NON-CONTACT MONITORING OF BREATHING FUNCTION USING INFRARED IMAGING

R. Murthy and I. Pavlidis
Department of Computer Science
University of Houston
Houston, TX, 77204, USA
http://www.cs.uh.edu

Technical Report Number UH-CS-05-09
April 09, 2005

Keywords: Breathing function, thermal imaging, Jeffrey's divergence measure, method of moments, contact-free polygraphy, sleep studies.

Abstract

We have developed a novel method for non-contact measurement of breathing function. The method is based on statistical modeling of dynamic thermal data captured through a highly sensitive infrared imaging system. The air that is breathed out has higher temperature than the typical background of indoor environments (e.g., walls). Therefore, the particles of the expired air emit at a higher power than the background, a phenomenon that is captured as a distinct thermal signature in the infrared imagery. There is significant technical difficulty in computing this signature, however, because the phenomenon is of very low intensity and transient nature. We use an advanced statistical algorithm based on multi-Normal data representation, the method of moments, and the Jeffreys divergence measure to address the problem. In experimental tests, we were able to compute correctly the breathing waveforms in eight (8) infrared video clips of three (3) subjects at distances ranging from 6-8 feet. The results were compared with ground-truth data collected concomitantly with a traditional contact sensor. Our experiments demonstrated the promise of this modality, which may find applications among others in the next generation contact-free polygraphy and in sleep studies.

[1]This work was supported by DARPA.

NON-CONTACT MONITORING OF BREATHING FUNCTION USING INFRARED IMAGING

R. Murthy and I. Pavlidis

Abstract

We have developed a novel method for non-contact measurement of breathing function. The method is based on statistical modeling of dynamic thermal data captured through a highly sensitive infrared imaging system. The air that is breathed out has higher temperature than the typical background of indoor environments (e.g., walls). Therefore, the particles of the expired air emit at a higher power than the background, a phenomenon that is captured as a distinct thermal signature in the infrared imagery. There is significant technical difficulty in computing this signature, however, because the phenomenon is of very low intensity and transient nature. We use an advanced statistical algorithm based on multi-Normal data representation, the method of moments, and the Jeffreys divergence measure to address the problem. In experimental tests, we were able to compute correctly the breathing waveforms in eight (8) infrared video clips of three (3) subjects at distances ranging from 6-8 feet. The results were compared with ground-truth data collected concomitantly with a traditional contact sensor. Our experiments demonstrated the promise of this modality, which may find applications among others in the next generation contact-free polygraphy and in sleep studies.

Index Terms

Breathing function, thermal imaging, Jeffrey's divergence measure, method of moments, contact-free polygraphy, sleep studies

I. INTRODUCTION

MONITORING of breathing function has applications among others in polygraphy, sleep studies, sport training, early detection of sudden infant death syndrome in neonates, and patient monitoring.

Various contact measurement methods have been developed for estimating the breathing rate of a subject. George B. Moody, et al. developed a contact modality in which numerous Electrocardiogram (ECG) electrodes and sensors are attached to the subject [1]. The principle of operation is based on the fact that the heart rate is typically modulated by breathing, a phenomenon known as sinus arrhythmia [2]. Therefore, a signal corresponding to the heart function contains breath information, which is filtered out using band-pass filters.

As an improvement over the ECG method, the BioMatt method [3] was developed in Finland by a group of researchers who were studying sleep disorders. BioMatt performs measurements of vital signs, such as breathing and cardiac activity without electrodes. Initially, BioMatt could not distinguish motion that was due to breathing versus cardiac activity or body movement. Later, Larson developed a signal processing technique to separate out the components of the BioMatt signal [4].

Photoplethysmography (PPG) is a variant method of the ECG, developed to measure blood volume changes in living tissues by absorption or scattering of near-infrared radiation. This modality consists of an infrared Light Emitting Diode (LED) and a photodiode which can be clamped to the ear lobes, thumbs, or toes. It is advantageous in that it is portable, compact, and needs very little maintenance. The measurement of blood volume changes by PPG depends on stronger absorption of near-infrared light by blood when compared to other superficial tissues [5]. The amount of backscattered light corresponds to the variation of the blood volume. As in ECG, the breath waveform is separated from the cardiac signal through various methods that have been developed [6], [7]. However, using heart function as a basis for acquiring the breathing waveform is unreliable since sinus arrhythmia is not present in all individuals. Control of cardiac activity by breathing depends on the age and medications administered to subjects.

Other contact modalities are capable of measuring directly the breathing signal. An example of such modality is the abdominal strain gauge transducer [8] that is strapped around the subject's chest and measures the change in thoracic or abdominal circumference while breathing. Another example is a thermistor measuring nasal air temperature variation as an indication of breathing [9].

The disadvantage of all the aforementioned technologies is that they require close contact with the subject, which in certain cases may be quite uncomfortable and awkward (e.g., abdominal transducer). A contact-free but

active technology called Radar Vital Signs Monitor (RVSM) [10] was developed in 1996 to monitor the performance of Olympic athletes. The RVSM detects breathing-induced movement of the chest based on the Doppler phenomenon. It measures breath at distances of up to 15 feet behind an 8 inch hollow concrete or wooden wall. A Radar Flashlight [11] was built to make use of this capability in assisting law enforcement personnel to detect individuals hidden behind walls. In 2000, RVSM was used in non-contact polygraphy [12]. The disadvantage of this technique is that motion artifacts corrupt breath signals and specialized frequency filters need to be used to separate them.

In 2000, infrared imaging proved its potential in deception detection when thermal image analysis was used by Pavlidis et al. to detect facial patterns of stress at a distance [13]. A little later Pavlidis et al. used infrared imaging to compute periorbital perfusion as a replacement of the corresponding polygraph channel that uses finger contact sensing [14], [15]. The proposed use of infrared imaging for computing breathing function may also replace the corresponding polygraph channel that uses abdominal transducer. Incremental replacement of contact channels with non-contact ones may prove very effective in the field of polygraphy, where it is essential that subjects feel as comfortable as possible during examination.

Moreover, highly automated, non-contact monitoring of breathing function may have a significant impact on certain biomedical applications. For example, in sleep studies this new methodology will enable monitoring of sleep apnea with minimal or no wiring of the subject, potentially at his\her home and not in the lab. This will not only improve the subject's comfort but also facilitate much more sustained monitoring than is currently feasible.

The use of infrared imaging for measuring breathing function is based on the fact that the exhaled air has temperature higher than the typical background of indoor environments. This creates a discriminating thermal signature that can be captured through an infrared imaging sensor. The phenomenon is quasi-periodic and can be quantified using either statistics or calculus. From the statistical point of view one can model the breathing cycles as multi-Normal distributions – one with "cold" temperatures corresponding to inhalation and one with "hot" temperatures corresponding to exhalation. From the Calculus point of view one can model the quasi-periodicity of breathing through Fourier analysis.

In the present article, we describe a statistically based methodology for quantifying breathing rate on infrared imaging data. Alternative methodologies, like Fourier analysis, can be used but are not addressed in our present work. Our goal is to open a new line of research by demonstrating the feasibility of monitoring breathing function in a highly automated and non-contact fashion.

In the rest of the paper we describe briefly the physiology of breathing (section II). Then, in section III we refer to the visual tracking mechanism that enables consistent breathing measurements in the presence of subject motion. In section IV we describe our breath visualization scheme, which is of paramount importance during the training phase of our measurement algorithm. In section V, we explain in great detail the statistical algorithm that performs the breathing measurement on the infrared imaging data. In sections VI and VII we outline our experimental design and results. We conclude our paper in sections VIII and IX, where we discuss the strong and weak points of our methodology, its prospects, and our planned work for the future.

II. BREATHING FUNCTION

Respiration in a man involves three well defined stages [16]. The first stage called breathing comprises of inspiration, which is the process of taking oxygenated air into the lungs and expiration, which is discharging out air rich in carbon dioxide. The second stage involves the transport of the oxygen to the cells of the body using the heart and the vascular system. The third stage is called cellular respiration where oxygen is used in the process of generating energy for physiological activities.

In our study, we are interested in monitoring breathing using infrared imaging. The breathing cycle consists of inspiration, expiration, and post-expiratory pause. During quiet breathing, inspiration begins due to negative pressure created inside the chest cavity by the contraction of the diaphragm. Expiration is a passive process where the air flow occurs due to the elastic recoil property of the lungs. The post-expiratory pause is caused when there is equalization of the pressures inside the lungs and the atmosphere.

