Hiding Identity, Preserving Respiration: Semantic-Decoupled Privacy Protection for mmWave Radar Sensing

XIN-YU LI, Hefei University of Technology, China
JINYANG HUANG*, Hefei University of Technology, China
MENG WANG, Hefei University of Technology, China
PENG ZHAO, Hefei University of Technology, China
FENG-QI CUI, University of Science and Technology of China, China
YUANHAO FENG, The University of Electro-Communications, Japan
ZHENG GONG, Tianjin University, China
FEIYU HAN, Nanjing University of Information Science and Technology, China
FUSANG ZHANG, Beihang University, Key Laboratory of Space-Air-Ground Integrated Network Security, Ministry of Industry and Information Technology, and Inspur Computer Technology Co., Ltd, China

Millimeter-wave (mmWave) radar-based respiratory sensing has attracted increasing attention due to its contactless, unobtrusive nature and high sensitivity to micro-motions. However, mmWave radar signals inevitably encode identity-sensitive information beyond respiration, raising serious privacy concerns. Existing privacy protection strategies are often coarse-grained or indiscriminately applied, leading to substantial degradation in respiratory sensing performance. To address these issues, we propose P^2CRM, an identity-preserving mmWave radar–based contactless respiratory sensing system that explicitly balances sensing performance and identity protection. Specifically, by disentangling respiration rhythm from rhythm-irrelevant components, we design a Wavelet Packet Decomposition–based Semantic Decoupling (*WSD*) method to transform radar signals into a structured semantic representation. Building on this representation, a Key-Controlled Semantic Mixing (*KCSM*) mechanism is introduced to selectively perturb rhythm-irrelevant semantics to suppress identity cues while preserving respiration-critical information. While such semantic-level perturbation reduces identity leakage, it inevitably introduces distortions that impair respiratory sensing accuracy. To compensate for this perturbation-induced degradation, we further propose a Cross-Security Respiratory Knowledge Transfer (*CRKT*) network that transfers respiration-relevant knowledge from semantically rich signals to their perturbed representations, enabling accurate respiratory sensing directly from perturbed signals. Extensive experiments show that P^2CRM reduces identity recognition accuracy from 83.38% to 28.47%, while maintaining a respiratory sensing MAE of 1.08 bpm. These results demonstrate that P^2CRM achieves a favorable balance between reduced

*Corresponding author. E-mail: hjy@hfut.edu.cn.

Authors' Contact Information: Xin-Yu Li, xinyuli@mail.hfut.edu.cn, Hefei University of Technology, Hefei, Anhui, China; Jinyang Huang, hjy@hfut.edu.cn, Hefei University of Technology, Hefei, Anhui, China; Meng Wang, mengw512@mail.hfut.edu.cn, Hefei University of Technology, Hefei, Anhui, China; Peng Zhao, ZhaoPengWork@mail.hfut.edu.cn, Hefei University of Technology, Hefei, Anhui, China; Feng-Qi Cui, fengqi_cui@mail.ustc.edu.cn, University of Science and Technology of China, Hefei, Anhui, China; Yuanhao Feng, fengyuanhao@uec.ac.jp, The University of Electro-Communications, Tokyo, Japan; Zheng Gong, gongzheng@tju.edu.cn, Tianjin University, Tianjin, China; Feiyu Han, fyhan@nuist.edu.cn, Nanjing University of Information Science and Technology, Nanjing, China; Fusang Zhang, Beihang University, Beijing and Key Laboratory of Space-Air-Ground Integrated Network Security, Ministry of Industry and Information Technology and Inspur Computer Technology Co., Ltd, China, zhangfusang@buaa.edu.cn.

ACM 2476-1249/2026/8-ART111
https://doi.org/XXXXXXX.XXXXXXX

identity leakage while preserving respiratory sensing performance, enabling trustworthy and identity-preserving contactless respiratory monitoring. Our code is available at: https://github.com/xiaozhu990724-lxy/Trusted-BR-monitoring-P2CRM-.

CCS Concepts: • **Human-centered computing** → *Ubiquitous and mobile computing systems and tools.*

Additional Key Words and Phrases: Respiratory Sensing, Privacy Protection, mmWave Radar, Knowledge Transfer

ACM Reference Format:
Xin-Yu Li, Jinyang Huang, Meng Wang, Peng Zhao, Feng-Qi Cui, Yuanhao Feng, Zheng Gong, Feiyu Han, and Fusang Zhang. 2026. Hiding Identity, Preserving Respiration: Semantic-Decoupled Privacy Protection for mmWave Radar Sensing. *Proc. ACM Meas. Anal. Comput. Syst.* 37, 4, Article 111 (August 2026), 28 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

Millimeter-wave (mmWave) radar has gained increasing attention for respiratory sensing due to its contactless, user-friendly, and high sensitivity to micro-motions [6, 7, 28, 33, 60]. By capturing the subtle chest wall movements caused by respiration, mmWave radar enables the extraction of respiratory biosignals. This approach has achieved notable progress in various respiratory monitoring applications, including respiratory rate estimation [6, 33], sleep apnea monitoring [47], and pulmonary function assessment [14, 55].

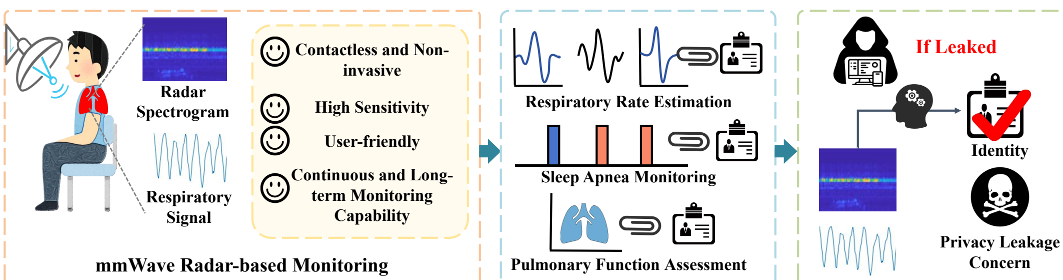

Fig. 1. Overview of mmWave radar-based respiratory monitoring. mmWave radar enables contactless, user-friendly, and long-term respiratory monitoring, supporting applications such as respiratory rate estimation, sleep apnea detection, and pulmonary function assessment. However, radar signals inherently encode fine-grained individual-specific information, thereby raising concerns about potential identity leakage.

Although mmWave radar enables high-accuracy respiratory sensing, the captured signals inevitably encode identity-sensitive information [21] that can be exploited for identity inference [5, 12, 26, 27, 30, 34, 36, 52, 59, 61], as shown in Fig. 1. Whether raw radar signals are temporarily stored on the device for window-based processing, uploaded to the cloud for advanced analytics, or released as public datasets for respiratory sensing research, they remain vulnerable to unauthorized access or leakage [3, 35]. This raises critical privacy concerns, as such exposure could lead to unintended identity disclosure. Although several existing studies have explored privacy-preserving techniques for radar-based sensing [29, 31], they typically adopt coarse-grained or indiscriminate protection strategies, which inevitably disrupt sensing-related semantics and consequently degrade sensing performance [1, 50]. Therefore, it is necessary to develop an identity-preserving mechanism that selectively suppresses identity-sensitive information while preserving sensing-critical semantics.

However, achieving effective identity protection without compromising sensing performance presents significant challenges. Firstly, disentangling sensing-critical features and identity-related features is inherently difficult due to the strong coupling between these components. Secondly, the identity-oriented perturbation mechanism must be carefully designed to preserve the integrity of the sensing-critical semantics while suppressing identity-relevant information. Finally, due to the inevitable temporal and spectral distortions introduced by the perturbation process, accurately extracting respiratory features from identity-preserved signals remains a significant challenge.

Our insight stems from the intrinsic priors in the signal structure, which consists of multiple components encoding distinct semantic meanings. First, mmWave radar signals inherently comprise both respiration-related and non-respiratory components during respiratory sensing. Within the respiration-related component, the signal can be further decomposed into two distinct aspects: the respiratory rhythm and the respiratory pattern. Specifically, the rhythm reflects the frequency and temporal periodicity of respiration, while the pattern exhibits individual-specific morphological characteristics and fine-grained waveform structures. These pattern-level variations can therefore be exploited as cues for identity recognition [5, 12, 26, 27, 30, 36, 52]. Second, although the rhythm is coupled with other signal components, they capture different aspects of the underlying signal content. This distinction motivates separating the rhythm from rhythm-irrelevant components and selectively perturbing the latter. When this separation is sufficiently accurate, identity-related cues can be anonymized while preserving the rhythm information required for respiratory sensing.

Building upon these foundational insights, we propose P^2CRM, an identity-preserving system for non-contact respiratory sensing. Specifically, P^2CRM integrates three key components: signal disentanglement, semantic perturbation, and robust respiratory sensing under identity constraints, to achieve both identity protection and efficient respiratory sensing. At the signal level, a Wavelet Packet Decomposition (WPD)-based semantic decoupling (*WSD*) is applied to transform raw radar measurements into a structured representation. Specifically, the signal is first decomposed into the respiratory-relevant signal, low-frequency drift, and high-frequency noise, and the respiratory-relevant signal is subsequently decomposed into a rhythm component and a pattern component. Building upon these decoupled semantic representations, a Key-controlled Semantic Mixing Scheme (*KCSM*) is introduced to enable selective identity perturbation in the semantic space. Specifically, it performs component-wise semantic perturbation with scheduled perturbation strength applied to rhythm-irrelevant components, to preserve respiratory sensing utility while suppressing identity-related cues beyond rhythm semantics. Although the respiration-dominant rhythm component is preserved, the perturbation process inevitably introduces temporal distortions and spectral irregularities across semantic components, which undermine the assumptions of conventional signal processing methods for respiratory rate estimation. To compensate for semantic distortions introduced by perturbation, a Cross-security Respiratory Knowledge Transfer (*CRKT*) strategy based on a teacher–student paradigm is employed. Specifically, a teacher model is trained on semantically rich signals to learn robust respiratory representations, and the learned knowledge is transferred to guide the student model to estimate respiratory rate directly from identity-preserved signals.

We implement P^2CRM on a commercial mmWave radar (TI AWR 1843 BOOST) and evaluate it on a dataset of 20 subjects, including 9 males and 11 females. Experimental results demonstrate that P^2CRM substantially reduces identity-recognition accuracy while maintaining accurate respiratory-rate estimation. Furthermore, experiments across a wide range of conditions (*e.g.*, sensing distances, respiratory states, body postures, body types, orientation angles, sample durations, environments, health conditions, and age groups, *etc.*) demonstrate the robustness of the proposed system. Our major contributions can be summarized as follows:

- We propose P^2CRM, a novel identity-preserving contactless respiratory sensing system based on commercial mmWave radar, which achieves a favorable balance between effective identity protection and accurate, robust respiratory monitoring.
- We propose a series of techniques to address the challenges of identity-preserving respiratory sensing. Specifically, we introduce a WPD-based Semantic Decoupling (*WSD*) method to decompose raw radar signals into meaningful semantic components, a Key-controlled Semantic Mixing (*KCSM*) mechanism to selectively perturb respiration rhythm-irrelevant components, and a Cross-security Respiratory Knowledge

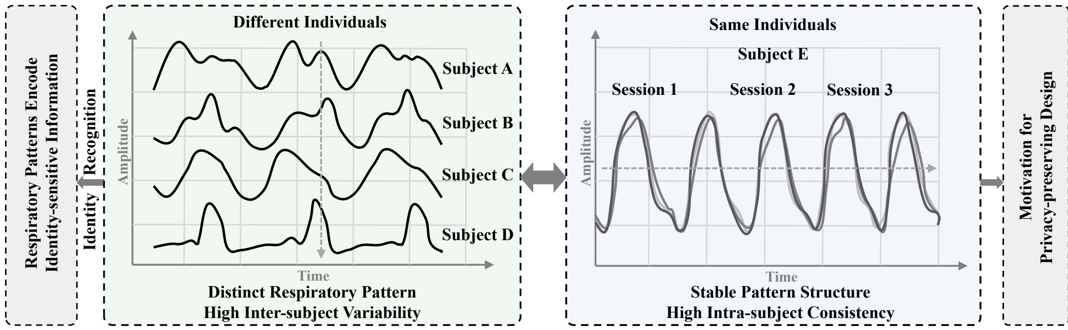

Fig. 2. Comparison of respiratory patterns. Respiratory patterns exhibit pronounced inter-subject variability, whereas those of the same subject recorded across different sessions remain highly consistent.

Transfer (*CRKT*) network to transfer respiration-dominant knowledge from semantically rich signals to identity-preserved signals, thereby enabling robust respiratory sensing under privacy constraints.

- Extensive experimental results demonstrate that the proposed system achieves remarkable performance in both identity protection and respiratory sensing, attaining a mean absolute error (MAE) of 1.08 bpm and a standard deviation (STD) of 0.31 bpm for respiratory sensing, while reducing identity recognition accuracy from 83.38% to 28.47%, indicating its potential for secure and accurate long-term health monitoring in real-world environments.

