Expiratory Flow and Volume Estimation through Thermal-CO_2 Imaging

Shane Transue, *Member, IEEE*, Se Dong Min, *Member, IEEE* and Min-Hyung Choi, *Member, IEEE*

***Abstract*— Objective: In this work, we introduce a quantitative non-contact respiratory evaluation method for fine-grain exhale flow and volume estimation through Thermal-CO_2 imaging. This provides a form of respiratory analysis that is driven by visual analytics of exhale behaviors, creating quantitative metrics for exhale flow and volume modeled as open-air turbulent flows. This approach introduces a novel form of effort-independent pulmonary evaluation enabling behavioral analysis of natural exhale behaviors. Methods: CO_2 filtered infrared visualizations of exhale behaviors are used to obtain breathing rate, volumetric flow estimations (L/s) and per-exhale volume (L) estimations. We conduct experiments validating visual flow analysis to formulate two behavioral Long-Short-Term-Memory (LSTM) estimation models generated from visualized exhale flows targeting per-subject and cross-subject training datasets. Results: Experimental model data generated for training on our per-individual recurrent estimation model provide an overall flow correlation estimate correlation of $R^2 = 0.912$ and volume in-the-wild accuracy of 75.65-94.44%. Our cross-patient model extends generality to unseen exhale behaviors, obtaining an overall correlation of $R^2 = 0.804$ and in-the-wild volume accuracy of 62.32-94.22%. Conclusion: This method provides non-contact flow and volume estimation through filtered CO_2 imaging, enabling effort-independent analysis of natural breathing behaviors. Significance: Effort-independent evaluation of exhale flow and volume broadens capabilities in pulmonological assessment and long-term non-contact respiratory analysis.**

***Index Terms*— Exhale visualization, CO2 exhale flow, Non-contact exhale analysis, Quantitative exhale analysis**

I. Introduction

IDENTIFYING abnormalities in continuous respiratory waveforms is one of the most prevalent diagnostic tools used to identify and evaluate pulmonary impairments ranging from acute respiratory failure [1] to the onset of chronic conditions [2]. Accurate quantification of these pulmonary indicators enables condition trajectory analysis, preventative solutions for in-hospital catastrophic events, and interventions for severe declines in pulmonary function. Metrics including Respiratory Rate (RR or BPM) and Pulmonary Function Test (PFT) volumetric flow measures (VC, FVC, ERV, TV, etc.) are used to provide feature descriptors of breathing behaviors that signify unique signatures of chronic conditions such as asthma, apnea in Sleep Disordered Breathing (SDB), and Chronic Obstructive Pulmonary Disease (COPD) [2]. These indicators are also critical for both the detection and prevention of severe respiratory conditions encountered in Intensive Care Units (ICUs) including cardiac and respiratory arrest, Pulmonary edema, Systematic Inflammatory Response Syndrome (SIRS), and Acute Respiratory Distress Syndrome (ARDS). To broaden the form of respiratory analysis used to monitor both acute and chronic conditions, we introduce a new form of non-contact, *effort-independent* evaluation that enables quantitative analysis (unit measures) of *natural* breathing through predictive CO_2 exhale flow waveform modeling.

An ongoing challenge in the improvement of abnormal behaviors that contribute to chronic and acute condition diagnosis is the quantification of exhale flow in *natural breathing*. The problem is that most respiratory monitoring methods record breathing that is *effort-dependent* or is subject to *device interference*. These factors incorporate how physical effort and usability alter how metrics are obtained. This incurs subtle differences in how the pulmonary behavior is quantified, leading metrics to exhibit secondary characteristics caused by the measurement itself. This may include limiting body movement, breathing through a tube, and strain from continuous unnatural breathing. This introduces implications for most existing methods in that: (1) they are not well adapted for capturing natural breathing and (2) are limited to short-term evaluation which is better suited for momentary function and severe symptom identification. Effort also interferes with reproducibility of specific behaviors that may indicate early traits of intermittent conditions due to symptoms that may not be exhibited during short monitoring durations. This limits the ability to detect condition progression and early intervention opportunities for broadening pulmonary pathology diagnostics.

To provide an effort-independent respiratory evaluation method, we present a novel vision-based approach for directly evaluating exhale flow and volume through *visual flow analytics*. The principle of our approach is to quantify exhaled CO_2 through thermography paired with spectral band-pass filtering (3-5μm) and exhale tracking to evaluate volumetric flow associated with pulmonological function [3] through Long-Short-Term-Memory (LSTM) neural network modeling

Manuscript received 23 February 2022; revised 21 July 2022; accepted 2 January 2023. This work is partially supported by NSF Partnerships for Innovation (PFI-TT) Grant: 1941221, BK21 FOUR (Fostering Outstanding Universities for Research): 5199990914048, Soonchunhyang University Research Fund, and Children's Hospital Colorado. (Corresponding authors: Min-Hyung Choi; Se Dong Min.)

Shane Transue and Min-Hyung Choi are with the University of Colorado Denver, CG and VR Lab, Denver, Colorado, USA (e-mails: {shane.transue, min.choi}@ucdenver.edu). Se Dong Min is with the Department of Medical IT Eng. and Software Convergence, Soonchunhyang University, Asan, South Korea (e-mail: sedongmin@sch.ac.kr).

[4]. In this approach, the aim is to measure and quantify visualized exhale CO_2 density and flow behaviors as both detailed flow fields and quantitative metrics compatible with standard PFTs. We build on the synergy between thermal CO_2 exhale visualization and apparent visual flow behavioral analysis to generate an accurate quantitative evaluation of breathing rate (BPM), flow rate (L/s), and exhale volume (L). Based on this, the proposed solution provides a novel step towards non-contact, *quantitative flow analysis* of exhale breathing behaviors. The significance of this approach is that it enables the development a new form of long-term respiratory evaluation that identifies subtle differences in natural breathing behaviors while enabling the foundation of long-term longitudinal studies where gradual changes in respiratory behavior can be quantified and correlated with condition progression. This has broad applicability to a wide variety of different conditions where specific signature traits are difficult to capture due to intermittent occurrences and short monitoring durations provided by existing flow measurement devices.

II. Related Work

Numerous techniques exist for both contact and non-contact respiratory analysis [5]. From the wide variety of available monitoring solutions, most devices can be categorized into one of four primary design spaces: (1) the device is a *contact* solution in which the user actively breathes on or through the device (pneumotach spirometer), (2) *wearable* solutions that track behaviors through movement or heat fluctuations (transducer belts, thermistors), (3) *non-contact* monitors that are based on radar, ultrasonic, wireless [6] signals or forms of imaging including depth [7], and thermal [8]–[11], and (4) *environmental* sensors (pressure, load cells) [12]. The primary differentiating factors of our approach is that we aim to obtain accurate physical measurements of flow through visual CO_2 density to volume analysis [3] as shown in Fig. 1. Of the existing related devices, the primary subset of solutions that directly measure air-flow and CO_2 pressure/concentrations are various forms of spirometry and capnography [13], [14]. The closest related non-contact exhale evaluation method is based on exhale characterization through Schlieren imaging [15] of exhale behaviors; however, while this similar form of analysis has been well studied, there has not been a direct quantification of exhale flow and volume based on visualized behaviors. Additionally, this method requires a precise experimental setup (large mirror, extensive camera setup) making it difficult to use with varied clinical settings. In contrast to prior efforts that evaluate exhale flow and tidal volume as indicators driving many PFT metrics, our method provides a first step towards *direct* analysis of exhale behaviors to obtain

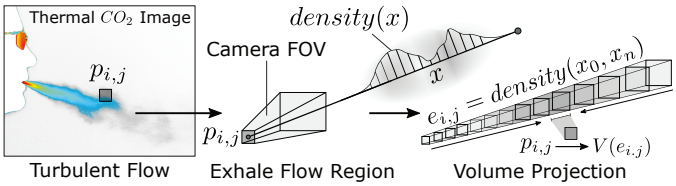

Fig. 1. Illustration of projected exhale volumetric turbulent flow as captured through the infrared (IR) thermal camera for each pixel $p_{i,j}$ to obtain the discrete volume density estimate $V(e_{i,j})$ [3].