Breathing cycle is defined as the time interval between the beginning of inspiration and the end of post-expiratory pause. During quiet breathing, the breathing rate may vary from 12-20 breaths/min and after physical activity, 30-40 breaths/min in healthy individuals. Figures 1(a) and 1(b) show typical duration of the three phases during quiet breathing and after physical activity respectively.

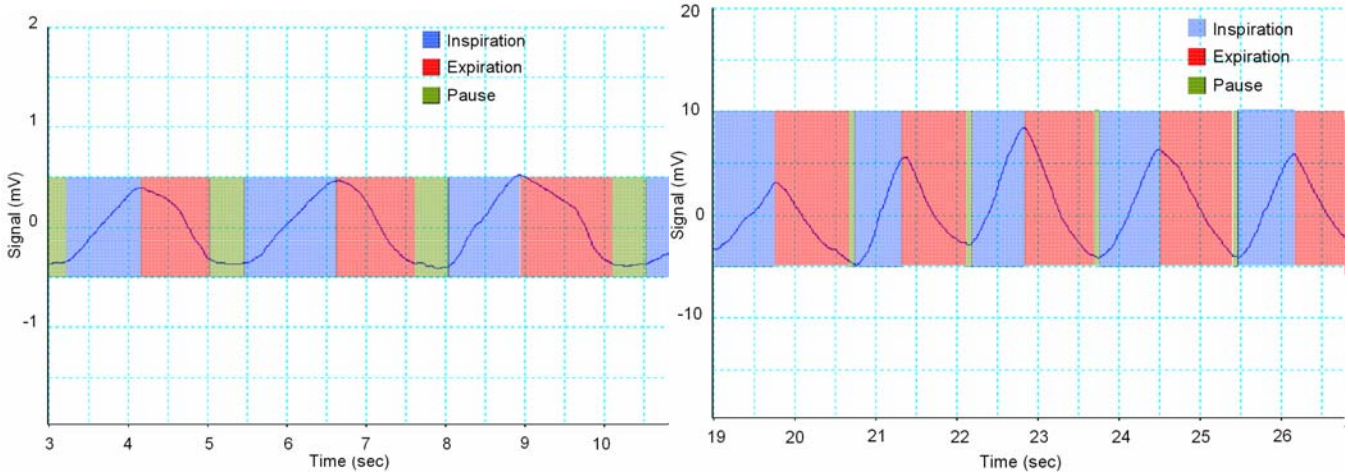

Fig. 1. Output from a piezo-respiratory belt transducer showing the three breathing phases during (a) quiet breathing and (b) after exercise.

During quiet breathing, the duration of post-expiratory pause is comparable to that of inspiration and expiration. After a person undergoes physical exertion, the post-expiratory duration reduces considerably and in some cases, this phase may even cease to exist.

III. TRACKING THE REGION OF INTEREST (ROI)

We define as the Region of Interest (ROI) *R* the region in the background, where there is possible presence of respiratory airflow. It is in this small image region that our statistical algorithm is applied. The ROI is characterized by its size, shape, and position. Over time the size and shape of ROI remain the same but its position changes to cope with the subject's motion (tracking). We have experimented with different ROI sizes and we will give more details about the optimal determination of this parameter in section VII. In this section, we will address the issues of ROI shape and dynamic positioning.

For simplicity, *R* was chosen to be a rectangular region. Typically, subjects are breathing through the nasal cavity, which results in a downward airflow profile. Breathing through the mouth is less prevalent and results in horizontal airflow profile. In our data set we observed downward airflow profiles (Figure 2(a)) in eight (8) video clips and horizontal profile (Figure 2(b)) in one (1) video clip. Hence, we chose a rectangular region *R* arranged in a longitudinal fashion to closely match the prevalent downward profile of airflow. Our experiments have also shown that this shape still worked quite well on the video clip featuring the horizontal airflow profile (accuracy of 97.74%).

Fig. 2. Respiratory airflow profiles: (a) downward and (b) horizontal.

The respective aims of the initial positioning and tracking algorithms are to provide the user with an approximate position of the ROI in the vicinity of the nasal-mandible region and follow it automatically throughout the breathing rate computation process. Technical details of the initial positioning and tracking algorithms can be found in [17]. In summary, the initial positioning algorithm locates the tip of the subject's nose and places the ROI underneath it, in a position that is between the nostrils and the mouth (Figure 3). This placement works for the

typical monitoring scenario of a subject imaged at a side view. For different monitoring scenarios the above heuristic approach tends to place *R* incorrectly. But, the Graphical User Interface (GUI) gives options to move the ROI around the image so that it can be placed within the airflow with just a mouse click. In such adjustments the user is aided by the breath visualization tool (see Section IV). Hence, the initial positioning of ROI is semi-automatic wherein, the algorithm gives the approximate position of *R* and then the user may need to move it to a better position, within the field of the airflow. By contrast, the tracking algorithm is completely automatic since it tracks *R* with respect to the tip of the nose in subsequent video frames without any feedback from the user.

IV. VISUALIZATION OF BREATH

Visualization is important among other reasons for adjusting the initial ROI position and for training the statistical algorithm. Since the thermal signature of breath is not very strong, we have to apply image processing techniques to visually perceive breath in infrared video frames. Specifically, we apply the following operations on the video clip frames:

1. Otsu's adaptive thresholding [18];
2. Differential Infrared Thermography (DIT) [19]; and
3. Image opening [20].

Otsu's adaptive thresholding is applied to segment the skin region from the background. Then, in the background region, we apply DIT to generate a breath mask of all pixels whose temperature has increased beyond a preset threshold. This operation makes sure that the color-map is applied only to expiration frames in which DIT senses an increase in ROI pixel temperatures above the preset value. The result is a highly contrasting effect of no color during inspiration versus vivid color during expiration (see Figure 3(a) and (b)). As a final step, an image open operation is applied on the output binary mask of DIT to improve breath visualization.

(a) (b)

Fig. 3. Visualization of breath during (a) non-expiration and (b) expiration. The ROI is anchored just under the subject's nose tip.

The visualization scheme works well only for a limited period, as the temperature distributions of inspiration and expiration tend to drift over time for physiological and other reasons. DIT cannot handle this distribution drift well. If it did, it could have also been used to measure breathing rate in the place of the more sophisticated statistical algorithm (see section V). However, the short time window of good visualization performance is adequate for adjusting the initial ROI position and training the statistical algorithm.

V. STATISTICAL METHODOLOGY

Integral to breathing rate computation using infrared imaging is the labeling of frames as expiratory and non-expiratory. For this purpose, we have adopted a statistical methodology based on multi-Normal distributions, the method of moments, and the Jeffreys' divergence measure [21]. We identify two phases in the statistical method: training and testing. We describe both phases in subsequent sections.

First, we ascertain the Normal nature of the temperature distributions in the ROI for the various breathing stages (Figure 4). Therefore, we can represent ROI distributions by their mean μ and variance σ^2 only. Our method lumps together inspiration and post-expiratory pause phases, since the thermal signatures of these two are almost identical. We designate the combined stage as non-expiratory phase.

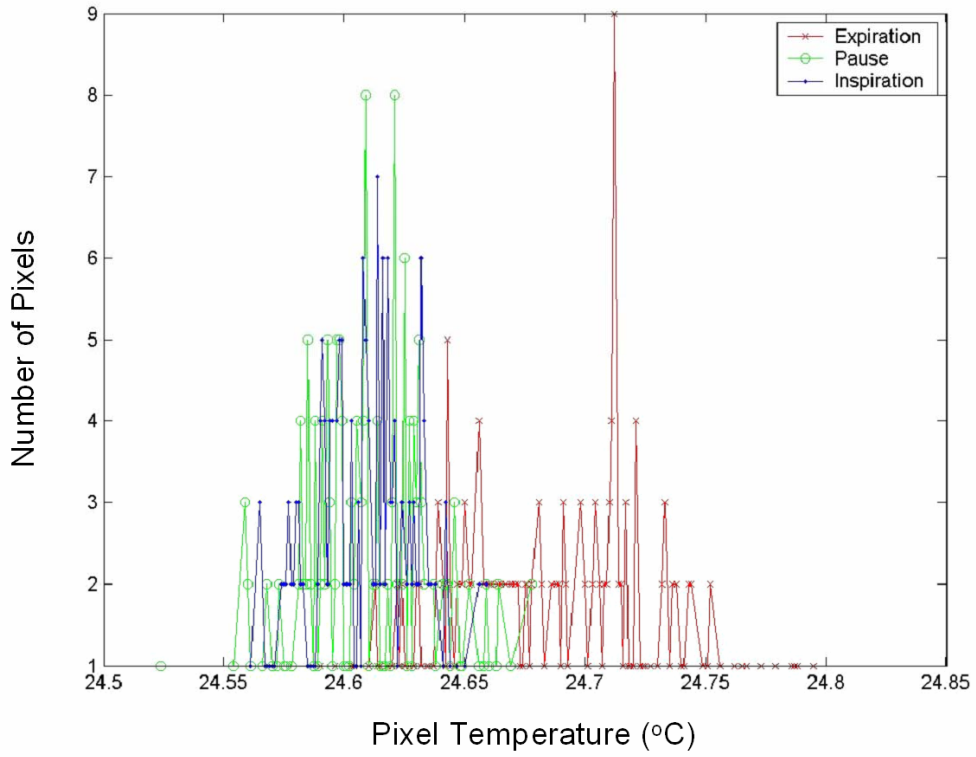

Fig. 4. Experimental temperature distributions for expiration, post-expiratory pause, and inspiration. The Normal nature of the distributions and overlapping between post-expiratory pause and inspiration are evident.

