Sense with Polyface Mirror: Enhancing Wi-Fi Sensing Diversity via Programmable Metasurfaces

Long Fan[1,*], Yinghui He[2,*], Lei Xie[1,†], Serene Zhang[3], Jun Luo[2,†]
[1]Nanjing University, China [2]Nanyang Technological University, Singapore [3]Raffles Institution, Singapore
Email: fanl@smail.nju.edu.cn, {yinghui.he, junluo}@ntu.edu.sg, lxie@nju.edu.cn

ABSTRACT

While gaining significant attention for device-free applications, Wi-Fi sensing still faces challenges in differentiating multiple targets; this stems from the design priorities of Wi-Fi systems that prioritize coverage and stability over sensing diversity. Existing proposals that either expand bandwidth or increase antennas to enhance *sensing diversity* can be confined by the limited access to Wi-Fi firm/hardware. To this end, we propose Mirror-Fi, a novel Wi-Fi sensing system that improves sensing diversity without modifying Wi-Fi firm/hardware. Exploiting the reconfigurability of *metasurfaces*, Mirror-Fi augments beamforming with spatially significant features, facilitating the construction of exclusive sensing signal links for individual targets. We innovate in an encoding scheme that equips each metasurface with a distinct phase coding sequence to mark link uniqueness. We then train a deep neural model to leverage prior coding sequences for decomposing non-linearly superimposed channel samples into mutually independent channels; it removes the need for complex channel matrix parameter estimation and mitigates hardware-related offsets inherent to Wi-Fi. Extensive evaluations demonstrate that, with a sufficient number of auto-configured metasurfaces, Mirror-Fi successfully achieves multi-target sensing.

CCS CONCEPTS

• **Human-centered computing** → **Ubiquitous and mobile computing**; • **Computing methodologies** → *Machine learning*.

KEYWORDS

Metasurfaces, Wi-Fi sensing, sensing diversity, multi-target sensing, spatial resolution.

ACM Reference Format:
L. Fan, Y. He, L.Xie, S. Zhang, J. Luo. 2026. Sense with Polyface Mirror: Enhancing Wi-Fi Sensing Diversity via Programmable Metasurfaces. In *ACM/IEEE International Conference on Embedded Artificial Intelligence and Sensing Systems (SenSys'26), May 11–14, 2026, Saint Malo, France.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3774906.3800469

* Both authors have contributed equally to this research, which is done when Long Fan works as a visiting scholar at NTU, under the funding support offered by China Scholarship Council (CSC).

† Lei Xie and Jun Luo are both the corresponding authors.

SenSys'26, May 11–14, 2026, Saint Malo, France

ACM ISBN 979-8-4007-2309-4/2026/05.
https://doi.org/10.1145/3774906.3800469

1 INTRODUCTION

Wi-Fi sensing has garnered considerable attention from both academia and industry due to its potential applications in various domains. Leveraging Wi-Fi signals (specifically channel state information, or CSI [16]), Wi-Fi sensing enables contact- and device-free sensing, including vital signs monitoring [41, 43, 76], gesture detection [22, 61], activity recognition [18, 31, 34], as well as localization and motion tracking [1, 9, 63, 74]. However, a significant challenge persists in the evolution of Wi-Fi sensing despite its promising advancements: while fine-grained sensing seeks to enhance the diversity of signal links related to sensing target, communication systems are inherently designed to prioritize broader signal coverage and stability. Consequently, enhancing sensing capabilities demands trading certain aspects (e.g., communication coverage, stability, and hardware constraints) within the current Wi-Fi infrastructure for *sensing diversity* in spectral, spatial, and temporal sense.

