Quantifying the Physical Separability of RF-based Multi-Person Respiration Monitoring via SINR

Shujie Zhang[1*] Tianyue Zheng[1,2*] Hongbo Wang[1] Zhe Chen[3] Jun Luo[1]
[1]School of Computer Science and Engineering, Nanyang Technological University (NTU), Singapore
[2]Energy Research Institute, Interdisciplinary Graduate Programme, Nanyang Technological University, Singapore
[3]AIWiSe, China-Singapore International Joint Research Institute, China
Email: {shujie002, tianyue002, hongbo001, junluo}@ntu.edu.sg, chenz@csijri.com

ABSTRACT

Recent years have witnessed a growing interest in contact-free respiration monitoring leveraging *radio-frequency* (RF) technologies. However, the proposed solutions mostly consider single-person scenarios, whereas a few multi-person monitoring proposals simply apply blind source separation to handle inter-person interference, without drawing a clear line between *physical* and *algorithmic* separability. In this paper, we set out to answer: *under what condition(s) one may physically separate multiple respiration signals sensed by diversified RF technologies*? Drawing inspiration from conventional signal processing, we propose *respiration-to-interference-plus-noise ratio* (RINR) as a novel metric, taking into account the impact from both background noise and various interfering sources. Instead of attenuation in Euclidean distance, RINR has to be evaluated upon *range/angle bin*s where physical separation actually take place. As signal attenuation has never been modeled in this manner, we rise to this challenge by levering a deep learning model to fit a *spread* function upon range/angle bins. The resulting RINR model allows us to concretely indicate the limit of physical separability of RF-based multi-person respiration monitoring. Our extensive experiments firmly validate the RINR model, thus evidently demonstrating the benefits of employing RINR model as a guideline for conducting respiration monitoring with different RF technologies.

CCS CONCEPTS

• **Human-centered computing** → **Ubiquitous and mobile computing systems and tools**; • **Hardware** → Sensor applications and deployments.

KEYWORDS

Contact-free sensing, RF-sensing, respiration monitoring, SINR.

ACM Reference Format:
S. Zhang, T. Zheng, H. Wang, Z. Chen, and J. Luo. 2022. Quantifying the Physical Separability of RF-based Multi-Person Respiration Monitoring via SINR. In *The 20th ACM Conference on Embedded Networked Sensor Systems (SenSys '22), November 6–9, 2022, Boston, MA, USA.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3560905.3568513

* Both authors contributed equally to this research.

SenSys '22, November 6–9, 2022, Boston, MA, USA

ACM ISBN 978-1-4503-9886-2/22/11…$15.00
https://doi.org/10.1145/3560905.3568513

1 INTRODUCTION

Since respiration monitoring plays a key role in reflecting human health status [5, 11, 16], it has enabled a wide range of healthcare applications, such as sleep status monitoring [31, 39, 61], fatigue detection [47, 57, 68], and medical care [16, 43, 45]. Traditional approaches for respiration monitoring depend on wearable devices to achieve high-quality measurement [15, 18, 19, 24, 35], but these contact-based sensing methods tend to cause discomfort and inconvenience, rendering them less practical for long-term monitoring. Therefore, both academia and industry have devoted efforts lately to developing contact-free respiration monitoring schemes. Essentially, contact-free sensing leverages the reflected signals of various media (e.g., light, acoustic, and radio waves) to capture the micro-motions of human chests during breathing [1, 7, 8, 25, 37, 38, 46, 49, 56, 61, 62, 68]. Among all the contact-free sensing media, *radio-frequency* (RF) [1, 8, 25, 37, 61, 68], appears to be the most appealing option, because it is less susceptible to ambient noise and interference suffered by vision- and acoustic-based solutions [7, 38, 46, 49, 51, 56, 62], and they are promoted by highly configurable commercial-grade devices [3, 10, 14].

Given all the projected benefits, it is a surprise that many existing RF respiration monitoring solutions only deal with *single-person* scenarios [8, 25, 33, 37, 40, 54, 65, 67], while realistic scenarios often involve *multiple persons* (e.g., sharing a couch). This situation is likely caused by the difficulty to handle the mutual interference among the respiration signals (i.e., reflected RF signals induced by respiration) from multiple persons (*subject*s): when multiple subjects are present within sensing range, the induced RF reflections superimpose onto each other, thus forming a signal mixture that demands a specific mechanism to handle (i.e., to separate individual respiration signals), as briefly illustrated in Figure 1. Existing proposals for multi-person respiration monitoring can be roughly classified into two categories. On one hand, *physical separation* method relies on large bandwidth or multiple antennas to achieve a high range or angular resolution, thus zooming in on individual subjects while excluding interference from others [1, 2, 50, 55, 68]. On the other hand, *algorithmic separation* method leverages statistical independence (e.g., distinct rates or phases) to drive either spectral estimation (only for rate estimation) [6, 31, 53] or one-time blind source separation schemes [63, 66].

Theoretically, there is no limit on the algorithmic separability, as it may always succeed as far as the assumption of statistical independence holds true. However, this independence assumption may

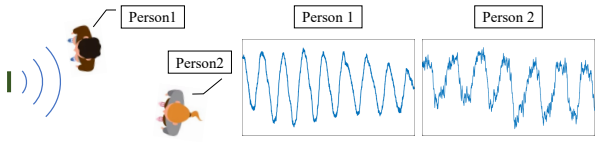

(a) Two persons are distantly located.

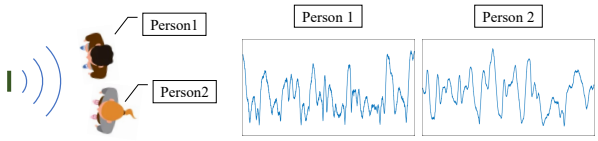

(b) Two persons are closely located.

Figure 1: Respiration signals from two persons form a mixture demanding potential separations. For subjects located far enough, the mutual-interference is minor and hence the mixture can be readily separated. As they get closer to each other, separating the mixed respiration signals becomes more and more difficult, till totally inseparable at a certain point. So the question is: *can we quantify the separability*?

not always be justified in reality because, for example, breaths of two close-by subjects can be somewhat "synchronized" due to a phenomenon called respiration coupling [17]. Moreover, algorithmic separation schemes tend to incur a very high computational complexity [30, 64] and hence hamper a real-time monitoring. Therefore, instead of quantifying the limit between separability and inseparability, we turn to a more reasonable objective of demarcating *physical separability* from *algorithmic separability*. In particular, we aim to answer the following questions:

(1) How to assess the physical separability of diversified RF technologies?
(2) Where is the limit of physical separability under different RF technologies, and what determine(s) the limits?
(3) When conducting multi-person respiration monitoring, what positions, ranges, and angular spacing among the subjects may achieve satisfactory monitoring performance?

To answer these questions, we first propose a new metric termed *respiration-to-interference-plus-noise ratio* (RINR) to measure the quality of the respiration signals under multi-person monitoring scenarios. Inspired by the definition of *signal-to-interference-plus-noise ratio* (SINR) in conventional signal processing, RINR takes into account the impact of both background noise and mutual interference among monitored subjects. Unlike SINR often evaluated on Euclidean distances, RINR is defined over the *range/angle bins* containing targeted respiration information. Since the definition over range/angle bins conforms to the data format of RF-sensing [13], RINR is more practical in evaluating the performance of respiration monitoring, yet this novel definition does entail a challenge: a new *spread function* has to be in place to characterize how respiration from one subject spreads across neighboring bins and interferes with others' respiration signals.

Due to the intrinsic complexity in physical environments and hardware implementations, the spread function may not admit a closed-form representation. Consequently, we resort to a deep learning model to fit the spread function given inputs such as the radar bin index, the bandwidth, and the number of antennas. After substituting the spread function into the RINR model, we are able to depict how multi-person respiration interferes with each other. By setting proper thresholds, the RINR further indicates the physical limit of RF-based multi-person respiration monitoring, hence provides practical guidance for the corresponding operations. Through extensive experiments, we validate the RINR model and also quantify the relationship between RINR and a variety of practical factors. To sum up, our main contributions are:

- We propose RINR as a new metric to not only quantify the quality of respiration signals but also characterize the physical separability; it is generalizable to diversified RF schemes for measuring their respective performance.
- To our best knowledge, this is the first work investigating the mechanism of multi-person respiration interference; it establishes guidelines for conducting multi-person respiration monitoring with different RF technologies.
- We evaluate the developed model with extensive experiments using three RF technologies: Wi-Fi, IR-UWB, and FMCW; the results evidently confirm the validity of our RINR model.

In fact, our metric and model are applicable to all types of contact-free motion/vibration sensing applications, yet we stick to respiration monitoring for the sake of focus in remaining of the paper. In particular, Section 2 surveys the related works of multi-person respiration monitoring; Section 3 explains the preliminaries of RF-based contact-free sensing; Section 4 introduces the RINR metric, and Section 5 completes the RINR model by fitting the spread function with a neural network; Section 6 presents the evaluation results to validate the RINR model, and Section 7 concludes the paper.

2 RELATED WORKS

In this section, we briefly review the state of the art in RF-based respiration sensing. We broadly group past works into two categories, i.e., Wi-Fi-based and radar-based respiration monitoring.

Solutions Leveraging Wi-Fi. Wi-Fi-based sensing mostly exploits Channel State Information (CSI) retrieved from received signals to estimate respiration rate and waveform in different scenarios [6, 31, 32, 34, 53, 65, 66]. Liu et al. [32] proposes a sleep monitoring system that can extract respiration patterns from CSI in a single-person scenario. Zeng et al. [65] investigated the complementary relationship between amplitude and phase of CSI to address the "blind spot" problem sensing by Wi-Fi signals. However, since the main function of Wi-Fi is communication, its designed bandwidth is narrow and does not provide sufficient range resolution for sensing multiple close-by respiration signals. As a result, researchers have to resort to advanced signal processing techniques to retrieve respiration information in multi-subject scenarios. Liu et al. [31] propose a method to simultaneously estimate respiration rates for a known number of subjects by analyzing the power spectral density of CSI. PhaseBeat [53] and TR-BREATH [6] apply the root-MUSIC algorithm [44] to resolve the respiration signals of multiple subjects. Yang et al. [59] utilize the Fresnel zone model to optimize the deployment of Wi-Fi transceivers, so that each target's respiration

signal is less interfered with by others. However, this method requires the accurate location of each person and fixed placements of the transceivers. MultiSense [66] views multi-subject respiration monitoring as a blind source separation problem, and applies the ICA (independent component analysis) algorithm [20] to extract respiration waveforms. Distinct in terms of the adopted commercial hardware, ViMo [50] leverages IEEE 802.11ad [42] with large bandwidth and a 32-antenna array for respiration monitoring; the powerful hardware makes ViMo comparable to the radar-driven approaches introduced below.

Radar-driven Approaches. Numerous commercial-grade radars have become available for near-field sensing in the past few years [4, 22, 48]. Compared with Wi-Fi, radar systems are dedicated to sensing tasks, so they provide a higher bandwidth, thus having bred a wealth of literature on high-accuracy and fine-grained respiration monitoring [1, 28, 55, 58, 63, 68]. Generally speaking, radars can be categorized into IR-UWB (impulse radio ultra-wideband) and FMCW (frequency modulated continuous wave). While IR-UWB radars transmit pulses in the time-domain [58, 68], FMCW radars send chirp signals whose frequency increases over time [2, 63]; they can be deemed as a time-frequency dual pair. Conventional ways to distinguish respiration signals from multiple subjects is to separate them on range bins. Vital-Radio [1] adopts an FMCW radar that sweeps from 5.46 GHz to 7.25 GHz to achieve satisfactory separation of respiration signals when the subjects are at least 1 m apart; V^{2}iFi [68] leverages the wide 1.4 GHz bandwidth of its IR-UWB radar to discriminate different subjects in a vehicle.