2 Preliminaries

2.1 mmWave Sensing

The mmWave radar works by transmitting frequency-modulated continuous waves (FMCW), *i.e.*, chirps, and detecting the reflections from surrounding objects. When the transmitted signals interact with the objects, they are reflected back to the radar, and the system measures the time delay of these reflections. The received signals are mixed with transmitted signals to obtain intermediate frequency (IF) signal X_{IF} in the equation [19]:

$$X_{IF}(i,t) \approx \alpha_i \exp[-j4\pi(f_c + Bt/T_c)d(t)/c], \tag{1}$$

where i represents the i-th reflection path channel. α_i corresponds to the complex path loss associated with the i-th channel. f_c indicates the initial frequency of the chirp. c is the speed of light, and B is the bandwidth of sweeping frequency. T_c refers to the duration of the chirp. $d(t)$ represents the distance from the detected object to the mmWave radar. By performing a Range-FFT operation [24] on the IF signals, the frequency and phase of these IF signals can be extracted:

$$f = 2Bd(t)/(cT_c), \phi(t) = 4\pi d(t)/(c/f_c). \tag{2}$$

Based on the above equations, the IF signal allows for object localization with a resolution of $\frac{c}{2B}$. In addition, it is capable of detecting micrometer-scale changes in distance via the phase $\phi(t)$ [24], which can be applied to monitor subtle chest-wall movements induced by respiration. We use the complex-valued representations of mmWave Range-FFT results to preserve information from both amplitude and phase aspects [19].

2.2 Privacy in mmWave Respiratory Sensing

Respiratory signals captured by mmWave radar encode rich physiological information manifested through structured temporal dynamics and diverse morphological patterns [5, 26], such as breathing cycle duration,

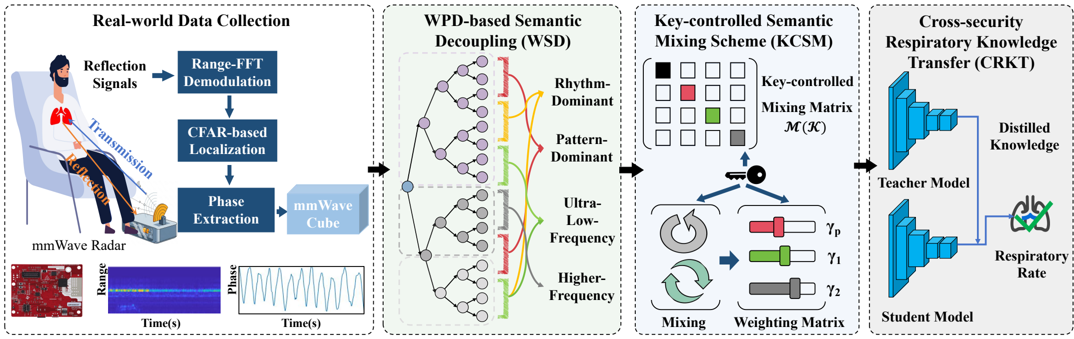

Fig. 3. Overview of P^2CRM. The radar signals are semantically decoupled, selectively protected, and exploited for robust respiratory sensing under identity constraints.

inhalation–exhalation asymmetry, amplitude modulation, waveform shape variation, *etc.* These signal characteristics arise from the combined influence of anatomical factors (*e.g.*, thoracic structure, lung capacity, *etc.*), neuromuscular control mechanisms, and habitual breathing behaviors [4]. Prior studies have shown that these characteristics often exhibit pronounced subject-specific traits, making them exploitable for identity recognition and biometric inference [5, 12, 23, 26, 30, 36, 52, 59, 61]. As shown in Fig. 2, substantial inter-subject variability can be observed in respiratory waveforms, while recordings from the same individual across different sessions remain highly consistent. This pronounced inter-subject distinctiveness and intra-subject stability indicate that respiratory signals inherently carry identifiable biometric signatures, posing potential privacy risks when utilized without adequate protection.

3 System Overview

P^2CRM is an identity-preserving contactless respiratory sensing system based on mmWave radar. As shown in Fig. 3, P^2CRM first collects mmWave reflections from the subject's chest region and performs Range-FFT demodulation to obtain slow-time signals. The extracted signals are then fed to the three core components of P^2CRM: the WPD-based Semantic Decoupling (*WSD*) method, Key-Controlled Semantic Mixing (*KCSM*) mechanism, and Cross-Security Respiratory Knowledge Transfer (*CRKT*) network.

The *WSD* aims to transform the demodulated phase signals into a semantically structured representation. By employing a WPD [20], the phase signals are decomposed into a set of uniformly partitioned subbands, which are then flexibly regrouped according to their physiological relevance to isolate respiration-dominant components from ultra-low-frequency drift and high-frequency residuals. Within the respiration-dominant band, an intra-band semantic decoupling strategy is further employed to explicitly separate respiration rhythm information and respiratory pattern information.

The *KCSM* is a lightweight yet effective identity-preserving mechanism designed to enable selective semantic protection. By applying key-controlled perturbations with component-wise strength scheduling directly in the semantic token space, *KCSM* selectively suppresses sensitive identity-related information, while preserving the respiration rhythm semantics that are critical for reliable respiratory sensing.

The *CRKT* is designed to enable robust respiratory sensing under identity-preserving constraints by transferring respiration-relevant semantics across security domains. Instead of directly training a model solely on identity-preserved signals, *CRKT* adopts a teacher–student paradigm [49], where respiration-dominant knowledge learned from semantically rich representations is used to guide the learning process on identity-preserved representations.

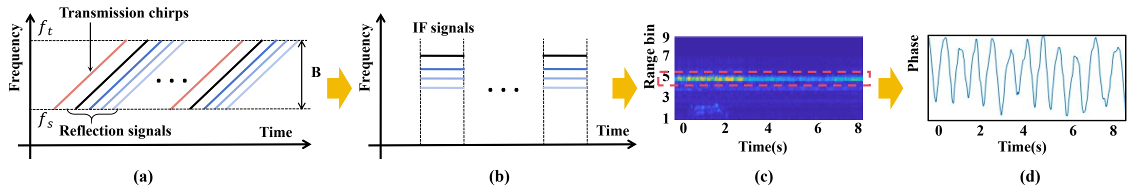

Fig. 4. Data collection process. (a). Transmission and reflection signals from the target; (b). The IF signals; (c). Targeted range bin localization; (d). Phase signal.

This cross-security knowledge transfer enables the student model to capture effective respiratory semantics from identity-preserved representations.

4 System Design

4.1 Data Collection

During data collection, we employ a commercial mmWave radar to transmit FMCW chirps (Fig. 4a) and obtain the corresponding IF signals (Fig. 4b). To extract range information from the IF signals, we apply a Range-FFT along the fast-time dimension, which transforms the signals into the range domain [19]. Subsequently, we adopt a Constant False Alarm Rate (CFAR) detection algorithm [37] to identify candidate regions associated with the human chest and then select the range bin with the largest reflection energy (Fig. 4c).

It is worth noting that the phase of the signal is highly sensitive to subtle movements, thereby offering higher precision and reliability in detecting fine-grained human body motions [15]. Therefore, we explicitly extract and process the phase information from the signals of the targeted range bin to enhance the sensitivity to respiration-induced micro-movements, as illustrated in Fig. 4d.

4.2 WPD-based Semantic Decoupling (*WSD*) Architecture

The phase signal extracted from mmWave radar data is a composite observation that integrates multiple heterogeneous components with distinct physical origins and semantic meanings. Directly applying coarse-grained or indiscriminate perturbation to such composite signals can inevitably distort sensing-relevant semantics, thereby severely degrading respiratory sensing performance.

In this section, we propose the *WSD* to enable multi-resolution decomposition of signals, thereby disentangling sensing-relevant semantics from heterogeneous signal components. Instead of conventional wavelet transforms or Fourier-based filtering that typically adopt fixed or predefined frequency partitions, we leverage WPD [20] to achieve uniform and fine-grained frequency decomposition, as illustrated in Fig. 5. This design enables flexible semantic decoupling and regrouping of subbands according to their semantic relevance.

4.2.1 Semantic Band Decoupling and Regrouping. To transform the composite signal into a form amenable to semantic manipulation, we first perform a multi-resolution decomposition using WPD [20]. Unlike fixed band-pass filtering, which imposes rigid frequency boundaries and is sensitive to non-stationarity, WPD recursively decomposes phase signal $X(t)$ into a complete binary tree of subband signals, thereby providing a uniform and fine-grained partition:

$$X(t) = \sum_{n=0}^{2^L-1} X_{L,n}(t) = \sum_{n=0}^{2^L-1} \sum_{k} c_{L,n}(k)\psi_{L,n,k}(t), \tag{3}$$

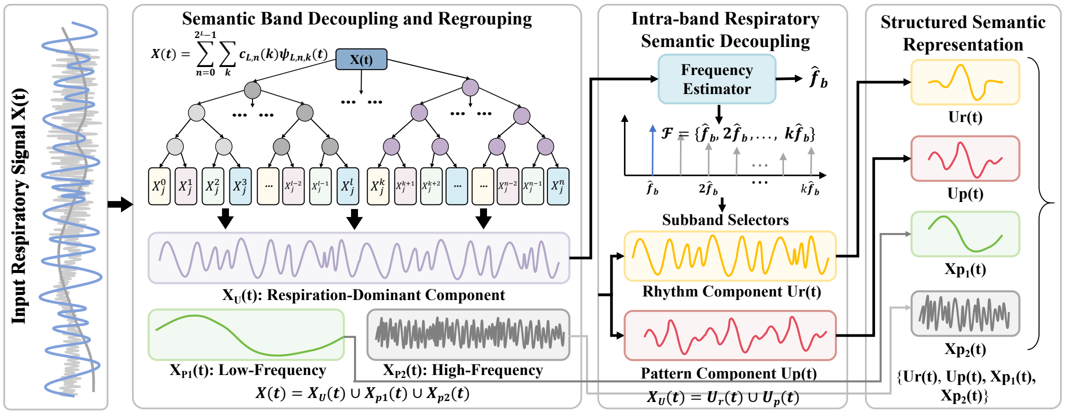

Fig. 5. Semantic decoupling architecture.

where $X_{L,n}(t)$ denotes the reconstructed subband signal at the n-th node of the L-level wavelet packet tree, $c_{L,n}(k)$ represents the wavelet packet coefficients, and $\psi_{L,n,k}(t)$ denotes the corresponding wavelet packet basis functions. In our implementation, we use the Daubechies 4 (db4) wavelet family and set the decomposition depth to $L = 5$.

Guided by the frequency characteristics exposed through WPD decomposition, the decomposed subbands are reorganized into three semantically interpretable components:

$$X_U(t) = \sum_{n \in S_U} X_{L,n}(t), \qquad X_{P1}(t) = \sum_{n \in S_{P1}} X_{L,n}(t), \qquad X_{P2}(t) = \sum_{n \in S_{P2}} X_{L,n}(t), \tag{4}$$

where $X_U(t)$ denotes the respiration-dominant component corresponding to the respiratory frequency range, $X_{P1}(t)$ captures ultra-low-frequency variations such as baseline drift, and $X_{P2}(t)$ aggregates higher-frequency residual components arising from non-respiratory signal variations and noise. The sets S_U, S_{P1} and S_{P2} denote the wavelet packet node indices assigned to each semantic band, respectively.

4.2.2 Intra-band Respiratory Semantic Decoupling. To enable finer-grained semantic separation within $X_U(t)$, we introduce an intra-band respiratory semantic decoupling mechanism to distinguish respiratory rhythm semantics from identity-sensitive respiratory pattern components. This is motivated by the fact that respiratory signals inherently comprise both respiratory rhythm information and cycle-level structural patterns, the latter of which demonstrate strong subject dependency [5, 12, 26, 27, 30, 36, 52].

Formally, $X_U(t)$ is modeled as the superposition of two latent semantic components:

$$X_U(t) = U_r(t) \cup U_p(t), \tag{5}$$

where $U_r(t)$ denotes the rhythm-dominant component encoding respiratory frequency information, and $U_p(t)$ represents the pattern-dominant component associated with waveform morphology and individual-specific variations.

Given the respiration-dominant signal $X_U(t)$, we first obtain a window-level respiratory frequency estimate $\hat{f}_b$ using a spectral peak-based method. Based on $\hat{f}_b$, we construct a target frequency set consisting of the fundamental frequency and its harmonics:

$$\mathcal{F} = \{\hat{f}_b, 2\hat{f}_b, \ldots, k\hat{f}_b\} \cap [0.1, 1.0], \tag{6}$$

where k denotes the number of harmonics considered, and the frequency range $[0.1, 1.0]$ Hz corresponds to the physiologically plausible band for respiratory-related components.

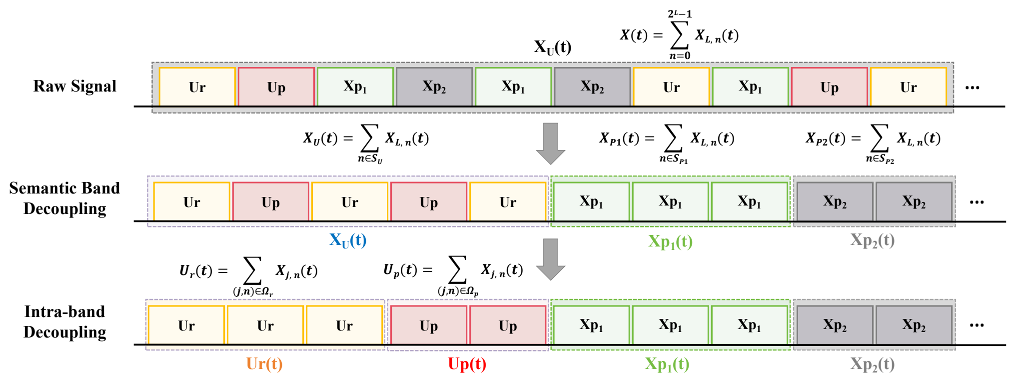

Fig. 6. Illustration of semantic component separation.