A. Training Phase

The algorithm runs through a training phase to generate estimates of the expiration and non-expiration distributions from the first few video frames. These estimates are then used to label pixels as expiratory or non-expiratory in the initial video frame of the testing phase.

We use a variant of the *K*-means clustering algorithm [22] to generate training data. Our objective is to form $K = 2$ representative distributions through an iterative process; a "hot" one for expiration and a "cold" one for non-expiration. Initially, we specify as the expiration distribution $D_{e,0}$ the one with the "hottest" mean temperature $\mu_{e,0}$ in the training set; respectively, we specify as the non-expiration distribution $D_{i,0}$ the one with the "coldest" mean temperature $\mu_{i,0}$ in the training set:

$$\mu_{e,0} = \max_{1 \le j \le M} \{\mu_j\}, \tag{1}$$

$$\mu_{i,0} = \min_{1 \le j \le M} \{\mu_j\}, \tag{2}$$

where, $N(\mu_j, \sigma_j^2)$, $1 \le j \le M$ is the set of temperature distributions for ROI R corresponding to the first $M = 100$ training frames of infrared video. We sort the distributions in ascending order with respect to their means, in order to facilitate the iterative process.

On every step $j, 1 \le j \le M$ we find the statistical distance of distribution $D_j \square N(\mu_j, \sigma_j^2)$ from $D_{e,j-1}$ and $D_{i,j-1}$. For this purpose, we use the Jeffreys' divergence measure as follows:

$$J(D_{e,j-1}, D_j) = \frac{1}{2}\left(\frac{\sigma_{e,j-1}}{\sigma_j} - \frac{\sigma_j}{\sigma_{e,j-1}}\right)^2 + \frac{1}{2}\left(\frac{1}{\sigma_j^2} + \frac{1}{\sigma_{e,j-1}^2}\right)\left(\mu_{e,j-1} - \mu_j\right)^2, \tag{3}$$

$$J(D_{i,j-1}, D_j) = \frac{1}{2}\left(\frac{\sigma_{i,j-1}}{\sigma_j} - \frac{\sigma_j}{\sigma_{i,j-1}}\right)^2 + \frac{1}{2}\left(\frac{1}{\sigma_j^2} + \frac{1}{\sigma_{i,j-1}^2}\right)\left(\mu_{i,j-1} - \mu_j\right)^2. \tag{4}$$

The Jeffreys' Divergence measure is a symmetric form of the Kullback-Leibler Divergence measure. It is a function of the means and standard deviations of the two distributions being compared. Hence, it is an appropriate distance measure for bivariate distributions. We choose the winner distribution $D_{w,j-1}$ ($w = e$ or i) at step j as the one whose Jeffreys' distance from the training population D_{j-1} is the smallest. The mean and variance of the winning distribution are then updated at each step as follows:

$$\mu_{w,j} = \frac{\mu_{w,j-1} + \mu_j}{2}, \tag{5}$$

$$\sigma^2_{w,j} = \frac{\sigma^2_{w,j-1} + \sigma^2_j}{2}. \tag{6}$$

The loser distribution retains its previous values.

We iterate this process for all the training populations except those that were marked as initial clusters. At the end of the process we obtain the estimates of the two distributions corresponding to the expiration and non-expiration phases of the breathing cycle. Figure 5 depicts our *K*-means training method.

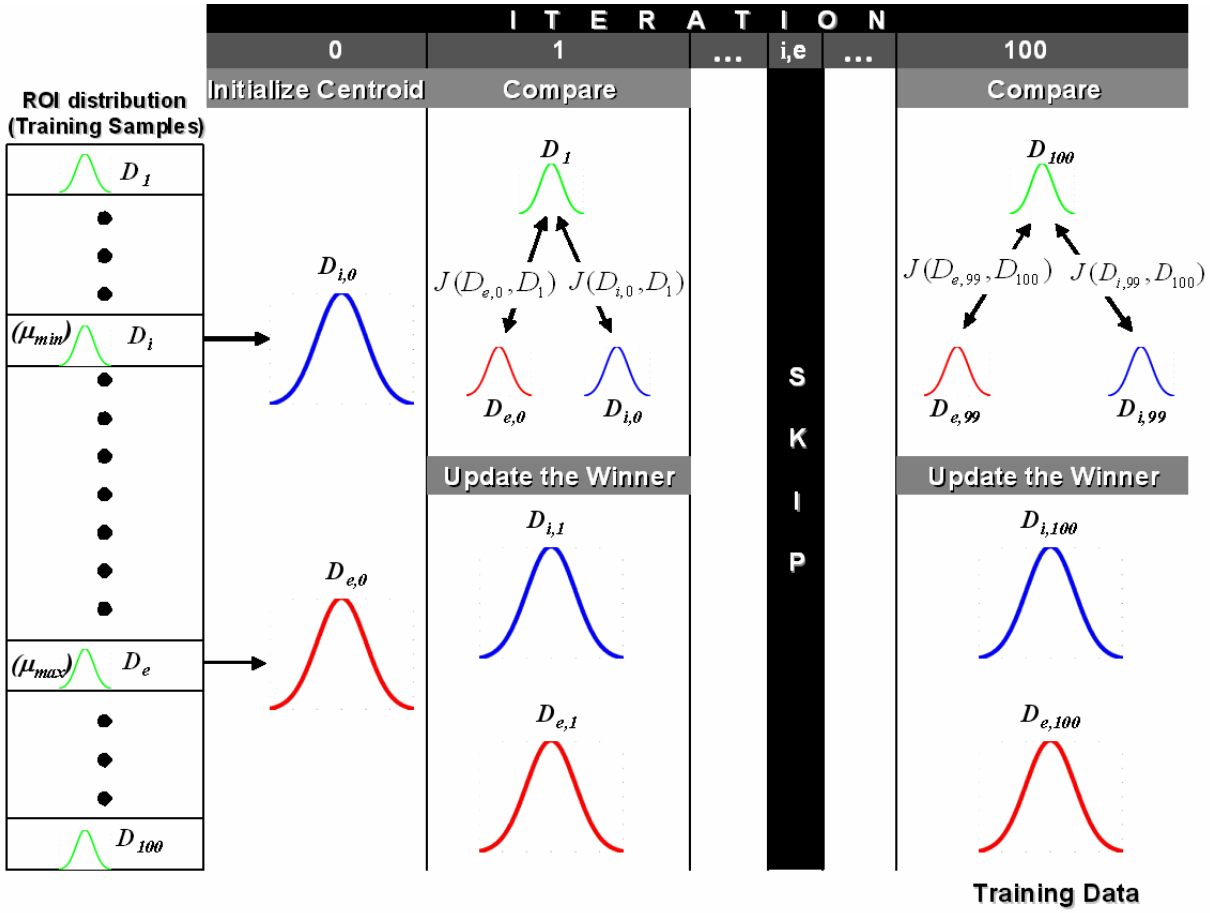

Fig. 5. Iterations in the K-means training data acquisition method. The resulting expiration and non-expiration distributions at the end of the M=100th iteration are used to jump start the testing phase.