Moreover, researchers have exploited beamforming techniques [2, 55] to further mitigate the interference. Xiong et al. [55] propose a 1×8 SIMO radar system to achieve an angular respiration separation of 15°. Ahmad et al. [2] adopt a 3×4 MIMO radar system to separate signals of 15° apart. In addition, Yang et al. [60] propose a multi-beam technique, where an FMCW radar mechanically steers its beam by a rotator, yet requiring a minimum separation of 12° and 1.67 m between two subjects still. Similar to Wi-Fi-based approaches, radar-based systems also explore various sophisticated signal processing techniques and algorithms. Lee et al. [28] proposed a parametric spectral estimation technique to mitigate the mutual interference among subjects, and overcome the theoretical range resolution to detect subjects that are 40 cm apart. DeepBreath [63] shows that ICA may help extract the respiration signals of multiple subjects even when they are separated by zero distance.

Summary. Despite extensive research efforts on RF-based respiration sensing, there is still an evident gap between research and real-life application. The root cause of this gap is that the underlying assumptions of the specially designed algorithms [20, 53, 66] may not always be robust in varying environments and scenarios. To avoid this issue, certain proposals demand that sensing subjects be placed in fixed positions, for example, DeepBreath [63] and V^{2}iFi [68] require their users to remain to lie in bed or seated in the car, respectively. Such stringent requirements make sensing conditions largely invariant across different environments and scenarios, yet they are not remedies for all real-life deployments. A recent proposal [52] advocates that theoretical analysis on sensing performance is required to guarantee a consistent sensing performance. They come up with a new metric termed SSNR (sensing-signal-to-noise-ratio) to aid the design of sensing systems that are robust to environment changes. Unfortunately, they only focus on the single-person case under the Wi-Fi scheme, totally disregarding common scenarios involving multiple subjects.

3 CAPTURING MULTI-PERSON RESPIRATION

In this section, we provide the theoretical background of sensing respiration with RF technologies. We carefully study and derive the multi-person respiration model, laying a foundation for further discussions on the metric selection and quantification of respiration signal quality in Section 4.

Existing proposals have employed various RF technologies (Wi-Fi, IR-UWB, and FMCW) for monitoring respiration. Although they adopt different central frequencies, bandwidths, and waveforms, the received signal model can be unified to facilitate exposition and further processing [13]. Basically, these proposals exploit the channel information contained in the reflected signals off the subjects in a give space to perform monitoring. The channel at the carrier frequency f_c can be represented as:

$$h(t) = \sum_{p=1}^{P} \alpha_p(t) e^{-j2\pi f_c \frac{2d_p(t)}{c}} + w(t), \tag{1}$$

where P is the number of indoor multipaths,[1] $\alpha_p(t)$ is the amplitude of the real channel gain of the p-th reflected signal, c is the speed of light, $d_p(t)$ is the length of the p-th path, and $w(t)$ is Gaussian noise. Given the transmitted signal $x(t)$, the received signal is $y(t) = x(t) * h(t)$ with $*$ denoting convolution. After I/Q downconversion, the demodulated baseband signal becomes a complex sequence $r(t) = r_I(t) + jr_Q(t)$, with the carrier phase of $r(t)$ representing fine displacements of human chest, as shown in Eqn. (1).[2]

In Figure 2a, we plot the amplitude of the received complex signal $r(t)$ with the peaks indicating the subjects being sensed. Due to limited bandwidth and sampling frequency, fast-time t is a discrete index resulting from binning its continuous counterparts. Since the index t denotes the propagation time related to distance, it is termed *range bin*, and the interval between consecutive bins is

[1] P indicates the number of reflectors for a radar, while this number should be increased by 1 for Wi-Fi to include the LoS (line-of-sight) path.

[2] The phase difference of the waveform is a linearly function of the difference in distance traveled by signal: $\Delta\phi = 2\pi \frac{\Delta d}{\lambda}$.

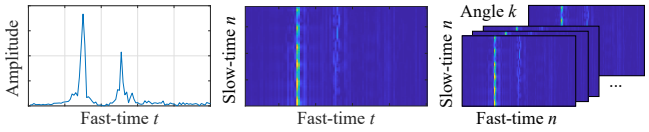

(a) S. frame $|r(t)|$. **(b) S. matrix $|r(t,n)|$.** **(c) S. tensor $|r(t,n,k)|$.**

Figure 2: Reflections from two breathing subjects captured by the radar. (a) shows a single *signal frame* indicating the distances (or range bins) of the two subjects. By adding a slow-time dimension, (b) shows the respiration *signal matrix* depicting how respiration evolves over time. (c) adds angle dimension to form a *signal tensor* involving signal matrices captured from different directions (or angle bins).

the range resolution[3] of the concerned RF technology. While the samples of IR-UWB radar directly correspond to range bins, those of Wi-Fi OFDM symbols or FMCW chirps have to be transformed with FFT to obtain corresponding range bins. With sufficiently wide bandwidth, subjects at different distances will fall into different bins and can be clearly differentiated as shown in Figure 2a.

However, a single *signal frame* $r(t)$ only allows one to differentiate subjects at distinct distances. To monitor the micro-displacement of the human body induced by respiration, one has to consecutively transmit multiple frames at a regular time interval; stacking N such frames forms a *signal matrix* $r(t,n) = [r_1(t), \cdots r_n(t), \cdots r_N(t)]^T$ in Figure 2b, with n being the *slow-time* indices. In other words, the slow-time dimension of the frame describes how the OFDM symbols, FMCW chirps, or IR-UWB pulses vary in a larger timescale. As shown in Figure 2b, the patterns of respiration of two subjects can be readily observed at distinct locations indicated by the corresponding range bins. Moreover, for RF devices equipped with multiple antennas, the signal matrix is extended to a signal tensor $r(t,n,k)$ in Figure 2c with k indicating the *angle bin* (representing the angle of arrival) as resolved by multiple antennas. To summarize, the signal tensor $r(t,n,k)$ containing the respiration of a single person can be represented as:

$$\begin{aligned} r(t,n,k) &= \beta_{\tau,\kappa}(t,n,k) \circ |r^{\mathrm{H}}(t,n,k)| e^{-j4\pi(\frac{\bar{d}+d^{\mathrm{r}}}{\lambda})} \\ &+ r^{\mathrm{LoS}}(t,n,k) + r^{\mathrm{N}}(t,n,k), \end{aligned} \tag{2}$$

where $\beta_{\tau,\kappa}(t,n,k)$ is the *spread function* parameterized by the subject's location (as determined by τ and κ) that describes how a subject's respiration signal spread to the neighboring range and angle bins, the operator $\circ$ denotes the Hadamard product, $r^{\mathrm{H}}(t,n,k)$ is the human respiration signal, $\bar{d}$ is the mean distance from the radar to the subject's body, d^{r} is the micro-displacement caused by respiration, λ is the wavelength at the carrier frequency, $r^{\mathrm{LoS}}(t,n,k)$ is the interference due to the LoS path for Wi-Fi, and $r^{\mathrm{N}}(t,n,k)$ denotes noise. Extending Eqn. (2) to cases involving M subjects in the environment, we have:

$$\begin{aligned} r(t,n,k) &= \sum_{m=1}^{M} \beta_{t_m,k_m}(t,n,k) \circ |r_m^{\mathrm{H}}(t,n,k)| e^{-j4\pi\left(\frac{\bar{d}_m+d_m^{\mathrm{r}}}{\lambda}\right)} \\ &+ r^{\mathrm{LoS}}(t,n,k) + r^{\mathrm{N}}(t,n,k). \end{aligned} \tag{3}$$

Let $m = i$ to be the subject of interest for respiration extraction, and $m = j$ to be other subjects that cause the interference, then r_i can be rewritten as a sum of the respiration signal of interest r_i^{H}, interference from other breathing subjects r_j^{I}, LoS signal r^{LoS} (only for Wi-Fi sensing), and noise r^{N} as follows:

$$\begin{aligned} r_i(t,n,k) &= \beta_{t_i,k_i}(t,n,k) \circ |r_i^{\mathrm{H}}(t,n,k)| e^{-j4\pi\left(\frac{\bar{d}_i+d_i^{\mathrm{r}}}{\lambda}\right)} \\ &+ \sum_{j\neq i}^{M} \beta_{t_j,k_j}(t,n,k) \circ |r_j^{\mathrm{H}}(t,n,k)| e^{-j4\pi\left(\frac{\bar{d}_j+d_j^{\mathrm{r}}}{\lambda}\right)} \\ &+ r^{\mathrm{LoS}}(t,n,k) + r^{\mathrm{N}}(t,n,k) \\ &= r_i^{\mathrm{H}}(t,n,k) + r_i^{\mathrm{I}}(t,n,k) + r^{\mathrm{LoS}}(t,n,k) + r^{\mathrm{N}}(t,n,k). \end{aligned} \tag{4}$$

[3]Range resolution indicates the ability of an RF sensing device to differentiate reflectors at different distances. It relies on the bandwidth (B) of the baseband signal and can be calculated as $\Delta r = \frac{c}{2B}$.

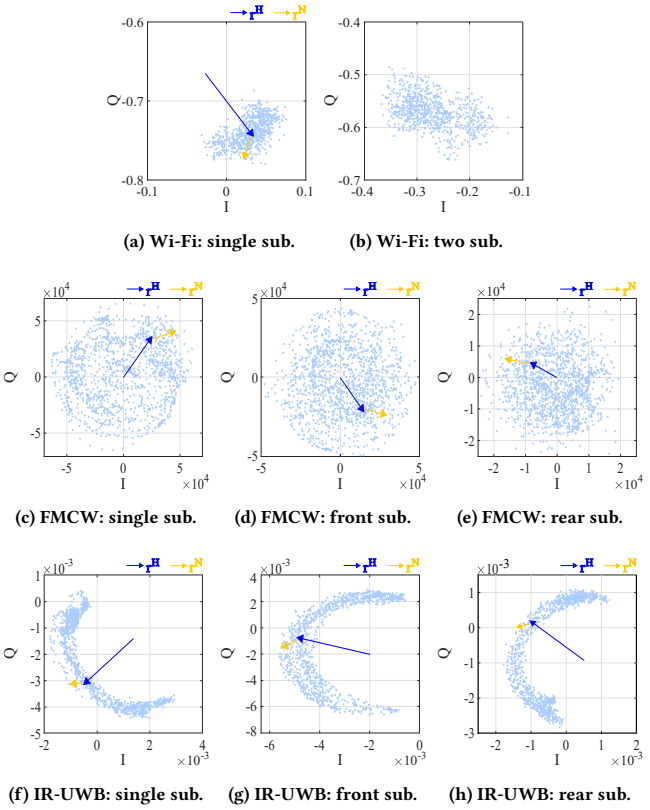

Figure 3: Constellation diagrams of respiration signals. While Wi-Fi performs poorly in differentiating multi-person respiration, both radars can differentiate the respiration signals of sufficiently separated subjects.