For each level-L wavelet packet node n contributing to $X_U(t)$, with frequency support $[f^{\ell}_{L,n}, f^{h}_{L,n}]$, we assign it to the rhythm set Ω_r if its band overlaps with a tolerance neighborhood of any target respiratory frequency in $\mathcal{F}$:

$$n \in \Omega_r \iff \exists f \in \mathcal{F} : f^{\ell}_{L,n} \leq f + \delta \wedge f^{h}_{L,n} \geq f - \delta, \tag{7}$$

where δ controls the allowable frequency deviation. All remaining subbands within the respiration-dominant range are assigned to the pattern set Ω_p.

The rhythm and pattern components are then reconstructed as:

$$U_r(t) = \sum_{n\in\Omega_r} X_{L,n}(t), \qquad U_p(t) = \sum_{n\in\Omega_p} X_{L,n}(t), \tag{8}$$

where $X_{L,n}(t)$ denotes the reconstructed time-domain signal corresponding to node n at decomposition depth L.

4.2.3 Semantic-driven Respiratory Reorganization. Based on the intra-band semantic decoupling described above, the respiration-dominant signal is represented by two explicitly disentangled components, denoted as $U_r(t)$ and $U_p(t)$, which correspond to respiratory rhythm and respiratory pattern information, respectively. As shown in Fig. 6, together with other components $\{X_{P1}, X_{P2}\}$, the radar signal can be reformulated into a semantically structured representation as:

$$s(t) = [U_r(t), U_p(t), X_{P1}, X_{P2}]^{\mathrm{T}}. \tag{9}$$

4.3 Key-controlled Semantic Mixing (*KCSM*) Scheme

Building upon the semantic decoupling and reorganization described in Section 4.2, the radar signal is reformulated as a set of explicitly disentangled semantic components, including respiratory rhythm, respiratory pattern, and non-respiratory residuals.

To reduce identity leakage while preserving respiratory sensing utility, we introduce a *KCSM* scheme that operates directly in the semantic space. Rather than discarding semantic components, *KCSM* retains the complete semantic structure during perturbation. In contrast to conventional signal-level perturbation strategies, which indiscriminately perturb all signal components and inevitably impair task-relevant information, *KCSM* leverages the functional roles of decoupled semantic components and performs selective, key-driven perturbations, as shown in Fig. 7. Specifically, the rhythm-related component, which is essential for reliable respiratory rate estimation, is explicitly preserved, whereas components encoding respiratory pattern and non-respiratory interference are mixed in a key-dependent manner to suppress their separability and identity-related cues.

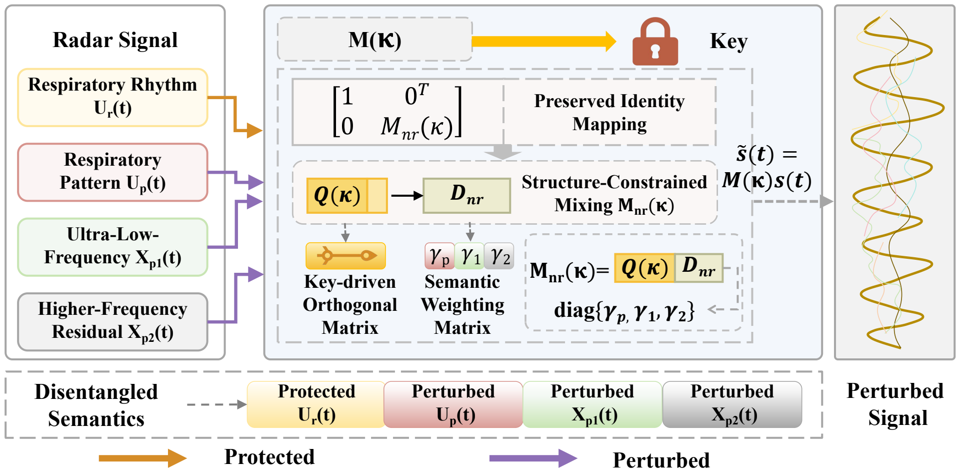

Fig. 7. Overview of *KCSM* scheme.

Given the semantic vector $\mathbf{s}(t)$, a key-dependent semantic mixing matrix $\mathbf{M}(\mathcal{K})$ is generated from a secret key $\mathcal{K}$:

$$\mathbf{M}(\mathcal{K}) \in \mathbb{R}^{4\times 4}, \tag{10}$$

where $\mathbf{M}(\mathcal{K})$ defines a linear transformation operating in the semantic space and remains fixed for a given key within each monitoring session.

To explicitly preserve respiration rhythm information while suppressing identity-sensitive semantics, $\mathbf{M}(\mathcal{K})$ is designed as a *structure-constrained mixing matrix* that selectively operates on non-rhythm components. Specifically, the released semantic representation is obtained as:

$$\tilde{\mathbf{s}}(t) = \mathbf{M}(\mathcal{K})\mathbf{s}(t), \qquad \mathbf{M}(\mathcal{K}) = \begin{bmatrix} 1 & \mathbf{0}^\top \\ \mathbf{0} & \mathbf{M}_{\mathrm{nr}}(\mathcal{K}) \end{bmatrix}, \tag{11}$$

the first dimension corresponds to the respiration rhythm component $U_r(t)$ and is preserved through an identity mapping, while $\mathbf{M}_{\mathrm{nr}}(\mathcal{K}) \in \mathbb{R}^{3\times 3}$ performs key-controlled mixing over the remaining non-rhythm semantic components, including respiration morphology and non-respiratory residuals.

To ensure numerical stability and controlled semantic contribution, the non-rhythm mixing matrix is further factorized as:

$$\mathbf{M}_{\mathrm{nr}}(\mathcal{K}) = \mathbf{Q}(\mathcal{K})\mathbf{D}_{\mathrm{nr}}, \tag{12}$$

where $\mathbf{Q}(\mathcal{K})$ is a key-driven random orthogonal matrix, ensuring energy preservation and preventing numerical amplification during semantic mixing. The diagonal matrix $\mathbf{D}_{\mathrm{nr}}$ is a semantic weighting matrix defined as:

$$\mathbf{D}_{\mathrm{nr}} = \mathrm{diag}\{\gamma_p, \gamma_1, \gamma_2\}, \tag{13}$$

where γ_p controls the weighting of the respiration pattern component $U_p(t)$, and γ_1 and γ_2 scale the low-frequency drift and high-frequency residual components, respectively. In our experiments, the baseline values are set to $\gamma_p = 0.35$, $\gamma_1 = 0.7$, and $\gamma_2 = 0.7$ to achieve a balanced trade-off between identity protection and respiratory sensing utility.

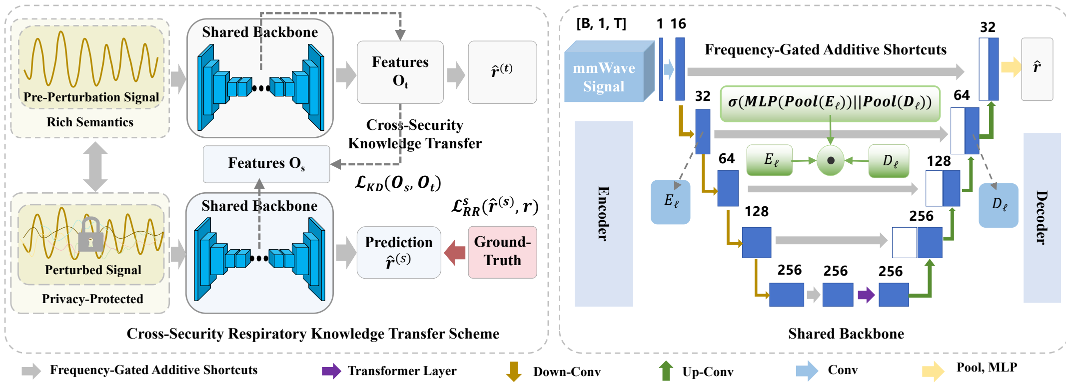

Fig. 8. Framework of cross-security knowledge transfer.

Finally, the mixed semantic representation is mapped back to a public time-domain signal through a fixed linear projection:

$$\tilde{x}(t) = \mathbf{w}^{\top}\tilde{\mathbf{s}}(t), \tag{14}$$

where $\mathbf{w}$ is a normalized weighting vector used to aggregate the released semantic components into a single observable signal.

Specifically, the secret key is generated and stored exclusively on the local radar device and is not transmitted to any third party. We assume that the perturbation construction is publicly known. The adversary may observe the released perturbed radar signals but does not possess the corresponding session key. A new independent key is generated for each monitoring session, avoiding key reuse and limiting the impact of key compromise to that session. Our security goal is empirical identity obfuscation rather than a formal cryptographic privacy guarantee.

4.4 Cross-security Respiratory Knowledge Transfer (*CRKT*) Scheme

It is worth noting that, although the rhythm-dominant component is preserved, the perturbation process inevitably introduces temporal distortions and spectral irregularities across semantic components, thereby undermining the assumptions of conventional signal processing methods for respiratory rate estimation.

To address this issue, we propose the *CRKT*, as illustrated in Fig. 8, which adopts a teacher–student learning paradigm [49] to transfer respiration-dominant knowledge from pre-perturbation radar representations to their perturbed radar representations. During the training stage, a teacher network is first trained on semantically rich pre-perturbation radar signals to learn a robust and discriminative respiratory representation. Knowledge extracted from its intermediate feature layers is then transferred to the student network through a cross-security distillation constraint, thereby regularizing the student to capture respiratory rate–invariant characteristics under perturbation-induced distortions. During the inference stage, the student model operates exclusively on perturbed radar signals, without access to the teacher network or any pre-perturbation information.

4.4.1 Shared Backbone. The *CRKT* is built upon a shared temporal backbone that is instantiated by both the teacher and the student networks. The backbone consists of two key components: a multi-resolution UNet–based time feature extractor for capturing respiration dynamics across different time scales, and a frequency-gated additive shortcut module that selectively propagates respiration-dominant information while suppressing perturbation-sensitive signal variations.

Multi-Resolution UNet-based Temporal Feature Extractor. To extract reliable respiratory rhythm information from radar signals under identity-preserving perturbation, we formulate perturbed respiratory sensing as a temporal regression task, in which the objective is to infer respiratory rate directly from radar-derived time sequences. It is worth noting that a robust temporal mapping should remain consistent across multiple time resolutions, as respiratory motion remains inherently periodic and manifests across different temporal scales.

Motivated by this temporal consistency requirement, we design a multi-resolution temporal feature extractor that explicitly captures respiratory dynamics across different time scales. To aggregate features from multi-resolution layers, we leverage a UNet architecture [39] with frequency-gated additive shortcuts to selectively propagate respiration-dominant information across scales while suppressing perturbation-sensitive distortions from the same resolution. As illustrated in Fig. 8, the encoder progressively encodes the input radar sequence using convolution operations with 5 blocks corresponding to different time resolutions. Each layer encodes the mmWave signals and downsamples them for the subsequent layer using convolutional kernels with a stride of $S = 2$. The final layer, with the largest receptive field, is further enhanced with Transformer layers [45] to leverage their superior capability in modeling time series relationships [19]. The decoder mirrors the encoder structure by progressively upsampling the downsampled embeddings from lower-resolution layers using transposed convolution kernels and combining them with the passed representations from the encoder at the same layer.

Frequency-Gated Additive Shortcuts. In the feature extractor, shortcut connections between the encoder and decoder play a critical role in concatenating feature maps to mitigate information loss across different time scales. Unlike typical UNet architectures that both the input and the output share the same dimension [39], the shortcuts in our feature extractor are designed to selectively transmit respiration-relevant information while suppressing residual non-respiratory components that may persist after identity-preserving perturbation.

To this end, we introduce frequency-gated additive shortcuts to regulate shortcut information flow and ensure effective cross-scale feature propagation. Specifically, given the encoder and decoder feature maps at the ℓ-th temporal resolution, denoted as $\mathbf{E}_\ell \in \mathbb{R}^{C_\ell \times T_\ell}$ and $\mathbf{D}_\ell \in \mathbb{R}^{C_\ell \times T_\ell}$, respectively, where C_ℓ and T_ℓ represent the number of feature channels and the temporal length at resolution level ℓ. The gating vector is calculated as follows:

$$\mathbf{g}_\ell = \sigma(MLP(Pool(\mathbf{E}_\ell)||Pool(\mathbf{D}_\ell))), \tag{15}$$

where $Pool(\cdot)$ denotes global temporal pooling that aggregates frequency-related responses over time, $||$ represents feature concatenation, and $\sigma(\cdot)$ is the sigmoid activation. The resulting gating vector $\mathbf{g}_\ell$ adaptively modulates the shortcut information flow in a channel-wise manner. The gated shortcut is then incorporated using an additive formulation:

$$\mathbf{D}_\ell \leftarrow \mathbf{D}_\ell + \mathbf{g}_\ell \odot \mathbf{E}_\ell, \tag{16}$$

where $\odot$ denotes element-wise multiplication. This additive gating mechanism enables selective reinforcement of respiration-dominant features while attenuating perturbation-sensitive residual components.

4.4.2 Teacher: The Pre-Perturbation Respiratory Model. The teacher model is designed to learn respiration-dominant representations from pre-perturbation radar signals, where the respiratory information is fully preserved. The teacher model adopts the shared backbone introduced in Section 4.4.1, without requiring any additional modules.