B. Testing Phase

During the testing phase we represent at frame *t* each pixel x_t in region *R* as a mixture of two distributions:

$$f(x_t) \sim \pi_{e,t} N(\mu_{e,t}, \sigma^2_{e,t}) + \pi_{i,t} N(\mu_{i,t}, \sigma^2_{i,t}) \tag{7}$$

$f_e(x_t) \sim N(\mu_{e,t}, \sigma^2_{e,t})$ is the Normal expiration distribution, $f_i(x_t) \sim N(\mu_{i,t}, \sigma^2_{i,t})$ is the Normal non-expiration distribution, and $\pi_{e,t}$ and $\pi_{i,t}$ are their respective weights in the mixture satisfying the criterion:

$$\pi_{e,t} + \pi_{i,t} = 1.$$

In the beginning of the testing phase ($t = 0$) the distributions for non-expiration and expiration are equiprobable with $\pi_{e,t} = \pi_{i,t} = 0.5$ and are parameterized by the respective means and variances that we computed during the training phase. Therefore, every pixel in region R is represented as having the following starting distribution:

$$f(x_0) \sim 0.5N(\mu_{i,0}, \sigma^2_{i,0}) + 0.5N(\mu_{e,0}, \sigma^2_{e,0}) \tag{8}$$

At time $t > 0$ and for pixel x_t we compare the incoming temperature value from the sensor with the estimated distribution from the previous frame at time $t - 1$. For this comparison to be effective we consider that the incoming temperature θ_{x_t} can be associated to a Normal distribution $g(\theta_{x_t}) \sim N(\mu_{g,t}, \sigma^2_{g,t})$, where $\mu_{g,t} = \theta_{x,t}$ and $\sigma_{g,t} = NEDT$. For the camera model that we use $NEDT = 0.01^0\,C$.

We compute the Jeffreys' divergence measures between the incoming distribution $g(\theta_{x_t})$ and the existing non-expiration $f_i(x_{t-1})$ and expiration $f_e(x_{t-1})$ distributions respectively. Specifically:

$$J(f_i(x_{t-1}), g(\theta_{x_t})) = \frac{1}{2}\left(\frac{\sigma_{i,t-1}}{\sigma_{g,t}} - \frac{\sigma_{g,t}}{\sigma_{i,t-1}}\right)^2 + \frac{1}{2}\left(\frac{1}{\sigma^2_{g,t}} + \frac{1}{\sigma^2_{i,t-1}}\right)\left(\mu_{i,t-1} - \mu_{g,t}\right)^2, \tag{9}$$

$$J(f_e(x_{t-1}), g(\theta_{x_t})) = \frac{1}{2}\left(\frac{\sigma_{e,t-1}}{\sigma_{g,t}} - \frac{\sigma_{g,t}}{\sigma_{e,t-1}}\right)^2 + \frac{1}{2}\left(\frac{1}{\sigma^2_{g,t}} + \frac{1}{\sigma^2_{e,t-1}}\right)\left(\mu_{e,t-1} - \mu_{g,t}\right)^2. \tag{10}$$

We consider that the incoming distribution is closer to the existing distribution that features the minimum Jeffreys' divergence measure. We call this the winning distribution $f_w(x_{t-1})$ and the other the losing distribution $f_l(x_{t-1})$. Based on this information we update the parameters of the mixture following the method of moments. Specifically, we update the weights for both distributions and the mean and variance of the winning distribution only. The mean and variance of the losing distribution remain the same.

The weights of the winning and losing distributions are updated as follows:

$$\pi_{w,t} = (1-\rho)\pi_{w,t-1} + \rho, \tag{11}$$

$$\pi_{l,t} = (1-\rho)\pi_{l,t-1}. \tag{12}$$

The mean and variance of the winning distribution are updated as follows:

$$\mu_{w,t} = (1-\rho)\mu_{w,t-1} + \rho\mu_{g,t}, \tag{13}$$

$$\sigma^2_{w,t} = (1-\rho)\sigma^2_{w,t-1} + \rho\sigma^2_{g,t} + \rho(1-\rho)(\mu_{g,t} - \mu_{w,t-1})^2. \tag{14}$$

The parameter ρ is a learning parameter that is computed from the following formula [23]:

$$\rho = e^{-\frac{1}{2}\left[\frac{\frac{1}{2}(\mu_{g,t} - \mu_{w,t-1})}{\sigma_{w,t-1}}\right]}. \tag{15}$$

Figure 6 is a visualization example of how the incoming distribution may affect the existing distribution. The updated data acts as the new estimate for the corresponding pixel in the next incoming frame.

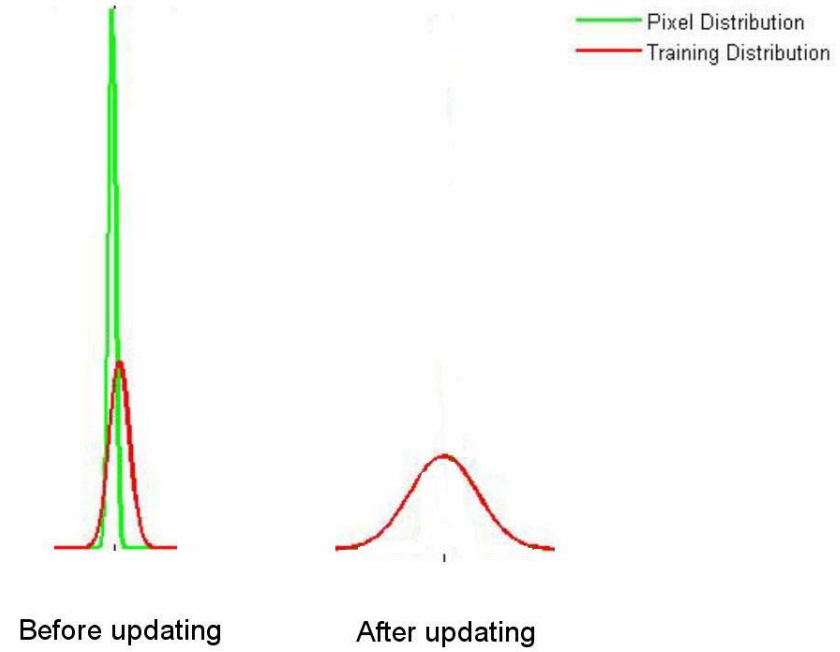

Fig. 6. Overview of the statistical methodology for labeling infrared video frames as expiration or non-expiration.

The pixel x_t is given the label of the distribution with the highest updated weight. A count is kept of the number of expiration C_e and non-expiration C_i pixels in region *R* at time *t*. Once all the pixels in region R are processed the frame gets the label of the most frequently occurring pixel label. For example, if $C_i > C_e$ the frame is labeled as non-expiration. Otherwise, the frame is labeled as expiration. Figure 7 shows the flow of control and data through the statistical algorithm.

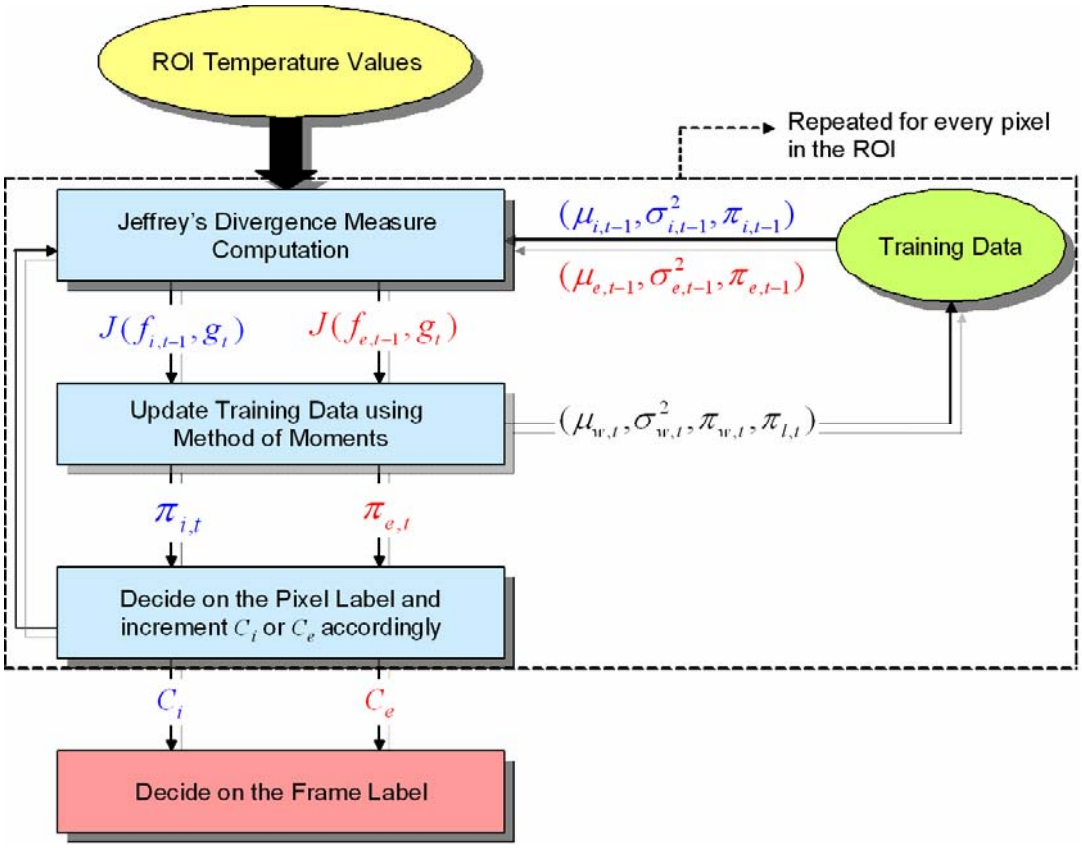

Fig. 7. Comparison of a training distribution before and after the application of the Method of Moments.