Clearly, both the interference and noise are integral parts of the received signals in Eqn. (4): they complicate the extraction of respiration signals. To demonstrate the damaging effects of the interference and noise, we plot the constellation diagrams of the received signals of different RF schemes under both single-person and dual-person scenarios in Figure 3. While the noise $r^{\mathrm{N}}(t,n,k)$ is Gaussian-like and it disperses the signal in the IQ space by a small amount, the interference $r_i^{\mathrm{I}}(t,n,k)$ (together with the LoS path $r_i^{\mathrm{LoS}}(t,n,k)$ for Wi-Fi), whose strength is much higher, may totally overwhelm the targeted respiration signal. Figures 3a and 3b show that Wi-Fi-based respiration monitoring is able to trace the respiration of a single person out as an arc, but it is susceptible to multi-person interference and can barely separate the mixed signal from close subjects (of 1 m distance), because each range bin of Wi-Fi spans across a large distance of 7.5 m. In Figures 3c, 3d, and 3e, FMCW radar is demonstrated to be more robust to multi-person interference, as it can differentiate the respiration signals of reasonably separated subjects of 1 m distance, though the front (closer) subject exhibits a clearer trace. Similarly, it can be observed in Figures 3g and 3h that IR-UWB can also distinguish two respiration signals from 1 m separated subjects, though the rear subject may be slightly more dispersed due to a further distance.

To find out the factors that affect the physical separability of the respiration signals, we perform a quantified analysis on the relationship between target and interfering respiration signals, given background noise. In the next two sections, we choose to conduct experiments and trace-driven analysis using an FMCW radar, because it is the most re-configurable one and capable of representing all RF sensing technologies: while FMCW and IR-UWB are a time-frequency dual pair [13], Wi-Fi introduces an extra (constant) LoS term that incurs only a scaling effect to the analysis.

4 QUANTIFYING SIGNAL QUALITY

In this section, we give a formal definition of our newly proposed metric RINR, which measures the quality of received respiration signal. We also present a thorough investigation into the RINR-affecting factors to facilitate the modeling of RINR.

4.1 RINR: A New Metric

The *signal-to-noise ratio* (SNR) is a widely used measure in science and engineering that measures the quality of a signal by comparing it to noise. Specifically, SNR can be defined as the ratio between the power of a signal and that of the background noise. According to Eqn. (4), we can express the SNR of a particular frame as:

$$\mathrm{SNR}(n) = \frac{P^{\mathrm{H}}(n)}{P^{\mathrm{N}}(n)} = \frac{\sum_k \sum_t |r^{\mathrm{H}}(t,n,k)|^2}{\sum_k \sum_t |r^{\mathrm{N}}(t,n,k)|^2}. \quad (5)$$

However, applying SNR directly to RF respiration sensing may not be exactly practical as it is only suitable for a single subject. Interference induced by other subjects and the LoS signal is not properly taken into consideration. Enhancing over SNR, *signal-to-interference-plus-noise ratio* (SINR) of the i-th subject further takes into account the interfering signals induced by other sources:

$$\mathrm{SINR}_i(n) = \frac{P_i^{\mathrm{H}}(n)}{P^{\mathrm{I}}(n) + P^{\mathrm{N}}(n)} = \frac{\sum_k \sum_t |r_i^{\mathrm{H}}(t,n,k)|^2}{\sum_k \sum_t (|r_i^{\mathrm{I}}(t,n,k)|^2 + |r^{\mathrm{LoS}}(t,n,k)|^2 + |r^{\mathrm{N}}(t,n,k)|^2)}. \quad (6)$$

Since SINR integrates power over all range and angle bins, the location information of the subject is lost in the metric. Consequently, physically separable interference from far-away subjects that would normally not interfere with a respiration signal will also contribute to a degraded SINR, thus making SINR less useful for assessing multi-person respiration monitoring.

To exclude physically separable interference, we propose a new metric RINR to quantify the respiration signal quality under interference and noise. Specifically, we define RINR as the ratio between concerned signal power and the sum of interference and noise power only over specific range and angle bins (t,k), as identified by the CFAR (constant false alarm rate) algorithm [29]. According to Eqn. (4), RINR can be defined as follows:

$$\mathrm{RINR}_i(t,n,k) = \frac{P_i^{\mathrm{H}}(t,n,k)}{P^{\mathrm{I}}(t,n,k) + P^{\mathrm{N}}(t,n,k)} = \frac{|r_i^{\mathrm{H}}(t,n,k)|^2}{|r_i^{\mathrm{I}}(t,n,k)|^2 + |r^{\mathrm{LoS}}(t,n,k)|^2 + |r^{\mathrm{N}}(t,n,k)|^2}. \quad (7)$$

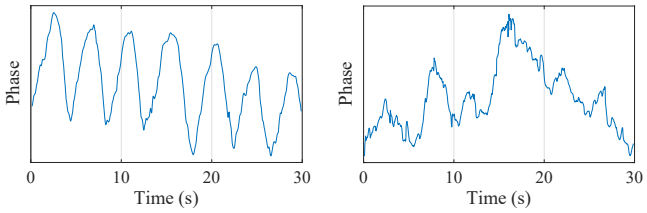

(a) Low SINR and high RINR. **(b) High SINR and low RINR.**

Figure 4: Respiration signals from the bin where the subject of interest is located. The comparison between (a) and (b) demonstrates that RINR is more practical in evaluating respiration signal quality.

This new metric zooms into the range bin t and angle bin k of the concerned subject i and exclude physically separable interference, thus correctly characterizing the quality of respiration signal.

To demonstrate that RINR is a better metric than SINR in characterizing the quality of respiration monitoring, we conduct the following experiment in an indoor environment with an FMCW radar and two subjects A and B. While B is the target subject, A acts as the interfering subject, and both of them directly face the radar. In the case shown in Figure 4a, subject A and B respectively sit at distances of 1.5 m and 2 m from the radar. Another case shown in Figure 4b keeps A intact but moves B closer to A. Comparing these two cases, the power of A's respiration signal remains unchanged but that of B's increases as it moves toward the radar. As a result, B's SINR gets increased but its RINR can be decreased. Plotting B's respiration waveforms in Figure 4, one may observe the distortion when B gets very close to A in Figure 4b: it contradicts the SINR implication but strongly validates the efficacy of RINR in measuring the quality of respiration monitoring.

4.2 Factors Affecting Signal Quality

In this section, we further investigate what factors may affect the quality of respiration signal monitoring leveraging RINR. To perform this experiment analysis, we adopt an FMCW radar [48] as it is the representative RF technology explained in Section 3. This radar is equipped with 2 transmitter (Tx) antennas (for azimuth), 4 receiver (Rx) antennas, and a tunable bandwidth up to 4 GHz theoretically. Consequently, it is deemed to have up to $2 \times 4 = 8$ *virtual* antennas [2].

Target and Interfering Range Bin Indices. In this experiment, we activate a pair of Tx-Rx antennas and a 2 GHz bandwidth to investigate the relationship between the respiration signal quality and range bin indices. We first sit subject A facing the radar while varying the subject-radar distance from 1 m to 4 m with a step size of 1 m, and plot the results in Figure 5a. Subsequently, we fix A at 2 m distance facing the radar (along the normal radiating direction), then sit another interfering subject B at 15° from the radar normal direction, and vary B's position along the radial direction from −0.4 m to 0.4 m with a step size of 0.2 m from A (a negative sign indicating B closer to the radar than A); the results are shown in Figure 5. It is clear from Figure 5a that increasing range has a negative impact on the quality of respiration monitoring, as the reflected

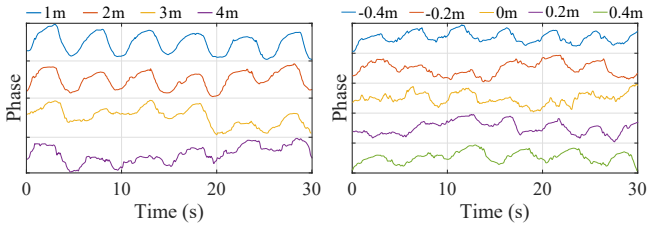

(a) Respiration without interferer. (b) Respiration with interferer.

Figure 5: Respiration signals obtained by changing range bin indices of target and interfering subjects.

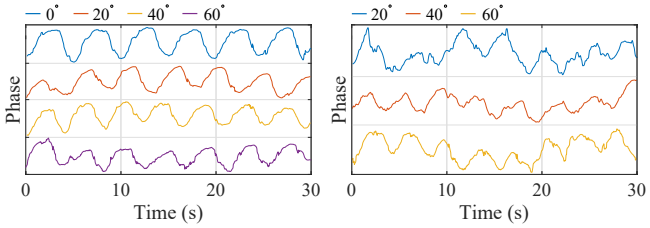

(a) Single subject. (b) Respiration with interferer.

Figure 7: Respiration signals obtained by changing angle bin indices of the target and interfering subjects.

respiration becomes weaker. Moreover, Figure 5b demonstrates that the interference induced by a close-by subject may severely degrade the quality of respiration monitoring.

Different Bandwidths. We further study the impact of bandwidths of an RF technology on the quality of respiration signals, since bandwidth is one of the key parameters that determine range resolution. We again activate a pair of Tx-Rx antennas of the radar, but sit two subjects respectively at distances of 2 m and 2.2 m from the radar, with azimuth angles 20° and -20° from the radar normal. To sense the respiration signal, we chose three radar bandwidths: 750 MHz, 1.5 GHz, and 3 GHz, and the resulting waveforms are shown in Figure 6. It can be observed in Figure 6a that the respiration signals can barely be separated under a narrow bandwidth of 750 MHz. As the bandwidth gradually increases to 1.5 and 3 GHz, the respiration signals achieve higher fidelity as they can be better separated under a lower mutual interference, as proven by Figures 6b and 6c.

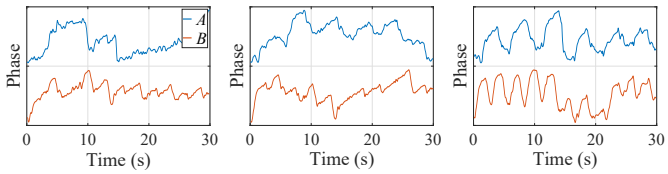

(a) 750 MHz. (b) 1.5 GHz. (c) 3 GHz.

Figure 6: Respiration signals from two subjects obtained by employing different radar bandwidths.

Target and Interfering Angle Bin Indices. We activate one Tx and four Rx antennas of the radar in this experiment and perform respiration monitoring of two subjects. We first fix the distance between the radar and subject A to be 2 m, and we vary A's azimuth angle (with respect to the radar normal) from 0° to 60° with a step size of 20°. For each setting, we plot a segment of the received respiration signal and show the results in Figure 7a. We further fix the range and angle of subject A to be respectively 2 m and 0°, and let another interfering subject B sit at the same distance, but vary B's angle so that the angle between A and B is changed from 20° to 60° with a step size of 20°. The results of A's respiration signals are shown as curves of different colors in Figure 7b. It is very evident that changing the angle of a single subject A does not have much impact on the quality of the respiration signals. However, an interfering subject B with the same range as A may clearly degrade A's signal quality, unless they have a rather large discrepancy in angles (e.g., till 60°). Note that, with $1 \times 4 = 4$ virtual antennas, each angle bin has on average roughly a span of $180°/4 = 45°$, so both 20° and 40° may not sufficiently separate the two subjects.