Given a pre-perturbation radar signal $\mathbf{x}$, the teacher network processes the input to produce a respiratory rate estimate $\hat{r}^{(t)}$, along with an intermediate feature representation $\mathbf{O}_t$ extracted from the bottleneck layer of the shared backbone. This representation encodes global respiratory rhythm characteristics and serves as the supervision source for subsequent cross-security knowledge transfer. To minimize the error of estimation results, the teacher model is optimized using a regression loss that quantifies the discrepancy between the predicted respiratory rate and the ground-truth label:

$$\mathcal{L}_{RR}^{t}(\hat{r}^{(t)}, r) = ||\hat{r}^{(t)} - r||_2^2, \tag{17}$$

where r denotes the ground-truth respiratory rate.

Table 1. Configuration of mmWave radar.

Parameters	Values	Parameters	Values
Starting Frequency	77 GHz	Bandwidth	4 GHz
RX Gain	50 dB	Idle Time	10 μs
Chirp Cycle Time	50 μs	Chirps/Frame	128
Frame Periodicity	50 ms	Samples/Chirp	256
ADC Sample Rate	8000K	Frequency Slope	80 MHz/μs

4.4.3 Student: The Perturbed Respiratory Model. The student model is responsible for estimating respiratory rate directly from perturbed radar signals and adopts the same backbone architecture as the teacher model.

Specifically, the student takes a perturbed radar signal $\tilde{\mathbf{x}} = \varepsilon(\mathbf{x})$ as input, where $\varepsilon(\cdot)$ denotes the proposed semantic decoupling and targeted perturbation process. During training, the student model is optimized under a cross-security knowledge transfer framework. For each paired sample $(\mathbf{x}, \tilde{\mathbf{x}})$, the student minimizes both the respiratory-rate prediction error and the distance between the intermediate representations extracted by the student and teacher models. The loss function $\mathcal{L}$ for the student model is:

$$\begin{aligned}
\mathcal{L}_{KD}(\mathbf{O}_s, \mathbf{O}_t) &= ||\mathbf{O}_s - \mathbf{O}_t||_2^2, \\
\mathcal{L}_{RR}^s(\hat{r}^{(s)}, r) &= ||\hat{r}^{(s)} - r||_2^2, \\
\mathcal{L}(\hat{r}^{(s)}, r, \mathbf{O}_s, \mathbf{O}_t) &= \alpha \mathcal{L}_{RR}^s(\hat{r}^{(s)}, r) + \beta \mathcal{L}_{KD}(\mathbf{O}_s, \mathbf{O}_t),
\end{aligned} \tag{18}$$

where $\mathbf{O}_s$ and $\mathbf{O}_t$ denote the intermediate representations extracted from the bottleneck layer of the student and teacher models, respectively. α and β are hyperparameters that balance the respiratory rate estimation accuracy and the strength of cross-security knowledge transfer.

4.4.4 Identity Recognition. To evaluate identity leakage before and after perturbation, we adopt multiple representative network architectures, including the 1D-CNN [22], ResNet-18 [16], CNN-LSTM [46], and Transformer [13]. Specifically, the 1D-CNN consists of several convolutional blocks, each containing a convolutional layer, batch normalization, and ReLU activation, to progressively extract features and capture identity-related representations. ResNet-18 effectively mitigates the vanishing gradient problem in deep networks through skip connections, enabling the extraction of richer high-level features. CNN-LSTM combines the local feature extraction capability of CNN with the temporal modeling ability of LSTM to capture long-term dependencies in radar signals. The Transformer uses a multi-head self-attention mechanism to model global relationships in the signal. All architectures are adapted for radar time-series classification and are trained separately on raw and perturbed signals to evaluate the identity leakage. Unless otherwise specified, the 1D-CNN is used as the default identity recognition model.

5 Implementation

Sensing Platform. We implement P^2CRM using a commercial off-the-shelf (COTS) mmWave radar, the TI AWR 1843 BOOST [43], together with a TI DCA1000EVM data capture board [44] for acquiring raw radar data, as illustrated in Fig. 9. The radar integrates three transmitting antennas and four receiving antennas, enabling multi-channel acquisition of reflected signals from the subject's chest region. The detailed radar configuration parameters used throughout the experiments are summarized in Tab. 1.

Ground Truth Platform. The ground truth respiratory information for P^2CRM is collected using a Bigrun Team respiratory belt, which provides a continuous measurement of respiratory activity by sensing thoracic

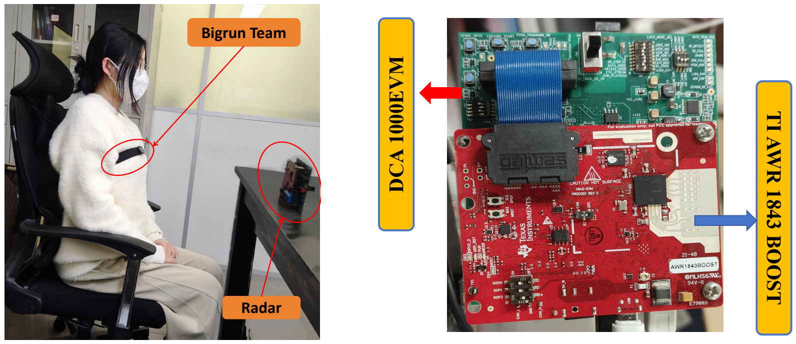

(a) Data Collection Setup.

(b) mmWave Radar Sensing Platform.

Fig. 9. Experimental setup for evaluating the proposed system.

expansion and contraction. The respiratory belt is worn around the subject's chest and does not interfere with the mmWave sensing process. The belt output data is recorded synchronously with the radar data via software-controlled simultaneous triggering, and serves as the reference for respiration rhythm evaluation.

Data Collection and Split. We collect a dataset to implement and evaluate P^2CRM. We recruit 20 volunteers, including 11 females and 9 males, with body weights ranging from 40 kg to 80 kg and ages between 18 and 60 years old. During the experiments, each participant sits on a chair in a relaxed posture and is instructed to avoid any large movements to minimize motion interference. The mmWave radar is placed on a desk directly in front of the participant, with the antenna oriented toward the chest to capture respiratory motion. For each participant, approximately 30 minutes of radar and respiratory belt data are collected across different days and at varying times. The data collection consists of five sessions: a 15-minute session of normal respiration at 0.5 m, followed by four 4-minute sessions covering deep respiration at 0.5 m, irregular respiration at 0.5 m, normal respiration at 1.0 m, and normal respiration at 1.5 m, respectively. In our experiments, a sliding window (window size equal to the sample duration, as specified in Sec. 6.4.7, with a stride of half the window size) is first applied to each participant's continuous radar signal to generate fixed-length samples. We then adopt a subject-dependent temporal split for both identity recognition and respiratory-rate estimation. Specifically, for each subject, the temporally ordered samples are divided into three partitions: the first 60% are used for training, the next 20% for validation, and the final 20% for testing. The temporal split is performed before key-controlled semantic mixing. For the identity-recognition evaluation, the training, validation, and testing partitions are perturbed using independently generated keys, such that no key or corresponding orthogonal matrix used for the test samples appears in the training set.

Training Settings. The shared temporal backbone is built upon a 5-layer multi-resolution 1D U-Net. We train the models using the Adam optimizer with a learning rate of $1e^{-3}$ and a batch size of 32. The first convolutional layer takes 1 input channel, produces 16 output channels, and uses a kernel size of 5. All other convolutional layers also adopt a kernel size of 5. The output channel dimensions of the UNet-convolution kernels, down-conv, and frequency-gated additive modules in the encoder are 16, 32, 64, and 128, respectively. The final encoder stage outputs 256 channels and is followed by 12 Transformer layers, each with 8 attention heads. The decoder symmetrically employs UNet convolution and Up-conv layers with output channel counts of 128, 64, 32, and 16, respectively. All nonlinear activations are implemented as LeakyReLU with a negative slope of 0.3. For the loss function, the weighting coefficients are empirically set to $\alpha = 1.0$, $\beta = 2.0$.

Table 2. Overall performance comparison with different baselines.

Method	MAE (bpm)	STD (bpm)	IRAC (%)
Raw (Non-perturbed Upper Bound)	0.90	0.30	83.38
All-Perturb (Naive Full-Signal Perturbation)	Disabled	Disabled	8.21
MuFilter	Disabled	Disabled	10.37
Ours	**1.08**	**0.31**	**28.47**

6 Evaluation

6.1 Evaluation Metrics

Metrics for Respiratory Sensing. To evaluate the performance of respiratory sensing, we adopt two commonly used error-based metrics, the Mean Absolute Error (MAE) and the Standard Deviation of the estimation error (STD), to jointly assess the accuracy and stability of respiratory sensing. Specifically, MAE measures the average deviation between the estimated respiratory rates and the corresponding ground-truth values, reflecting the overall sensing accuracy, while STD quantifies the dispersion of the estimation errors across different samples, indicating the robustness and consistency of the sensing performance. All metrics are measured in beats per minute (bpm).

- MAE: Measures the average magnitude of the estimation error, $\mathrm{MAE} = \frac{1}{J}\sum_{j=1}^{J}\left|\widehat{b}_j - b_j\right|$.
- STD: Reflects the stability and dispersion of the absolute estimation errors, $\mathrm{STD} = \sqrt{\frac{1}{J}\sum_{j=1}^{J}\left(\left|\widehat{b}_j - b_j\right| - \mu\right)^2}$.

where $\widehat{b}_j$ denotes the estimated respiratory rate of the j-th sample, b_j is the corresponding ground-truth value, J represents the total number of samples, and μ denotes the MAE.

Metrics for Identity Recognition. Identity recognition is formulated as a multi-class classification task, and identity recognition accuracy (IRAC) measures how accurately the evaluated classifier infers subject identities from radar signals. A lower IRAC indicates that fewer identity-related cues are extracted by the evaluated classifier.

- IRAC: Measures identity recognition accuracy, $\mathrm{IRAC} = \frac{1}{J}\sum_{i=1}^{J}\mathbb{I}\left(\hat{c}_i = c_i\right)$.

where c_i and $\hat{c}_i$ denote the true label and the predicted label of the i-th sample, respectively, and $\mathbb{I}(\cdot)$ is the indicator function.

To comprehensively evaluate the performance of our work, we report 95% confidence intervals (CIs) for both MAE and IRAC. Overall, an effective identity-preserving respiratory monitoring system is expected to achieve low MAE and STD for respiratory rate estimation, while simultaneously yielding a low IRAC for identity recognition, thereby balancing sensing utility and identity protection.

6.2 Overall Performance

We evaluate the proposed system against three baselines: raw data (no encryption, serving as the utility upper bound), a naive full-signal perturbation strategy (uniformly perturbing the entire radar signal without semantic awareness), and MuFilter [31], a state-of-the-art privacy-preserving scheme that adds random inter-frame phase fluctuation to vital signs. The raw data-based model performs direct respiratory rate extraction and identity recognition on unmodified radar signals and serves as a reference upper bound due to its access to complete, unperturbed information. As shown in Table. 2, our method achieves competitive respiratory sensing performance

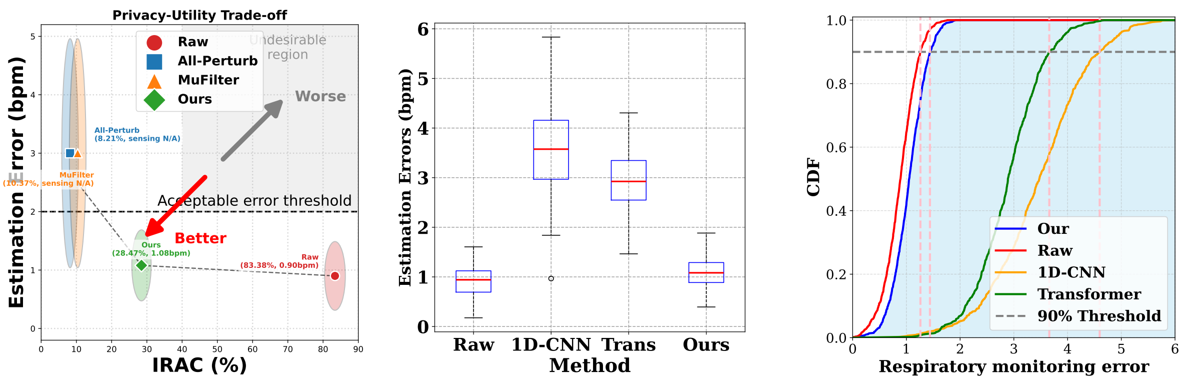

(a) Identity–utility trade-off. (b) Respiratory error distribution. (c) CDF of respiratory estimation error.

Fig. 10. Visualization of sensing performance and identity protection capability.

Table 3. Comparison of identity recognition performance.

Model	IRAC based on perturbed data (%) (95% CI)	IRAC based on raw data (%) (95% CI)
1D CNN	28.47 (26.78, 30.05)	83.38 (81.98, 84.76)
ResNet-18	29.39 (27.82, 30.90)	89.91 (88.59, 91.25)
CNN-LSTM	29.62 (28.06, 31.07)	90.85 (89.53, 92.20)
Transformer	29.97 (28.47, 31.24)	93.31 (91.92, 94.63)

with an MAE of 1.08 bpm and an STD of 0.31 bpm, while reducing the IRAC to 28.47%. These results highlight the model's ability to selectively preserve respiratory rhythm information while perturbing identity-specific fingerprints, striking a practical balance between utility and privacy. To provide a more intuitive visualization of the privacy-utility trade-off, we present the results in Fig. 10a. Compared with Raw, our method demonstrates competitive performance: the IRAC drops sharply from 83.38% to 28.47%, while the respiratory sensing error increases only slightly (from 0.90 bpm to 1.08 bpm). A paired bootstrap test on the absolute errors of the same test samples confirms that this increase in MAE is statistically significant ($p < 0.01$). In contrast, the naive full-signal perturbation strategy and MuFilter yield lower IRAC values but at the cost of completely unusable respiratory sensing.

