Demo: Statistical Acoustic Sensing For Real-Time Respiration Monitoring and Presence Detection

Sheng Lyu, Ruiming Huang, Yuemin Yu, Chenshu Wu
Department of Computer Science
The University of Hong Kong
{shenglyu,huangruiming,yuyuemin}@connect.hku.hk,chenshu@cs.hku.hk

ABSTRACT

In this demo, we present an all-in-one real-time system for breathing monitoring and presence detection using statistical acoustic sensing. By applying Auto-Correlation Function (ACF) to the Channel Frequency Response (CFR), our system captures both motion statistics and breathing rates. We devise novel weight combining schemes to enhance the SNR of the weak sensing signals. We then enable human presence detection by integrating both motion statistics and breathing rate as vital indicators. Our system operates using a single microphone without relying on a bulky microphone array. Our demo functions in real-time and supports any device that is equipped with a commodity microphone and speaker. Our demo can be accessed through https://youtu.be/JUB6yQ1rQUo

CCS CONCEPTS

• **Computer systems organization** → **Embedded and cyberphysical systems**; • **Human-centered computing** → **Ubiquitous and mobile computing**.

1 INTRODUCTION

The vital nature of monitoring daily human breathing has garnered immense attention. The ability to accurately capture and analyze such data is critical in healthcare, in-car child presence detection [1], and sleep monitoring. There is a pressing need for reliable noncontact sensing for capturing breath. Acoustic sensing is a potential solution due to its widespread availability and practicality. The rationale lies in that chest and abdomen motions induced by breath can be captured by the acoustic channel, usually represented as
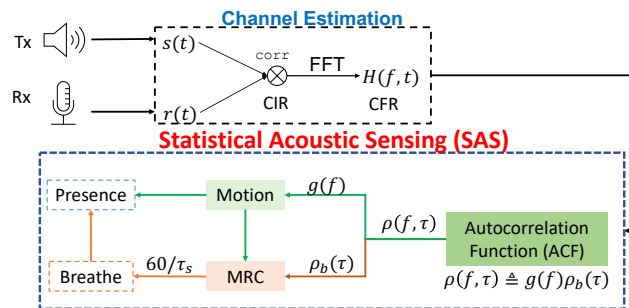
Figure 1: System Framework

Channel Impulse Response (CIR) or its frequency-domain counterpart, Channel Frequency Response (CFR). To estimate the acoustic channel, certain sensing signals are modulated and transmitted via a speaker, which, during propagation, will reflect off the human body before being recorded by the microphone. Traditional approaches usually work by detecting the subject's reflection at a single tap in the measured CIR profile, and breathing rate estimation is done by periodicity analysis, *e.g.*, via FFT, of a time series of the CIR taps. These approaches usually suffer from limited sensing coverage, in addition to being vulnerable to subject locations and orientations.

In this demo, we present a novel breath monitoring and presence detection system via statistical acoustic sensing. Specifically, we employ the Auto-Correlaton Function (ACF) on acoustic CFR to acquire presence and breath information in a unified model of statistical acoustic sensing introduced by VeCare in [1]. The statistical acoustic sensing model leverages all multipath reflections that are potentially distorted by the subject for sensing, therefore promising an enlarged sensing coverage and enhanced robustness.

We demonstrate our demo in a real-time mode, as shown in Fig. 1. The acoustic transmitter will first broadcast the Kasami sequence and signals will be bounced back when encountering reflectors. Our system implements a statistical acoustic sensing framework to capture both motion and breathing statistics simultaneously, with a single speaker-microphone channel. To boost the accuracy of breathing detection, we perform Maximal-ratio Combining (MRC) to enhance the SNR of the ACF and leverage a dedicated peak-finding algorithm. To lower the false alarm rate and enhance the robustness of presence detection, we employ a decision tree that integrates breathing and motion data, alongside a duty-ratio algorithm that optimizes presence reporting. This versatile system is compatible with any mobile or endpoint device equipped with a speaker and microphone.
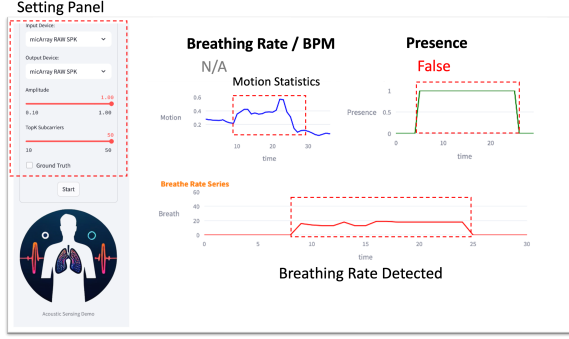
Figure 2: GUI of Our Demo