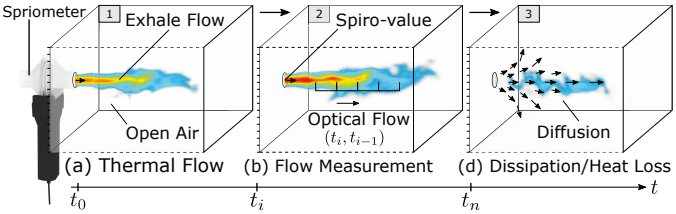

Fig. 2. Conceptual design. Visualized exhale behaviors are measured quantitatively through *gold-standard* (spirometer) and correlated visual *optical-flow* measurements integrated into a trained predictive model.

quantitative metrics based on visualized airflow. In contrast to complex 3D simulation models formed by flow observation constraints that closely replicate Navier-Stokes fluid dynamics (flow reconstruction), our aim is to formulate a computationally efficient *data-driven* approach suitable for real-time estimates in clinical settings. This enables our method to be used in a wide variety of clinical use-cases currently underaddressed by existing technologies used in sleep studies [16], pediatrics, and neonatal care [17].

III. Quantitative Analysis of Visual Exhale Flows

Visual analysis of breathing behaviors for obtaining quantitative metrics represents a form of *non-contact* respiratory monitoring method that directly extracts breathing behaviors from exhaled CO_2 turbulent flows. The main objective of this method is to expand the availability of new measurements used in pulmonary evaluation based on generating behavioral metrics from visualized CO_2 flow patterns. This form of analysis can lead to measurements of breathing rate, flow, exhale volume, and the natural distribution between *nose and mouth* breathing [3]. By capturing effort-independent, natural breathing, the visual representation of projected turbulent flows is quantified directly from *open-air* flows. We then formulate a method of correlating open-air *visual flow measurements* representing exhaled CO_2 concentrations with volumetric flow measurements obtained using a *gold-standard*. This enables quantitative measurements of flow by correlating measured volumetric flow with the visual flow from imaged CO_2 thermal exhale distributions as illustrated in Fig. 2.

A. Linking to a Quantitative Ground-truth

Standard solutions for *directly* measuring respiratory flow or CO_2 density enforce a *controlled flow* that is captured by breathing through a *tube* connected to a measurement device (spirometry, plethysmography, capnography). These devices capture flow characteristics based on a variety of different methods (pressure, ultrasonic) that sample the rate of flow as it passes through an enclosed region (the *tube* of the device). We consider this an instantaneous measurement of the flow as it passes through the measurement region of the device. These measurements are straight-forward to compute and are generally accurate due to how the direct measurement is performed. Unlike prior non-contact respiratory analysis methods that indirectly infer breathing measurements, our approach *measures the quantity of exhaled air* through the direct analysis of CO_2 flow. This inherits the benefits of a direct method, but enables the analysis of natural breathing and long-term monitoring. To provide quantifiable metrics for the

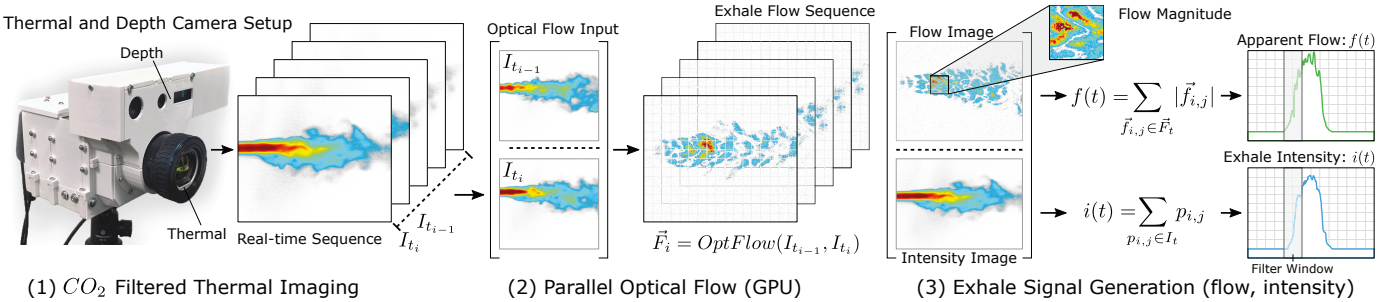

Fig. 3. Signal generation pipeline for measuring flow and intensity of exhales within $\boldsymbol{CO_2}$ filtered thermal images (1). We use optical flow (2) to isolate exhale behaviors and then convert the dense flow-field representations into 1D waveforms (3) that will be used inputs used to estimate exhale behavior measurements including flow in $(\boldsymbol{L/s})$ and volume in $(\boldsymbol{L})$. This represents the exhale flow extraction processing pipeline.

measurement of the observed behaviors, our method employs a predictive model to correlate our image-generated waveforms with the physical units provided by a *gold-standard*. We establish a tight correlation between this metric and the visual flow that *links* measurements of the physical quantity (flow in L/s) and the signals extracted from our image analysis. This provides a prediction model that establishes the link between visual analysis of apparent flow and a quantitative ground-truth for method validation, model training, and evaluation.

B. Measuring Unconstrained Exhale Behaviors

Sensing and quantitatively measuring respiratory behaviors through open-air analysis presents unique challenges in representing pulmonary health. Recent advances in inexpensive wireless technologies estimate flow [6], but only address a portion of our overall objective in vision-based pulmonary analysis. Our objective is to provide flow and volume measurements as a baseline, but also drastically expand respiratory diagnostic methods to include direct behavioral analysis to identify effort-to-flow relations, nose-mouth separation, and complex secondary behaviors as trace indicators for various pulmonary conditions. In our approach, our analysis generates a measure (SI units) of the apparent fluid dynamics representing flow, heat dissipation, and diffusion of carbon dioxide into open-air that we define as a *mixed*-signal representing a combination of the thermal energy and CO_2 concentration. This presents several challenges including: flow behaviors are spatially ambiguous, data is lost due to volume-to-plane projection (Fig. 1), viewing angle, and environmental factors alter flow behaviors (air movement, temperature, humidity). Due to this complexity, we focus on how to efficiently extract, track, and translate visual flow measurements [18], [19] with *unconstrained* open-air exhale behaviors. To address this, we build upon isolating detailed flow behaviors [3] to build a diagnostic viable for various forms of pulmonary evaluation, including polysomnography (PSG) studies in SDB (Fig. 4) that include visual flow measures [8] used for apnea event detection [20] as well as clinical out-patient PFT evaluations.

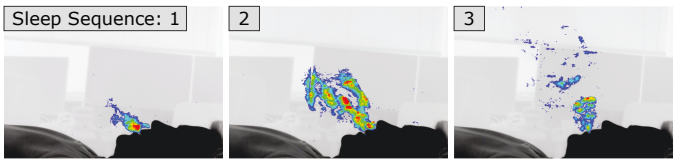

Fig. 4. Sequence illustration of quantified natural respiratory behavior tracking through our direct exhale tracking for SDB pulmonary analysis.

IV. Method

The principle of our approach is based on establishing a correlation between *visual flow* behaviors extracted from CO_2 thermal image sequences through *flow-field* analysis and *volumetric flow* measured as the flow rate of exhaled CO_2. We exploit the periodic behavior of the respiratory cycle to pair this correlation with a recurrent prediction model that provides an estimate of the flow as physical measurement in (L/s). This semi-continuous waveform prediction is then numerically integrated to obtain an estimated measurement of total exhale volume measured in (L). We then broaden the capabilities of this approach to *inferred* inhale predictions, breathing abnormality detection, and long-term monitoring. The proposed method utilizes a multi-stage pipeline where we define five primary steps: (1) real-time CO_2 paired thermal/depth image recording or streaming, (2) extraction of flow behaviors through flow-field analysis, (3) conversion of detailed flow vector-fields into continuous waveforms, (4) generation of datasets on a per-individual and cross-subject basis, and (5) the creation of data-driven models that predict continuous flow behaviors. The construction and training of these models provides the basis for predicting the rate waveforms of the observed flow which are then numerically integrated to obtain individual volume estimates for each exhale episode (ξ). The trained models and volume calculations can then be used to directly monitor and measure open-air exhale behaviors.

To establish our flow prediction models, we develop a new set of image-based datasets used to generate waveforms of *apparent flow* that quantify exhale behaviors as shown in Fig. 3. In the constructed datasets, we employ two strategies: (1) the construction of *per-subject* models that can accurately replicate individualized behaviors unique to each subject and (2) a generalized *cross-subject* model that incorporates all recorded data into a consolidated prediction model. Unlike the per-subject model which provides accurate predictions but requires training for each individual, the cross-subject approach provides a practical solution by eliminating the need for per-patient training at a slight cost to accuracy due to generalization. Both models predict the flow rate as compared with the ground-truth, in which we evaluate the characteristics of the predicted signals, correlation with spirometer readings, and directly compare measurements. Evaluation of our method is performed based on three primary metrics including: breathing rate (BPM), flow (L/s), and volume (L) on a preliminary study size of 12 subjects for individual and cross-subject models.