6.2.1 Evaluation of Identity Recognition. To evaluate the identity protection capability of the proposed system, we compare the identity recognition performance of multiple classifiers trained separately on raw and perturbed radar signals. As shown in Table. 3, on raw signals, all models achieve accuracies above 80%. Notably, ResNet-18, CNN-LSTM, and the Transformer achieve higher accuracies (89.91%, 90.85%, and 93.31%, respectively) than the 1D-CNN (83.38%), confirming their superior ability to capture discriminative identity features when present. In contrast, on perturbed signals, the 1D-CNN achieves an identity recognition accuracy of 28.47%, while ResNet-18, CNN-LSTM, and the Transformer achieve 29.39%, 29.62%, and 29.97%, respectively. These accuracies are consistently low and comparable across all architectures, despite their large differences in model capacity. For all identity-recognition

Table 4. Comparison of respiratory sensing performance.

Configuration	MAE (bpm)	STD (bpm)
Traditional Signal Processing (FFT-based)	Disabled	Disabled
1D-CNN	3.54 (3.11, 3.97)	0.87
Transformer	2.91 (2.57, 3.21)	0.64
Ours (*CRKT*)	**1.08 (0.94, 1.22)**	**0.31**

Table 5. Validation of Semantic Decoupling.

Configuration	IRAC(%) (95% CI)
Original Respiratory Signal	81.95 (79.85, 84.15)
Pattern Component	74.58 (72.92, 76.30)
Rhythm Component	22.10 (20.55, 23.67)

models, the test samples are perturbed using keys that are not used to generate the training samples. Therefore, the reported IRAC evaluates identity recognition under an unseen-key setting rather than a same-key train/test setting. The consistently reduced recognition accuracy across different attacker architectures suggests that key-controlled semantic mixing empirically suppresses identity-related cues. However, it does not eliminate identity-related information. With 20 subjects, the chance-level IRAC is 5%, whereas the proposed method yields an IRAC of 28.47%, approximately 5.7 times the chance level. This residual identity information may partly arise from weak person-specific statistical tendencies retained in the deliberately preserved respiratory rhythm component.

6.2.2 Evaluation of Sensing Performance. To evaluate the respiratory sensing performance of the proposed framework on perturbed radar signals, we compare our *CRKT* framework against three representative methods: a traditional signal processing approach (FFT) and two deep learning models (a compact 1D-CNN and a shallow Transformer). As shown in Table. 4, our method achieves the lowest MAE (1.08 bpm) and STD (0.31 bpm), indicating superior accuracy and stability in respiratory sensing under perturbed conditions. To provide a more intuitive assessment of sensing performance, we further visualize the distribution of respiratory estimation errors, as shown in Fig. 10b and Fig. 10c. Specifically, since the traditional signal processing method fails to support reliable respiratory sensing under perturbation, it is excluded from the visualization for clarity. Compared with the 1D-CNN and Transformer, the proposed method exhibits a clear shift toward lower error regions in the error distributions, demonstrating its superior ability to capture respiration-related semantics from identity-preserved signals. Compared with the Raw (non-perturbed) signal, our method exhibits a slight degradation in sensing performance, which is expected due to the applied identity-preserving perturbation that inevitably introduces some loss of signal fidelity. Nevertheless, the overall sensing error remains well controlled. This performance advantage is primarily attributed to the proposed *CRKT* scheme. By distilling respiration-dominant knowledge from the teacher network trained on semantically rich signals, the student model inherits strong respiration-specific priors and is guided to focus on rhythm-relevant semantics that remain invariant under perturbation, while suppressing sensitivity to perturbation-induced distortions.

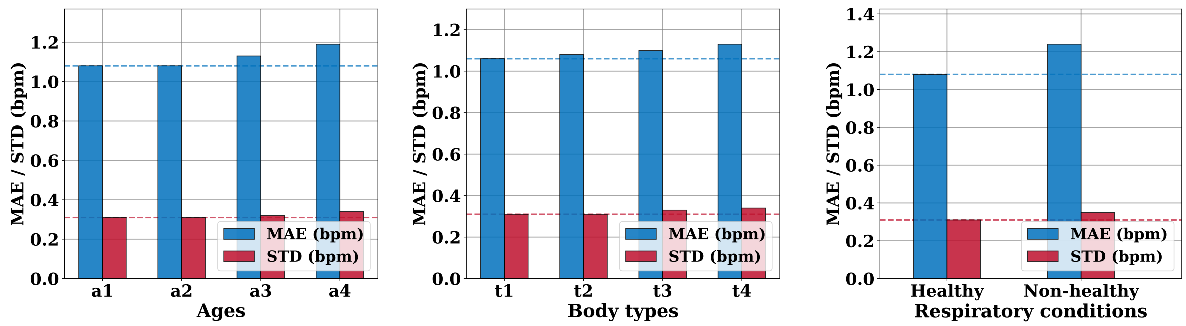

(a) Ages.

(b) Body types.

(c) Health conditions.

Fig. 11. Evaluation on demographic subgroups.

6.2.3 Verification of Semantic Decoupling Assumption. Additionally, to verify our premise that the respiratory frequency band contains identity information and that such information is primarily encoded in the pattern component, we conduct three supplementary experiments. As shown in Table. 5, the original respiratory signal achieves 81.95% accuracy (slightly lower than the 83.38% on the full raw radar signal, due to minor identity cues in non-respiratory components), the pattern component achieves 74.58%, and the rhythm component achieves only 22.10%. These results demonstrate that identity cues are concentrated in the pattern component, while the rhythm component retains less but non-negligible identity information.

In summary, our system demonstrates superior performance in achieving an effective identity–utility trade-off, enabling accurate respiratory frequency estimation while substantially suppressing identity-related information in mmWave signals.

6.3 Evaluation on Demographic Subgroups

To investigate the influence of demographic factors on performance, we provide detailed statistics on the age, body types (body mass index, BMI), and health conditions of the participants, and report the respiratory sensing performance for each subgroup. Specifically, for age grouping, we adopt the following intervals: < 21 (a1, $n = 3$), 21–35 (a2, $n = 11$), 36–50 (a3, $n = 4$), and > 50 (a4, $n = 2$) years. For body type grouping, participants are classified into four categories: underweight (t1, BMI < 18.5, $n = 3$), normal weight (t2, BMI 18.5–24, $n = 12$), overweight (t3, BMI 24–27, $n = 2$), and obese (t4, BMI ≥ 27, $n = 3$). For health conditions, participants are divided into two groups: healthy individuals ($n = 15$) and those with mild respiratory conditions (controlled asthma or chronic bronchitis, $n = 5$).

As shown in Fig. 11, we evaluate the model's sensing performance within each age, body type, and health condition subgroup. Apart from a slightly larger error for the >50 years group, the other age groups (<21, 21–35, 36–50) show consistently low sensing errors, indicating that the model exhibits strong robustness across different age groups. For body type subgroups, a marginal increase in error is observed for obese participants, while underweight, normal, and overweight groups show consistently low sensing errors, indicating that even under different chest wall motion amplitude and breathing mechanics induced by body type, the model maintains stable sensing performance. For health conditions, healthy participants achieve lower MAE and STD, whereas participants with mild respiratory conditions exhibit slightly higher errors, which may be due to subtle irregularities in their breathing patterns affecting signal consistency. Despite this slight increase, the errors remain well within the acceptable range for respiratory monitoring. These results confirm that the proposed system is also robust to mild

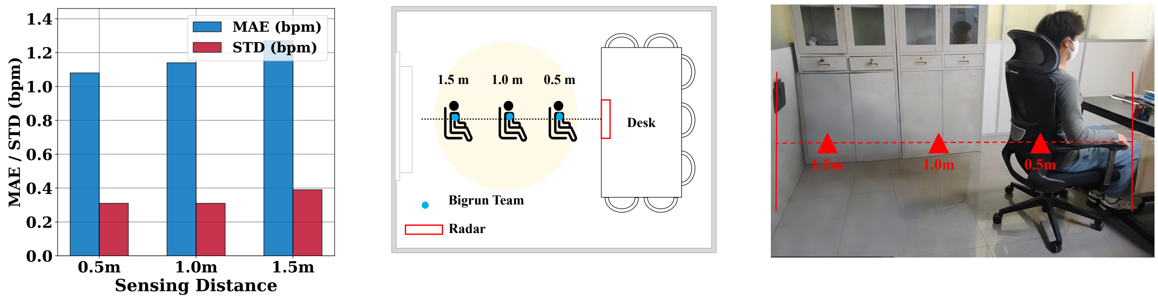

Fig. 12. Impact of sensing distance.

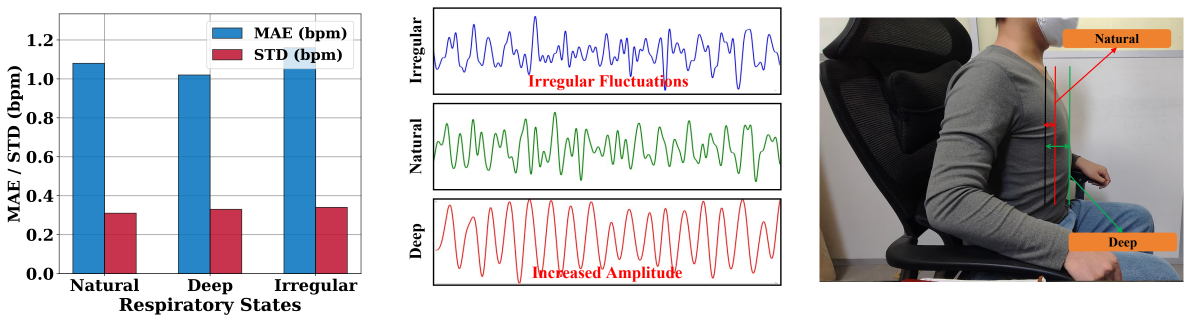

Fig. 13. Impact of respiratory states.

respiratory abnormalities. These results suggest that the proposed method is generally robust to demographic variations, demonstrating strong generalization across different ages, body types, and health conditions.

6.4 Robustness Study

6.4.1 Impact of Sensing Distance. Sensing distance is a critical factor for the practical deployment of contactless respiratory sensing systems. To evaluate the robustness of our system with respect to distance variations, we conduct experiments at different distances by adjusting the distance between the radar and the target from 0.5 m to 1.5 m with a fixed orientation 0°. During each trial, subjects are instructed to remain seated at the designated position and to maintain a natural breathing state. As shown in Fig. 12, the proposed system demonstrates stable respiratory sensing performance across all evaluated distances. Although a slight performance degradation is observed as the sensing distance increases, both the MAE and STD remain within acceptable ranges across all distance settings. These results indicate that, despite the gradual attenuation of signal strength with increasing distance, our system is able to reliably capture rhythm-relevant semantics from the signals.

6.4.2 Impact of Respiratory States. As respiratory characteristics may vary substantially across different breathing states, it is essential to evaluate the robustness of the proposed system under diverse respiratory states. To this end, we ask subjects to perform three representative patterns: deep respiration, natural respiration, and irregular respiration. These patterns are deliberately selected to represent a broad spectrum of respiratory conditions that may occur in real-world scenarios. Specifically, to control other factors, the sensing distance and orientation are fixed at 0.5m and 0° to ensure a high signal-to-noise ratio and stable measurements. As shown in Fig. 13,

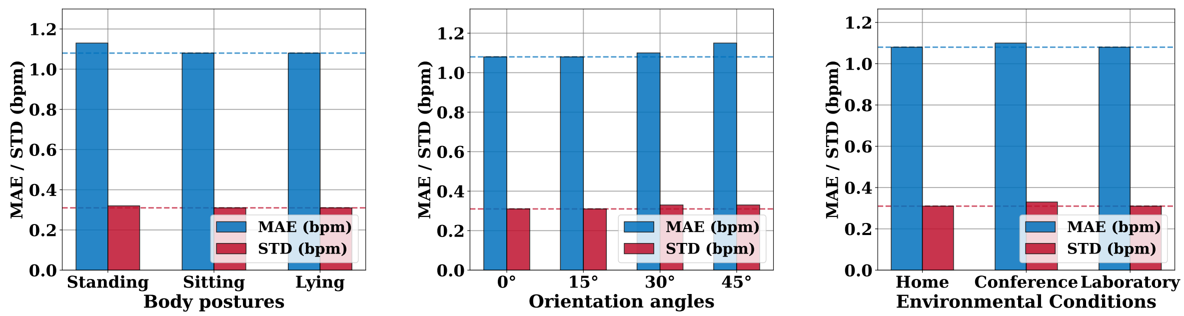

(a) Body postures. (b) Orientation angles. (c) Environmental Conditions.

Fig. 14. Impact of body postures, angles and environments.