Number of Antennas. We then study the effect of deploying different numbers of the virtual antennas in the radar. In the experiment, both subjects are 1.5 m away from the radar, and their azimuth angles from the radar normal are -20° and 20°, respectively. We vary the number of antennas to be 2, 4, and 8 and perform FFT beamforming to focus on signals coming from different directions (i.e., angle bins). Regardless of the number of antennas, the respiration signals of A and B always reside in the two middle angle bins encompassing $-20°$ and 20°, respectively. For example, given 8 angle bins enabled by 8 antennas, A and B respectively reside in the 4-th and 5-th bins. According to the received waveforms of A and B shown in Figure 8, one may readily observe that a larger number of antennas helps with resolving mixed signals and improves the quality of respiration signals: while 2 antennas can barely separate the two respiration signals, 4 antennas can already capture the signal of one subject, and the results with 8 antennas clearly showcase that respiration signals reflected from different angles can be perfectly separated.

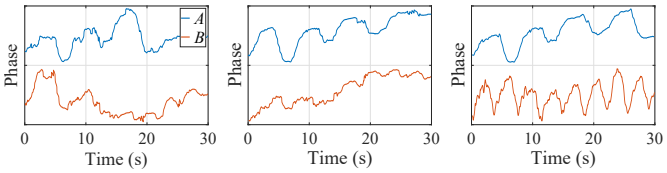

(a) 2 antennas. (b) 4 antennas. (c) 8 antennas.

Figure 8: Respiration signals obtained by deploying different numbers of virtual antennas: the more the better.

4.3 Summary

Based on the experiment analysis discussed above, we conclude that RINR is dependent on all the factors including range/angle bin indices, bandwidth, and the number of antennas. Although we have investigated other potential factors, such as center frequency and waveform of the RF technologies, we find out that the

quality of respiration signals is largely insensitive to the change of these factors. Therefore, we simply omit to present the related results. Motivated by the preliminary results, we aim to establish a comprehensive characterization on how respiration signals spread across range/angle bins, so as to associate the above-mentioned four factors with our RINR model.

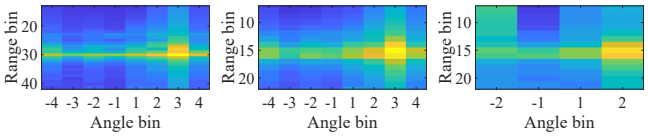

(a) 3 GHz, 8 antennas. (b) 1.5 GHz, 8 antennas. (c) 1.5 GHz, 4 antennas.

Figure 9: Visualization of $\beta_{t,k,B,N^{\text{ant}}}$ of subject sitting at 1.2 m and 45°. The respiration signal attenuates across both neighboring range and angle bins, manifesting differently under distinct settings.

5 MODELING THE INTERFERENCE

In this section, we first model the relationship between signal attenuation and the aforementioned four factors (i.e., range bin, bandwidth, angle bin, and antenna number), with the help of a deep learning network. Leveraging the resulting spread functions of both the target and interfering subjects, we can further derive the RINR model of the target subject's respiration. Finally, we summarize the key properties of the RINR model as a guide for conducting real-life multi-person respiration monitoring.

5.1 Spread Function

According to Eqn. (4) and (7), one needs $r_i^{\text{H}}(t,n,k)$ and $r_i^{\text{I}}(t,n,k)$ to calculate RINR, which further requires modeling the spread function $\beta_{\tau,\kappa}(t,n,k)$. This function measures the interference from the source "location" (in terms of range and angle bins) (τ,κ) to another "location" (t,k). Conventionally, interference strength is determined by the "distance" between (τ,κ) and (t,k), but the spread function has to characterize this distance in the "bin" space, hence becoming dependent on bin-related parameters: bandwidth B, antenna number N^{ant}, and range/angle bin indices (τ,κ). Since respiration signal is cyclostationary, the spread function does not change along slow-time n. Consequently, we could expect the actual form of the spread function to be $\beta_{\tau,\kappa,B,N^{\text{ant}}}(t,k)$, which is visualized by our experiment data in Figure 9, as a "bell" shape around (τ,κ) whose influence may spread to more than 20 range bins and 8 angle bins. However, we also notice a few challenges in obtaining a closed-form for the spread function: i) the function does not have a unified shape across range and angle bins, ii) it may not be monotonic, but rather displays ripple-like sidelobes. For example, in Figure 9a, signal strengths in angle bins -4 and -2 are greater than those in bins -3 and -1, which are closer to the source.

To tackle these challenges, we intend to fit $\beta_{\tau,\kappa,B,N^{\text{ant}}}(t,k)$ via a deep neural network trained by our experiment data. Our SPD-Fit aims to overcome three issues not readily handled by pure analytical models. First, human respiration cannot be modeled as an ideal point reflector in space, it is instead the overall result of a reflective continuum in motion. Second, the complex physics of a human body (e.g., scattering, refraction, and attenuation) further complicate the model. Last but not least, instead of being evaluated on Euclidean distance, the spread function has to be evaluated upon the range and angle bins where physical separation among RF signals actually takes place. However, the discrete nature of the bins (discretized due to intervals specified by the range and angle resolutions) makes it impossible to obtain a closed-form solution. Fortunately, the universal approximation ability [12] of deep learning networks enables us to recover a complex function relationship, so we leverage SPD-Fit to fit $f(\tau,\kappa,B,N^{\text{ant}}) = \beta_{\tau,\kappa,B,N^{\text{ant}}}(t,k), \forall t,k$. In other words, f takes in τ, κ, B, and N^{ant} and outputs a matrix across all possible (t,k). The Encoder-Decoder architecture of SPD-Fit, arguably the best for fitting a complex function, is shown in Figure 10. It first performs one-hot positional encoding to transform the information of τ and κ to an 8×120 matrix whose value is 1 at (τ,κ) and 0 elsewhere. Meanwhile, all 4 input parameters are transformed to a vector whose length is 960 with a fully connected layer, then reshaped to matrices of size 8×120. The 5 matrices are combined to form an 8×120 "image" with 5 channels. SPD-Fit then employs an encoder consisting of three 2D convolutional layers, followed by a fully connected layer to encode the transformed image to a 32×1 vector in the latent space. Finally, a decoder consisting of four transposed 2-D convolutional layers and one convolutional layer is invoked to upsample the latent vector as the output spread function matrix. Specifically, inside SPD-Fit, kernel size is set to 3×3,

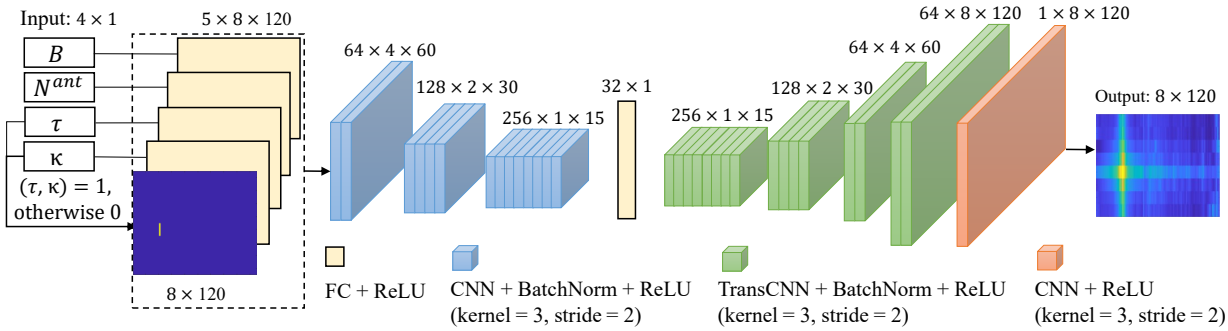

Figure 10: The architecture of our SPD-Fit model for fitting $\beta_{\tau,\kappa,B,N^{\text{ant}}}(t,k)$.

Table 1: Parameter settings for data collection.

Angle bin index	-4, -3, -2, -1, 1, 2, 3, 4
Range bin index	from 1 to 100
Bandwidth (MHz)	150, 750, 1000, 1500, 2000, 3000
Number of (virtual) antennas	2, 4, 8

and stride is set to 2. Multiple batch normalization layers [23] are used for rescaling and centering intermediate values, while ReLU layers [27] are used as the activation function to add nonlinearity as indicated in Figure 10.

To train the SPD-Fit model, we first collect a dataset of spread functions under different parameter combinations specified in Table 1, and we traverse all 8,400 possible combinations. Since the spread function describes a time-averaged effect of how respiration signals spread to neighboring bins, we take a 20 seconds average over the slow-time axis to include the effect of at least 4 complete respiration cycles. It is worth noting that not all parameter combinations are possible; for example, 2 antennas would invalidate angle bin indices of 3 to 8. To guarantee a consistent 8×120 input data format for SPD-Fit, we pad unavailable bin indices with 0's to indicate the absence of sensing signal. During the training process, we force the output spread function to be as close to the ground truth as possible by minimizing the MSE (mean squared error) loss. The training is conducted with Pytorch on a GeForce RTX 2080 Ti GPU with batch size set to 64. We train the network for 500 epochs using Adam optimizer [26] with an initial learning rate of 0.0001.

To demonstrate the efficacy of using SPD-Fit network for generating spread functions, we feed several sets of parameters into the network and plot the generated results in Figure 11. It can be observed the respiration signals spread to both neighboring range and angle bins just as the ground truth spread function does in Figure 9: i.e., the spread function resembles a "bell" in the 2-D space constituted by the range and angle bins, and conforms to the following rules. First, as the bandwidth of the radar increases, the number of affected range bins also increases. This can be explained by the fact that higher bandwidth results in a decreased range bin width, leading to more range bins within a certain Euclidean distance. Similarly, we find that a larger number of antennas results in more affected angle bins due to the decreased angle bin intervals. Last but not least, it can be observed that while the spread function attenuates monotonically in the range bin, it displays sidelobes causing fluctuations along the angle bin dimension.

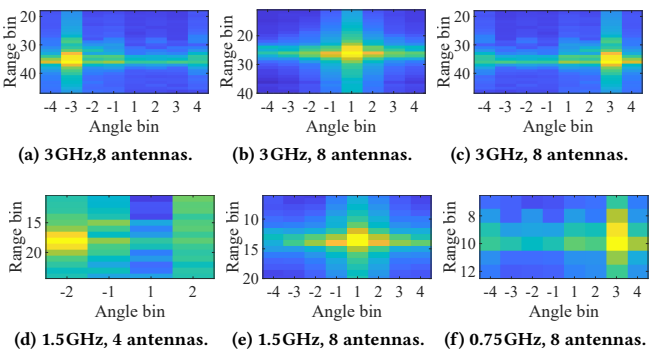

(a) 3GHz,8 antennas. (b) 3GHz, 8 antennas. (c) 3GHz, 8 antennas.

(d) 1.5GHz, 4 antennas. (e) 1.5GHz, 8 antennas. (f) 0.75GHz, 8 antennas.

Figure 11: Visualization of $\beta_{\tau,\kappa,B,N^{\text{ant}}}(t,k)$ of a subject generated by SPD-Fit with different parameter settings.