The breathing rate computation algorithm keeps track of the frame labels and continuously updates the time of the current breathing cycle T_c by using the current timestamp T_n and the previous timestamp T_{n-1}. The initial run of similar frame labels (Figure 8) is skipped since the testing phase might have started in the middle of a cycle. The algorithm keeps track of the two sub-cycles in the breathing cycle and once it detects the beginning of the next cycle, the breathing rate of the current cycle is computed in cycles/min using the formula:

$$Rate = \frac{60}{T_c} cycles / \min. \tag{16}$$

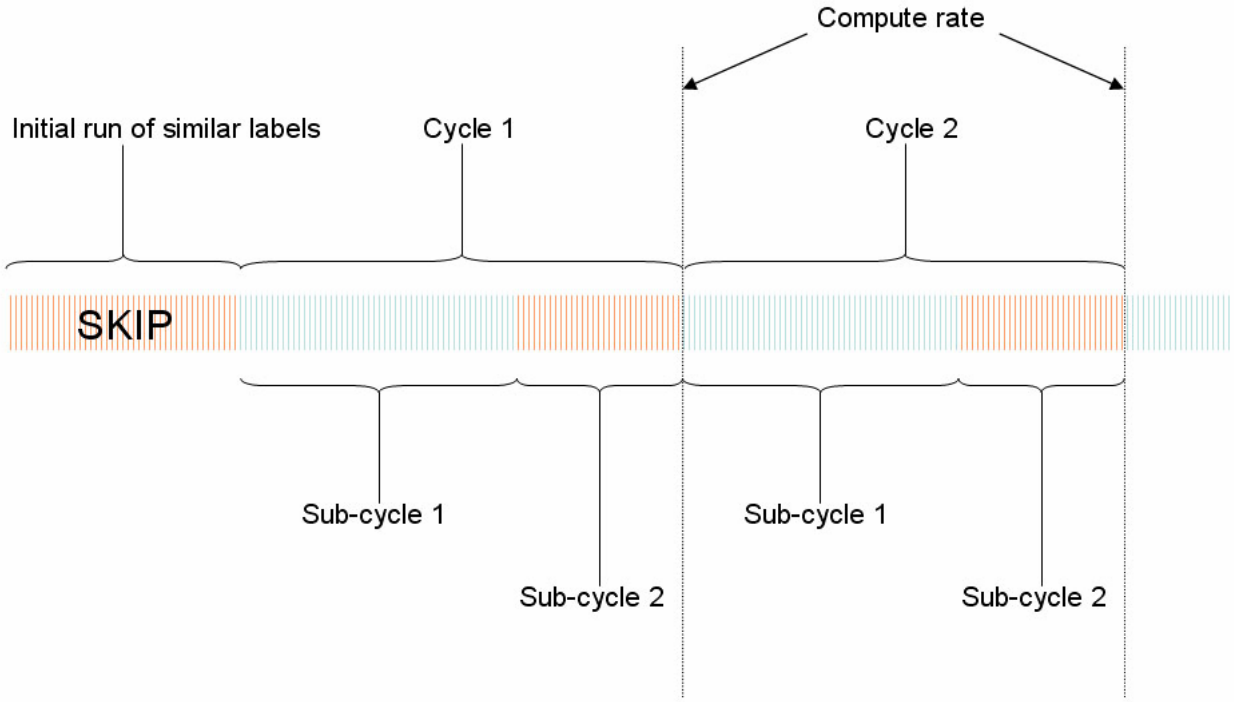

Fig. 8. Example labeling of the video frame line and computation of the breathing rate.

VI. EXPERIMENTAL SETUP

We used a cooled mid-wave infrared Phoenix camera with a spectral range of $3.0-5.0\mu m$ (Indigo Systems, Goleta, CA) equipped with a $50mm$ lens (Figure 9). The focal plane array of the camera is $FPA = 640 \times 512$ pixels in size and has thermal sensitivity $NEDT = 0.01^0 C$. An external black body (Santa Barbara Infrared, Santa Barbara, CA) was used to calibrate the camera. Infrared video frames were acquired at a rate of 31 frames per second.

Fig. 9. Phoenix mid-wave infrared camera equipped with a $50mm$ lens.

We captured the profile view of the subjects' face and respiratory airflow from a distance of 6-8 feet (Figure 10). A piezo-strap transducer (Figure 11(a)) wrapped around the subject's diaphragm was measuring the thoracic circumference during expiration and non-expiratory phase. The transducer was sending its signal to a PowerLab Data Acquisition System (ADI Instruments, Australia – Figure 11(b)). This was the gold standard that we used for ground truthing the infrared imaging measurements.

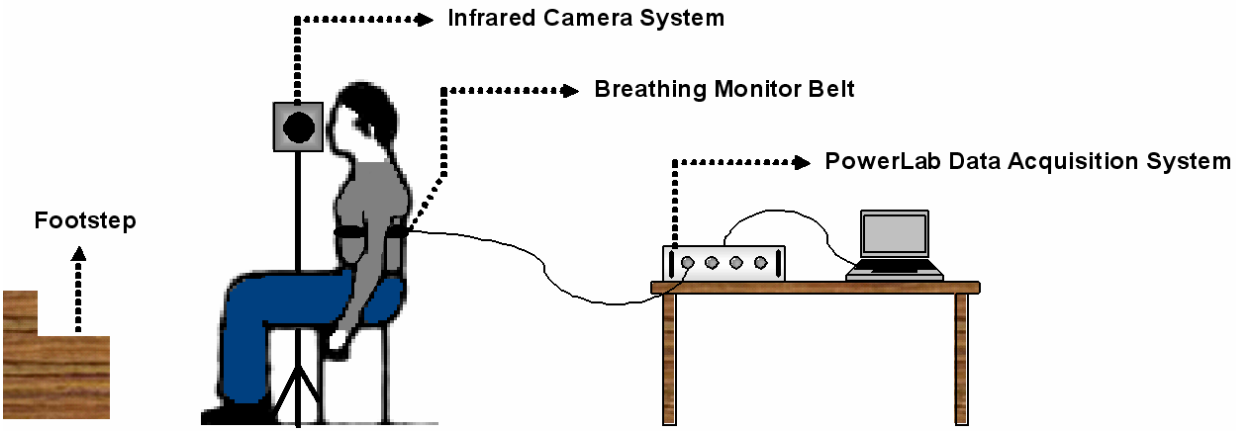

Fig. 10. Experimental setup with the profile view of the subject's face captured by a mid-wave IR Phoenix camera and ground truth data recorded concomitantly using a piezo-strap transducer and a PowerLab Data Acquisition System.

(a) (b)

Fig. 11. (a) Piezo-strap transducer. (b) PowerLab Data Acquisition System.

VII. EXPERIMENTAL RESULTS

In this section we investigate experimentally different aspects of the breathing function and our method's parameters. We also describe the performance of our non-contact methodology against ground-truth measurements taken by the PowerLab/4SP ADI instrument.

When air is breathed in, it gets warmed up during its passage into the respiratory system and during its brief stay in the lungs. Figure 12 shows the plot of mean ROI temperature along the timeline for one of the subjects in our data set. From the plot, we observe that the ROI temperature increases by around 0.1^{o}C during expiration. Figure 13 shows the plot of ROI variance along the timeline for one of the subjects in our data set. From the plot we observe that the ROI variance during expiration is quadrupled. This is due to the fact that within ROI there are clusters of "hot" air molecules interspersed with "cold" air resident in the room. The "hot" air molecules are the ones recently expired through the nostrils.

Our experimental protocol called for measurements during the following phases:

1. Breathing while the subject rests in a chair;
2. Breathing after the subject has stepped on and off a foot step 30 times in 30 sec;
3. Breathing after the subject rested for 10 min so that his physiology reverts back to baseline; and
4. Breathing after the subject has stepped on and off a foot step 60 times in 60 sec.

