R and C

We have selected our records solely from patients whose ventilators operate in *constant-flow* mode during forced inhalation. As noted in Section II, the airflow model simplifies in any interval during which the airflow is constant at some value $-F < 0$. If pleural pressure does not vary much in such an interval, we expect a linear ramping up of airway pressure with slope F/C in the interval, allowing C to be computed from a least-mean-square-error linear fit to the ramp. The total insufflated volume V_{ins} during the transient, before constant airflow sets in, can also be computed; dividing this by the estimated C determines the pressure rise in the compliance above

its starting value, before its linear rise. (See Supplementary Material for figures embodying the preceding computations.) The total gradient across the compliance before it begins its linear rise is thus $\bar{P} + (V_{\text{ins}}/C)$. Subtracting this from airway pressure at the beginning of the linear rise gives the gradient FR across the resistance, which therefore allows R to be estimated.

The product RC , if estimated during an exhalation, would be expected to correspond to the time-constant τ that governs the decay of airflow during exhalation. However, since lung characteristics can and do differ between inhalation and exhalation, the RC product estimated during inhalation may only correspond approximately to τ . Similarly, to the extent that the R estimated from inhalation data applies to exhalation, we might obtain an approximate estimate of the inertance using the relation $L = \delta R$.

IV. RESULTS

In the first part of this section we present a summary of the results from our validation studies for the Vcap and Tcap models in Section II, using the ventilator data and the approaches outlined in Section III. The last part of this section presents the results of COPD–CHF discrimination using our Tcap-based analog of a standard spirometry test.

A. Fitting Vcap, Tcap and Airflow Models to Ventilator Data

The fits of our models to Vcap, Tcap, and airflow waveforms in the ventilator data can be assessed via the associated root-mean-square-error (rmse) values.

a) Vcap: The example in Fig. 4, with rmse 0.18 mmHg, provides a benchmark for assessing the fits on other exhalations. The rmse across all 310 exhalations in our records spanned 0.10–0.70 mmHg (and 0.4–2.2% of end-tidal partial pressure), with mean rmse 0.36 mmHg and standard deviation 0.14 mmHg. The largest within-record standard deviation for rmse was 0.07 mmHg, across the 17 records. Note that these various rmse values are obtained with instrumentation that has a 0.075 mmHg quantization interval for CO₂ measurement, and for exhalations that have end-tidal CO₂ values in the range of 23.7 mmHg to 45.6 mmHg.

b) Tcap: For Tcap, a benchmark is in the upper set of traces in Fig. 5, where the rmse was 0.3 mmHg. The rmse across all 310 exhalations spanned 0.14–1.11 mmHg (and 0.5–3.5% of end-tidal partial pressure), with mean rmse 0.57 mmHg and standard deviation 0.25 mmHg. The largest within-record standard deviation for rmse was 0.12 mmHg.

c) Airflow: A benchmark for the airflow estimates obtained by scaling Tcap-derived normalized airflows for least-mean-square-error fit to the corresponding measured airflows is in the lower set of traces in Fig. 5, where the rmse was 30 mL/s. The rmse of the airflow fits across all 310 exhalations ranged from 14 mL/s to 79 mL/s, with mean rmse 42 mL/s and standard deviation 13 mL/s. The largest within-record standard deviation for rmse was 10 mL/s.

B. Parameter Estimates From Ventilator Data

The key parameters for our model are:




Fig. 6. Cross-record comparison of estimates for ν , V_L , and V_M estimated directly from Vcap, and V_M estimated in two ways from Tcap: (i) using the relation $V_M = V_T/\alpha$, and (ii) by scaling the estimated airflow waveform $\dot{v}(t) = \dot{V}(t)/V_M$ to match the measured airflow. The symbol in each case indicates the mean value across all exhalations in the corresponding record, and the line segment marks one standard deviation on either side. Quantities in red are derived from Tcap alone.