5.2 RINR Model

Fitting the spread function allows us to derive a more concrete RINR model. Since the spread function describes a time-average effect, the RINR model should also be evaluated in a time-average sense. Therefore, we take the expectation of the right side of Eqn. (7) so as to absorb the slow-time dimension n; the resulted notation of RINR (only with (t,k) as parameters) can be written as:

$$\text{RINR}_i(t,k) = \frac{\mathbb{E}_n(|r_i^{\text{H}}(t,n,k)|^2)}{\mathbb{E}_n(|r_i^{\text{I}}(t,n,k)|^2 + |r^{\text{LoS}}(t,n,k)|^2 + |r^{\text{N}}(t,n,k)|^2)}, \quad (8)$$

where $\mathbb{E}_n(|r^{\text{LoS}}(t,n,k)|^2)$ and $\mathbb{E}_n(|r^{\text{N}}(t,n,k)|^2)$ are constant powers of the LoS path and noise, according to [52]. $\mathbb{E}_n(|r^{\text{LoS}}(t,n,k)|^2) > 0$ for Wi-Fi-based respiration monitoring; $\mathbb{E}_n(|r^{\text{LoS}}(t,n,k)|^2) = 0$ otherwise. Assuming the same amplitude for the respiration signals of all subjects, we can simplify Eqn. (8) using the newly fitted spread function $\beta_{\tau,\kappa,B,N^{\text{ant}}}(t,k)$:

$$\text{RINR}_i(t,k) = \frac{\beta^2_{t_i,k_i,B,N^{\text{ant}}}(t,k)}{\sum_{j\neq i}\beta^2_{t_j,k_j,B,N^{\text{ant}}}(t,k) + C}, \quad (9)$$

where $C = \mathbb{E}_n(|r^{\text{LoS}}(t,n,k)|^2) + \mathbb{E}_n(|r^{\text{N}}(t,n,k)|^2)$; it is empirically set to 2×10^{-8} and 7×10^{-5} for IR-UWB and FMCW radar, respectively.[4] Moreover, for Wi-Fi-based respiration monitoring, C depends not only on the noise strength, but also on the strength of LoS path (to be evaluated in Section 6.4). Eqn. (9) indicates that RINR is in the form of a 2-D matrix whose two dimensions are indices of range and angle bins, respectively.

Fixing the radar with 3GHz bandwidth and 8 antennas, we show two cases of deriving RINR from the spread functions of two subjects[5] in Figure 12. The first case has the two subjects distantly located (with 1m distance), while the second case has them closely located (with 0.4m distance); we show them in the upper and lower panels of Figure 12, respectively. In particular, the sum of squares of their spread functions $\sum_{i=1}^{2}\beta^2_{t_i,k_i,B,N^{\text{ant}}}(t,k)$ (the first term in the denominator of Eqn. (9)) is shown in Figure 12a, where the actual bin locations of subjects A and B are marked with black and red crosses, respectively. Dividing the power of one subject's spread function by the sum of noise power and that shown in Figure 12a, we get the RINR matrices for both subjects, and show them in Figure 12b and 12c, respectively. One may readily observe that, whereas the optimal RINR location coincides with the actual bin and angle bins in the first case, this coincidence can be slightly violated in the second case, the actual peaks are "pushed away" by nearby interfering subjects. The inconsistency between the bins of an RINR peak and the actual location of a subject necessitates a two-step procedure, i.e., calculate the RINR matrix first and then select the bins of RINR peak for retrieving respiration signals.

[4]The empirical value of C is obtained by a two-hour statistics.
[5]The derivation can be readily extended to more subjects, but we stick to two-subject scenarios here for better clarity and brevity.

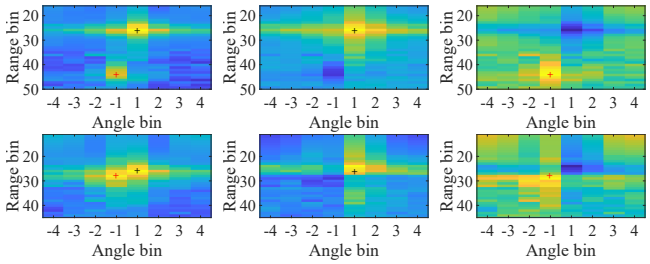

(a) $\sum_{i=1}^{2} \beta^2_{t_i,k_i,B,N^{\text{ant}}}(t,k)$. (b) RINR_A. (c) RINR_B.

Figure 12: Visualization of $\beta^2_{\tau,\kappa,B,N^{\text{ant}}}(t,k)$ of subject A and B and their RINR; the positions of subjects are labeled by black and red crosses respectively.

We can further derive an RINR heatmap to describe the spatial distributions of the achievable RINR of multi-person respiration monitoring. We set the radar at the origin of the coordinate system, and align x and y axes with the parallel and normal directions of the radar, respectively. We fix the x coordinate of subject A (marked as a red cross) to be 0 m and 1 m in Figure 13a and 13b, respectively, with the top, middle, and bottom panels having the A's y coordinate set as 1 m, 2 m, and 3 m, respectively. Given a fixed position of A, we move subject B to scan across the x-y plane. For each A-B position pair, we calculate RINR matrices and extract peak values respectively for A and B, and we put the smaller RINR value (termed *achievable* RINR) at the midpoint between them to obtain RINR heatmaps characterizing the spatial distribution of quality of respiration monitoring. Noted that each subfigure is independently normalized to fit the color scale, so the same color scale does not necessarily represent the same RINR value across subfigures.

The top panel of Figure 13a clearly depicts a triangular "forbidden" area around A; it indicates that, when B gets too close to A, the achievable RINR becomes too small to enable physical separation of their respiration signals. One may also observe a "wheel spoke"-like pattern at different angles; it is caused by the sidelobes induced by FFT beamforming, as discussed in Section 5.1 (e.g., Figure 9). As A moves to 2 m along the y-axis, the middle panel of Figure 13a shows that the forbidden area enlarges to a trapezoid, since the width along the x-axis spanned by a single angle bin increases. Moreover, further moving A to 3 m, as shown in the bottom panel of Figure 13a, makes the forbidden area grow boundlessly along the radial direction from the radar, where the signal reflected from B becomes too weak to be detected. While Figure 13a illustrates a symmetric case, we change the x-coordinate of subject A to 1 m to show an asymmetric case in Figure 13b. It is clear that similar observations made in Figure 13a (e.g., a forbidden area around A) still apply here; the major differences in patterns can be attributed to the asymmetric deployment. Moreover, the conclusions definitely hold true for negative x-coordinates of A due to space symmetry. Therefore, the results shown in Figure 13 are generalizable to arbitrary subject positions. To summarize the key insights suggested by the RINR heatmap, we further quantify the relationship between RINR and the distance between the subjects in Section 5.3.

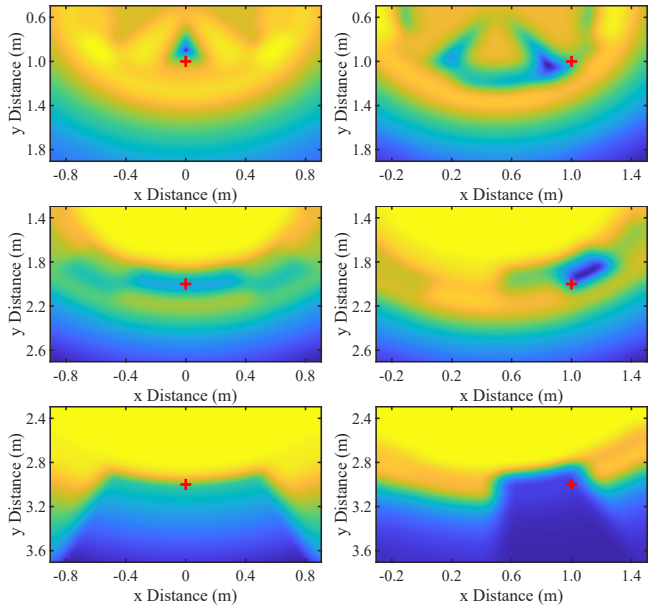

(a) $x = 0$, $y = 1, 2, 3$ m. (b) $x = 1$, $y = 1, 2, 3$ m.

Figure 13: RINR heatmaps generated for a 3 GHz and 8 antenna radar. The x distance of subject A is 0 m and 1 m in (a) and (b), respectively. The top, middle, and bottom panels show that his y distance is 1 m, 2 m, and 3 m, respectively.

5.3 Key Takeaways from the Model

Given the radar settings and the bin indices of the subjects, the RINR model in Section 5.2 outputs the worst-case (among several subjects) achievable RINR that indicates the quality of multi-person respiration monitoring. However, we are often interested in reverse tasks in practice, e.g., obtaining the minimum distance between two subjects (in terms of their gravity centers) to guarantee a desired RINR. To find such a minimum distance for the scenarios studies in Section 5.2, we maintain the same coordinates of subject A, fix the x-coordinate to be 0 m for subject B but let its y-coordinate varying from 1.3 m to 2.1 m, so as to obtain the resulted RINR in Figure 14a (with legend indicating the bandwidth and antenna number). We also fix the y-coordinate of B to be 1 m and move its x-coordinate from 0.5 m to 1.3 m to obtain the RINR results in Figure 14b.

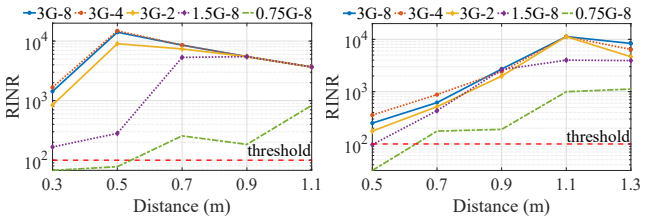

(a) B moves in the y direction. (b) B moves in the x direction.

Figure 14: The smaller RINR between subjects A and B as a function of their Euclidean distance.

One may readily observe that, given an RINR threshold of 100, which is shown as red dashed lines in the figures[6], the minimum distances between the subjects under 0.75 GHz and 8 antennas (the worst resolution given this specific type of radar [48]) are slightly below 0.6 m and 0.7 m along the y and x directions, respectively. The more stringent requirement along the x axis is actually reflected in the top panel of Figure 13a: the dark-blue forbidden area spans a larger distance along the x-axis. Moreover, distances along both x- and y-axis decrease with the radar bandwidth; they respectively drop below 0.3 m and 0.5 m eventually. As these are the closest distances two subjects can ever reach (given the dimension of a human body), *physical separability alone is proven sufficient*. It is noticeable in Figure 14a that the smaller RINR between two subjects first increases with the distance due to weaker interference then drops slightly because the respiration power attenuates with one subject moving further from the radar. Therefore, to achieve high-quality monitoring for both subjects, we have to consider both the subjects' positions and their separation. Furthermore, one may readily observe from Figure 14b that the RINR under different antenna numbers is similar, but shows distinctions under different bandwidth of 3 GHz and 0.75 GHz. This later observation implies that, though both range and angle resolutions may help improve the physical separability of multi-person respiration signals, the former plays a more prominent role than the latter.

6 EVALUATION

In this section, we first demonstrate the efficacy of employing our RINR model as a guideline for conducting multi-person respiration monitoring, by studying the relationship between RINR and two key respiration indicators. Then we investigate how practical factors (e.g., multipath and LoS) affect the performance of the RINR model. Finally, we show how proper radar placement, guided by RINR, can help improve the quality of multi-person respiration monitoring. Though only three of the authors (healthy adults) are involved as the subjects, we still strictly follow the approved IRB procedure of our institute. During the evaluation, all subjects sit statically and breathe naturally without extra interference. For each scenario, we conduct corresponding experiments 30 times.

6.1 Respiration Rate Error vs RINR

To demonstrate the effectiveness of RINR in reflecting the performance of respiration monitoring, we verify the relation between *respiration rate error* and RINR generated by our model. We conduct experiments in a 30 m^2 empty room, and use the same radar configuration and placement as in Section 5.2. We employ a NeuLog respiration monitor belt logger sensor NUL-236 [36] to collect ground truth respiration signals; its sampling rate is set to the same 50 fps as the radar and the Precision Time Protocol [21] is used to synchronize the clocks so as to align the ground truth and radar-sensed signal. Furthermore, to estimate the respiration rate from the waveform, we apply a sliding window of 1000 sampling points, then perform FFT with interpolation to find the frequency peaks in the desired respiration range of 0.16 Hz to 0.6 Hz [68]. Finally, we obtain the respiration rate error as the absolute difference between the estimated and actual respiration rates.