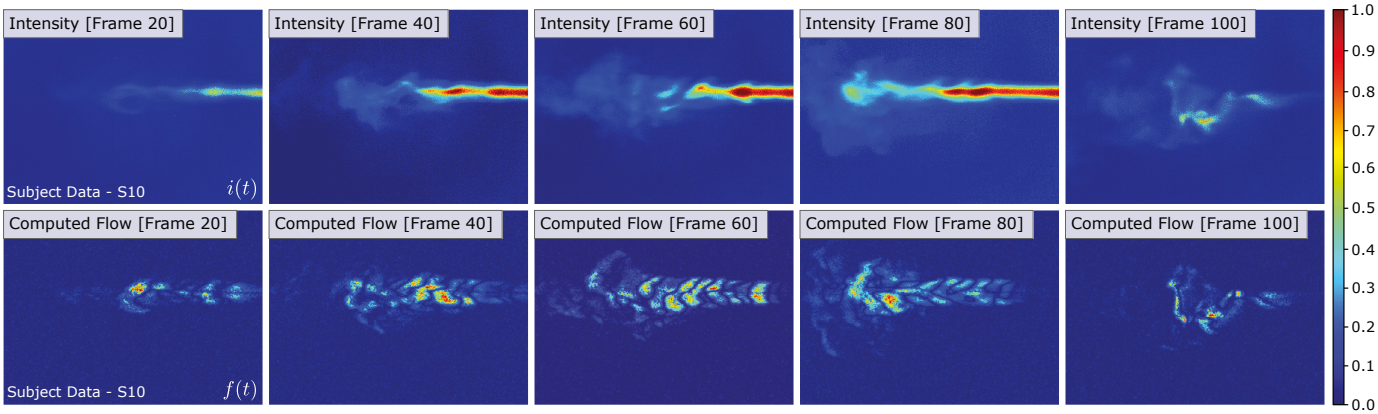

Fig. 5. Signal generation pipeline for measuring flow and intensity of exhales within $\boldsymbol{CO_2}$ filtered thermal images (1). We use optical flow (2) to isolate exhale behaviors and then convert the dense flow-field representations into 1D waveforms (3) that will be used inputs used to estimate exhale behavior measurements including flow in $(\boldsymbol{L/s})$ and volume in $(\boldsymbol{L})$. This represents the exhale flow tracing/extraction processing pipeline.

A. CO_2 Imaging and Respiratory Signals

To enable a real-time stream of quantitative metrics based on the sequence of thermal images, we integrate the dense optical flow result into a waveform generation pipeline. From the observed optical flow, we can isolate exhale behaviors based on vector classification and extract two measurements to define *intensity* and *flow* measurements over time. The thermal image stream provides sequential images used to compute the optical flow vector field $\vec{F}_t$, which is then used to isolate detailed flow behaviors ($|f_{i,j}| \in \vec{F}_t > \varepsilon$). This produces two *discontinuous* wave-forms from the two real-time image streams that contain: (1) a CO_2 image I_t for an intensity signal $i(t)$ and (2) the paired field $\vec{F}_t$ representing quantitative measurement of the visualized exhale flow signal as $f(t)$. We perform a parallel variant of dense Optical Flow [19]. This fulfills two requirements: (1) isolation of flow behaviors from background heat sources/movement and (2) enables identification and quantification of minute flow details from which we can build respiratory metrics. We minimize standard intensity and smoothness constraints [19] of the form:

$$\min \iint (\nabla I \cdot v + \frac{\partial I}{\partial t})^2 + \alpha(\|\nabla v_x\|^2 + \|\nabla v_y\|^2)\, dxdy \quad (1)$$

Where $\vec{v} \in \mathbb{R}^2$ and $\alpha \in [0,1]$; we select small values: $\alpha = 0.15$ to retain fine flow movements for detailed behavioral analysis. For the minimization problem, we employ an iterative optimization scheme executed in parallel. This leads to a formulation based on two input images: I_{t_i} and $I_{t_{i-1}}$, two intermediate images U and V, and the time derivative image I_t. We compute $\partial I/\partial t$ using backward differences and compute both $\partial I/\partial x$ and $\partial I/\partial y$ using x and y filters respectively. Illustrated results are shown in Fig. 5.

$$v_x^{i+1} = \bar{v}_x^i - \left[\frac{\frac{\partial I}{\partial x}\bar{v}_x^i + \frac{\partial I}{\partial y}\bar{v}_y^i + \frac{\partial I}{\partial t}}{\alpha + \frac{\partial I}{\partial x}^2 + \frac{\partial I}{\partial y}^2}\right]\frac{\partial I}{\partial x} = U_i \quad (2)$$

$$v_y^{i+1} = \bar{v}_y^i - \left[\frac{\frac{\partial I}{\partial x}\bar{v}_x^i + \frac{\partial I}{\partial y}\bar{v}_y^i + \frac{\partial I}{\partial t}}{\alpha + \frac{\partial I}{\partial x}^2 + \frac{\partial I}{\partial y}^2}\right]\frac{\partial I}{\partial y} = V_i \quad (3)$$

Where the update to the v_x and v_y flow vectors are based on the discrete iterative formulation [21] based on Equation 1 where all filter processing is done on the GPU [22].

Flow Signal. To provide an estimate of the flow obtained by an instantaneous measurement at each captured frame, we convert the sum of the magnitudes of the segmented optical flow vectors into a semi-continuous waveform. This formulates the changes in the observed exhale patterns as an estimate of the *apparent flow*, which represents the changes in CO_2 movement throughout each exhale ξ. For a given optical flow vector field $\vec{F}$ at time t computed from optical flow using images $I_{t_{i-1}}$ and I_{t_i}, where the discrete form of the flow magnitude is computed in Equation 4.

$$f(t) = \sum_{\vec{f}_{i,j} \in \vec{F}_t} |\vec{f}_{i,j}| \quad ; \quad \vec{F}_t = OptFlow(I_{t_{i-1}}, I_{t_i}) \quad (4)$$

Each optical flow field $\vec{F}_t$ also performs the pre-processing required for segmenting the exhale flow from body movement and other thermal objects within the monitored area. This provides segmentation for both flow and intensity behaviors from the surrounding environment. The flow vectors $\vec{f}_{i,j}$ can be classified (movement, noise) and then used to generate a weighted mask applied to the intensity image ($p_{i,j} > 0$) to isolate exhale contributions for each quantitative measurement. A consequence of this approach is that the flow segmentation introduces bias the intensity behavior towards the characteristics exhibited by the flow signal, but is required to effectively isolate exhale behaviors from the surrounding environment.

Intensity Signal. The intensity of the thermal signature in the filtered wavelength provides a snap-shot representation of the CO_2 density as a function of detected CO_2 emission and thermal intensity. This density contributes to the overall exhale volume due to the higher sensor activations from the mixed thermal-CO_2 signature. To compute the total amount of exhale within each frame, we consider the flow-segmented intensity image and generate the sum of the non-zero pixel values $p_{i,j}$ of image I_{t_i} as shown in Equation 5.

$$i(t) = Intensity\,(I_{t_i}) = \sum_{p_{i,j} \in I_t} p_{i,j} \quad \forall p_{i,j} > \varepsilon_{flow} \quad (5)$$

This generates two signals that measure approximations of the visual changes in air movement as a representation of flow signal $f(t)$ and CO_2 density over time as the intensity signal $i(t)$. These signals represent our recurrence model inputs.

Signal Processing. Turbulent exhale flows naturally exhibit complex fluid dynamics that are lost in the imaging process. This is due to complex flow behaviors and the projection of the $3D$ phenomena to the $2D$ image plane, heat distribution changes, and the chaotic dissipation of the CO_2 into the open air. To address this, we implemented two forms of signal filtering: (1) windowed average filtering ($n = 12$ @ 30Hz) and (2) Gaussian weight filtering. The filtering reduces the high-variance regions of the signal create large fluctuations that exist within the peaks of exhale signals due to chaotic dissipation behaviors. To minimize the variance of the signal spikes during each exhale, we weighting higher flow magnitudes in the Gaussian window that softens high-variance peaks. The aim is to align the waveform characteristics of the generated signals with ground-truth measurements. The resulting raw and filtered signals are shown in Fig. 6.