Existing solutions to enhance sensing diversity typically focus on expanding bandwidth, adding antennas, or increasing links. Expanding bandwidth, though promising, faces several challenges, including issues with channel availability [65, 68], frequency response boundary effects, insufficient coherence time, interference with communication [41], and limited applicability to general sensing tasks. Alternatively, adding antennas to Wi-Fi devices does improve spatial diversity, yet such approaches [67, 72] significantly increase deployment costs and present integration challenges due to the firm/hardware limitations of commodity Wi-Fi devices. Finally, whereas increasing links directly contributes to sensing diversity [10, 22, 26, 52], its reliance on user devices (e.g., smartphones) [10, 22] or firmware restricts its application [26, 52] in device-free scenarios and raises privacy and security concerns [21, 34, 44]. Therefore, *can we devise a solution that enhances sensing diversity without touching Wi-Fi firm/hardware and disrupting communication?*

Fortunately, metasurfaces are widely deemed as a powerful extension to commodity wireless devices. Starting with *passive* versions leveraging fixed dielectric properties to control signal reflection [45, 50] and transmission [11, 54], they already exhibit potential to improve diversity. However, their immutable propagation nature (e.g., fixed reflection paths [53, 62]) hinder adaptation to environmental dynamics. Consequently, *active* (programmable) metasurfaces have emerged to enable real-time control on reflected signals [14, 37, 42]. As Wi-Fi sensing faces unknown target variations (e.g., activities), only active metasurfaces may effectively counteract these disturbances through dynamic signal optimization techniques (e.g., beamforming [15]). Therefore, leveraging active metasurfaces as external "plug-in" components to explore *reflection path diversity* (RPD) holds great potential for improving Wi-Fi sensing performance without disrupting communication or modifying

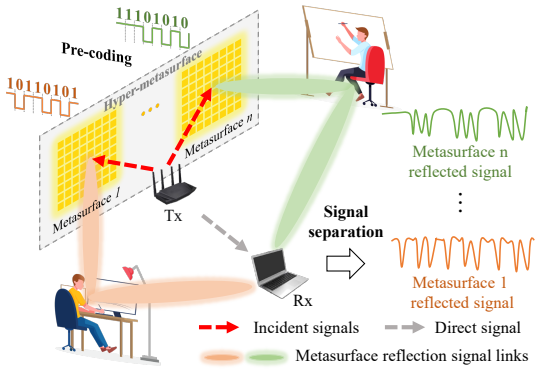

Figure 1: Enhancing sensing diversity via reflection path diversity (RPD) enabled by metasurfaces.

firm/hardware. Here, RPD refers to exploiting multiple metasurfaces to dynamically establish unique sensing links for individual targets, as shown in Figure 1.

Although the basic concept appears straightforward, this novel diversity presents three non-trivial challenges. First, constructing sensing links through active beamforming with metasurfaces requires knowledge of target locations, but commodity Wi-Fi devices cannot accurately acquire these locations, again due to the inherent limitations in bandwidth and antenna resources. Second, the homogeneous nature of metasurfaces causes their reflected signals to exhibit high similarity, complicating the identification of individual metasurfaces associated with distinct sensing links; yet this identification is crucial to realizing RPD. Third, achieving RPD also necessitates separating the reflected signal paths from all metasurfaces at a Wi-Fi receiver (Rx), yet this is highly non-trivial given intrinsic coupling between metasurface-reflected signals and random phase shifts.

To tackle these challenges, we propose Mirror-Fi as a novel wireless sensing system that leverages the reconfigurability of external metasurfaces to enhance sensing diversity; it operates seamlessly with communications without requiring modifications to existing Wi-Fi firm/hardware, as shown in Figure 1. Our approach begins with a metasurface-assisted multi-target azimuth estimation method enabled by a large-aperture dynamic beamformer that integrates multiple metasurfaces into a hyper-metasurface. This method extracts angular power spectrum through dynamic searching, facilitating accurate identification of position candidates. To distinguish the individual sensing links constructed by distinct metasurfaces, we design a pre-coding scheme that modulates the phase shifts of respective metasurfaces so that each metasurface is equipped with a unique phase coding sequence. At the Rx side, the challenge of signal separation persists even with our pre-coding scheme: since existing signal processing techniques fail to achieve the goal, we introduce a novel model-based deep neural network (mDNN) to tackle this challenge. mDNN exploits the pre-coding sequence unique to each metasurface to train its network in separating superimposed samples into mutually independent channels. Finally, we implement Mirror-Fi prototype involving customized metasurfaces and commodity Wi-Fi devices, and conduct extensive experiments to evaluate its performance. In summary, our main contributions are:

- We innovate in constructing Mirror-Fi to enhance Wi-Fi sensing diversity (via RPD) while maintaining full compatibility with existing Wi-Fi hardware.
- We integrate multiple metasurfaces into a large-aperture beamformer and leverage a dynamic search strategy to achieve accurate target detection.
- We propose a pre-coding scheme that assigns unique phase coding sequences to individual metasurfaces, aiming to facilitate sensing links differentiation.
- We design a scenario-agnostic mDNN architecture at Wi-Fi Rx to handle the superposed and phase-distorted metasurface signals.
- We implement a prototype of Mirror-Fi and validate its multi-target sensing performance through extensive real-world experiments.

The remainder of this paper is organized as follows. Section 2 provides background on Wi-Fi sensing with metasurfaces and the motivation of Mirror-Fi. Section 3 presents the design of Mirror-Fi, including azimuth estimation using metasurfaces, phase pre-coding scheme, and separation of individual sensing links. Sections 4 and 5 describe the implementation and evaluate the performance of Mirror-Fi, respectively. We briefly discuss related work in Section 6. Finally, Section 7 concludes our paper.

2 BACKGROUND AND MOTIVATION

In this section, we introduce reconfigurable metasurface and its sensing model, analyze the spatial diversity it enables, and demonstrate the feasibility of implementing CSI phase pre-coding for Wi-Fi signals through metasurfaces.

2.1 Wi-Fi Sensing with Metasurface

A reconfigurable metasurface is an emerging artificial electromagnetic material technology composed of numerous programmable sub-wavelength units, referred to as meta-atoms. By adjusting the voltage applied to the meta-atom, the phase shift to the reflected wireless signal can be controlled, enabling the reconfiguration of wireless signals [46, 70, 75]. Consequently, when deployed in a Wi-Fi sensing system, metasurfaces may purposefully affect the CSI measured at the Rx side. Specifically, considering a Wi-Fi sensing system assisted by N metasurfaces, the received signal consists of two kinds of components: one that reaches the Rx directly without affected by any metasurfaces, denoted by h^{D}, and another that is reflected by certain metasurfaces before arriving at the Rx, denoted by h_n^{M} for the n-th metasurface. Since our goal is to leverage the

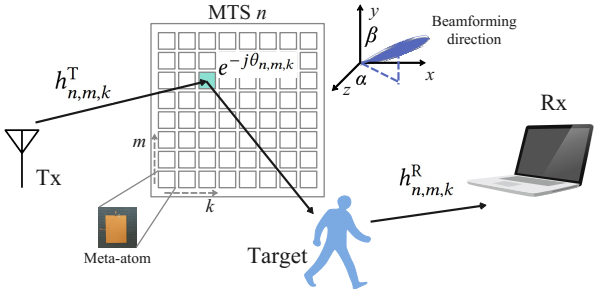

Figure 2: Channel reflected by the n-th metasurface.

diversity introduced by the metasurfaces, we focus on the reflected component. As shown in Figure 2, for the reflective meta-atom (m,k) on the n-th metasurface, the reflected channel h_n^{M} is determined by three factors: i) the channel from the transmitter (Tx) to the reflective meta-atom $h^{\mathrm{T}}_{n,m,k}$, ii) the phase shift $\theta_{n,m,k}$ introduced by the meta-atom; and iii) the channel from the reflective meta-atom to the Rx antenna, $h^{\mathrm{R}}_{n,m,k}$, which includes the path associated with the target. Thus, the CSI between the Tx and Rx can be modeled as:[1]

$$h = \left(h^{\mathrm{D}} + \sum_n \left(\sum_{m,k} h^{\mathrm{T}}_{n,m,k} e^{-j\theta_{n,m,k}} h^{\mathrm{R}}_{n,m,k} \right) \right) e^{j\theta^{\mathrm{offset}}}, \tag{1}$$

where θ^{offset} is the time-varying phase offset brought by the Rx.