[6]This RINR threshold will be further explained in Sections 6.1 and 6.2.

We employ two subjects and present the target's error heatmaps in Figures 15a, 15b, and 15c, while depicting the close error-RINR correlation in Figure 15d for each subject in a more generalized 3-subject scenario. The heatmaps differ from Figure 13a in two aspects: i) only the target errors are shown here by moving an interfering subject to scan across the x-y plane, and ii) the color scale is chosen to maintain a consistent relationship between color and monitoring quality (i.e., blue and yellow indicate high and low errors, respectively). It is intuitive to observe from the three heatmaps that the high-error (blue) zones are all surrounding the target, yet their shapes keep changing when the target moves from $(0, 1)$ m to $(0, 3)$ m. This change in shape can be explained by the fact that the resolution of the angle bin becomes more dominating than the range bin, since the same angle bin covers a larger distance along the x-axis with an increasing distance between the target and the radar. It is also worth noting that these high-error zones are not symmetric along the y-axis: since an interfering subject closer to the radar causes stronger interference, its distance from the target subject should be larger to maintain the same level of error.

To further quantify the relation between respiration rate error and RINR, we perform a more general experiment by employing three mutually interfering subjects, meanwhile recording their respective rate errors and randomly changing their positions after each recording. We show the measured respiration rate error as a function of the corresponding RINR in Figure 15d. One may readily observe that the error and its variance are both decreasing functions of RINR. When RINR is at 30, the median respiration rate error is 3.5 bpm, which is too large to render the respiration monitoring useful. However, as RINR increases, the median error steadily decreases and it drops below 1 rpm (sufficiently accurate for respiration monitoring purpose) after RINR reaches 100: this value can hence be deemed as a threshold of physical separability, which will be corroborated later. The results evidently confirm that i) RINR can effectively indicate respiration rate error and ii) our RINR model (albeit mostly illustrated for two subjects) can be readily extended to multi-subject scenarios.

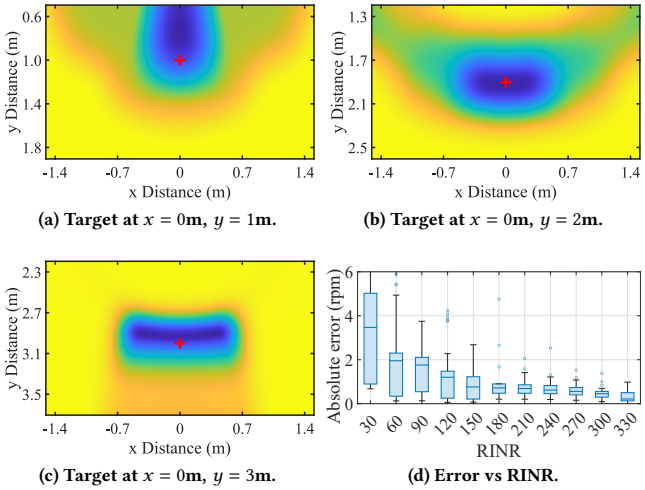

(a) Target at $x = 0$m, $y = 1$m. (b) Target at $x = 0$m, $y = 2$m.

(c) Target at $x = 0$m, $y = 3$m. (d) Error vs RINR.

Figure 15: Respiration rate error heatmaps of a target (a-c) and the relation between errors and RINR (d).

6.2 Waveform Similarity vs RINR

We further investigate whether RINR can be used for reflecting the waveform similarity, which is a more fine-grained respiration indicator than respiration rate. Given the same experiment setup as in Section 6.1, we first employ a Savitzky-Golay filter [41] to process the captured RF respiration signal $r(n)$ and obtain the smoothed signal $r'(n)$. To measure the similarity between $r'(n)$ and $r_{\text{gt}}(n)$, we first map both signals to the range $[0, 1]$, and then leverage the cosine similarity metric $\mathcal{S}(r', r_{\text{gt}})$ to quantify how much the two signals resemble each other:

$$\mathcal{S}(r', r_{\text{gt}}) = \frac{r' \cdot r_{\text{gt}}}{\|r'\|\|r_{\text{gt}}\|} = \frac{\sum_{n=1}^{N} r'(n) r_{\text{gt}}(n)}{\sqrt{\sum_{n=1}^{N} r'^2(n)} \sqrt{\sum_{n=1}^{N} r_{\text{gt}}^2(n)}}. \quad (10)$$

Essentially, the cosine similarity is calculated by measuring the cosine of the angle between two vectors $r'(n)$ and $r_{\text{gt}}(n)$, and then determines to what extent the two vectors point to the same "direction" in a high dimensional space. A higher cosine similarity indicates better fidelity of the estimated waveform.

Similar to the experiments in Section 6.1, we employ two subjects, and move the interfering subject to scan across the plane and obtain three similarity heatmaps of the target subject; the results are presented in Figures 16a, 16b, and 16c, where yellow and blue indicate high and low similarities, respectively. It can be observed that the patterns of the heatmaps roughly follow those in Figure 15, yet the relatively larger blue areas (especially in Figure 16b and 16c) seem to suggest that the performance in terms of waveform similarity is much worse than that of respiration rate. However, the seeming enlargement of the blue area is actually caused by the separate normalizations: the same color scale does not necessarily represent the same value across subfigures. To indicate the actual sizes of the low-similarity zones, we draw red contours in the heatmaps corresponding to a similarity of 0.8. One may readily observe that as the target subject moves away from the radar, the size of the low-similarity zone first decreases and then increases. The size change can be explained as follows. First, when the target is near the radar, the stronger interference nearby enlarges the area, as shown in Figure 16a. Moreover, when the subject moves further to 3 m (Figure 16c), the target signal quality is degraded, leading to an enlarged low-similarity zone again.

To further quantify the relation between waveform similarity and RINR, we perform a more general experiment by employing three mutually interfering subjects, and meanwhile recording their respective respiration waveform similarities. We show waveform similarity as a function of the corresponding RINR in Figure 16d, where the two quantities clearly exhibit a monotonic relation, i.e., the waveform similarity increases with RINR. When RINR takes on a small value of 30, the median cosine similarity can be as small as 0.62, indicating a failed recovery of the respiration waveform. The similarity increases above 0.8 (suggesting a strong positive correlation between the estimated and ground truth waveform) after RINR exceeds 100: this value corroborates the same threshold of physical separability identified in Section 6.1, and again verifies that our RINR model can be readily applied as a guideline for conducting multi-subject respiration monitoring.

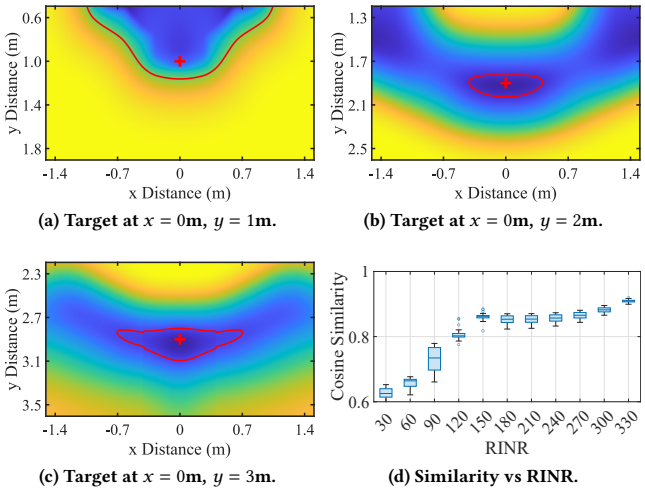

(a) Target at $x = 0$m, $y = 1$m. (b) Target at $x = 0$m, $y = 2$m.

(c) Target at $x = 0$m, $y = 3$m. (d) Similarity vs RINR.

Figure 16: Waveform similarity heatmaps of a target (a-c) and the relation between similarities and RINR (d).

6.3 Does Multipath Matter?

The RINR model in Eqn. (9) is developed without taking multipath into account, because we conjecture that it holds true even in a multipath environment given that typical multipath reflections are much weaker and with longer paths to be well separated in range bins. To prove this point, we repeat the experiments in Section 6.1 and 6.2 in another two rooms, i.e., a small room of 5 m^2 cluttered with furniture and a medium-sized room of 10 m^2 with a desk and a cabinet, to represent different levels of multipath severity. We plot the respiration rate error and waveform similarity of the target in the three rooms in Figure 17. As shown in Figure 17a, the respiration rate errors start respectively for large, medium, and small rooms at 3.2, 3.5, and 3.9 rpm, when RINR is at its minimum of 30. As RINR increases to 330, the rate errors gradually drop to three close values of 0.2, 0.2, and 0.45. Similarly, in Figure 17b, it can be seen that when the RINR is 30, the waveform similarity starts at three nearby values of 0.61, 0.63, and 0.66, respectively. As RINR increases, the similarities gradually reach their respective saturated values of 0.82, 0.88, and 0.92. To sum up, we conclude that although the same RINR manifests as slightly different respiration rate errors and waveform similarities in different environments (i.e., better in a spacious room and worse in a cramped room), the two quantities as functions of RINR follow environment-independent trends and are largely insensitive to the existence of multipath.

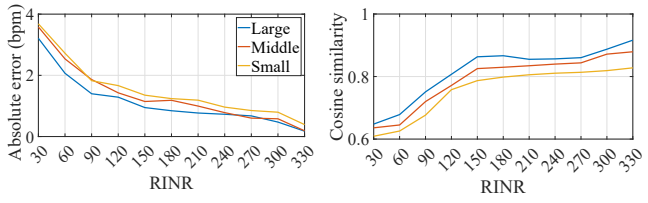

(a) Rate error. (b) Waveform similarity.

Figure 17: Effect of multipath on respiration rate error (a) and waveform similarity (b).

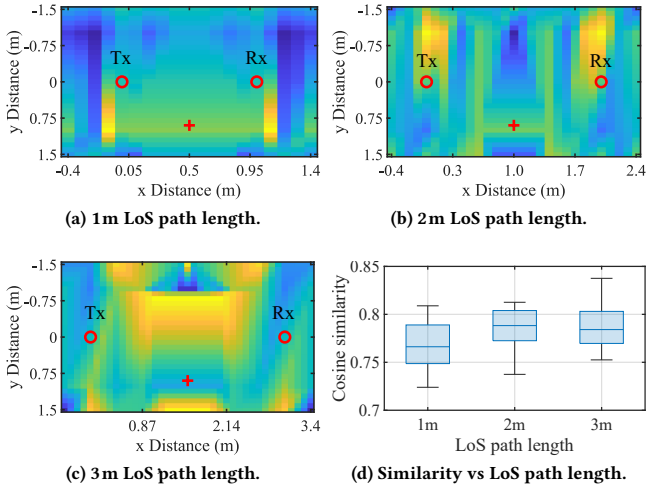

(a) 1m LoS path length. (b) 2m LoS path length.

(c) 3m LoS path length. (d) Similarity vs LoS path length.

Figure 18: Waveform similarity heatmaps of a target under different LoS path lengths (a-c) and the relation between similarities and LoS path length (d).