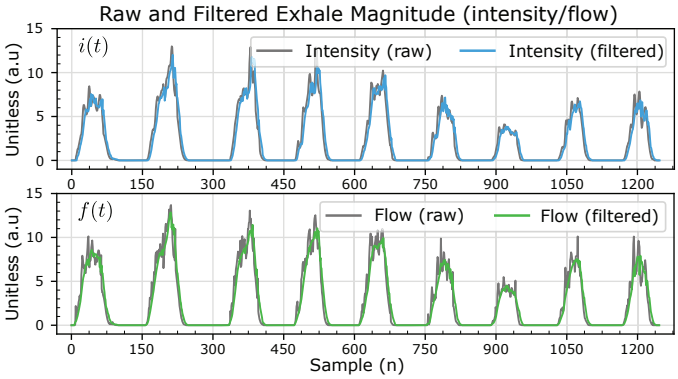

Fig. 6. Raw and filtered intensity $\boldsymbol{i(t)}$ and flow $\boldsymbol{f(t)}$ signals. Zero regions within the generated signals correspond to inhale segments. The exhale behavior has been varied to illustrate different exhale sizes.

B. Volumetric Flow Estimation Model

The aim of our approach is to provide an efficient method for estimating the volumetric flow which represents the rate of exhaled CO_2 measured over time. The model accounts for the similarity of the data characteristics in the periodic exhale signal and is formulated as a *recurrent* model that takes (s) prior steps into account for each prediction. We define this model as $\mathcal{F}(t) = \mathcal{M}(i(t), f(t), s)$ where $i(t)$ and $f(t)$ represent the input signals and $\mathcal{F}(t)$ is the predicted waveform. The structure of the model is presented in Fig. 7.

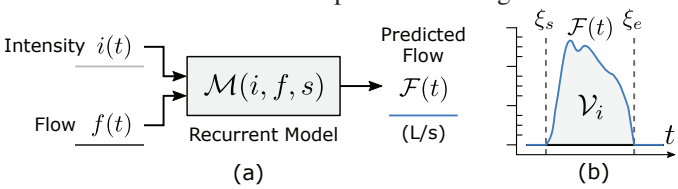

Fig. 7. Model overview. (a) The inputs of the model $\boldsymbol{i(t)}$ and $\boldsymbol{f(t)}$ are used to estimate the rate of volumetric flow representing the amount of $\boldsymbol{CO_2}$ exhaled breath per second measured in (L/s). The model prediction is then shown in (b) as an estimate of the total volume $\boldsymbol{\mathcal{V}_i}$ from exhale at index $\boldsymbol{i}$ from associated (end - start) duration $\boldsymbol{\xi_e - \xi_s}$.

The prediction of the model represents the conversion from *visual flow* measured as an arbitrary unit (a.u) to the *volumetric flow* defined as the rate of exhaled CO_2 over time, measured in liters per second (L/s). The error of this prediction is evaluated as the deviation from the ground-truth flow waveform provided by the reference spirometer acting as our gold standard.

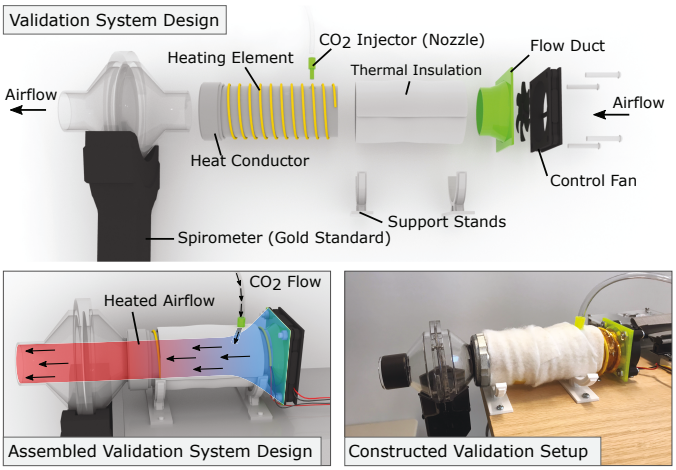

Fig. 8. Exhale flow validation device design. This device provides a reoccurring control for evaluating synthesized behaviors. The validation system emulates the heat, $\boldsymbol{CO_2}$ concentration, and airflow of a typical expiratory episode as measured by both our method and spirometry.

C. Volume Estimation

The flow estimation model $\mathcal{M}(i, f, s)$ provides an estimation of the volumetric flow $\mathcal{F}(t)$ representing the quantity of CO_2 during each exhale episode based on the input intensity $i(t)$ and flow $f(t)$ signals. To obtain the i^{th} per exhale volume in liters as $\mathcal{V}_i$, we integrate the predicted signal from the start (ξ_s) to the end (ξ_e) of the i^{th} exhale episode as shown in Fig. 7 (b). From the model prediction, the volume is computed through integration as shown in Equation 6.

$$\mathcal{V}_i = \int_{\xi_s}^{\xi_e} \mathcal{F}(t)\, dt \quad \text{where} \quad \mathcal{F}(t) = \mathcal{M}(i, f, s) \tag{6}$$

The discrete form of the predicted waveform is integrated using Simpson integration. This represents the total volume exhaled within a single episode measured in liters (L).

V. Method Validation

Unconfined turbulent exhale flows exhibit high levels of variance due to several factors including: exhale and ambient temperature, moisture contained in the exhale, and $C0_2$ density, all of which determine the visual flow behavior and dissipation rate of each expiratory episode. Establishing an accurate flow prediction model capable of potentially representing quantitative measurements requires a experimental control for establishing the validity of the relationship between visualized exhale behaviors and accurate flow measurements.

A. Visual Flow Measurement Control

We validate this approach by creating consistent and repeatable exhale behavior analogue for measurement verification. To do this, we have designed an experimental control setup (Fig. 8) for accurately emulating exhale behaviors that can be measured with both an existing gold standard as well as the proposed model. The premise of this validation is based on heating controlled CO_2 concentrations that are forced through the validation system at a fixed interval using a programmatically controlled fan. This simulates expiratory episodes that are simultaneously captured by both the imaging system and

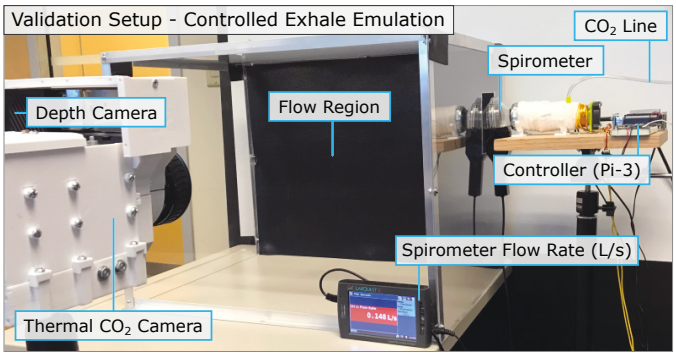

Fig. 9. Validation setup. This setup has been created to provide an *idealized flow region* where the flow measurements of turbulent exhale behaviors can be calibrated for controlled timing, temperature, distance, and $\boldsymbol{CO_2}$ concentrations compared with spirometer measurements.

spirometer measurements of flow (L/s) and volume (L). The experimental setup for validating our control measurements is shown in Fig. 9. To maintain an accurate level of CO_2 density within the flow, we exhale into an external reserve that holds exhaled breath and limit the CO_2 flow through an injector nozzle to ensure a constant release that will mix with the air drawn in from a $40mm$ control fan. To precisely control the fan speed and duration, we utilized a simple motor controller (TI-SN754410) connected to a Raspberry Pi [23]. Through Pulse Width Modulation (PWM) we alter the duty cycle of the fan to provide varying exhale strengths and durations generated $2.0s$ apart. This setup is then used to vary experimental parameters including: (1) fan speed for exhale strength, (2) duration for volume, and (3) camera-to-flow distance, while fixing CO_2 mixture rate and temperature.

B. Flow and Volume Measurement Validation

To establish a baseline of the variance incurred by quantifying turbulent flow behaviors, we used the validation setup to demonstrate the accuracy of the relational model. We formed a preliminary validation dataset to train an initial prediction model to estimate flow magnitude and account for the differences in waveform characteristics between the flow, intensity, and spirometer measurements. The model was then used to predict the flow waveform for 12 simulated exhale flows, with the results of the predictions shown in Fig. 10.

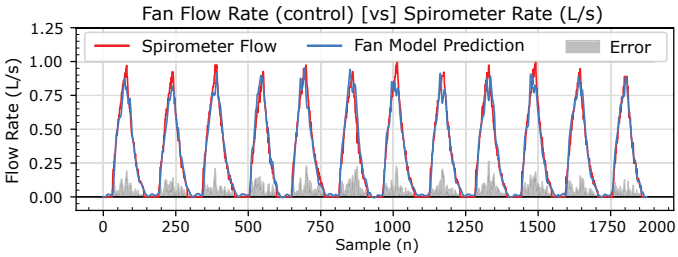

Fig. 10. Flow validation: Predicted $\boldsymbol{\mathcal{F}(t)}$ versus Spirometer flow $\boldsymbol{\mathcal{S}(t)}$.