We measured the performance of our method on eight (8) thermal clips captured from three (3) subjects. The thermal clips were 1000 frames long. The first 100 frames of each clip were used for training and the remaining 900 frames for testing. Only for subject #2 we were able to perform the infrared imaging measurements for all four

phases of the experimental protocol. For the other two subjects, we had to discard some of the measurements, because of technical problems in synchronizing the camera with the ADI vital signs monitor.

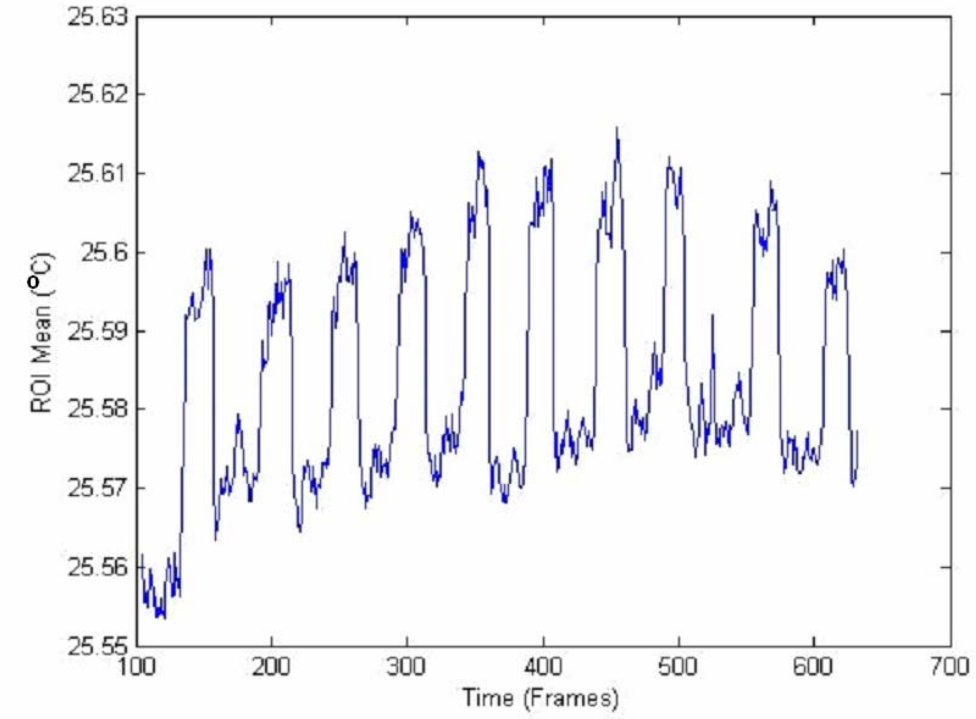

Fig. 12. Plot of the mean pixel temperature in the ROI along the timeline.

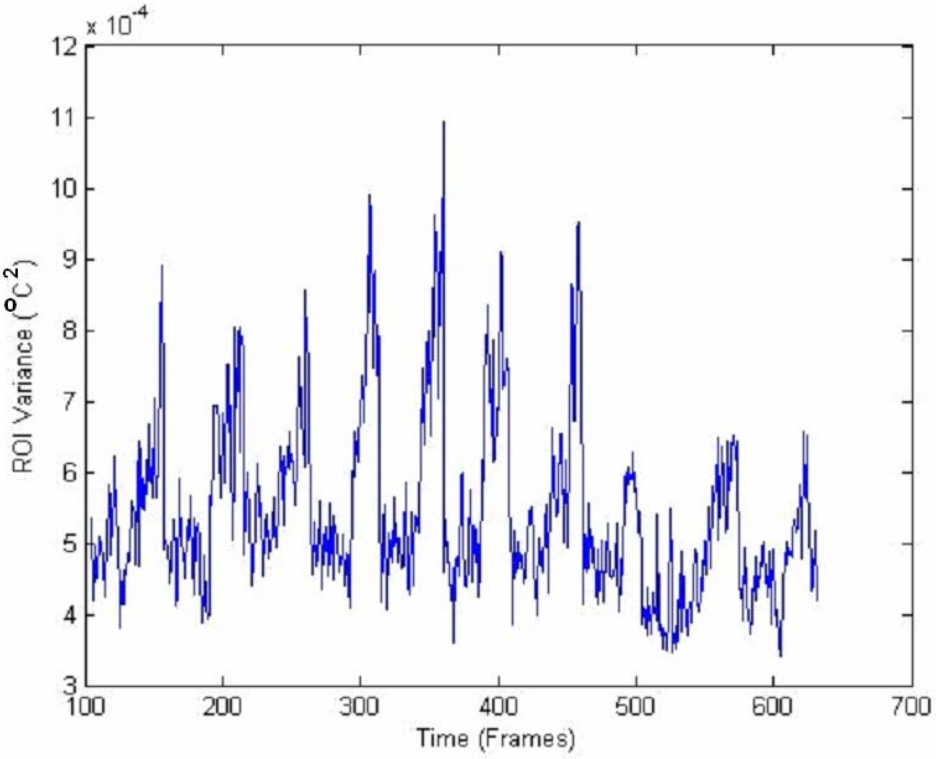

Fig. 13. Plot of the variance of pixel temperature in the ROI along the timeline.

Since the performance of the method clearly depends on the size of the ROI where the statistical computation is taking place, we have experimented with three different ROI sizes: small (7×3 pixels), medium (21×9 pixels), and

large (63×27 pixels). From the experimental results in Table 1 we observe that the medium size ROI outperforms the other two sizes. The interesting fact is that there is a clear breakdown in performance when the ROI size is getting large. In such a case a significant number of the ROI pixels are background and not expiratory pixels. As a result they bias the ROI labeling towards non-expiration and the accuracy drops. The absolute ROI sizes are of course dependent on the optics. In our specific experimental scenario we recorded from a distance of 6-8 feet with a 50*mm* lens and the ROI sizes we cite are commensurate to this optical arrangement.

TABLE I
EXPERIMENTAL RESULTS

Subject #	Video Clip #	Small ROI	Medium ROI	Large ROI
1	1	98.19	94.32	57.71
	2	92.59	96.14	83.03
	3	97.50	97.74	97.31
2	1	94.76	96.70	67.34
	2	97.91	93.36	78.26
	3	99.23	99.31	56.65
	4	98.06	99.05	97.95
3	1	87.71	94.82	67.62
Average % Accuracy		**95.74**	**96.43**	75.73

Accuracy of breathing rate measurements based on infrared imaging. Results are shown for three (3) different ROI sizes. The accuracy was determined by ground-truthing the imaging measurements against concomitant measurements performed with the ADI vital signs monitor.

Figure 14, Figure 15, Figure 16, and Figure 17 show the variation of breathing rate computed through the infrared imaging method for the four (4) phases of the experimental protocol for subject #2. They also show the comparison between the mean breathing rate computed through the infrared imaging method and the mean ground truth rate measured by the ADI vital signs monitor. From the plots we observe that the breathing rate increases from 12-20 cycles/min during rest to 30-40 cycles/min after the brief exercise. The accuracy of the infrared imaging method appears to be consistent during rest as well as during active periods.

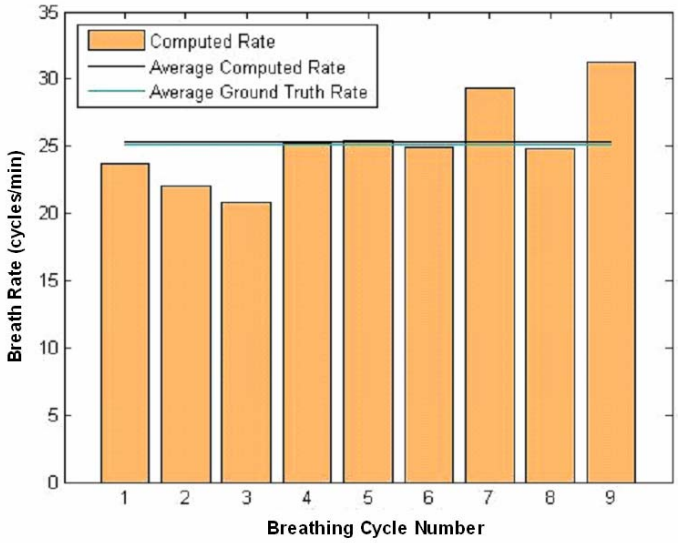

Fig. 14. Breathing rate measurements during initial resting phase for subject #2.