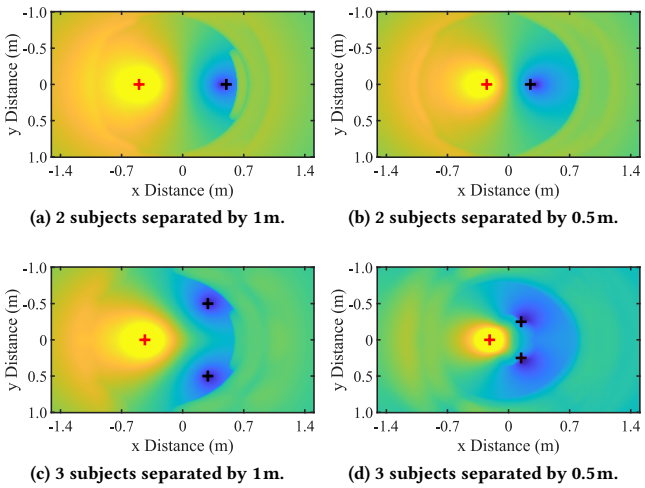

(a) 2 subjects separated by 1m. (b) 2 subjects separated by 0.5m.

(c) 3 subjects separated by 1m. (d) 3 subjects separated by 0.5m.

Figure 19: RINR heatmaps of a target obtained by moving a radar across the plane, under one (a-b) or two (c-d) interfering subjects separated by different distances.

6.4 How about LoS Interference under Wi-Fi?

As stated in Section 5.2, the denominator in the RINR expression for Wi-Fi involves not only the noise and interference from other subjects, but also the strength of the LoS path signal. To study how LoS path affects the performance of Wi-Fi-based multi-person monitoring, we conduct similar experiments as in Section 6.1, but place the Wi-Fi receiver (with 4 Rx antennas) instead of the radar at the origin. In the meantime, the Wi-Fi transmitter (with 2 Tx antennas) is placed at different LoS distances along the x-axis. Moreover, we fix the distance between the target subject (marked as a red cross in the heatmap) and the Tx-Rx center to be 1 m, and their connecting line to be perpendicular to the x-axis. By moving the interfering subject and varying the length of the LoS path from 1 m to 3 m at a step size of 1 m, we measure waveform similarities under different LoS path lengths and plot their respective heatmaps in Figures 18a, 18b, and 18c. One may readily observe that the heatmaps do not display regular ribbon-like patterns of high- and low-RINR regions as in Figure 15, since the respiration signals are basically interfered with everywhere due to the low range and angle resolutions of Wi-Fi. Moreover, due to the bistatic nature of Wi-Fi-based respiration sensing, the heatmaps become more complicated and cannot be directly compared to the ones in Figure 16.

To further quantify the effect of LoS path, we show the measured waveform similarities under different LoS path lengths as a boxplot in Figure 18d. Apparently, the relation between waveform similarity and LoS path length is not monotonic anymore, as a LoS path length of 2 m seems to result in the best performance in similarity. In fact, the increase of the similarity from 1 m LoS path length to 2 m is evidently the result of a reduced LoS signal strength as expressed in Eqn. (9), but the latter similarity drop has to be the joint effect of i) a reduced LoS signal strength and ii) a reduced reflection signal (representing respiration) strength due to a diminishing cross-section (of a human body). In general, the current sensing mode offered by Wi-Fi is always subject to the LoS interference. This can be readily addressed by switching to the monostatic radar sensing or explore our recent proposal to integrate monostatic sensing with Wi-Fi communications [9].

6.5 Guidance to Radar Placement

In the previous sections, we have discussed how the subject distance can be adjusted to obtain a sufficient RINR, thus guaranteeing sufficient physical separability indicated by low respiration rate error and high waveform similarity. However, we are often not allowed to change the subject positions in practice due to the specific environment layout (e.g., two subjects sleeping in bed and passengers sitting in a vehicle). Essentially, we are more concerned about the reverse problem, i.e., given fixed subject positions, how should we place the radar to mitigate the interference in multi-person respiration monitoring. We perform the following experiment to investigate the optimal placement of the radar. In the experiment, we fix the positions of 2 to 3 subjects (the target and interfering subject(s) are shown as red and black crosses, respectively), move the radar to scan the RINR value of the target subject, and illustrate the resulted heatmaps in Figure 19. We omit the plots for corresponding rate errors and waveform similarities as they have been depicted in Figures 15d and 16d.

It is noticeable that RINR takes on the maximum and minimum values when the radar is placed very close to the target subject and interfering subject(s), respectively. Specifically, there is a bright yellow high-RINR parabolic region around the target where the radar can be placed to achieve high RINR. By comparing the top and bottom panels of Figure 19, it is evident that increasing the number of subjects will suppress the high RINR region by decreasing the eccentricity of contours. Similarly, by comparing the results in the left and right panels of Figure 19, it becomes clear that a closer distance between the subjects further shrinks the high-RINR region, and pushes the optimal position for radar placement to the opposite directions of the interfering targets. Furthermore, most of

the time we aim to perform respiration monitoring for all subjects simultaneously. Consequently, the radar placement has to take the RINR of multiple subjects into account. By inspecting the space symmetry of the subjects, it is intuitive that the radar should not be placed in the vicinity of any single subject (and not within the convex hull defined by all subjects due to the directionality), as doing so will result in a low RINR for other subjects. All in all, the radar position optimizing RINR for all subjects should i) fall in the yellowish-green region outside by the strong-interference region roughly bounded by an ellipse, and ii) not too far from all subjects to maintain sufficient reflected signal strengths for each of them.

7 CONCLUSION

We study the conditions under which multiple respiration signals can be physically separated by RF sensing in this paper. To this end, we propose a novel metric RINR to quantify not only the quality of respiration signal but also the physical separability of mutually interfering respiration signals. Employing a carefully designed deep learning model, we fit the spread function of respiration signals in order to numerically characterize the RINR model; this in turn leads to a concrete indication on the limit of physical separability of multi-person respiration signals. As a by-product, we have also found out that, with an adequate radar, physical separability alone is sufficient for separating respiration signals from two side-by-side subjects. With extensive experiments, we have firmly validated the RINR model, and demonstrated the potential of employing this model as a guideline for conducting multi-person respiration monitoring. We foresee that integrating the RINR metric and physical separability concept into future system design will substantially push RF-based respiration monitoring towards real-life deployments.

ACKNOWLEDGEMENT

We are grateful to anonymous reviewers for their constructive suggestions. This research was support in part by National Research Foundation (NRF) Future Communications Research & Development Programme (FCP) grant FCP-NTU-RG-2022-015.

REFERENCES

[1] Fadel Adib, Hongzi Mao, Zachary Kabelac, Dina Katabi, and Robert C. Miller. 2015. Smart Homes that Monitor Breathing and Heart Rate. In *Proc. of the 33rd ACM CHI*. 837–846.

[2] Adeel Ahmad, June Chul Roh, Dan Wang, and Aish Dubey. 2018. Vital Signs Monitoring of Multiple People Using a FMCW Millimeter-wave Sensor. In *Proc. of IEEE Radar Conference*. 1450–1455.

[3] Kiarash Amiri, Yang Sun, Patrick Murphy, Chris Hunter, Joseph R. Cavallaro, and Ashutosh Sabharwal. 2007. WARP, A Unified Wireless Network Testbed for Education and Research. In *2007 IEEE International Conference on Microelectronic Systems Education (MSE'07)*. 53–54.

[4] Novelda AS. 2017. Single-Chip Radar Sensors with Sub-mm Resolution - XETHRU. https://www.xethru.com/. Accessed: 2022-04-24.

[5] David Blumenthal, Elizabeth Malphrus, and J. Michael McGinnis. 2015. *Vital Signs: Core Metrics for Health and Health Care Progress*. National Academies Press.

[6] Chen Chen, Yi Han, Yan Chen, Hung-Quoc Lai, Feng Zhang, Beibei Wang, and K. J. Ray Liu. 2018. TR-BREATH: Time-Reversal Breathing Rate Estimation and Detection. *IEEE Transactions on Biomedical Engineering* 65, 3 (2018), 489–501.

[7] Weixuan Chen and Daniel McDuff. 2018. DeepPhys: Video-Based Physiological Measurement Using Convolutional Attention Networks. In *Proc. of the 15th IEEE ECCV*. 349–365.

[8] Zhe Chen, Tianyue Zheng, Chao Cai, and Jun Luo. 2021. MoVi-Fi: Motion-robust Vital Signs Waveform Recovery via Deep Interpreted RF Sensing. In *Proc. of the 27th ACM MobiCom*. 392—-405.

[9] Zhe Chen, Tianyue Zheng, Chao Hu, Hangcheng Cao, Yanbing Yang, Hongbo Jiang, and Jun Luo. 2022. Integrating Monostatic Sensing with Communication for IoT. In *Proc. of the 1st ACM MobiCom Workshop on ISAC Systems*. 43–48.

[10] Zhe Chen, Tianyue Zheng, and Jun Luo. 2021. Octopus: A Practical and Versatile Wideband MIMO Sensing Platform. In *Proc. of the 27th ACM MobiCom*. 1–14.

[11] Jennifer Gonik Chester and James L Rudolph. 2011. Vital Signs in Older Patients: Age-related Changes. *Journal of the American Medical Directors Association* 12, 5 (2011), 337–343.

[12] Balázs Csanád Csáji et al. 2001. Approximation with Artificial Neural Networks. *Faculty of Sciences, Etvs Lornd University, Hungary* 24, 48 (2001), 7.

[13] Shuya Ding, Zhe Chen, Tianyue Zheng, and Jun Luo. 2020. RF-Net: A Unified Meta-Learning Framework for RF-Enabled One-Shot Human Activity Recognition. In *Proc. of the 18th ACM SenSys*. 517–530.

[14] Ettus Research. 2014. USRP X310 High Performance Software Defined Radio - Ettus Research. https://www.ettus.com/all-products/x310-kit/. Online; accessed 28 April 2022.

[15] Biyi Fang, Nicholas D. Lane, Mi Zhang, Aidan Boran, and Fahim Kawsar. 2016. BodyScan: Enabling Radio-based Sensing on Wearable Devices for Contactless Activity and Vital Sign Monitoring. In *Proc. of the 14th ACM MobiSys*. 97–110.

[16] Mia Folke, Lars Cernerud, Martin Ekström, and Bertil Hök. 2003. Critical Review of Non-Invasive Respiratory Monitoring in Medical Care. *Medical and Biological Engineering and Computing* 41, 4 (2003), 377–383.

[17] Pavel Goldstein, Irit Weissman-Fogel, and Simone G. Shamay-Tsoory. 2017. The Role of Touch in Regulating Inter-Partner Physiological Coupling During Empathy for Pain. *Scientific Reports* 7, 1 (2017), 1–12.

[18] Firat Güder, Alar Ainla, Julia Redston, Bobak Mosadegh, Ana Glavan, T.J. Martin, and George M. Whitesides. 2016. Paper-based Electrical Respiration Sensor. *Angewandte Chemie International Edition* 55, 19 (2016), 5727–5732.

[19] Tian Hao, Chongguang Bi, Guoliang Xing, Roxane Chan, and Linlin Tu. 2017. MindfulWatch: A Smartwatch-based System for Real-Time Respiration Monitoring during Meditation. In *Proc. of the 19th ACM UbiComp*. 1–19.

[20] Aapo Hyvärinen and Erkki Oja. 2000. Independent Component Analysis: Algorithms and Applications. *Neural Networks* 13, 4-5 (2000), 411–430.

[21] IEFT. 2017. Precision Time Protocol Version 2 (PTPv2). Accessed: 2022-04-30.