Errors between the spirometer and predicted flow are a product of the variance exhibited by the turbulent flow, even in the instance where every episode is similar in temperature, duration, and strength. This is also exhibited at the peak magnitudes of each exhale. Capturing this variance presents the primary challenge of establishing an adequate neural network model for accurate predictions, motivating the use

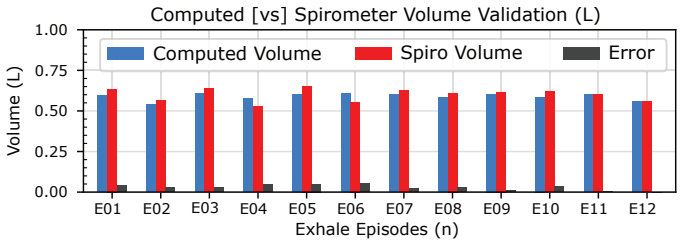

Fig. 11. Model flow estimations integrated per exhale validating control volume measurements. Variance of turbulent flow measurements are illustrated in the volume differences across the $\mathbf{12}$ exhale episodes.

of a recurrent model to capture these behaviors. Validation of the per-episode (ξ_i) volumes generated from the predicted waveforms are obtained by integrating each exhale instance from Fig 10, with the computed volumes shown in Fig. 11.

VI. Exhale Model Generation and Training

The proposed model must provide a transformative relationship between the observed thermal distribution of CO_2 and volumetric flow measurements provided a gold-standard. This requires capturing quasi-periodic but volatile exhale flow behaviors as a continuous recurrent waveform representing flow predictions over time. To achieve this, we employ an LSTM network architecture to provide a predictive model architecture that can estimate both waveform characteristics and physical quantified metric from apparent flow. The continuous prediction of flow based on the input signals then forms the basis of a discontinuous non-linear regression encoded into the trained model predicted using the last s (look_back) samples. To evaluate the loss within the model based on the observed value provided from the spirometer sample $\mathcal{S}(t)$ and the predicted flow generated by our model $\mathcal{F}(t)$, we use *Mean Squared Error* where $MSE_{loss} = 1/n \sum_{i=1}^{n} (\mathcal{S}_i - \mathcal{F}_i)^2$. In our approach we formulate two models: (1) an *individual* and (2) *cross-subject model*. The individual model is per-subject and accurately predicts waveforms with unique breathing behaviors. The cross-subject model provides an inter-subject prediction that only requires initial training on a larger dataset to generalize across expected flow behaviors, increasing practical use and utility. An overview of the data collection and model separation is shown in Fig. 12. We form three dataflow training procedures: (a) randomized cross-subject (mixed), (b) individual datasets, and (c) selective data as shown in Fig. 13. We employ parallel (GPU) LSTM layers [24] provided by *Cuda* [25] in Keras [26] with TensowFlow back-end [27]. We select a recurrent architecture to capture the time dependency across input sequences; however, other architectures (1D CNN/etc.) are also viable alternatives.

A. Dataset Generation

High variance within turbulent exhale data requires that we account for significant fluctuations within the observations included within the training data. To establish an accurate correlation between visual flow measurements and exhale volume, we select *best candidate* exhale episodes to form training datasets. This is achieved by selecting exhale measurements that closely represent the same area under the spirometer flow curve as primary candidates for the training set. First,

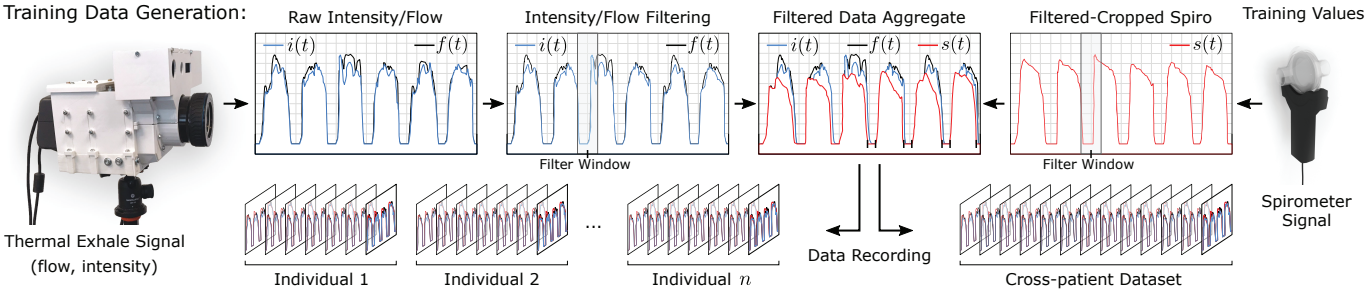

Fig. 12. Model training datasets. We formulate *individual* and *cross-subject* datasets used to train different measurement objectives including per-subject accuracy (individual) and train towards generalization through mixing randomized exhale behaviors within the (cross-subject) model.

both signals are aligned through re-sampling to account for data acquisition rates. Then each individual dataset $\mathcal{D}$ is composed of at least 25 exhale episodes collected in 4 recorded sessions for each subject S01 - S12 ($n = 12$) to generate a training dataset defined as: $\mathcal{D}_t = \{\mathcal{D}_0, ..., \mathcal{D}_4\}$. Each dataset is composed of an average range of 9,000-20,000 samples. As there are portions of $null$ measurements between exhales representing the inspiratory phase ($\mathcal{I}$), the valid sample size of each dataset is approx. 6,000 non-zero measured samples. In our datasets, depending on subject, each exhale typically contains 50-70 filtered samples (an average ξ duration of $2s$ @ $\sim 30Hz$). After acquisition, each exhale is separated as an individual episode, each of which can be randomized within the training data to avoid individual behavioral bias.

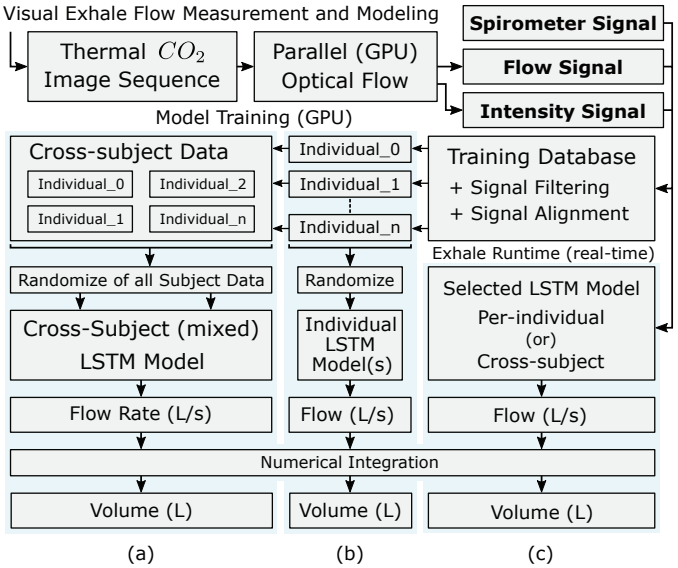

Fig. 13. Dataflow overview: Flow behaviors are extracted from $\boldsymbol{CO_2}$ image sequences and correlated with spirometer readings to generate training databases for predictions of flow ($\boldsymbol{L/s}$) and volume ($\boldsymbol{L}$).

To identify exhales that will be added to the training set $\mathcal{D}_t$, we select a subset for training based on relational area constraints. Each exhale ξ is bound by non-zero start ξ_s and end ξ_e to form each individual episode. For each exhale in the filtered thermal exhale signal, we define the average area of the flow $f(t)$ and intensity $i(t)$ signals as $\mathcal{T}_i$. Similarly, the expiratory area of the spirometer for exhale i is defined as $\mathcal{S}_i$.

$$\mathcal{T}_i = \frac{1}{2}\int_{\xi_s}^{\xi_e} \big[f(t) + i(t)\big]\, dt \quad ; \quad \mathcal{S}_i = \int_{\xi_s}^{\xi_e} s(t)\, dt \tag{7}$$

Where the ratio of the overlapped signal area of each exhale is evaluated as the primary selection criteria. This is defined as $\mathcal{O}_i = \mathcal{T}_i/\mathcal{S}_i$. These overlap values are computed for all exhales in each dataset. The average and standard deviation of these overlaps are computed as $d = mean(O_i)$ and $\varepsilon = std(O_i)\ \forall i \in \mathcal{D}$. For the secondary selection criteria, we evaluate the similarity of the waveform. We compute the mean correlation of both signals as $\gamma = \frac{1}{2}[r^2_{f,s} + r^2_{i,s}]$ where $r^2_{f,s}$ and $r^2_{i,s}$ are the correlations between the flow, intensity, and spirometer waveforms respectively. This provides the formulation of two numerical constraints C_1 and C_2 that define the per-exhale training data given the following selection criteria:

$$\begin{aligned} C_1 &= (d - \varepsilon) < \mathcal{O}_i < (d + \varepsilon) \\ C_2 &= r_i^2 > \gamma \end{aligned} \tag{8}$$

where r_i^2 is the correlation of exhale ξ_i. Exhale episodes satisfying both of these constraints are selected for the construction of each training dataset $\mathcal{D}_t$. That is: $\xi_i \in \mathcal{D}_t$ iff $C_1 \wedge C_2$.