2.2 Diversity of Metasurfaces

Multi-target sensing with Wi-Fi devices presents challenges due to the low resolution caused by the limited number of antennas and bandwidth. Metasurfaces may overcome this limitation by enabling beamforming through tuning of the phase shift $\theta_{n,m,k}$, allowing Wi-Fi signals to be steered toward specific directions [15]. As shown in Figure 2, for a desired beamforming direction with azimuth α and elevation β, the phase shift of each meta-atom is given by:

$$\theta_{n,m,k} = (\angle h^{\mathrm{T}}_{n,m,k} - \frac{2\pi}{\lambda}(x_k \cos\beta + y_m \sin\beta)\sin\alpha), \tag{2}$$

where λ is the wavelength, and x_k and y_m denote the respective distances of the k-th and m-th meta-atoms from the center, respectively. Moreover, $\angle h^{\mathrm{T}}_{n,m,k}$, i.e., the phase of $h^{\mathrm{T}}_{n,m,k}$, is given by $\frac{2\pi}{\lambda}\sqrt{x_k^2 + y_m^2 + \bar{d}^2}$, with $\bar{d}$ being the distance between the Tx (i.e., feed) and the center of the metasurface. In practice, cost considerations often limit metasurfaces to have only 1-bit meta-atoms, where $\theta_{n,m,k}$ can take only two discrete values: 0 and π. Therefore, the phase shift $\theta_{n,m,k}$ derived from Eqn. (2) needs to be binarized as follows:

$$\theta^{\mathrm{bin}}_{n,m,k} = \begin{cases} \pi, & \text{if } |\theta_{n,m,k}| > \frac{\pi}{2}, \\ 0, & \text{otherwise,} \end{cases} \quad \forall n,k,m. \tag{3}$$

To validate the effectiveness of metasurfaces in enhancing spatial diversity for Wi-Fi sensing, we conduct experiments by configuring the phase shifts of meta-atoms, as described in Eqn. (2), to direct the beam toward desired directions. Specifically, Figure 3(a) illustrates the configuration where the beamforming direction of the metasurface, consisting of an 8×8 grid of meta-atoms, is set to $-30°$ and $20°$. The received CSI is then measured at a distance of 3 m from the metasurface, within a field of view of $120°$. As shown in Figure 3(b), the experiment results exhibit a 6 dB increase in the CSI signal strength in the target direction when using the metasurface, consistent with the simulation outcomes. These findings highlight the ability of metasurfaces to enhance spatial diversity by steering Wi-Fi signals, offering the potential for multi-target localization and sensing.

[1] Since metasurface exhibits similar effect on distinct subcarriers, this expression applies to all subcarriers used in Wi-Fi communications.

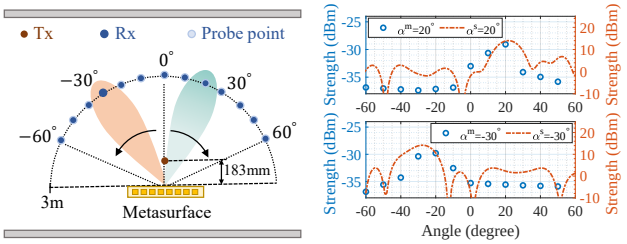

(a) The experiment layout. (b) Beamforming with metasurface.

Figure 3: Verifying the spatial diversity of metasurfaces: (a) layout and (b) signal strength at various azimuth directions given elevation $\beta = 0$, where the dashed and dot lines represent the simulation and experiment results, respectively.

2.3 Feasibility of CSI Phase Pre-coding

Allocating a metasurface for each target to establish a unique sensing link does offer RPD that may conceptually address multi-target sensing in practice. However, owing to the homogeneous design of metasurfaces and the inherent limitations of Wi-Fi Rx antennas, the Rx encounters difficulties in directly distinguishing reflected signals originating from different metasurfaces. To address this issue, we propose a phase pre-coding that uniquely encodes each metasurface, allowing it to serve as a distinct identifier for its respective propagation path. It implements phase pre-coding by adding an extra shared phase, denoted by $\varphi_n(t)$ to all $\theta_{n,m,k}$ of the n-th metasurface and at each time slot t; $\varphi_n(t)$ is updated as:

$$\varphi_n(t) = \begin{cases} 0, & \text{if } \varphi_n(t-1) = \pi \\ \pi, & \text{if } \varphi_n(t-1) = 0 \end{cases}. \tag{4}$$

According to Eqn. (1), a uniform phase flip preserves phase differences among meta-atoms, enabling global phase shifts for controlled signal modulation and enhanced RPD without altering beam geometry or orientation. For example, flipping the phase values of all meta-atoms, from 0 to π (encoding a logical '1') at time slot $t-1$, and from π to 0 (encoding a logical '0') at time slot t, enables the link reflected by the metasurface to carry phase-encoded information.