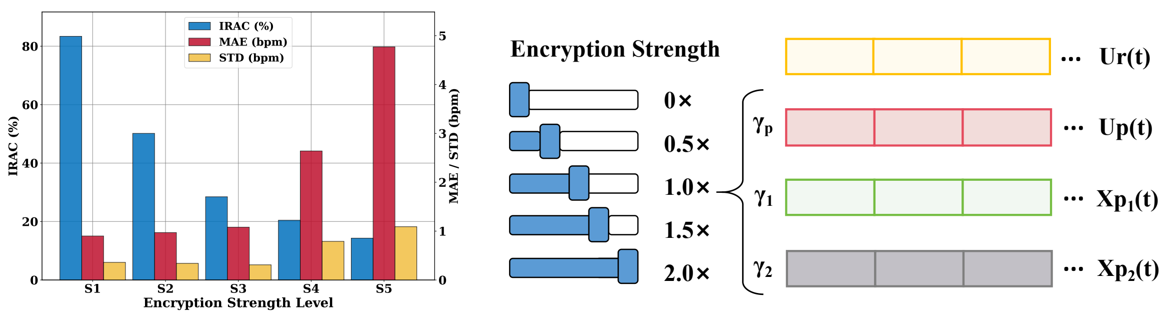

Fig. 15. Impact of perturbation strength.

the proposed system maintains stable respiratory sensing performance across various respiratory patterns. This robustness can be attributed to the model's ability to focus on respiration-related frequency semantics, enabling reliable respiratory monitoring under diverse and unconstrained breathing behaviors.

6.4.3 Impact of Body Postures. To assess the impact of body postures, we select three representative postures: standing, sitting (default), and lying supine. The sensing distance and orientation are fixed at 0.5 m and 0°, respectively. As shown in Fig. 14a, the standing posture results in a slight decline in sensing performance, with MAE and STD increasing to 1.13 bpm and 0.32 bpm. We attribute this to subtle involuntary body movements when subjects are standing, which can introduce artifacts that partially affect the respiratory signal. However, even in the standing posture, the system is able to provide acceptable results, indicating that it retains practical utility in less ideal scenarios.

6.4.4 Impact of Orientation Angles. To investigate the impact of orientation angle, we conduct comparative experiments at different orientation angles. Considering that the radar's beam pattern and the chest's radar cross-section maintain reasonable sensitivity within a moderate angular range, and that extreme orientations (*e.g.*, 90°) are uncommon in typical monitoring scenarios, we select a practical set of orientation angles for evaluation: 0° (default), 15°, 30°, and 45°. As shown in Fig. 14b, MAE and STD increase only slightly as the orientation angle increases, with the worst case at 45° (MAE = 1.15 bpm, STD = 0.33 bpm), corresponding to a gradual decline in sensing performance. This decrease is likely because the effective reflective area of the chest decreases at larger

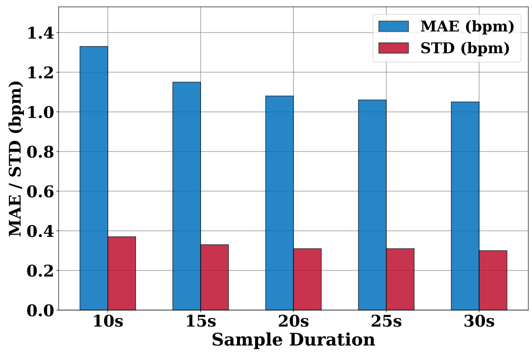

Fig. 16. Impact of sample duration.

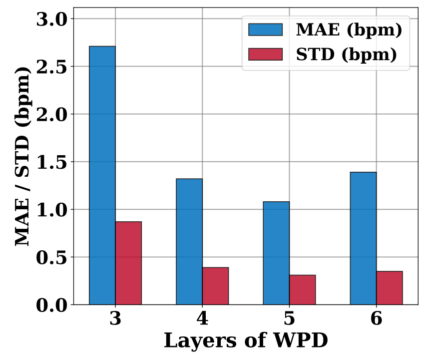

Fig. 17. Impact of layers of WPD.

angles, reducing the signal-to-noise ratio. Nevertheless, the system maintains reliable sensing performance even at the largest tested angle, demonstrating its robustness to moderate orientation variations.

6.4.5 Impact of Environmental Conditions. To evaluate the environmental adaptability and robustness of the proposed system, we perform experiments in three different scenarios: a quiet home environment, a conference room, and a laboratory. As shown in Fig. 14c, the system maintains stable respiratory sensing performance across all three environments, with MAE values of 1.08 bpm, 1.10 bpm, and 1.08 bpm, and STD values of 0.31 bpm, 0.33 bpm, and 0.31 bpm, respectively, indicating that the system experiences no noticeable performance degradation when the environment changes. Therefore, the proposed system exhibits stable performance across different environments.

6.4.6 Impact of Perturbation Strength. To investigate the impact of perturbation strength, we evaluate the proposed method under different levels by jointly scaling $\gamma_p, \gamma_1, \gamma_2$ while keeping all other settings unchanged. Five scaling factors are considered: $\alpha \in \{0, 0.5, 1, 1.5, 2\}$. As shown in Fig. 15, with the increase of scaling factor, the respiratory sensing performance gradually degrades. This is because stronger mixing introduces greater semantic distortion, which may overwhelm sensing-related information and degrade respiratory sensing performance. Within the evaluated moderate perturbation range, the method maintains a low estimation error while reducing identity recognition accuracy.

6.4.7 Impact of Sample Duration. To investigate the impact of sample duration, we conduct experiments with window lengths of 10 s, 15 s, 20 s, 25 s, and 30 s. As shown in Fig. 16, both MAE and STD consistently decrease as the sample duration increases, indicating that longer data sequences enable the proposed system to extract more stable and representative respiratory rhythm features. However, when the duration exceeds 20 s, the performance gains become negligible. Therefore, a monitoring duration of approximately 20 s provides a well-balanced trade-off among sensing performance, system efficiency, and user experience.

6.4.8 Impact of Layers of WPD. The number of WPD layers L determines the granularity of frequency partitioning and the quality of semantic separation. To investigate the impact of decomposition depth, we conduct experiments by varying L from 3 to 6. As shown in Fig. 17, insufficient decomposition ($L = 3$) causes coarse subbands and incomplete separation, increasing MAE and STD, indicating degraded sensing performance. Moderate depths ($L = 4, 5$) significantly improve both metrics by enabling effective semantic regrouping. Over-decomposition ($L = 6$) slightly degrades performance due to fragmented semantics and noise sensitivity. Therefore, we adopt $L = 5$ as the default decomposition depth.

Table 6. Perturbation parameter analysis.

Parameter	Value	MAE (bpm) (95% CI)	STD (bpm)	IRAC (%) (95% CI)
	0	0.90 (0.76, 1.04)	0.30	82.08 (79.91, 84.23)
γ_p	0.35 (ours)	1.08 (0.94, 1.22)	0.31	28.47 (26.78, 30.05)
	0.7	1.91 (1.70, 2.12)	0.35	24.71 (23.08, 26.35)
	0.35	1.08 (0.94, 1.22)	0.31	29.72 (27.87, 31.75)
γ_1	0.7 (ours)	1.08 (0.94, 1.22)	0.31	28.47 (26.78, 30.05)
	1.0	1.14 (0.95, 1.33)	0.35	28.41 (26.69, 30.23)
	0.35	1.08 (0.94, 1.22)	0.31	29.85 (27.99, 31.90)
γ_2	0.7 (ours)	1.08 (0.94, 1.22)	0.31	28.47 (26.78, 30.05)
	1.0	1.12 (0.97, 1.34)	0.34	28.38 (26.13, 29.97)

6.5 Perturbation Parameter Analysis

To further evaluate the performance of the proposed model under different perturbation parameter settings, we conduct a sensitivity analysis on the three key semantic weighting parameters: γ_p, γ_1 and γ_2. In this experiment, we vary one parameter at a time while keeping the others fixed at their default values to isolate and measure its individual impact on model performance.

Table. 6 summarizes the performance variations under different parameter configurations. Increasing γ_p strengthens the perturbation applied to the pattern component, thereby achieving better privacy protection. However, this inevitably degrades sensing performance because the pattern component carries fine-grained waveform details that are helpful for accurate respiratory rate estimation. Moderately increasing γ_1 and γ_2 can improve privacy protection without noticeably affecting sensing performance, as the low-frequency drift and high-frequency residual components contain limited respiratory information. Nevertheless, excessively strong perturbation on these components introduces additional noise that may overwhelm useful information, which unavoidably harms sensing performance. Therefore, the default configuration ($\gamma_p = 0.35$, $\gamma_1 = 0.7$, $\gamma_2 = 0.7$) strikes a practical balance between privacy protection and sensing utility.

6.6 Ablation Study

The P^2CRM consists of three key components: *WSD*, *KCSM*, and *CRKT*. Starting from raw radar signals, the *WSD* module performs structured semantic decoupling through two successive stages, namely WPD-based Semantic Band Decoupling and Regrouping and Intra-band Respiratory Semantic Decoupling. The former conducts multi-resolution decomposition and semantic regrouping to isolate respiration-dominant components from interference, while the latter further disentangles respiratory rhythm and pattern within the respiration-dominant band. Building upon the structured semantic representation produced by *WSD*, the *KCSM* module operates directly in the semantic space to selectively conceal identity-sensitive information while explicitly preserving respiration rhythm semantics that are critical for reliable respiratory sensing. Finally, the *CRKT* adopts a teacher–student learning paradigm to transfer respiration-dominant knowledge learned from semantically rich pre-perturbation radar representations to their perturbed counterparts, thereby improving the robustness of respiratory sensing

Table 7. Ablation study. The symbol "/" indicates that under this configuration, we focus solely on respiratory sensing performance and do not evaluate identity recognition capability.

Method	MAE (bpm) (95% CI)	STD (bpm)	IRAC (%) (95% CI)
W/O Intra-band Decoupling (C1)	0.90 (0.76, 1.04)	0.30	81.95 (79.85, 84.15)
W/O *WSD* (C2)	Disabled	Disabled	8.21 (5.40, 10.87)
W/O *KCSM* (C3)	0.90 (0.76, 1.04)	0.30	83.38 (81.98, 84.76)
W/O *CRKT* (C4)	3.07 (2.70, 3.45)	0.71	/
W/O Frequency-Gated Additive Shortcut (C5)	2.12 (1.78, 2.47)	0.64	/
W/O Pattern (Direct Removal) (C6)	1.98 (1.80, 2.18)	0.36	22.10 (20.55, 23.67)
Ours	**1.08 (0.94, 1.22)**	**0.31**	**28.47 (26.78, 30.05)**

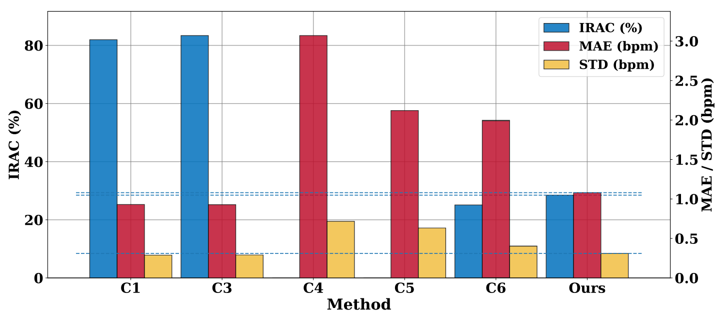

Fig. 18. Visualization of P^2CRM ablation experiment.

under identity-preserving constraints. To validate the effectiveness of the proposed modules in the overall architecture, we conduct a systematic ablation study, as shown in Tab. 7.

- **W/O Intra-band Decoupling:** It refers to removing the intra-band respiratory semantic decoupling within the respiration-dominant frequency band, treating the entire respiratory component as a single semantic unit. Consequently, the subsequent perturbation is applied only to the respiration-irrelevant components.
- **W/O *WSD*:** It refers to removing the entire semantic decoupling module. Consequently, the subsequent perturbation is directly applied to the full signal without semantic awareness.
- **W/O *KCSM*:** It refers to removing the key-controlled semantic mixing module and directly releasing the decoupled semantic components without applying any identity-preserving perturbation.
- **W/O *CRKT*:** It refers to removing the teacher–student knowledge transfer scheme and training the student model solely on perturbed representations.
- **W/O Frequency-Gated Additive Shortcut:** It refers to removing the frequency-gated additive shortcuts in the backbone, thereby disabling frequency-aware feature modulation during multi-scale feature aggregation.
- **W/O Pattern (Direct Removal):** It refers to directly removing the pattern component $U_p(t)$ without any mixing, retaining only the rhythm component and residuals.

Fig. 18 provides the visualized comparison of the ablation study (C2 loses respiratory sensing capability and is therefore excluded from the visualization). Specifically, when the intra-band respiratory semantic decoupling module is removed, the respiration-dominant component is no longer further disentangled and perturbed. Therefore, the respiratory sensing performance is almost unaffected. However, since perturbation is still applied to the remaining non-respiratory components, which may contain residual identity-related cues, the identity recognition accuracy exhibits a slight reduction. When the *WSD* module is removed, perturbation is directly applied to the raw radar signals. Although the IRAC is substantially reduced, the respiratory sensing performance is severely compromised. Furthermore, removing the *KCSM* module disables semantic-level identity protection, allowing identity-related information to remain exposed and resulting in severe identity leakage. Without the *CRKT* scheme, the student model is trained solely on perturbed representations without guidance from the teacher network. This severely limits its ability to extract respiratory rhythm from perturbed inputs, leading to degraded sensing performance. In addition, disabling the frequency-gated additive shortcut weakens the model's capability to selectively emphasize rhythm-relevant components during feature propagation, thereby leading to inferior respiratory sensing performance. Finally, to evaluate the privacy–utility benefit of controlled mixing, we compare directly removing the pattern component with our system. When the pattern component is directly removed, identity recognition accuracy drops further, but respiratory sensing accuracy also suffers due to the loss of fine-grained waveform details that are beneficial for accurate respiratory sensing. Paired bootstrap tests confirm that our system achieves significantly lower MAE than Direct Removal, whereas Direct Removal achieves significantly lower IRAC than our system ($p < 0.01$ for both comparisons). In contrast, our system retains these fine-grained details via controlled mixing, providing a different trade-off between sensing accuracy and identity leakage. Consequently, these components jointly support the reported privacy–utility trade-off.