- for Vcap, the laminar-flow deadspace volume V_L , the mixing deadspace volume V_M , the parameter ν that governs the discharge of alveolar gas into V_M , and the mean alveolar gas partial pressure \bar{p}_A ;
- for Tcap, the time constants δ (governing the fast initial transient on exhalation) and τ (governing the slower subsequent settling), also α (the V_M -normalized tidal volume V_T), ϵ to describe alveolar discharge, and \bar{p}_A — and we also showed that matching the Tcap-derived normalized airflow and tidal volume estimates to the actual values could yield two additional estimates of V_M ;
- from the inhalation data, R and C .

We present cross-record comparisons here for many of these parameters. For other parameters we simply indicate the range of mean values obtained for the estimates across all records. To avoid clutter, and because they are not as interesting or novel, we omit presenting the estimates of \bar{p}_A .

Figure 6 shows ν , V_M , and V_L values estimated from Vcap as specified in Section III-B. Each row is labeled by the patient record identifier in Capnabase [8]. Also shown are the values of V_M estimated as the scale factors that cause the Tcap-derived estimates of normalized airflow and normalized tidal volume to respectively match the measured $V(t)$ (with least-mean-square-error) and tidal volume V_T (exactly).

Figure 7 similarly shows the cross-record comparison of estimates for δ and τ obtained from Tcap, along with the small and large time constants, δ' and τ' respectively, estimated through least-mean-square-error fit of the sum of two exponentials to the airway-pressure waveform during exhalation (see Supplementary Material). The estimated ϵ values are in the same range as the estimated δ 's, but omitted to avoid clutter.

Also incorporated for comparison in Fig. 7 is the product $RC = \tau''$ constructed from the estimates of R and C obtained from the airflow and airway-pressure waveforms during the




Fig. 7. Cross-record comparison of estimates for: δ and τ obtained from Tcap; small and large time constants δ' and τ' from least-mean-square-error fit of sum of exponentials to airway-pressure waveform during exhalation; and $RC = \tau''$ from estimates of R and C using airflow and airway-pressure waveforms during forced inhalation. The mean value across all exhalations and the corresponding standard deviation on either side are shown. Quantities in red are from Tcap alone.

forced inhalations. We omit a detailed presentation of the estimates obtained for R and C individually, except to note that the mean value of R estimated for each of the 17 records ranged from 3.45 to 18.74 mmHg \cdot L $^{-1}$ s $^{-1}$, and the mean value of C ranged from 49.0 to 101.1 mL \cdot mmHg $^{-1}$. These values are consistent with values in the literature (for instance [32]).

Finally, the mean value of α (the V_M -normalized tidal volume) ranged from 2.61 to 4.59 across the 17 records.

C. COPD–CHF Discrimination Using Tcap Airflow Estimate

Evaluation of our Tcap-based COPD–CHF classifier using the UEVI/TV ratio $u = w(T)/w(1)$ is done on the training and test data from standard sidestream capnometers (with no airflow or air pressure measurements) described in Section III-A (2). We use the same models and procedures as above (but with $\epsilon = 0$, so only 4 parameters) to fit each of the (*thousands* of) exhalations in the 109 records. The sample $w(t)$ plots in Fig. 8(a), chosen from exhalations in our test set, suggest that the ratio u might indeed differ systematically between COPD and CHF patients. This is further borne out by the histograms of u in Figs. 8(b) and (c), determined from the first 15 validated exhalations in each COPD and CHF record from the training and test sets. This suggests a simple threshold test on u , which we examine next. (A fuller evaluation more closely meeting STARD 2015 criteria [31], and comparison to other simple tests — such as a threshold test on τ , motivated by results in [19] — is left to future work.)