[22] Infineon Technologies AG. 2022. 24GHz Radar Sensors. https://www.infineon.com/cms/en/product/sensor/radar-sensors/radar-sensors-for-iot/24ghz-radar/. Accessed: 2022-04-24.

[23] Sergey Ioffe and Christian Szegedy. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *Proc. of ICML*. PMLR, 448–456.

[24] Zhenhua Jia, Amelie Bonde, Sugang Li, Chenren Xu, Jingxian Wang, Yanyong Zhang, Richard E. Howard, and Pei Zhang. 2017. Monitoring a Person's Heart Rate and Respiratory Rate on a Shared Bed using Geophones. In *Proc. of the 15th ACM SenSys*. 1–14.

[25] Ossi Kaltiokallio, Hüseyin Yiğitler, Riku Jäntti, and Neal Patwari. 2014. Non-Invasive Respiration Rate Monitoring using a Single COTS TX-RX Pair. In *Proc. of the 13th ACM IPSN*. IEEE, 59–69.

[26] Diederik P Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *Proc. of ICLR*.

[27] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Proc. of NIPS*, F. Pereira, C.J. Burges, L. Bottou, and K.Q. Weinberger (Eds.), Vol. 25.

[28] Hyunjae Lee, ByungHyun Kim, JinKwan Park, and JongGwan Yook. 2019. A Novel Vital-Sign Sensing Algorithm for Multiple Subjects Based on 24-GHz FMCW Doppler Radar. *Remote Sensing* 11, 10 (2019).

[29] Nadav Levanon. 1988. *Radar Principles*. Wiley.

[30] Qiu-Hua Lin, Yong-Rui Zheng, Fu-Liang Yin, Hualou Liang, and Vince D. Calhoun. 2007. A Fast Algorithm for One-unit ICA-R. *Information Sciences* 177, 5 (2007), 1265–1275.

[31] Jian Liu, Yan Wang, Yingying Chen, Jie Yang, Xu Chen, and Jerry Cheng. 2015. Tracking Vital Signs during Sleep Leveraging Off-the-shelf WiFi. In *Proc. of the 16th ACM MobiHoc*. 267–276.

[32] Xuefeng Liu, Jiannong Cao, Shaojie Tang, and Jiaqi Wen. 2014. Wi-Sleep: Contactless Sleep Monitoring via WiFi Signals. In *2014 IEEE Real-Time Systems Symposium*. 346–355.

[33] Xuefeng Liu, Jiannong Cao, Shaojie Tang, Jiaqi Wen, and Peng Guo. 2016. Contactless Respiration Monitoring Via Off-the-Shelf WiFi Devices. *IEEE Transactions on Mobile Computing* 15, 10 (2016), 2466–2479.

[34] Yongsen Ma, Gang Zhou, and Shuangquan Wang. 2019. WiFi Sensing with Channel State Information: A Survey. *ACM Computing Surveys (CSUR)* 52, 3 (2019), 1–36.

[35] Se Dong Min, Yonghyeon Yun, and Hangsik Shin. 2014. Simplified Structural Textile Respiration Sensor based on Capacitive Pressure Sensing Method. *IEEE Sensors Journal* 14, 9 (2014), 3245–3251.

[36] NeuLog. 2017. Respiration Monitor Belt Logger Sensor NUL-236. https://neulog.com/respiration-monitor-belt/. Accessed: 2022-04-28.

[37] Phuc Nguyen, Xinyu Zhang, Ann Halbower, and Tam Vu. 2016. Continuous and Fine-Grained Breathing Volume Monitoring from Afar Using Wireless Signals. In *Proc. of the 35th IEEE INFOCOM*. 1–9.

[38] A Pai, A Veeraraghavan, and A. Sabharwal. 2021. HRVCam: Robust Camera-based Measurement of Heart Rate Variability. *J. Biomed Opt* 26 (2021), 1–23.

[39] Kwang Suk Park and Sang Ho Choi. 2019. Smart Technologies toward Sleep Monitoring at Home. *Biomedical Engineering Letters* 9, 1 (2019), 73–85.

[40] Neal Patwari, Lara Brewer, Quinn Tate, Ossi Kaltiokallio, and Maurizio Bocca. 2014. Breathfinding: A Wireless Network That Monitors and Locates Breathing in a Home. *IEEE Journal of Selected Topics in Signal Processing* 8, 1 (2014), 30–42.

[41] William H. Press and Saul A. Teukolsky. 1990. Savitzky-Golay Smoothing Filters. *Computers in Physics* 4, 6 (1990), 669–672.

[42] Qualcomm Technologies, Inc. . 2022. Qualcomm 802.11ad 60GHz WiFi. https://www.qualcomm.com/products/features/80211ad. Accessed: 2022-05-28.

[43] A. Raji, P. Kanchana Devi, P. Golda Jeyaseeli, and N. Balaganesh. 2016. Respiratory Monitoring System for Asthma Patients based on IoT. In *IC-GET*. 1–6.

[44] B.D. Rao and K.V.S. Hari. 1989. Performance Analysis of Root-Music. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 37, 12 (1989), 1939–1949.

[45] Syed Tauhid Ullah Shah, Faizan Badshah, Faheem Dad, Nouman Amin, and Mian Ahmad Jan. 2019. Cloud-assisted IoT-based Smart Respiratory Monitoring System for Asthma Patients. In *Applications of Intelligent Technologies in Healthcare*. Springer, 77–86.

[46] Xingzhe Song, Boyuan Yang, Ge Yang, Ruirong Chen, Erick Forno, Wei Chen, and Wei Gao. 2020. SpiroSonic: Monitoring Human Lung Function via Acoustic Sensing on Commodity Smartphones. In *Proc. of The 26th ACM MobiCom*. 1–14.

[47] Shigeyuki Tateno, Xia Guan, Rui Cao, and Zhaoxian Qu. 2018. Development of Drowsiness Detection System based on Respiration Changes using Heart Rate Monitoring. In *57th SICE*. 1664–1669.

[48] Texas Instruments. 2020. IWR1843BOOST. https://www.ti.com/store/ti/en/p/product/?p=IWR1843BOOST. Accessed: 2022-05-30.

[49] Anran Wang, Jacob E. Sunshine, and Shyamnath Gollakota. 2019. Contactless Infant Monitoring Using White Noise. In *Proc. of the 25th ACM MobiCom*. 52:1–16.

[50] Fengyu Wang, Feng Zhang, Chenshu Wu, Beibei Wang, and K. J. Ray Liu. 2021. ViMo: Multiperson Vital Sign Monitoring Using Commodity Millimeter-Wave Radio. *IEEE Internet of Things Journal* 8, 3 (2021), 1294–1307.

[51] Tianben Wang, Daqing Zhang, Yuanqing Zheng, Tao Gu, Xingshe Zhou, and Bernadette Dorizzi. 2018. C-FMCW based Contactless Respiration Detection using Acoustic Signal. In *Proc. of the 20th ACM UbiComp*. 170:1–20.

[52] Xuanzhi Wang, Kai Niu, Jie Xiong, Bochong Qian, Zhiyun Yao, Tairong Lou, and Daqing Zhang. 2022. Placement Matters: Understanding the Effects of Device Placement for WiFi Sensing. In *Proc. of the 24th ACM UbiComp*, Vol. 6. 1–25.

[53] Xuyu Wang, Chao Yang, and Shiwen Mao. 2017. PhaseBeat: Exploiting CSI Phase Data for Vital Sign Monitoring with Commodity WiFi Devices. In *Proc. of the 37th IEEE ICDCS*. 1230–1239.

[54] Chenshu Wu, Zheng Yang, Zimu Zhou, Xuefeng Liu, Yunhao Liu, and Jiannong Cao. 2015. Non-Invasive Detection of Moving and Stationary Human with WiFi. *IEEE Journal on Selected Areas in Communications* 33, 11 (2015), 2329–2342.

[55] Junjun Xiong, Hong Hong, Hongqiang Zhang, Ning Wang, Hui Chu, and Xiaohua Zhu. 2020. Multitarget Respiration Detection With Adaptive Digital Beamforming Technique Based on SIMO Radar. *IEEE Transactions on Microwave Theory and Techniques* 68, 11 (2020), 4814–4824.

[56] Xiangyu Xu, Jiadi Yu, Yingying Chen, Yanmin Zhu, Linghe Kong, and Minglu Li. 2019. BreathListener: Fine-Grained Breathing Monitoring in Driving Environments Utilizing Acoustic Signals. In *Proc. of the 17th ACM MobiSys*. 54–66.

[57] Chao Yang, Xuyu Wang, and Shiwen Mao. 2020. Respiration Monitoring with RFID in Driving Environments. *IEEE Journal on Selected Areas in Communications* 39, 2 (2020), 500–512.

[58] Yanni Yang, Jiannong Cao, Xiulong Liu, and Xuefeng Liu. 2019. Multi-Breath: Separate Respiration Monitoring for Multiple Persons with UWB Radar. In *Proc. of the IEEE 43rd COMPSAC*, Vol. 1. 840–849.

[59] Yanni Yang, Jiannong Cao, Xuefeng Liu, and Kai Xing. 2018. Multi-person Sleeping Respiration Monitoring with COTS WiFi Devices. In *IEEE 15th MASS*. 37–45.

[60] Zhicheng Yang, Parth H. Pathak, Yunze Zeng, Xixi Liran, and Prasant Mohapatra. 2016. Monitoring Vital Signs Using Millimeter Wave. In *Proc. of the 17th ACM MobiHoc*. 211–220.

[61] Zhicheng Yang, Parth H. Pathak, Yunze Zeng, Xixi Liran, and Prasant Mohapatra. 2017. Vital Sign and Sleep Monitoring using Millimeter Wave. *ACM Transactions on Sensor Networks* 13, 2 (2017), 1–32.

[62] Zitong Yu, Wei Peng, Xiaobai Li, Xiaopeng Hong, and Guoying Zhao. 2019. Remote Heart Rate Measurement from Highly Compressed Facial Videos: an End-to-End Deep Learning Solution with Video Enhancement. In *Proc. of the IEEE/CVF ICCV*. 151–160.

[63] Shichao Yue, Hao He, Hao Wang, Hariharan Rahul, and Dina Katabi. 2018. Extracting Multi-Person Respiration from Entangled RF Signals. In *Proc. of the 20th ACM UbiComp*. 86:1–22.

[64] Pong C. Yuen and Jian-Huang Lai. 2002. Face Representation using Independent Component Analysis. *Pattern Recognition* 35, 6 (2002), 1247–1257.

[65] Youwei Zeng, Dan Wu, Ruiyang Gao, Tao Gu, and Daqing Zhang. 2018. FullBreathe: Full Human Respiration Detection Exploiting Complementarity of CSI Phase and Amplitude of WiFi Signals. *Proc. of the 20th ACM UbiComp* (2018), 148.

[66] Youwei Zeng, Dan Wu, Jie Xiong, Jinyi Liu, Zhaopeng Liu, and Daqing Zhang. 2020. MultiSense: Enabling Multi-Person Respiration Sensing with Commodity WiFi. In *Proc. of the 22th UbiComp*. Article 102, 29 pages.

[67] Jin Zhang, Weitao Xu, Wen Hu, and Salil S. Kanhere. 2017. WiCare: Towards in-Situ Breath Monitoring. In *Proc. of the 14th EAI MobiQuitous*. 126–135.

[68] Tianyue Zheng, Zhe Chen, Chao Cai, Jun Luo, and Xu Zhang. 2020. V^2iFi: in-Vehicle Vital Sign Monitoring via Compact RF Sensing. In *Proc. of the 22th ACM UbiComp*. 70:1–27.