7 Related Work

7.1 Respiratory Monitoring Methods

Long-term and continuous respiratory monitoring is critical for the early diagnosis and effective health management of chronic respiratory diseases [57]. Traditional respiratory monitoring methods typically rely on direct contact between sensors and patients, determining respiratory status by measuring physical parameters, *e.g.*, chest and abdominal movements [40], breath sounds, and airflow [8] generated during respiration. However, due to their complex operational procedures and the considerable discomfort caused by prolonged sensor attachment, these methods are unsuitable for continuous daily health monitoring. The emergence of portable devices, *e.g.*, smart watches [9], has extended respiratory monitoring from clinical settings to daily life, enabling continuous tracking of key respiratory parameters, like breathing rate [57]. Nevertheless, since these devices require direct skin contact, their measurements are susceptible to variations in wearing position, tightness, and users' daily activities, leading to low long-term compliance.

Owing to its non-invasive and contact-free characteristics, non-contact respiratory sensing has attracted considerable attention in recent years. Existing research mainly focuses on camera-based monitoring methods [38, 42, 56] and wireless sensing-based monitoring methods, *e.g.*, Wi-Fi-based sensing [18, 41, 51] and mmWave radar-based sensing [2, 17, 25, 53, 54]. Camera-based methods typically rely on cameras to capture subtle chest and abdominal movements for respiratory sensing [11, 58, 62]. Although they can achieve relatively high monitoring accuracy, they inevitably rely on favorable lighting conditions, which limits their robustness in real-world environments. In contrast, wireless sensing-based methods perceive respiration by analyzing changes in signal reflections or scattering induced by respiratory motions, making them inherently robust to illumination conditions [48]. Specifically, mmWave radar has emerged as a research hotspot in wireless sensing for respiratory monitoring due to its contactless, user-friendly, high sensitivity to micro-motions, and strong signal penetration capability.

7.2 MmWave radar-based Respiratory Sensing Security

Although mmWave radar enables high-accuracy respiratory sensing, the captured signals inevitably encode identity-sensitive information [26, 30, 36, 52, 59, 61]. This is because respiratory-induced chest motions and micro-Doppler patterns exhibit subject-specific characteristics related to physiological traits, body posture, and habitual breathing behaviors, which can be learned by data-driven models. Consequently, mmWave-based respiratory sensing may unintentionally reveal personal identity [10, 30, 52] or biometric attributes, raising significant privacy risks in real-world deployments [32].

To address privacy leakage concerns, several studies employed physical-layer attacks, utilizing specific devices to interfere with or manipulate the location information of the target [31]. For example, mmFilter [29] protected the privacy of individual users through targeted dimensional perturbations, while MuFilter [31] extended this approach to multi-user scenarios, enabling broader privacy protection by enabling subspace-level privacy protection policies on demand. However, while these methods provided privacy protection, they often introduced uncontrollable interference to critical sensing features, which inevitably impaired perception tasks and limited their practical applicability in scenarios including medical monitoring and smart home environments [63].

8 Discussion and Future Work

8.1 Usage Scenario

Our system assumes that the subject sits in front of the radar to sense chest micro-movements induced by respiration. Existing contact-based solutions including respiratory belts or wearable sensors often require tight attachment to the body or prolonged skin contact, which may introduce discomfort and are unsuitable for subjects with sensitive skin [19]. As the first step toward identity-preserving contactless respiratory sensing, P^2CRM operates on short-duration signal segments, making it suitable for daily scenarios. While the robustness study demonstrates stable respiratory sensing performance across various sensing distances, respiratory states, body postures, body types, orientation angles, environments, and sample durations, extending the system to more challenging conditions, such as strong motion interference, multi-person scenarios, older adults above 80 years, and individuals with severe respiratory diseases, remains an important direction for future work.

8.2 User Experience Study

P^2CRM aims to eliminate the burden of contact-based sensing through mmWave signals for identity-preserving respiratory monitoring. Extensive experiments have demonstrated its effectiveness and practicality. We would like to discuss user experience from the experiments. First, the proposed approach requires the subject to remain relatively still during sensing to ensure reliable extraction of respiration-related micro-motions. In future work, lightweight motion detection or activity recognition modules can be incorporated to automatically assess motion conditions and guide users during the sensing process, thereby improving system robustness and user interaction. Second, there still exists space to further optimize the system for deployment on extremely resource-constrained edge devices. The current system achieves a total latency of under one second per 20-second window, including preprocessing, semantic decoupling, key-controlled mixing, and student-model inference. In addition, the deployed student model requires approximately 65 MB to store its weights under 32-bit floating-point representation (FP32). This moderate model-weight footprint, together with the low processing latency, supports real-time operation in typical monitoring scenarios. In the future, we can further optimize and simplify the pipeline to achieve even lower latency and more responsive feedback on extremely resource-constrained devices. In addition, the form factor of the radar can be optimized. Despite the flexibility offered by the radar platforms used in this work (TI AWR 1843BOOST and DCA1000EVM), some users perceive such devices as uncommon in household environments and potentially distracting during daily activities such as watching TV [19]. We envision that future work will explore more compact radar modules or investigate the feasibility of embedding the sensing hardware into everyday objects in a more natural and unobtrusive manner.

8.3 Generalizability and Transferability

Although this work focuses on identity-preserving respiratory monitoring using mmWave radar, the proposed framework is not limited to a specific sensing task or hardware configuration. By adjusting sensing parameters, the approach can be readily adapted to other sensing devices without requiring fundamental modifications to the model architecture. Moreover, the system is transferable to a broader range of wireless sensing applications beyond respiratory monitoring. Tasks that rely on fine-grained physiological or motion-induced micro-movements, such as sleep stage analysis, respiratory abnormality detection, and other contactless health monitoring scenarios, can directly benefit from the proposed sensing paradigm. The selected WPD depth $L = 5$ controls frequency resolution rather than scaling directly with thoracic volume, and it should be revalidated for populations with substantially different or highly irregular respiratory profiles.

8.4 Scope of Preserved Respiratory Semantics

Although the proposed framework preserves the rhythm semantics required for accurate respiratory-rate estimation, the current evaluation focuses primarily on respiratory-rate estimation rather than advanced respiratory pattern recognition. Clinically meaningful respiratory pattern information, such as waveform morphology, inhalation/exhalation asymmetry, and other high-level respiratory patterns, may be important for diagnostic applications. However, these pattern-related semantics may also carry identity-sensitive information. Therefore, the publicly released perturbed signals are designed and evaluated mainly for privacy-preserving respiratory rate estimation. Their ability to support advanced pattern recognition remains an important direction for future work. Additionally, we assume that a discernible dominant respiratory frequency exists within each analysis window. Moderate frequency variations can be accommodated by the tolerance parameter δ. However, for transient or strongly non-periodic breathing with no stable dominant frequency, the estimated $\hat{f}_b$ and the resulting rhythm–pattern separation may become unreliable, thereby degrading respiratory-rate estimation. Handling such profiles through adaptive time-frequency analysis remains an important direction for future work.

9 Conclusion

This paper presents an identity-preserving contactless respiratory sensing system using mmWave radar, aiming to enable trusted, high-accuracy, and secure health sensing. The proposed system consists of three key components. First, the *WSD* module transforms demodulated phase signals into a structured semantic representation by isolating respiration-dominant components and further separating respiration rhythm and pattern information. Second, the *KCSM* selectively suppresses rhythm–irrelevant semantics through key-controlled perturbations in the semantic token space, while preserving respiratory rhythm information critical for accurate respiratory sensing. Finally, the *CRKT* scheme adopts a teacher–student paradigm to transfer respiration-relevant knowledge from semantically rich signals to perturbed representations, enabling robust respiratory sensing under identity-preserving constraints. Extensive experimental evaluations demonstrate that the proposed system achieves accurate and robust respiratory sensing while significantly reducing identity leakage, consistently outperforming existing baseline methods under diverse sensing conditions. These results indicate that the proposed approach holds strong potential for enabling trustworthy and privacy-preserving contactless respiratory monitoring in daily health sensing scenarios, such as long-term and sleep-related respiratory assessment.

10 Acknowledgements

This work is supported by National Natural Science Foundation of China (Grant No. 62302145, No. 62422213, No. 62372149, and No. U23A20303), Beijing Nova Program (20240484641), Major Scientific and Technological Project of Anhui Provincial Science and Technology Innovation Platform (Grant No. 202305a12020012), Key Laboratory of Space-Air-Ground Integrated Network Security, Ministry of Industry and Information Technology, Shandong Key Laboratory of Advanced Computing, and Fundamental Research Funds for the Central Universities (Grant No. JZ2025HGTB0225).

References

[1] Mohadese Aali, Fatemeh Sadat Lesani, and Faranak Fotouhi-Ghazvini. 2025. Smartphone continuous user identification via hybrid transfer learning: leveraging common semantic space and pre-trained models. *CCF Transactions on Pervasive Computing and Interaction* 7 (2025), 494 – 509. https://api.semanticscholar.org/CorpusID:280325563

[2] Shahzad Ahmed and Sung Ho Cho. 2023. Machine learning for healthcare radars: Recent progresses in human vital sign measurement and activity recognition. *IEEE Communications Surveys & Tutorials* 26, 1 (2023), 461–495.

[3] Jian An, Siyuan Wu, Xiaolin Gui, Xin He, and Xuejun Zhang. 2023. A blockchain-based framework for data quality in edge-computing-enabled crowdsensing. *Frontiers of Computer Science* 17, 4 (2023), 174503. doi:10.1007/s11704-022-2083-8

[4] Sufyan Ashhad, Kaiwen Kam, Christopher A Del Negro, and Jack L Feldman. 2022. Breathing rhythm and pattern and their influence on emotion. *Annual review of neuroscience* 45 (2022), 223–247.

[5] Humberto Basilio. 2025. BREATHING PATTERNS CAN IDENTIFY PEOPLE LIKE FINGERPRINTS. *Nature* 642 (2025), 851.

[6] Zhaoxin Chang, Xinyu Xue, Fusang Zhang, Jie Xiong, Weiyan Chen, Badii Jouaber, and Daqing Zhang. 2024. Robust Respiration Monitoring Under Body Motion Interference. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*. 1713–1715.

[7] Zhaoxin Chang, Fusang Zhang, Xujun Ma, Pei Wang, Weiyan Chen, Duo Zhang, Badii Jouaber, and Daqing Zhang. 2024. Mmecare: Enabling fine-grained vital sign monitoring for emergency care with handheld mmwave radars. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 8, 4 (2024), 1–24.

[8] Xing-Yu Chen, Qiu-Shi Zhu, Jie Zhang, and Li-Rong Dai. 2022. Supervised and self-supervised pretraining based COVID-19 detection using acoustic breathing/cough/speech signals. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 561–565.

[9] Gloria Cosoli, Luca Antognoli, Luna Panni, and Lorenzo Scalise. 2024. Indirect Estimation of Breathing Rate Using Wearable Devices. *IEEE Transactions on Instrumentation and Measurement* 73 (2024), 1–8.

[10] Feng-Qi Cui, Jinyang Huang, Sirui Zhao, Kun Li, Zhi Liu, Meng Li, Ziyu Jia, Dan Guo, and Meng Wang. 2026. PersoMoni: A Comprehensive Video-Based Benchmark Dataset for Fine-grained Personality Assessment with 15 Trait Dimensions. *IEEE Transactions on Affective Computing* (2026).

[11] Feng-Qi Cui, Anyang Tong, Jinyang Huang, Jie Zhang, Meng Li, Xin Yan, Linsheng Huang, Dan Guo, and Meng Wang. 2026. Towards Trustworthy Dynamic Facial Expression Recognition via Information Bottleneck Modeling. *IEEE Transactions on Information Forensics and Security* (2026).

[12] Zhengxin Guo, Xu Zhu, Linqing Gui, Biyun Sheng, and Fu Xiao. 2023. BreathID: Respiration Sensing for Human Identification Using Commodity Wi-Fi. *IEEE Systems Journal* 17, 2 (2023), 3059–3070.

[13] Kai Han, An Xiao, Enhua Wu, Jianyuan Guo, Chunjing Xu, and Yunhe Wang. 2021. Transformer in transformer. *Advances in neural information processing systems* 34 (2021), 15908–15919.

[14] Shijie Han, Dongheng Zhang, Jinbo Chen, Haoyu Wang, Jinli Zhang, Qibin Sun, and Yan Chen. 2023. Fine-grained lung function sensing based on millimeter-wave radar. In *2023 International Conference on Wireless Communications and Signal Processing*. IEEE, 471–476.

[15] Zhanjun Hao, Yue Wang, Fenfang Li, Guozhen Ding, Kai Fan, and Yifei Gao. 2025. Detection of vital signs based on millimeter wave radar. *Scientific Reports* 15, 1 (2025), 28112.