B. Model Architecture

To formulate the layer architecture of the proposed model, we perform parallel feature extraction that leverages features from both intensity $i(t)$ and flow $f(t)$ signals to establish an accurate flow estimate defined as the expected volumetric flow provided by the spirometer. The motivation for this architecture is to extract characteristics of the signals that include CO_2 density and flow behaviors that contribute to the volumetric flow measured by the spirometer. Using subsequent LSTM layers, features are extracted in parallel streams which are then combined in a layer concatenation that consolidates these features into a single prediction $\mathcal{F}$ that represents the expected volumetric flow. An overview of the model architecture, layers with unit sizes (individual, cross-subject), activations (layer_activation, recurrent_activation), and input (batch_size, look_back, 1) is shown in Fig. 14. In training, the model is trained for a total of 144 epochs with a batch size of

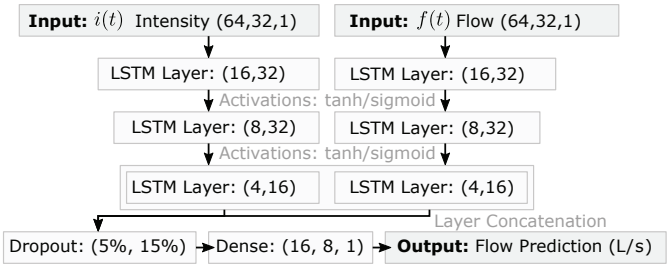

Fig. 14. Model architecture. Parallel feature extraction for both inputs providing an estimation of the expected spirometer flow measured in L/s. Each RNN layer is implemented on the GPU through *Cuda*.

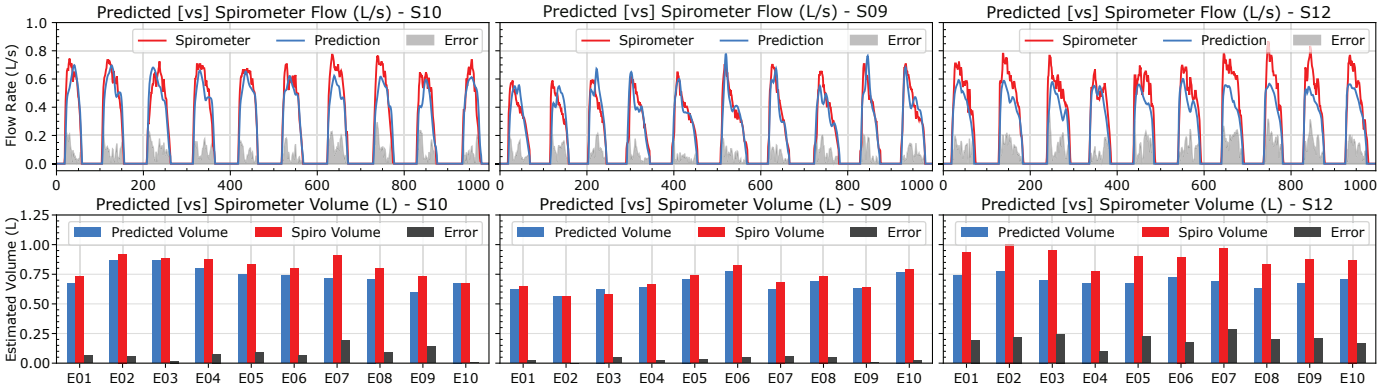

Fig. 15. Independent model results for three select subjects (S10, S09, S12). The flow plots (top) provide the predicted flow rate directly compared to the spirometer measurements and the volume plots (bottom) show the volume for each of the 10 episodes (E01 - E10) with associated error.

64. Each training dataset for both the individual and cross-subject models is set to a $80\%/20\%$ training/validation split. For the recurrent *look back* (input length s) the motivation is to capture typical episode behavior and was empirically obtained as $s = 32$ (@30Hz). For model optimization, we use *ADAM* [28] with parameters: $learning\_rate = 0.002$, $\beta_1 = 0.9$, $\beta_2 = 0.99$, and $decay = 0.001$. Dropout rate is defined as 5% and 15% for individual and cross-subject models respectively. The model architecture has been implemented in Keras using the functional model API to account for the parallel input and layer concatenation. For the cross-subject model, the model unit sizes are increased to aid in accounting for the higher potential variance in exhale waveform characteristics.

VII. Experimental Design

Experimental evaluation of the proposed method is based on two primary metrics: (1) the LSTM prediction of the volumetric flow waveform measured in (L/s) and (2) the resulting estimation of the per-episode exhale volume (L). This was completed through the collection of subject data used to train both individual and cross-subject LSTM models to evaluate their performance on data obtained from $n = 12$ healthy subjects within the normative range of lung function. The overall experimental setup is similar to the monitoring design shown in Fig. 9, adapted for each subject to sit comfortably and breathe naturally for the required monitoring period. The setup is composed of: (1) the depth/thermal camera device, (2) background (projector screen), and (3) the subject sitting 1.5m from the camera. Subject data collection included the collection of 4 to 6 trials, each recorded for 120 seconds, resulting in an average of 24.67 exhales obtained for each trial with an average of 14626 samples. Environmentally we ensured limited surrounding airflow, constant humidity, and interior ambient temperature ($22°C - 26°C$), which require initial calibration. Due to the direct visualization of exhale flow, any form of covering that prevents visual analysis of flow (tubes, masks, etc.) are not included within the study.

From the paired depth/thermal imaging system, the thermal stream is captured using a FLIR-A6788sc (640x512 @ $30Hz$) thermal camera with an actively cooled band-pass spectral filter 3-5μm generating 16-bit raw signal through the Gigabit Ethernet Vision Protocol (GigE Vision). These are then paired with registered depth images (512x424 @ $30Hz$) and stored as encoded video files. The gold-standard measurements for flow are provided by a Vernier spirometer, calibrated with a sample rate 30Hz, running concurrently with the recording. The overall experimental design was then separated into two training/evaluation sets based on the individual and cross-subject models respectively. Each exhale is then randomized on a per-exhale basis for model training. Extending the train/test split in the training phase, we also evaluate the accuracy using unseen data that evaluates new subject data trained models have not encountered. Human subject evaluation approved as part of the Colorado Multiple Institutional Review Board (COM-IRB FB F490): Non-contact Remote Breathing Analysis through Visualization of Thermal and CO2 Flow (VTCF).

VIII. Results

Evaluation of the proposed method is divided into two primary results: (1) the individual model and (2) the cross-patient model for both predicted flow and computed volume. Representing a selection of the computed results, Fig. 15 presents select predicted exhale waveforms and associated volumes. These present the advantage of an individual model that captures subtle characteristics of each subject's exhale breathing patterns can be captured and identified within the visual flow. This is due to the training data exhibiting the same characteristics as newly provided unseen data from the same subject, demonstrating the possibility of creating longitudinal studies to identify variance trends through long-term monitoring. Based on our evaluation of the individual model accuracy we obtained an average correlation $R^2 = 0.912$ and average volume accuracy of 90.92 within the $80/20$ training split with an overall accuracy of 88.02% for newly collected unseen data. The absolute error in flow between the prediction and spirometer $|\mathcal{F}(t) - \mathcal{S}(t)|$ associated with the individual model for each subject is shown in Fig. 16 with an overall analysis of the individual subject model presented in Table I.

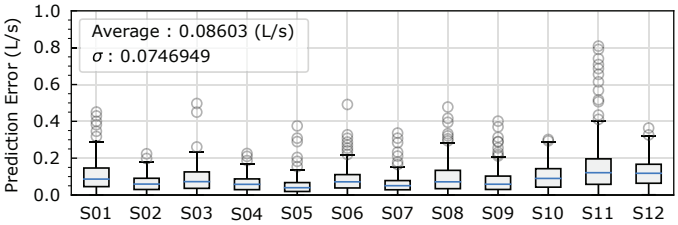

Fig. 16. Predicted flow error as measured by the absolute difference between the model prediction and spirometer $\epsilon = |\mathcal{F}(t) - \mathcal{S}(t)|$.