2 SYSTEM DESIGN

2.1 Stastical Acoustic Sensing

Normally, the room environment creates a rich multipath effect. The acoustic signals get reflected by scatterers in the environment. In general, these scatterers can be then decomposed as dynamic scatterers and static scatterers, *i.e.*

$$H(f,t) = \sum_{i \in R_D} H_i(f,t) + \sum_{j \in R_S} H_j(f,t) + N(f,t), \quad (1)$$

where R_D and R_S denote dynamic and static scatters respectively. $N(f,t)$ represents Additive Gaussian White Noise (AWGN) with variance σ_N^2. Practically, the static part can be removed by subtracting the mean. According to [1], the ACF of $H(f,t)$ is

$$\rho_H(f,\tau) \triangleq g(f)\rho_b(\tau) + n(f,\tau) \quad (2)$$

$g(f)$ is the *motion statistics* and $\rho_b(\tau)$ is the ACF of a periodic signal, which may observe the peak at the periodic point. Thus, once there is a breath signal, we detect the first prominent peak of $\rho_H(f,\tau)$ and get the time lag τ_s. The breath rate can then be estimated as $60/\tau_s$ BPM.

2.2 Breathing Monitoring

We apply sliding window on $H(f,t)$ and calculate an ACF matrix $\tilde{\rho}_H(f,\tau,t)$ for each window to obtain continuous measurements. Due to the notably weak reflections from chest movements, the received signal is not consistently strong, and the peaks are often not prominent. Consequently, it is necessary to develop a method to enhance these signals. We adopt normalized motion statistics to perform Maximal-ratio Combining (MRC). We then get a boosted version of ACF with the combination of different subcarriers. Afterward, we leverage a dedicated peak-finding algorithm to locate the first peak in the ACF. To alleviate the influence of real-world fluctuations, we add several constraints for time series, including using the Hampel filter to remove the outliers and applying interpolation to smooth the breathing signals. We depict the breathing signal on the GUI and demonstrate the breathing rate in real time.

2.3 Presence Detection

$g(f)$ is the motion statistics, which have been proven to be a robust indicator for profile motion dynamics. We use a combination of breathing rate and motion statistics to detect human presence. Specifically, we adopt an adaptive threshold filter to determine a
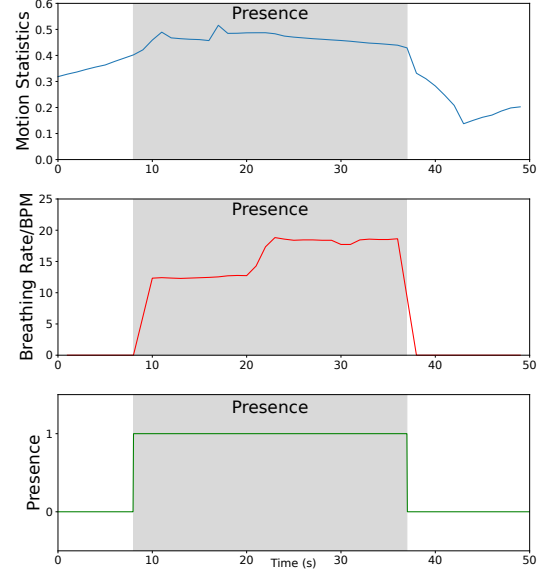

Figure 3: Experiment Result

threshold ϵ for the motion statistics. If $g(f) > \epsilon$, it means the motion is prominent, hereby present. Otherwise, it means the motion is not that prominent, which can be attributed to two reasons: there is weak breathing or there are no people present. Therefore, we employ a decision tree to combine the information of both breathing rate and motion statistics to judge whether a human is present. Furthermore, to boost the stability of the system and decrease the false alarm rate, we adopt a duty-ratio scheme, where we count the frequency that $g(f)$ is larger than the threshold or the breathing rate is detected in each sliding window.

3 DEMONSTRATION DETAILS

We use Streamlit to set up a web-based GUI for our system, as shown in Fig. 2. We use 2-s window to calculate $g(f)$, a 5-s window to detect presence and a 10-s window to compute breathing rate. The result would be rendered simultaneously on the screen. We implement our demo with a MiniDSP UMA-8-SP USB microphone array and PUI AS05308AS-R speaker. We only use one single microphone of the array in our experiments. Fig. 3 shows the experiment plot. Initially, there is no participant. Shortly after, a participant appears, and starts with a breathing rate of 13 BPM, which is later elevated to 19 BPM. The experiment concludes as the participant rises and departs from the scene.

ACKNOWLEDGEMENTS

REFERENCES

[1] Yi Zhang, Weiying Hou, Zheng Yang, and Chenshu Wu. 2023. VECARE: Statistical Acoustic Sensing for Automotive In-Cabin Monitoring. In *USENIX NSDI*.