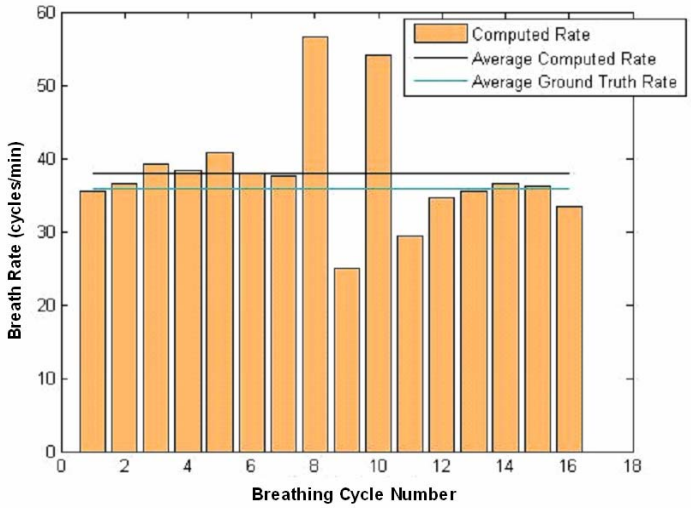

Fig. 15. Breathing rate measurements after 30 sec of exercise for subject #2.

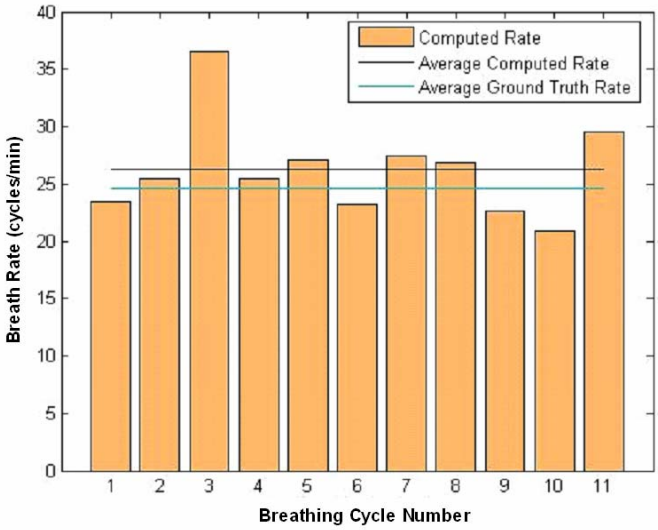

Fig. 16. Breathing rate measurements after 10 min rest for subject #2.

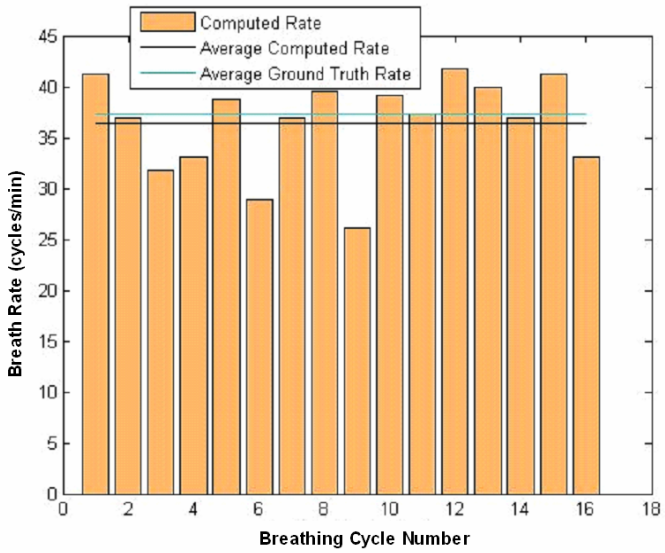

Fig. 17. Breathing rate measurements after 60 sec of exercise for subject #2.

The ground truth signal has output proportional to the expansion (zero level signal) and relaxation (positive level signal) of the breathing monitor belt during non-expiration and expiration respectively. The signal computed from the infrared imaging method has output labeled either as non-expiration or expiration. To make comparison between ground truth data and algorithmic results easier, we have digitized both signals by assigning a zero level signal to non-expiration and a positive level signal to expiration. In addition, we have assigned a negative signal level to frames used for acquiring training data. In Figure 18, we can observe that the cycles detected by the infrared imaging method are slightly out of phase with the ground truth cycles. This accounts for the small discrepancy that exists between the measurements of the infrared imaging method (middle ROI) and the ground-truth instrument.

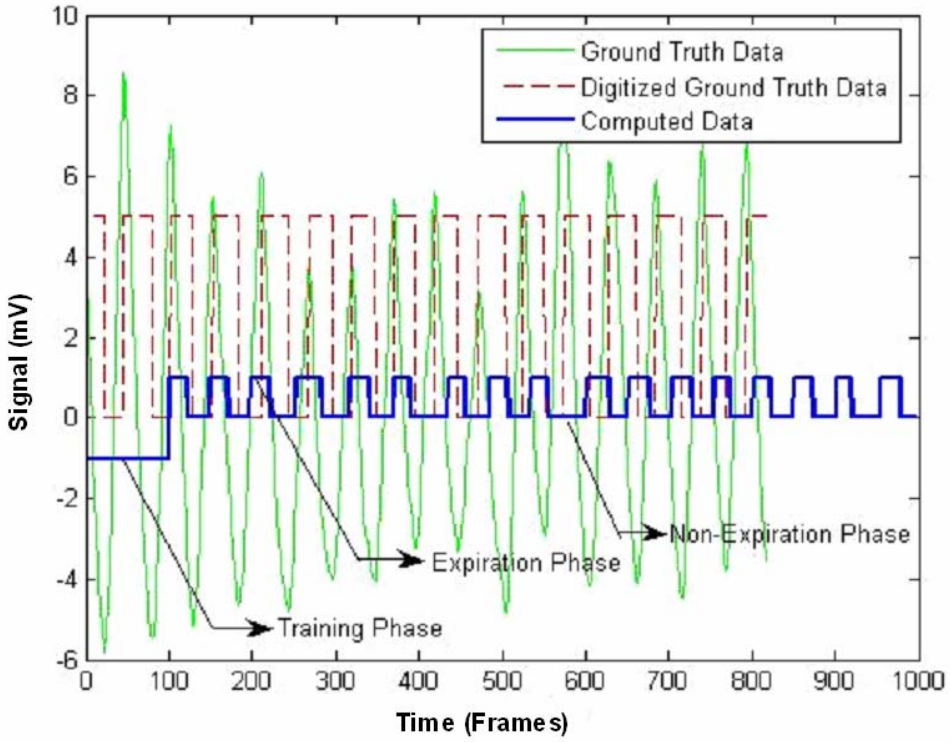

Fig. 18. This plot shows the phase shift between the breathing signal (subject #2, clip #4) computed from the infrared imaging method and the corresponding ground truth signal

Three primary factors account for this phase shift:

1. Imperfect (manual) synchronization of the beginning of the two recordings (infrared video and monitor belt).
2. Mismatch of recording frequencies. Our infrared camera records at 31 frames per second while the monitor belt samples at 100 times per second.
3. The monitor belt records ground truth data at the diaphragm level while our infrared imaging method classifies air flow at the nasal-mandible level.

The first factor can be addressed by developing a hardware trigger. The second factor can be addressed by down-sampling the ground-truth signal. The amount of phase shift due to the third factor can be determined and taken into account by performing further research in this direction.

VIII. CONCLUSION

Breathing function is one of the major indicators of an individual's health. It can be used to predict various life threatening disorders like sudden infant death syndrome and heart attacks. It is also used in sleep studies to detect sleep apnea. Finally, it is one of the psycho-physiological channels in polygraph examinations. Various modalities have been developed to measure breathing rate. Almost all the legacy methods require contact and hence they compromise the subject's comfort and mobility. Moreover, measurements by these methods are corrupted

either by movement artifacts or by their dependence on other physiological variables, like heart rate. We have proposed a method that is based on infrared imaging and statistical computation to measure passively, breathing rate at a distance. This method achieved an accuracy of *96.43*% on a small set of subjects during rest and after brief exercise. It has the potential to provide a unique capability for sustained monitoring of chronic or acute breathing problems and in sleep studies by overcoming the deficiencies of the existing measurement modalities. It also opens the way for the next generation contact-free polygraphy that will not affect the subject's psycho-physiology.

IX. FUTURE WORK

Since our method is contact-free, it has significant advantages over contact modalities like ECG, photo-plethysmography, nasal temperature probe, and breathing monitor belt in terms of comfort. With the use of a simple tracking algorithm, our method has overcome the drawback of active non-contact modalities, like the Radar Vital Signs Monitor, whose output gets corrupted by motion artifacts. But, our tracking algorithm cannot deal with situations wherein the ROI fails to remain in the field of respiratory airflow, which occur when:

1. The subject rotates his/her head towards or away from the camera (Figure 19).
2. The source of airflow (either the nose or the mouth) changes.