We used the first 15 exhalations from each record in the training database as our training set. A threshold of 0.8 on u maximized the proportion of correct decisions on the exhalations in this set, where u below threshold are deemed to have come from COPD records, and values above from CHF. We then carried out two COPD–CHF classification experiments:

- 1) Choosing 15 exhalations randomly from each record in the *trimmed* test set, Section III-A (2), we declared a record to be from a COPD patient if the threshold test declared more than half (≥ 8) of these exhalations to be COPD — and repeated this 100 times. The accuracies of COPD–CHF classification in these experiments were in the range 74.1%–86.2%, with a mean of 80.6%, stdev 2.1% (see Supplementary Material for a histogram of obtained accuracies).
- 2) For each record in the full test set, we declared the record to be COPD if more than a fraction f of *all* the exhalations were decided as COPD, varying f between 0 and 1 to generate a “receiver operating characteristic” (ROC) (the test in (1) corresponds to $f = 0.5$). This ROC (see Supplementary Material) has area under it (AUROC) of 0.84. The corresponding “equal-error-rate” (EER) value, where specificity = sensitivity [24], is 77% for our test set, with $f = 0.2$. The EER gives a rough lower bound for obtainable accuracy as the relative prevalences in the test set are varied.

V. DISCUSSION

Our results in Section IV have established that Vcap and Tcap waveforms can be modeled quantitatively and with good fidelity using low-order mechanistic models that reflect the underlying respiratory anatomy and the physiological mechanisms, though in a necessarily aggregated manner. Some reinforcement of the postulated anatomical basis of the model comes from the fact that the values of V_M and V_L estimated from Vcap are consistent with the nominal values of 53 mL and 115 mL respectively, derived from data quoted in [21], and indicated in Fig. 2. It should be kept in mind, however, that some part of the mixing volume V_M may be in the tubing that leads to the sensor, and hence capnometer-dependent.

The listed rmse values for the Vcap, Tcap and airflow fits to the ventilator data show that the particular set of waveforms selected as a running example to support the development in the paper falls in the best quartile of our data, from the point of view of model fits. However, even the poorest fits are not significantly worse.

We have previously demonstrated good fits to Tcap using similar models, [18], [19], but assuming a step increase in $p_A(t)$ at the start of exhalation, rather than an exponential settling to \bar{p}_A with time constant ϵ as in (7), and only [19] incorporated the inertance L (or equivalently δ) in the airflow model. Our models here thus involved one or two more parameters, permitting closer fits, and these parameters were robustly and consistently identifiable as well. (See Supplementary Material for comparison of fits to Tcap and airflow in Fig. 1, using different models.) These additional parameters, and similarly ν in the case of Vcap, are also potentially of clinical interest.

The 17 ventilator records analyzed here have allowed us to compare, on the same data, the performance of our Tcap and Vcap models, as well as models for the inspiratory phase on each breath. Although the individual parameter estimates in Figs. 6 and 7 vary across the records, each standard deviation

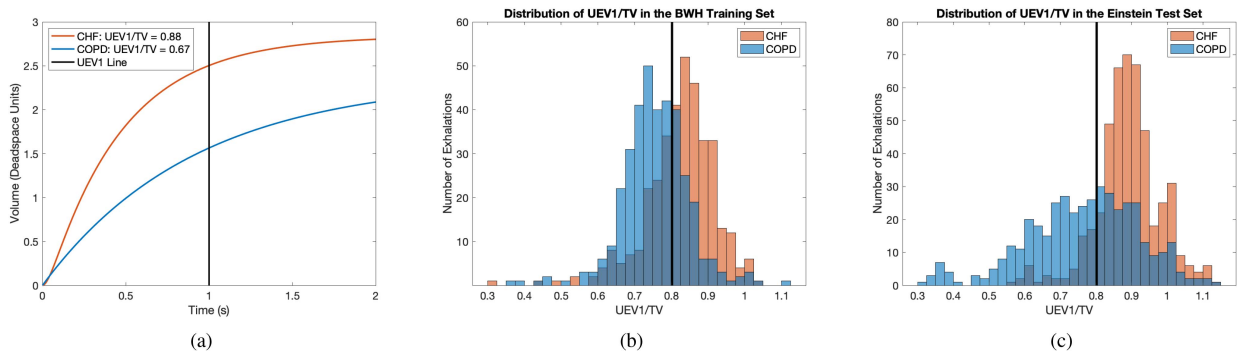


Fig. 8. (a) Plots of normalized airflow $w(t)$ on a representative exhalation from a CHF patient (red) and a COPD patient (blue) from test set. (b), (c) Histograms of $u = w(1)/\alpha$ for first 15 exhalations on each record in training and test databases respectively.