TABLE I

INDIVIDUAL SUBJECT MODEL FLOW AND VOLUME RESULTS - LEFT: PREDICTED FLOW (L/S), RIGHT: ESTIMATED VOLUME (L)

Subject Data			Raw Signal RMSE		Raw Signal Correlation R^2		LSTM Prediction-Spiro Results (L/s)					20% Split Test (Vol. Diff. Error in L)			Unseen Test Data (Result)		
SID	**SIS**	**Age**	**Flow-Spiro**	**Int-Spiro**	**Flow-Spiro**	**Int-Spiro**	**Mean**	**S.D.**	**RMSE**	R^2	**Improvement (%)**	**Error (L)**	**S.D.**	**Accuracy (%)**	**Error (L)**	**S.D.**	**Accuracy (%)**
S01	F	32	0.184	0.174	0.602	0.625	0.614	0.016	0.098	0.891	**27.75%**	0.096	0.077	91.09%	0.174	0.139	**84.55%**
S02	F	30	0.155	0.171	0.698	0.691	0.695	0.005	0.059	0.925	**23.03%**	0.087	0.801	91.26%	0.044	0.029	**90.55%**
S03	F	46	0.065	0.120	0.807	0.841	0.824	0.024	0.029	0.966	**14.20%**	0.049	0.039	91.56%	0.047	0.037	**92.64%**
S04	M	40	0.153	0.169	0.740	0.740	0.740	0.001	0.036	0.913	**17.28%**	0.052	0.035	90.53%	0.037	0.019	**93.94%**
S05	M	29	0.135	0.172	0.822	0.829	0.826	0.005	0.023	0.963	**13.75%**	0.029	0.020	93.42%	0.036	0.024	**91.09%**
S06	M	25	0.110	0.117	0.797	0.808	0.803	0.008	0.095	0.845	**42.50%**	0.131	0.100	83.62%	0.174	0.130	**76.44%**
S07	F	66	0.157	0.164	0.521	0.520	0.520	0.001	0.042	0.910	**38.98%**	0.037	0.033	94.04%	0.060	0.038	**90.15%**
S08	M	75	0.217	0.221	0.603	0.617	0.610	0.010	0.127	0.811	**20.10%**	0.221	0.149	80.69%	0.248	0.198	**82.30%**
S09	M	44	0.157	0.165	0.630	0.637	0.634	0.005	0.083	0.926	**29.25%**	0.067	0.047	95.72%	0.369	0.131	**93.51%**
S10	M	31	0.132	0.128	0.742	0.744	0.743	0.001	0.072	0.940	**19.73%**	0.088	0.060	92.34%	0.289	0.111	**94.44%**
S11	F	31	0.344	0.338	0.463	0.476	0.470	0.010	0.161	0.939	**46.95%**	0.084	0.064	94.66%	0.159	0.125	**91.27%**
S12	M	25	0.153	0.150	0.545	0.576	0.560	0.022	0.054	0.941	**38.09%**	0.041	0.037	92.11%	0.239	0.070	**75.65%**
Mean		39.5	0.164	0.174	0.664	0.675	0.670	0.009	0.073	0.912	**24.45%**	0.082	0.064	90.92%	0.163	0.089	**88.02%**
S.D.		16.1	0.068	0.059	0.121	0.121	0.120	0.008	0.042	0.046	**12.28%**	0.053	0.038	4.423%	0.124	0.060	**6.659%**

Breaths Per Minute (BPM). Peak detection is used on smoothed predicted waveforms to estimate the total number of breaths per minute. Due to the large number of exhale samples within the training datasets, we randomly select one trial from each subject and compute the BPM. The Bland-Altman plot for the BPM for the selected set of trials is shown in Fig. 17.

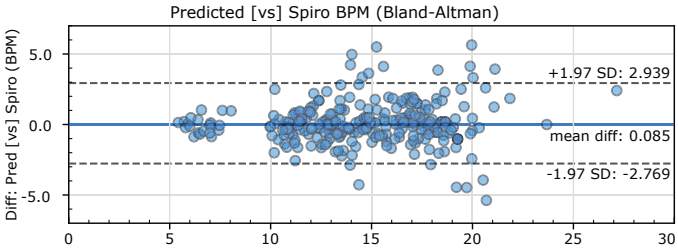

Fig. 17. Bland-Altman plot for all subjects measured as peak-to-peak time comparing predicted and spirometer BPM (one trial per subject).

Estimated Exhale Volume. Across all collected data, the flow associated with each exhale episode is numerically integrated to get a per-episode estimate of the total exhaled CO_2. Typically exhale volume at rest ranges between 0.5 and 1.0L for most adults [5]. The Bland-Altman plot for all exhale volumes for all subjects is shown in Fig. 18.

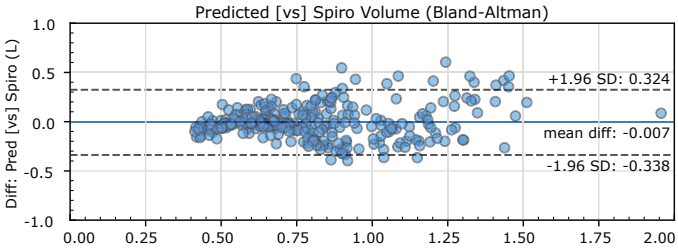

Fig. 18. Bland-Altman plot for predicted volume [vs] spirometer volume computed by integration for all subjects based on the individual model.

To establish the overall accuracy of the correlation between the BPM and volume measurements compared to the recorded spirometer measurements, we provide the correlation plots in Fig. 19. This demonstrates the overall correlation for the selected BPM measurements as $R^2 = 0.95$ and overall correlation of volume measurements as $R^2 = 0.867$. The strong correlation in these estimations provides a valid foundation for the development of new breathing metrics established on the independent subject model while maintaining a basis for mapping the results with standard PFT evaluations.

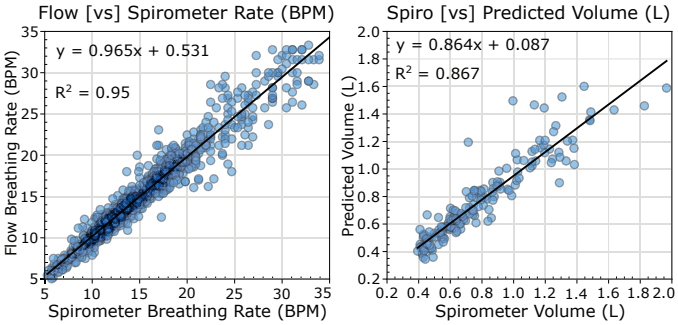

Fig. 19. Correlation of BPM and Volume for all subject data.

Cross-subject Model Flow and Volume. Training exhale prediction models for each subject presents a practical limitation within a clinical setting where time and healthcare resources are limited. To address this, we introduce a *cross-subject* model defined by a dataset composed of randomized exhale episodes from all subjects. The aim of the cross-subject model is to remove the requirement for training for each new subject. For the twelve subjects, an average of 24 episodes were collected for an average of 14626 samples per subject. From the cross-subject model, predicted flow exhibits lower correlation $\bar{R^2} = 0.804$ and higher average $RMSE = 0.121$ presented in Table II. Due to generalization, we see a drop in volume accuracy from 88.02% to 81.05%.

Inferred Inhale Estimation. A consequence of directly measuring the visual signature for respiratory analysis is that the inspiratory portion of the respiratory period cannot be directly measured; however, it can still be predicted through constrained *inference*. To extend the flow waveform generated by our approach to incorporate the inhale phase of the respiratory cycle, we construct an estimate of the complete respiratory waveform through *dual models* that predict both the inhale and exhale portions independently to construct the full waveform. This is done by delaying the prediction by a half-cycle and using the current exhale to infer flow characteristics of the inhale (exhale-inhale inference) where we link exhale volume with the volume of the prior inhale. This is not a direct relationship, but provides an estimate of the complete waveform under regular, natural breathing. Using this inference, the inhale model is trained on the relationship between the current exhale ξ_t and the preceding inhale $\mathcal{I}_{t-1}$.