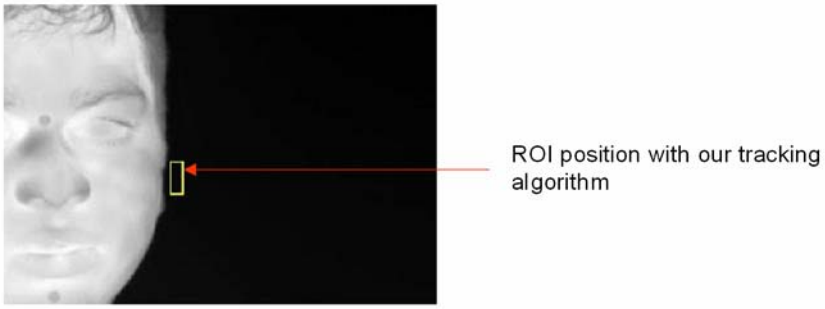

Fig. 19. ROI wrongly positioned when the subject turns his head towards the camera due to the loss of the reference point.

The first problem may be addressed by developing an advanced nasal-mandible tracking algorithm (Figure 20) along with further research on detecting the respiratory airflow signal in frontal views. The second problem can be addressed by using two ROIs, one each for nasal and mandible airflow (Figure 21).

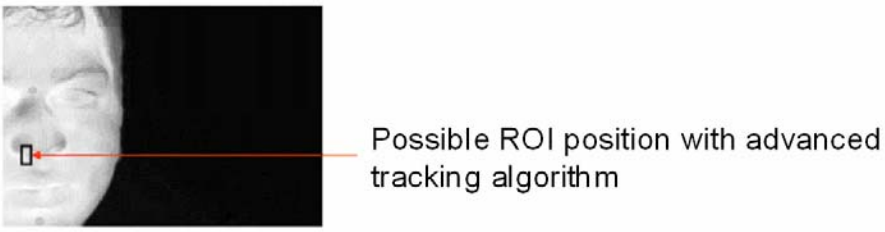

Fig. 20. Likely position of ROI if an advanced tracking algorithm were employed.

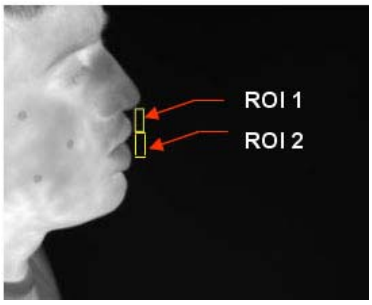

Fig. 21. The problem of change in source of respiratory airflow can be solved by using an ROI each at the nose and the mouth.

In our algorithm we have made use of the Gaussian nature of the thermal signature of breath to develop a statistical algorithm which classifies the frames as expiration or non-expiration. An alternative approach would be to consider the quasi-periodic nature of the thermal signal (Figure 12, Figure 13), which renders itself naturally to Fast Fourier Transforms.

X. ACKNOWLEDGEMENTS

This research was supported by DARPA grant #N00014-03-1-0622 and by the University of Houston start-up funds of Professor I. Pavlidis. The views expressed in this article do not necessarily reflect the views of the funding Agencies. We are thankful to Dr. Ralph Chatham at DARPA for his support. We would also like to thank Dr. Arcangelo Merla for his valuable help in the human experiments and Dr. Panagiotis Tsiamyrtzis for many valuable discussions on issues of advanced statistics.

REFERENCES

[1] G.B. Moody, R.G. Mark, M.A. Bump, J.S. Weinstein, A.D. Berman, J.E. Mietus, and A.L. Goldberger, "Clinical validation of the ECG-derived respiration (EDR) technique," *Computers in Cardiology*, vol. 13, pp. 507-510, 1986.

[2] T. Kim and M.C.K. Khoo, "Estimation of cardio respiratory transfer under spontaneous breathing conditions: A theoretical study," *The American Journal of Physiology*, vol. 273, no. 2, pp. H1012-H1023, August 1997.

[3] M. Partinen, J. Alihanka, and J. Hasan, "Detection of sleep apneas by the static charge-sensitive bed," in Proc. 6th European Congress on Sleep Research, Zurich, March 1982.

[4] B.H. Larson, "Signal processing techniques for non-invasive monitoring of respiration and heart rate," M.S. Thesis, Department of Electrical Science, University of Houston, May 1987.

[5] D. Barschdorff and W. Zhang, "Respiratory rhythm detection with plethysmographic methods," *Engineering Advances: New Opportunities for Biomedical Engineers*, vol. 2, pp. 912-913, November 1994.

[6] L.M. Vicente, A.B. Barreto, and A.M. Taberner, "DSP removal of respiratory trend in photoplethysmographic blood volume pulse measurements," in Proc. IEEE Southeastcon '96 - 'Bringing Together Education, Science and Technology', April 11-14, 1996, pp. 96 - 98, 98a.

[7] A. Johansson, L. Nilsson, S. Kalman, and P.A. Oberg, "Respiratory monitoring using photoplethysmography - Evaluation in the postoperative care unit," in Proc. 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, vol. 20, no. 6, 1998.

[8] K. Nepal, E. Biegeleisen, and T. Ning, "Apnea detection and respiration rate estimation through parametric modeling," in Proc. 28th IEEE Annual Northeast Bioengineering Conference, Philadelphia, Pennsylvania, April 20 -21, 2002.

[9] K. Storck, M. Karlsson, P. Ask, and D. Loyd, "Heat transfer evaluation of the nasal thermistor technique," *IEEE Transactions on Biomedical Engineering*, vol. 43, no. 12, pp. 1187-1191, December 1996.

[10] E.F. Greneker, "Radar sensing of heartbeat and respiration at a distance with applications of the technology," *RADAR*, vol. 97, no. 449, pp. 150-154, October 1997.

[11] E.F. Greneker and J.L. Geisheimer, "The RADAR flashlight three years later: an update on developmental progress," in Proc. IEEE 34th Annual International Carnahan Conference on Security Technology, October 23-25, 2000, pp. 170 - 173.

[12] J. Geisheimer and E.F. Greneker III, "A non-contact lie detector using radar vital signs monitor (RVSM) technology," in Proc. IEEE 34th Annual 2000 International Carnahan Conference on Security Technology, October 23-25, 2000, pp. 257 - 259.

[13] I. Pavlidis, J. Levine, and P. Baukol, "Thermal imaging for anxiety detection," in Proc. 2000 IEEE Workshop in Computer Vision Beyond the Visible Spectrum: Methods and Applications, Hilton Head Island, South Carolina, June 16, 2000, pp. 104-109.

[14] I. Pavlidis, N.L. Eberhardt, and J. Levine, "Human behavior: Seeing through the face of deception," *Nature*, vol. 415, no. 6867, January 3, 2002.

[15] I. Pavlidis and J. Levine, "Thermal image analysis for polygraph testing," *IEEE Engineering in Medicine and Biology Magazine*, vol. 21, no. 6, pp. 56-64, November/December 2002.

[16] D.U. Silverthorn, "Respiratory physiology," *Human Physiology: An Integrated Approach*, 2nd ed., Prentice-Hall, 2001, pp.498-508.

[17] R. Murthy, "Feasibility study of breathing rate computation using infrared imaging, M.S. Thesis," Department of Computer Science, University of Houston, December 2004.

[18] N. Otsu, "A threshold selection method from gray-level histograms," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 9, no. 1, pp. 62-65, 1979.

[19] G.C. Holst, *Common Sense Approach to Thermal Imaging*, SPIE Press, JCD Publishing, 2000, pp.144.

[20] M. Sonka, V. Hlavac, and R. Boyle, *Image Processing, Analysis, and Machine Vision*, 2nd ed., Brooks/Cole Publishing Company, 2001.

[21] I. Pavlidis, V. Morellas, P. Tsiamyrtzis, and S. Harp, "Urban surveillance systems: from the laboratory to the commercial world," *IEEE Proceedings*, vol. 89, no. 10, pp. 1478-1497, October 2001.

[22] J.T. Tou and R.C. Gonzalez, *Pattern Recognition Principles*, Addison-Wesley Publishing Company, 1974.

[23] A. Pednekar, I.A. Kakadiaris, and U. Kurkure, Adaptive fuzzy connectedness-based medical image segmentation, *Proceedings of the 2002 Indian Conf. on Computer Vision, Graphics, and Image Processing (ICVGIP'02)*, pp. 457-462, Ahmedabad, India, December 16-18, 2002.