within a record is relatively small, suggesting the parameters are characteristic of the patient’s current state. Also, when there are two records for a given patient, the two sets of estimates are quite close, suggesting the parameters in our models can be robustly identified in a patient-specific way from capnography alone. The estimated parameters are aligned with values in the literature, consonant between Tcap and Vcap on exhalation (particularly in the case of time-constants associated with the airway model), and with estimates extracted from the recordings during inhalation. Our results suggest one can have confidence in the use of these models and the associated parameter estimates extracted from Tcap and Vcap.

Models of greater fidelity to the anatomy and physiology can be constructed, and enhance understanding, e.g., in simulation studies [15], but robust identification from the available data becomes progressively harder. As our models closely fit the measured capnograms, there is not much information in the residual error to support additional modeling refinements. Though the parameters in our models represent aggregate effects, we anticipate that future work will discover useful associations of particular parameter ranges, combinations, or clusters with particular physiological states and clinical pathologies.

For Tcap, specific numerical results on the accuracy of our model fits, and the robustness and consistency of the resulting parameter estimates, were presented here only for the ventilator data. The accuracy and robustness of the Tcap fits for the thousands of exhalations in all the other Tcap data (used for the COPD–CHF study) were comparable.

Our models revealed, somewhat surprisingly, that the entire airflow profile can be determined to within a scale factor V_M using Tcap alone, see Section III-D. This accounts for why the time constants of airflow and airway pressure are well estimated using Tcap alone, see Fig. 7. The ventilator data have allowed us to verify that the scaled fit is good, see Fig. 5, despite having to assume zero initial conditions for (9) in order to generate a normalized airflow profile using Tcap alone. However, the scale factor V_M estimated by combining Tcap with a measurement of airflow or tidal volume V_T ended up larger (often 2–3 times larger) than the value estimated from Vcap, perhaps partly because of the assumption of zero initial conditions. Also, because of how we have defined the time origin for Tcap, the exhaled volume at this time is systematically lower than the V_L estimated as in Fig. 4 (by around 29 mL in the example in Fig. 1), perhaps

leaving this gap to be captured in Tcap’s estimate of V_M . These issues bear further investigation.

A related challenge is to more closely model and fit the portion of Vcap leading up to the volume we have chosen as our estimate of V_L . Our construction for estimating V_L in Fig. 4 — while plausible, and in the style of other constructions in the capnography literature — is somewhat *ad hoc*.

The ability to estimate a scaled airflow profile from Tcap alone motivated our proposal for a Tcap-based analog of a standard spirometry test for COPD–CHF discrimination, which we have evaluated here with encouraging results. The test may be more directly interpretable in a clinical setting than the discriminant tests described in our earlier work [13], [18], [19], which involved partitioning the multi-dimensional parameter space of our models into parameter clusters associated with COPD versus those associated with CHF. Those tests, using essentially the same database as in this article, have obtained test-set accuracies in the range of 70–82%. It would be interesting, following up on observations in [19], to explore how a simple threshold test on the estimated exhalation time-constant τ compares.

Even restricting ourselves to COPD–CHF discrimination tests based on estimated scaled airflow profile, the particular threshold test on u is only one of several possibilities, and others might be explored in future work. A more detailed view of pulmonary function is often obtained in spirometry by plotting airflow $\dot{V}(t)$ versus volume $V(t)$. An analog to this plot for the expiratory phase, and obtainable from Tcap measurements alone, is a plot of $\dot{w}(t) = \dot{V}(t)/V_M$ versus $w(t) = V(t)/V_M$, and would be interesting to explore.

It will also be worthwhile in future work to examine the relationship between the deadspace volumes V_L , V_M used in our models and those in more classical descriptions, for example the relation of $V_L + V_M$ to classical estimates of anatomical deadspace, [4], [5], [6], [7].