To validate this proposal, we employ a metasurface to encode the signal transmitted by a Wi-Fi Tx while utilizing two Rx antennas to capture CSIs. As shown in Figure 4(a), the phase ratio between the

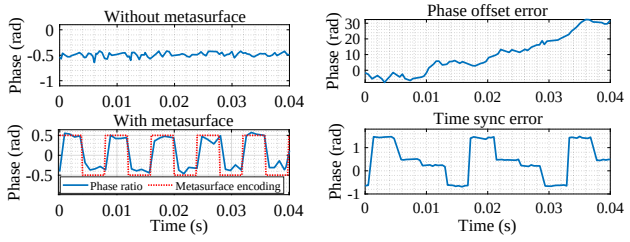

(a) Pre-coding one metasurface. (b) Pre-coding two metasurfaces.

Figure 4: Feasibility and challenges of CSI phase encoding. (a) Pre-coding a metasurface makes it identifiable, but (b) links reflected by two pre-coded metasurfaces may not be readily differentiable.

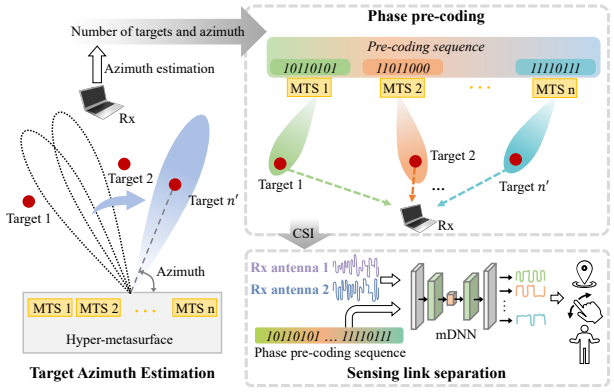

Figure 5: Overview of Mirror-Fi.

two antennas changes when our metasurface pre-coding is applied, indicating variations consistent with the pre-coding sequence of the metasurface. While phase pre-coding effectively identifies the reflected link by each metasurface, extending this technology to multiple metasurfaces introduces additional challenges, notably random phase offsets and time synchronization. As shown in Figure 4(b), when two metasurfaces encode signals simultaneously, the phase observed at a single Rx antenna is influenced by phase offsets, complicating the recovery of the individual phase-coded signals from the two independent metasurfaces. Due to these challenges, a joint pre-coding and separation design shall become one of the main design objectives of our Mirror-Fi.

3 SYSTEM DESIGN

Our Mirror-Fi is designed to enhance sensing diversity (i.e., RPD) by using metasurfaces to construct unique sensing links for individual targets, seamlessly operating with communication systems without requiring modifications to existing Wi-Fi hardware. As illustrated in Figure 5, Mirror-Fi consists of the following key modules:

- Target azimuth estimation: A hyper-metasurface configuration mechanism with phase coordination to narrow the main lobe of beamforming, thereby enabling accurate estimation of multiple target azimuths.
- Phase pre-coding: A metasurface pre-coding scheme for high sensing link diversity (or RPD) and precise signal synchronization across the metasurfaces.
- Sensing link separation: a multi-link separation neural model mDNN that leverages metasurface pre-coding sequences to separate mixed CSI signals and to isolate independent channel responses efficiently.

With the above design, Mirror-Fi supports independent sensing of multiple individuals and different types of motions simultaneously.

3.1 Target Azimuth Estimation

The first requirement in constructing unique sensing links for individual targets is to be aware of the target azimuth estimation. This is certainly not achievable via conventional Wi-Fi sensing given its limited bandwidth and antennas [41]. To this end, we propose to enhance spatial resolution (e.g., azimuth) and thus accurately estimate multiple targets azimuth, by leveraging a *hyper-metasurface configuration*. It consists of two key components: i) a configuration strategy employing phase coordination to expand the effective beamforming aperture and ii) a radio direction estimation algorithm that constructs angular power spectrum (APS) to accurately estimate target quantity and their respective azimuths.