TABLE II
CROSS-SUBJECT MODEL FLOW (LEFT) AND VOLUME (RIGHT)

	Experimental Data		Prediction-Spiro		Volume Difference (L)		Volume
SID	**Exhales**	**Samples**	**RMSE**	R^2	**Error (L)**	**S.D.**	**Accuracy (%)**
S01	23	14310	0.088	0.867	0.071	0.068	**92.59%**
S02	22	20814	0.073	0.880	0.041	0.038	**94.22%**
S03	26	15568	0.099	0.780	0.072	0.042	**89.04%**
S04	30	15882	0.119	0.726	0.165	0.090	**73.82%**
S05	19	8138	0.148	0.691	0.160	0.040	**67.06%**
S06	28	16546	0.162	0.774	0.299	0.132	**66.58%**
S07	30	13712	0.098	0.760	0.060	0.044	**87.96%**
S08	23	24004	0.174	0.752	0.387	0.227	**75.89%**
S09	23	10764	0.073	0.890	0.035	0.028	**94.99%**
S10	24	9382	0.081	0.873	0.038	0.022	**94.62%**
S11	24	12910	0.183	0.854	0.497	0.129	**62.32%**
S12	24	13484	0.152	0.799	0.203	0.062	**73.50%**
Mean	24.67	14626	0.121	0.804	0.169	0.077	**81.05%**
S.D.	3.284	4493.25	0.041	0.067	0.152	0.060	**12.39%**

The limitation is that this only provides an *estimated description* of the inhale behavior, which is only valid when current exhale flow is directly proportional to the prior inhale. Based on this modeling approach, we impose two constraints: (1) natural, regular breathing and (2) $\mathcal{I}_{t-1} \propto \xi_t$. The dual-model prediction that combines both the expiratory and inspiratory phases from the measured visible exhale is shown in Fig. 20.

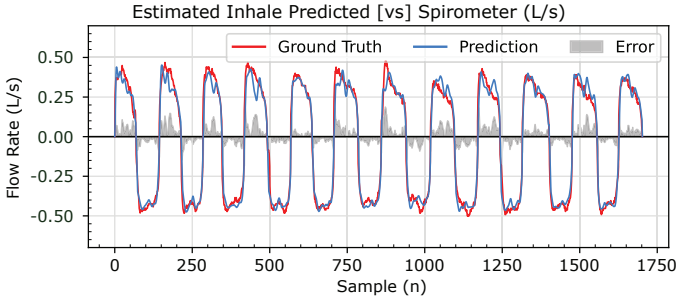

Fig. 20. Inferred inhale: Dual model estimation that predicts $\mathcal{I}_{t-1}$ and ξ_t portions independently to provide a complete interleaved waveform.

IX. LIMITATIONS AND FUTURE WORK

Visual exhale analysis is subject to various vision-based limitations including: field-of-view, hardware sensitivity/resolution, image-to-signal conversion, and model generalization. Trained model correlation of visual flow also has challenges related to the volatility of rapid dissipation, environmental factors, training domain selection, model architecture, and hyperparameter optimization. Future work incorporates broadening training heterogeneity and image-based evaluation to improve this form of diagnostic in pulmonary pathology.

X. CONCLUSION

In this work, we presented a computationally efficient method for the *direct* measurement of exhale flow and volume through spectral-filtered thermal CO_2 imaging for clinical pulmonary evaluation. Our method presents a direct vision approach for evaluating pulmonary function to enable long-term monitoring, natural breathing, and the identification of subtle variance in breathing behaviors that are difficult to capture using exiting devices. For flow estimations, we formulated two regression model datasets that account for individual and cross-patient characteristics obtaining an average flow correlation of $R^2 = 0.912$ and $R^2 = 0.804$ with average volume estimation accuracy of 88.02% and 81.05% respectively.

ACKNOWLEDGMENT

We would like to recognize Dr. Ann Halbower's contributions to both clinical insight and COM-IRB management and approval in coordination with Children's Hospital Colorado.

REFERENCES

[1] X. Li and X. Ma, "Acute respiratory failure in covid-19: is it "typical" ards?" *Critical Care*, vol. 24, no. 198, 2020.
[2] I. Tomasic *et al.*, "Continuous remote monitoring of copd patients—justification and explanation of the requirements and a survey of the available technologies," in *Medical and Biological Engineering and Computing*, vol. 56, 2018, pp. 547–569.
[3] S. Transue, S. M. Reza, A. C. Halbower, and M.-H. Choi, "Behavioral analysis of turbulent exhale flows," in *IEEE EMBS International Conference on Biomedical Health Info. (BHI'18)*, March 2018, pp. 42–45.
[4] S. Baker *et al.*, "Determining respiratory rate from photoplethysmogram and electrocardiogram signals using respiratory quality indices and neural networks," *PLoS ONE*, vol. 16, p. e0249843, 04 2021.
[5] A. R. Fekr *et al.*, "Design and evaluation of an intelligent remote tidal volume variability monitoring system in e-health applications," in *IEEE J. of Biomedical and Health Informatics*, 2015.
[6] A. Adhikari *et al.*, "mmflow: Facilitating at-home spirometry with 5g smart devices," in *IEEE (SECON'18)*, 2021, pp. 1–9.
[7] S. Transue, P. Nguyen, T. Vu, and M. Choi, "Real-time tidal volume estimation using iso-surface reconstruction," in *IEEE/ACM Connected Health: Applications, Sys. and Eng. Tech. (CHASE)*, 2016, pp. 209–218.
[8] J. Fei and I. Pavlidis, "Analysis of breathing air flow patterns in thermal imaging," in *IEEE EMBS'06*, 2006, pp. 946–952.
[9] J. Fei, Z. Zhu, and I. Pavlidis, "Imaging breathing rate in the co_2 absorption band," in *Engineering in Medicine and Biology*, 2005.
[10] A. Kwasniewska *et al.*, "Improving accuracy of respiratory rate estimation by restoring high resolution features with transformers and recursive convolutional models," in *IEEE CVPR Workshop*, 2021, pp. 3852–3862.
[11] D. Hanawa *et al.*, "Basic study on non-contact measurement of human oral breathing by using far infra-red imaging," in *Int. Conf. on Telecommunications and Signal Processing (TSP'16)*, 2016, pp. 681–684.
[12] S. Nizami *et al.*, "Comparing time and frequency domain estimation of neonatal respiratory rate using pressure-sensitive mats," in *2017 IEEE International Symposium (MeMeA)*, 2017, pp. 239–244.
[13] A. Abid *et al.*, "Model-based estimation of respiratory parameters from capnography, with application to diagnosing obstructive lung disease," *IEEE Trans. on Biomed Eng*, vol. 64, no. 12, pp. 2957–2967, 2017.
[14] X. Sun *et al.*, "Diagnostic value of volumetric capnography in patients with chronic cough variant asthma," *Clinics (Sao Paulo)*, vol. 75, 2020.
[15] C. Xu *et al.*, "Human exhalation characterization with the aid of schlieren imaging technique," *Building and Environment*, vol. 112, pp. 190 – 199, 2017.
[16] P. Jakkaew and T. Onoye, "Non-contact respiration monitoring and body movements detection for sleep using thermal imaging," *Sensors*, vol. 20, p. 6307, 2020.
[17] C. B. Pereira *et al.*, "Noncontact monitoring of respiratory rate in newborn infants using thermal imaging," *IEEE Transactions on Biomedical Engineering*, vol. 66, no. 4, pp. 1105–1114, 2019.
[18] J. L. Barron *et al.*, "Performance of optical flow techniques," *Intl. Journal on Comp. Vision*, vol. 12, no. 1, pp. 43–77, 1994.
[19] Berthold K.P. Horn *et al.*, "Determining optical flow," vol. 17, 1981.
[20] J. Y. An *et al.*, "Non-contact diagnosis of sleep breathing disorders using infrared optical gas imaging: a prospective observational study," *Scientific Reports*, vol. 12, no. 1, p. 21052, Dec 2022.
[21] Y. Wu, "Optical flow and motion analysis," Advanced Computer Vision Notes Series 6 (3-7), 2001.
[22] J. Marzat *et al.*, "Real-time dense and accurate parallel optical flow using cuda," in *The 17th Intl. Conf. (WSCG '2009)*, 2009, pp. 105–112.
[23] R. Foundation. (2021) Raspberry pi. [Online]. Available: raspberrypi.org
[24] J. Appleyard, T. Kociský, and P. Blunsom, "Optimizing performance of recurrent neural networks on gpus," *CoRR*, vol. abs/1604.01946, 2016.
[25] NVIDIA. Corporation. (2018) CuDNN. [Online]. Available: https://docs.nvidia.com/deeplearning/cudnn/
[26] F. Chollet *et al.* (2015) Keras. [Online]. Available: https://github.com/fchollet/keras
[27] M. Abadi *et al.*, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015, software available from tensorflow.org. [Online]. Available: https://www.tensorflow.org/
[28] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *International Conference on Learning Representations*, 12 2014.