VI. CONCLUSION

Our low-order mechanistic models for capnography reflect the underlying lung anatomy and physiology, though in a necessarily aggregated manner, using just 4 or 5 parameters. They closely and robustly characterize — in quantitative detail — the capnogram traces in Vcap and Tcap. The associated parameters

are thus potentially of clinical value in cardiorespiratory screening, diagnosis, and monitoring, beyond the specific example of COPD–CHF discrimination illustrated in this article and in our earlier work.

We have shown that the full exhaled airflow profile can be inferred from Tcap alone, to within a scale factor. Related to this is the fact that time-constants associated with airflow and airway pressure can be estimated well from Tcap.

ACKNOWLEDGMENT

Partial support has come from Philips Healthcare, and we are particularly grateful to Joseph Frassica and Francesco Vicario for their interest and encouragement, and to the latter for helpful feedback as this article took shape. Comments from reviewers were instrumental in improving the paper. EKM and CXY received support from the MIT EECS Department’s SuperUROP funds. We are indebted to the CapnoBase team [9] for its generosity in creating and disseminating this valuable database, and specifically to Walter Karlen and Mark Ansermino for helpful details. Our work has built on the efforts of prior students in the Computational Physiology and Clinical Inference Group at MIT’s Research Laboratory of Electronics, specifically Rebecca Mieloszyk, Abubakar Abid, Ekin Karasan, and Rhian Chavez.