3.1.1 Hyper-metasurface Configuration Strategy. A straightforward way for improving target azimuth estimation accuracy is to configure multiple metasurfaces to perform beamforming in the same direction simultaneously, leveraging the superposition of signals reflected by metasurfaces to enhance signal strength. However, due to the misalignment among these metasurfaces, multiple metasurfaces perform beamforming *independently* may not effectively widen effective aperture (i.e., narrowing the main lobe width), as shown by the baseline of Figure 6(b), thereby limiting its ability to enhance spatial diversity in multi-target azimuth estimation. To expand the effective aperture for beamforming, we control multiple metasurfaces deployed horizontally (i.e., along x-axis) in a *coordinated* manner to achieve large-aperture beamforming, as shown in Figure 6(a); this coordinated control allows us to tackle the main cause of the misalignment among metasurfaces that hampers a coherent superposition of their signals. As the misalignment stems from both geometric sense and signal space, we handle the former here but leave the latter to Section 4 where actual implementation is discussed. In particular, we propose a phase coordination method to mitigate the grating lobe effect [60] caused by inconsistent meta-atom spacing, especially the gap between adjacent metasurfaces.

Specifically, Figure 6(a) indicates that, even if both the n-th and $(n-1)$-th metasurfaces adopt the same beamforming direction α, the gap between them prevents coherent signal superposition in the target direction. To achieve coherent signal superposition between neighboring metasurfaces in the target direction, we define the phase compensation between them as: $\Delta\phi(\alpha) = \frac{2\pi}{\lambda} d^{\text{gap}} \sin\alpha$. By applying the same phase compensation $\Delta\phi(\alpha)$ to all meta-atoms on the n-th metasurface, we obtain the desired phase values for each meta-atom as follows:

$$\theta^{\text{des}}_{n,m,k}(\alpha) = \theta^{\text{des}}_{n,m,k}(\alpha) + (n-1)\Delta\phi(\alpha). \tag{5}$$

After binarizing $\theta^{\text{des}}_{n,m,k}$ in the same manner as Eqn. (3) and applying it to the metasurface, the desired beamforming direction can be achieved. The experiment results in Figure 6(b) demonstrate that, with a gap of 14.3mm, the beamforming performance (including both simulation and measurement) of the hyper-metasurface after

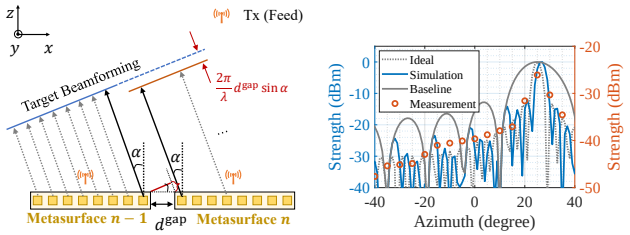

(a) Positional relationship. (b) Beamforming performance.

Figure 6: Hyper-metasurface configuration: (a) positional relationship and (b) beamforming performance when $\alpha = 25°$, $\beta = 0°$, and $d^{\text{gap}} = 14.3$mm.

phase coordination is consistent with the ideal scenario (i.e., without gaps). Moreover, the same figure shows that the half-power beamwidth of the hyper-metasurface is reduced by a factor of 4 (exactly corresponding to the number of metasurfaces) compared with the case without phase coordination (i.e., the baseline scenario).

3.1.2 Radio Direction Estimation. After phase configuration, we can now leverage the narrow main lobe achieved by the hyper-metasurface to enhance the beamforming spatial scanning resolution, thereby enabling multi-target azimuth estimation. Unlike traditional multi-target estimation methods (e.g., MUSIC [31, 55]) only passively processing Rx signals, Mirror-Fi scans different directions sequentially leveraging its hyper-metasurface and determines the presence of targets by analyzing the received CSI at Rx.