[16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.

[17] Chieh-Hsun Hsieh and Po-Hsuan Tseng. 2025. Multiperson Localization and Vital Signs Estimation Using mmWave MIMO Radar. *IEEE Transactions on Microwave Theory and Techniques* 73, 2 (2025), 1222–1234.

[18] Jingyang Hu, Hongbo Jiang, Tianyue Zheng, Jingzhi Hu, Hongbo Wang, Hangcheng Cao, Zhe Chen, and Jun Luo. 2024. M 2-Fi: Multi-person respiration monitoring via handheld WiFi devices. In *IEEE INFOCOM 2024-IEEE Conference on Computer Communications*. IEEE, 1221–1230.

[19] Qingyong Hu, Qian Zhang, Hao Lu, Shun Wu, Yuxuan Zhou, Qianyi Huang, Huangxun Chen, Ying-Cong Chen, and Ni Zhao. 2024. Contactless arterial blood pressure waveform monitoring with mmwave radar. *Proceedings of the ACM on Interactive, Mobile, Wearable*

and Ubiquitous Technologies 8, 4 (2024), 1–29.

[20] Dajian Huang, Wen-An Zhang, Fanghong Guo, Weijiang Liu, and Xiaoming Shi. 2023. Wavelet Packet Decomposition-Based Multiscale CNN for Fault Diagnosis of Wind Turbine Gearbox. *IEEE Transactions on Cybernetics* 53, 1 (2023), 443–453.

[21] Jinyang Huang, Yuanhao Feng, Feng-Qi Cui, Xiang Zhang, Zhi Liu, Xin Liu, Jianchun Liu, Fusang Zhang, and Meng Li. 2026. Identifying Who You Are No Matter What You Write through Abstracting Handwriting Style. *IEEE Transactions on Dependable and Secure Computing* (2026), 1–15. doi:10.1109/TDSC.2026.3668275

[22] Shuzhan Huang, Jian Tang, Juying Dai, and Yangyang Wang. 2019. Signal status recognition based on 1DCNN and its feature extraction mechanism analysis. *Sensors* 19, 9 (2019), 2018.

[23] Shekh MM Islam, Olga Borić-Lubecke, and Victor M Lubecke. 2022. Identity authentication in two-subject environments using microwave Doppler radar and machine learning classifiers. *IEEE Transactions on Microwave Theory and Techniques* 70, 11 (2022), 5063–5076.

[24] Chengkun Jiang, Junchen Guo, Yuan He, Meng Jin, Shuai Li, and Yunhao Liu. 2020. mmVib: micrometer-level vibration measurement with mmwave radar. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*. 1–13.

[25] Wei Kang, Chenwei Zhou, and Wen Wu. 2024. Respiration monitoring of all occupants in a vehicle using time-division multiplexing FMCW radar based on metasurface technology. *IEEE Transactions on Microwave Theory and Techniques* 72, 8 (2024), 4960–4974.

[26] Sangdong Kim, Bongseok Kim, Youngseok Jin, and Jonghun Lee. 2020. Human identification by measuring respiration patterns using vital FMCW radar. *Journal of Electromagnetic Engineering and Science* 20, 4 (2020), 302–306.

[27] Chi-Wei Lien and Sudip Vhaduri. 2023. Challenges and opportunities of biometric user authentication in the age of iot: A survey. *Comput. Surveys* 56, 1 (2023), 1–37.

[28] Fulong Liu, Jinbo Chen, Haoyu Wang, Hanqin Gong, Dongheng Zhang, Qibin Sun, Yang Hu, and Yan Chen. 2025. Finding Order in Chaos: Learning Disentangled Features for mmWave Cardiac Sensing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 9, 4 (2025), 1–33.

[29] Hankai Liu, Xiulong Liu, Xin Xie, Xinyu Tong, Tuo Shi, and Keqiu Li. 2023. Application-oriented privacy filter for mmWave radar. *IEEE Communications Magazine* 61, 12 (2023), 168–174.

[30] Jian Liu, Yingying Chen, Yudi Dong, Yan Wang, Tiannming Zhao, and Yu-Dong Yao. 2020. Continuous User Verification via Respiratory Biometrics. In *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications*. 1–10.

[31] Xiulong Liu, Hankai Liu, Jiaqi Zhang, Xin Xie, and Keqiu Li. 2025. Multi-user Behavioral Privacy Filtering for mmWave Radar Sensing. *IEEE Transactions on Mobile Computing* (2025).

[32] Zhusen Liu, Zhenfu Cao, Xiaolei Dong, Xiaopeng Zhao, Haiyong Bao, and Jiachen Shen. 2022. A verifiable privacy-preserving data collection scheme supporting multi-party computation in fog-based smart grid. *Frontiers of Computer Science* 16, 1 (2022), 161810. doi:10.1007/s11704-021-0410-0

[33] Zhaoyu Liu, Yuyong Xiong, Wendi Tian, Yingjie Gou, and Zhike Peng. 2024. Large-scale body movement cancellation in FMCW radar respiration detection. *IEEE Transactions on Microwave Theory and Techniques* (2024).

[34] Li Lu, Lingshuang Liu, Muhammad Jawad Hussain, and Yongshuai Liu. 2017. I sense you by breath: Speaker recognition via breath biometrics. *IEEE Transactions on Dependable and Secure Computing* 17, 2 (2017), 306–319.

[35] Ruotian Peng, Yang Gao, and Zhanpeng Jin. 2024. Gazenum: unlock your phone with gaze tracking viewing numbers for authentication. *CCF Transactions on Pervasive Computing and Interaction* 7 (2024), 1 – 14. https://api.semanticscholar.org/CorpusID:272300868

[36] Ashikur Rahman, Victor M Lubecke, Olga Boric-Lubecke, Jan H Prins, and Takuya Sakamoto. 2018. Doppler radar techniques for accurate respiration characterization and subject identification. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems* 8, 2 (2018), 350–359.

[37] Mark A Richards et al. 2005. *Fundamentals of radar signal processing*. Vol. 1. Mcgraw-hill New York.

[38] Chiara Romano, Emiliano Schena, Sergio Silvestri, and Carlo Massaroni. 2021. Non-contact respiratory monitoring using an RGB camera for real-world applications. *Sensors* 21, 15 (2021), 5126.

[39] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*. Springer, 234–241.

[40] Sibi C Sethuraman, Pranav Kompally, Saraju P Mohanty, and Uma Choppali. 2021. MyWear: A novel smart garment for automatic continuous vital monitoring. *IEEE Transactions on Consumer Electronics* 67, 3 (2021), 214–222.

[41] Wenchao Song, Zhu Wang, Yifan Guo, Zhuo Sun, Zhihui Ren, Chao Chen, Bin Guo, Zhiwen Yu, Xingshe Zhou, and Daqing Zhang. 2024. FinerSense: A Fine-Grained Respiration Sensing System Based on Precise Separation of Wi-Fi Signals. *IEEE Transactions on Mobile Computing* (2024).

[42] Xudong Tan, Menghan Hu, Guangtao Zhai, Yan Zhu, Wenfang Li, and Xiao-Ping Zhang. 2024. Lightweight Video-Based Respiration Rate Detection Algorithm: An Application Case on Intensive Care. *IEEE Transactions on Multimedia* 26 (2024), 1761–1775.

[43] Texas Instruments. 2024.04.10. TI AWR1843BOOST Evaluation Board. https://www.ti.com/tool/AWR1843BOOST.

[44] Texas Instruments. 2024.04.10. TI DCA1000EVM. https://www.ti.com/tool/DCA1000EVM.

[45] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).

[46] A Helen Victoria, SV Manikanthan, HR Varadaraju, Muhammad Alkirom Wildan, and Kakarla Hari Kishore. 2022. Radar based activity recognition using cnn-lstm network architecture. *International Journal of Communication Networks and Information Security* 14, 3 (2022), 303–312.

[47] Jinliang Wang, Zhiping Zhang, Qiuli Wang, Li Wen, Jiaqi Liu, Changzhan Gu, Peiyao Xiao, Ziyang Fang, Yikai Li, Xianfeng Yao, et al. 2025. Comparative Accuracy of Millimeter-Wave Radar and Polysomnography in Sleep Apnea Detection. *Medical science monitor: international medical journal of experimental and clinical research* 31 (2025), e948079.

[48] Kevin I-Kai Wang, Xiaokang Zhou, Wei Liang, Zheng Yan, and Jinhua She. 2022. Federated Transfer Learning Based Cross-Domain Prediction for Smart Manufacturing. *IEEE Transactions on Industrial Informatics* 18, 6 (2022), 4088–4096. doi:10.1109/TII.2021.3088057

[49] Lin Wang and Kuk-Jin Yoon. 2022. Knowledge Distillation and Student-Teacher Learning for Visual Intelligence: A Review and New Outlooks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44, 6 (2022), 3048–3068.

[50] Lu-Yao Wang, Zhu-Sen Liu, Qi Feng, Wei-Bin Wu, and Lu Zhou. 2026. Enhancing data quality with effective feature selection and privacy protection. *Frontiers of Computer Science* 20, 3 (2026), 2003607. doi:10.1007/s11704-025-41074-0

[51] Meng Wang, Jinyang Huang, Xiang Zhang, Zhi Liu, Meng Li, Peng Zhao, Huan Yan, Xiao Sun, and Mianxiong Dong. 2024. Target-oriented wifi sensing for respiratory healthcare: from indiscriminate perception to in-area sensing. *IEEE Network* 39, 5 (2024), 201–208.

[52] Xiaoyang Wang, Fan Li, Yadong Xie, Song Yang, and Yu Wang. 2022. Gait and Respiration-Based User Identification Using Wi-Fi Signal. *IEEE Internet of Things Journal* 9, 5 (2022), 3509–3521.

[53] Yao Wang, Tao Gu, Tom H Luan, and Yong Yu. 2022. Your breath doesn't lie: multi-user authentication by sensing respiration using mmwave radar. In *2022 19th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*. IEEE, 64–72.

[54] Yingqi Wang, Zhongqin Wang, Jian Andrew Zhang, Haimin Zhang, and Min Xu. 2024. Vital Sign Monitoring in Dynamic Environment via mmWave Radar and Camera Fusion. *IEEE Transactions on Mobile Computing* 23, 5 (2024), 4163–4180.

[55] Wei Xu, Zhu Wang, Yifan Guo, Zhihui Ren, Yandi Xu, Bin Guo, Zhiwen Yu, and Xingshe Zhou. 2025. Acoustic sensing mechanisms, technologies, and applications: a survey. *CCF Transactions on Pervasive Computing and Interaction* 7 (2025), 522 – 549. https://api.semanticscholar.org/CorpusID:280324891

[56] Yongshen Zeng, Dongfang Yu, Xiaoyan Song, Qiqiong Wang, Liping Pan, Hongzhou Lu, and Wenjin Wang. 2024. Camera-Based Cardiorespiratory Monitoring of Preterm Infants in NICU. *IEEE Transactions on Instrumentation and Measurement* 73 (2024), 1–13.

[57] Lei Zhang, Rong Bao, Chen Jiahao, and Yonghong Zhu. 2024. Smart City Healthcare: Non-Contact Human Respiratory Monitoring With WiFi-CSI. *IEEE Transactions on Consumer Electronics* 70, 3 (2024), 5960–5968.

[58] Xinxin Zhang, Xudong Tan, Yan Zhu, Mei Zhou, Menghan Hu, Zhanzhan Cheng, Nengfeng Qian, Changyin Wu, Guangtao Zhai, and Xiao-Ping Zhang. 2025. Optimizing video-based respiration monitoring: Motion artifact reduction and adaptive ROI selection. *IEEE Transactions on Multimedia* (2025).

[59] Haibo Zhao, Yongtao Ma, and Xiaofeng Wang. 2024. MN-UIV: Multimodal Neural Network Enabling User Identity Verification Based on Millimeter Wave Radar. *IEEE Internet of Things Journal* 11, 15 (2024), 26304–26313.

[60] Langcheng Zhao, Rui Lyu, Hang Lei, Qi Lin, Anfu Zhou, Huadong Ma, Jingjia Wang, Xiangbin Meng, Chunli Shao, Yida Tang, et al. 2024. AirECG: Contactless electrocardiogram for cardiac disease monitoring via mmWave sensing and cross-domain diffusion model. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 8, 3 (2024), 1–27.

[61] Yimeng Zhao, Zirui Zhang, Wei Huang, Zhao Li, Jianqi Wang, Chuantao Li, and Fugui Qi. 2025. Microwave HeartPrint: Non-contact Continuous Authentication Based on Multi-View Dual-Domain Feature of Heart Micro-motion Using UWB Bio-Radar. *IEEE Internet of Things Journal* (2025).

[62] Xiaokang Zhou, Wei Liang, Jianhua Ma, Zheng Yan, and Kevin I-Kai Wang. 2022. 2D Federated Learning for Personalized Human Activity Recognition in Cyber-Physical-Social Systems. *IEEE Transactions on Network Science and Engineering* 9, 6 (2022), 3934–3944. doi:10.1109/TNSE.2022.3144699

[63] Xiaokang Zhou, Wei Liang, Jinhua She, Zheng Yan, and Kevin I-Kai Wang. 2021. Two-layer Federated Learning With Heterogeneous Model Aggregation for 6G Supported Internet of Vehicles. *IEEE Transactions on Vehicular Technology* 70, 6 (2021), 5308–5317. doi:10.1109/TVT.2021.3077893