REFERENCES

- [1] M. B. Jaffe, “Infrared measurement of carbon dioxide in the human breath: ‘Breathe-through’ devices from Tyndall to the present day,” *Anesth. Analg.*, vol. 107, no. 3, pp. 890–904, Sep. 2008.
- [2] J. S. Gravenstein et al., Eds., *Capnography*, 2nd ed., Cambridge, U.K.: Cambridge Univ. Press, 2011.
- [3] C. Bohr, “Ueber die lungenathmung,” *Skandinavisches Archiv Für Physiologie*, vol. 2, no. 1, pp. 236–268, 1891.
- [4] W. S. Fowler, “Lung function studies II: The respiratory dead space,” *Amer. J. Physiol.*, vol. 154, no. 3, pp. 405–416, Sep. 1948.
- [5] R. Fletcher et al., “The concept of deadspace with special reference to the single breath test for carbon dioxide,” *Brit. J. Anaesth.*, vol. 53, no. 1, pp. 77–88, Jan. 1981.
- [6] S. Verscheure et al., “Volumetric capnography: Lessons from the past and current clinical applications,” *Crit. Care*, vol. 20, no. 1, Jun. 2016, Art. no. 184.
- [7] G. Tusman and S. Bohm, “Clinical monitoring by volumetric capnography,” in *Cardiopulmonary Monitoring*, S. Magder et al., Eds., Berlin, Germany: Springer, 2021, pp. 601–617.
- [8] W. Karlen, “CapnoBase In Vivo dataset,” 2021, doi: [10.5683/SP2/7JTWZDZ](https://doi.org/10.5683/SP2/7JTWZDZ).
- [9] W. Karlen et al., “CapnoBase: Signal database and tools to collect, share and annotate respiratory signals,” in *Proc. Annu. Meeting Soc. Technol. Anesth.*, 2010, Art. no. 27.
- [10] G. Tusman et al., “Model fitting of volumetric capnograms improves calculations of airway dead space and slope of phase III,” *J. Clin. Monit. Comput.*, vol. 23, no. 4, pp. 197–206, Jun. 2009.
- [11] M. B. Jaffe, “Using the features of the time and volumetric capnogram for classification and prediction,” *J. Clin. Monit. Comput.*, vol. 31, no. 1, pp. 19–41, Feb. 2017.
- [12] B. You et al., “Expiratory capnography in asthma: Evaluation of various shape indices,” *Eur. Respir. J.*, vol. 7, no. 2, pp. 318–323, Feb. 1994.
- [13] R. J. Mieloszyk et al., “Automated quantitative analysis of capnogram shape for COPD-Normal and COPD-CHF classification,” *IEEE Trans. Biomed. Eng.*, vol. 61, no. 12, pp. 2882–2890, Dec. 2014.
- [14] D. Bhagya and S. Manikandan, “Speed of sound-based capnographic sensor with second-generation CNN for automated classification of cardiorespiratory abnormalities,” *IEEE Sensors J.*, vol. 19, no. 19, pp. 8887–8894, Oct. 2019.
- [15] H. Benallal and T. Busso, “Analysis of end-tidal and arterial PCO₂ gradients using a breathing model,” *Eur. J. Appl. Physiol.*, vol. 83, no. 4, pp. 402–408, Nov. 2000.
- [16] C. E. W. Hahn and A. D. Farmery, “Gas exchange modelling: No more gills, please,” *Brit. J. Anaesth.*, vol. 91, pp. 2–15, 2003.
- [17] J. O. D. Buijs et al., “Bayesian tracking of a nonlinear model of the capnogram,” in *Proc. IEEE Eng. Med. Biol. Conf.*, 2006, pp. 2871–2874.
- [18] A. Abid et al., “Model-based estimation of respiratory parameters from capnography, with application to diagnosing obstructive lung disease,” *IEEE Trans. Biomed. Eng.*, vol. 64, no. 12, pp. 2957–2967, Dec. 2017.
- [19] E. Karasan et al., “An enhanced mechanistic model for capnography, with application to CHF–COPD discrimination,” in *Proc. IEEE Eng. Med. Biol. Conf.*, 2018, pp. 5267–5272.
- [20] E. R. Weibel, *Morphometry of the Human Lung*. Berlin, Germany: Springer, 1963.
- [21] C. Hsia et al., “Lung structure and the intrinsic challenges of gas exchange,” *Comprehensive Physiol.*, vol. 6, pp. 827–895, 2016.
- [22] C. Kleinstreuer and Z. Zhang, “Airflow and particle transport in the human respiratory system,” *Annu. Rev. Fluid Mech.*, vol. 42, no. 1, pp. 301–334, 2010.
- [23] T. Gemci et al., “Computational model of airflow in upper 17 generations of human respiratory tract,” *J. Biomech.*, vol. 41, no. 9, pp. 2047–2054, Jan. 2008.
- [24] A. V. Oppenheim and G. C. Verghese, *Signals, Systems and Inference*. Englewood Cliffs, NJ, USA: Prentice-Hall, 2017.
- [25] L. Ball et al., “Respiratory mechanics during general anaesthesia,” *Ann. Transl. Med.*, vol. 6, pp. 379–384, 2018.
- [26] C. Young et al., “Lung-protective ventilation for the surgical patient: International expert panel-based consensus recommendations,” *Brit. J. Anaesth.*, vol. 123, pp. 898–913, 2019.
- [27] T. Woo et al., “A comparison of various respiratory system models based on parameter estimates from impulse oscillometry data,” in *Proc. IEEE Eng. Med. Biol. Conf.*, 2004, pp. 3828–3831.
- [28] F. Vicario et al., “Noninvasive estimation of respiratory mechanics in spontaneously breathing ventilated patients: A constrained optimization approach,” *IEEE Trans. Biomed. Eng.*, vol. 63, no. 4, pp. 775–787, Apr. 2016.
- [29] B. L. Graham et al., “Standardization of spirometry,” *Amer. J. Respir. Crit. Care Med.*, vol. 200, pp. e70–e88, 2019.
- [30] K. Torén et al., “Vital capacity and COPD: The Swedish CardioPulmonary bioImage study (SCAPIS),” *Int. J. Chronic Obstructive Pulmonary Dis.*, vol. 11, pp. 927–933, 2016.
- [31] P. M. Bossuyt et al., “STARD 2015: An updated list of essential items for reporting diagnostic accuracy studies,” *BMJ*, vol. 351, 2015, Art. no. h5527.
- [32] D. R. Hess, “Respiratory mechanics in mechanically ventilated patients,” *Respir. Care*, vol. 59, no. 11, pp. 1773–1794, Nov. 2014.