Since the scan can be interfered with by the *direct signal* (see Figure 1) and common (RF) clutter, probing signals reflected by the hyper-metasurface need to be further processed. Given a base h_{t0} probed under no-target scenario, it is intuitive to subtract it out of the current probe outcome h_t to remove such background interference. However, as shown in Eqn. (1), the different phase offsets in the two probes at the receiver disrupt this subtraction, preventing effective interference removal via subtraction. To address this, we instead compare the absolute values of CSI from two probes. A significant variation in these values suggests the emergence of a new reflector (i.e., a target) along the probe direction. Based on this principle, we define the sum of squared differences across two receiving antennas for direction α as

$$\rho(\alpha) = (|h_t^{\mathrm{A1}}(\alpha)| - |h_{t0}^{\mathrm{A1}}(\alpha)|)^2 + (|h_t^{\mathrm{A2}}(\alpha)| - |h_{t0}^{\mathrm{A2}}(\alpha)|)^2, \quad (6)$$

where $h_{t0}^{\mathrm{A1}}(\alpha)$ and $h_{t0}^{\mathrm{A2}}(\alpha)$ denote the base CSI at two Rx antennas, while $h_t^{\mathrm{A1}}(\alpha)$ and $h_t^{\mathrm{A2}}(\alpha)$ represent the current CSI. To further enhance estimation accuracy, we non-coherently accumulate the real-valued metric $\rho(\alpha)$ across all subcarriers. Since the gain of the main lobe in the beamforming is significantly higher than that of the sidelobes, the target is expected to appear at the peak of the $\rho(\alpha)$ curve. Therefore, we first apply a predefined threshold to filter out part of the directions without targets and then perform a peak search within the remaining directions to determine the number of targets and their corresponding azimuths. Figure 7 shows the experiment results in a four-target scenario, with the probe azimuth range from $\alpha^{\min}$ to $\alpha^{\max}$ and a step size of α^{s} (e.g., $\alpha = \pm 60°$ and $\alpha^{\mathrm{s}} = 5°$). Figure 7(b) illustrates that threshold detection effectively

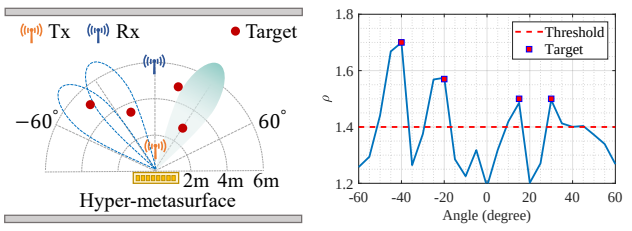

(a) Layout. (b) APS.

Figure 7: Multi-target azimuth estimation of Mirror-Fi: (a) layout, where the real azimuths of four targets are −45°, −25°, 15°, and 30°, respectively, and (b) the angular power spectrum with the estimated azimuths of the four targets are −40°, −20°, 15°, and 30°, respectively.

filters out insignificant directions, with the azimuth error remaining within 5° after peak search. This level of errors virtually has no impact on downstream (indoor) sensing applications, as they stay within the human body width.

3.2 Phase Pre-coding Scheme

Given the results obtained in Section 3.1, we are now ready to associate metasurfaces with individual targets. As pointed out by Section 2.3, the challenge faced by the association process is the homogeneity among all metasurfaces. To this end, we propose a phase pre-coding scheme that assigns a pre-coding sequence to each metasurface, enabling its unique identification. We also address the design considerations for pre-coding length and time synchronization. A byproduct of pre-coding scheme is that subsequent sensing tasks can focus solely on obtaining the corresponding sensing links, with no need to account for the background interference.

3.2.1 Sensing Model with Phase Pre-coding. Since metasurfaces' homogeneity makes reflected signals indistinguishable, the Rx cannot directly recognize which metasurface is logically associated with a certain target. Therefore, to enhance the distinguishability of reflected signals from different metasurfaces, we assign a unique phase pre-coding for each metasurface by introducing the additional shared phase $\varphi_n(t)$, as described in Section 2.3. This phase pre-coding introduces distinct variations in the received CSI over time, enabling the Rx to differentiate signals from different metasurfaces. Specifically, let $\boldsymbol{c}_n = [\varphi_n(1), \varphi_n(2), ..., \varphi_n(T)]$ denote the pre-coding sequence for the n-th metasurface, where T is the length of the pre-coding sequence (i.e., the number of time slots). At time slot t (with $t = 1, ..., T$), the received CSI can be rewritten as:

$$h(t) = \left(h^{\mathrm{D}} + \sum_{n=1}^{N} \left(h_n^{\mathrm{M}} e^{-j\varphi_n(t)} \right) \right) e^{j\theta^{\mathrm{offset}}(t)}, \quad (7)$$

where $h_n^{\mathrm{M}} = \sum_{m,k} (h_{n,m,k}^{\mathrm{T}} e^{-j\theta_{n,m,k}} h_{n,m,k}^{\mathrm{R}})$ is the CSI component that reflected by the n-th metasurface without $\varphi_n(t)$. By collecting measurements over T time slots, the received CSI vector, $\boldsymbol{h}^{\mathrm{En}}$, with phase pre-coding can be expressed as:

$$\boldsymbol{h}^{\mathrm{En}} = [h(1), h(2), ..., h(T)]'. \quad (8)$$

3.2.2 Pre-coding Sequence Design. Based on the above model, the distinguishability of reflected signals from different metasurfaces can be enhanced by increasing the length of the pre-coding sequence, $\boldsymbol{c}_n$, but comes at the expense of reduced phase continuity of the reflected signals. To this end, we propose a method to determine the pre-coding length T that ensures high-fidelity separation of metasurface sensing links at Rx. This method is achieved by minimizing T while guaranteeing a unique separation based on Eqns. (7) and (8), which in turn requires that the total number of unknown parameters on all sensing links is less than the number of observations collected by the receiving antennas in T time slots. Specifically, the Rx collects CSI from N^{Rx} antennas, where each CSI measurement is a complex number comprising real and imaginary components. According to Eqn. (7), at a given time slot t, the unknown parameters to be estimated include: i) the direct signal channel contributing $2N^{\mathrm{Rx}}$ unknown parameters; ii) the N

metasurface-reflected channels introducing $2N^{\mathrm{Rx}}N$ unknown parameters; and iii) an additional unknown parameter corresponding to the phase offset. Given that the total number of independent observations at the Rx over the measurement interval of T time slots is $2N^{\mathrm{Rx}}T$, the minimum pre-coding sequence length T for recovering the metasurface-reflected channel from the received CSI needs to satisfy:

$$T \geq \frac{2N^{\mathrm{Rx}}(N+1)}{2N^{\mathrm{Rx}}-1}. \tag{9}$$

Let us consider that Mirror-Fi employs $N^{\mathrm{Rx}} = 2$ Rx antennas and $N = 4$ metasurfaces, the length of the pre-coding sequence should thus be set to $T = 7$. Consequently, we have the design of a 7-bit pre-coding sequence for each metasurface based on the Gray code concept [6]. Specifically, during two consecutive time slots, the extra shared phase, $\varphi_n(t)$, of only one metasurface is reversed, while all others remain unchanged. The pre-coding sequences for the four metasurfaces over the time slot period are presented in Table 1. Although phase pre-coding can effectively enhance the distinguishability of metasurface's reflected signal, the periodicity of the pre-coding sequence, making it difficult to accurately identify the start time of pre-coding (i.e., $\varphi_n(1)$'s position), as shown in Figure 8(a). To mitigate potential time alignment issues (as discussed in Section 2.3), we add a synchronization bit after the 7-bit pre-coding sequence by repeating $\varphi_n(6)$ (i.e., 'bit-6') of the 7-bit pre-coding sequence, finally resulting in an 8-bit pre-coding sequence, as shown in Table 1. The underlying rationale is that adjacent synchronization bits are less error-tolerant, while bits repeated at longer intervals are more susceptible to desynchronization (i.e., different patterns in the received CSI). As illustrated in Figure 8(b), the introduction of the synchronization bit serves as a marker, enabling accurate identification of the starting position of the pre-coding sequence within a pre-coding period. Finally, to fit the sequence within CSI coherence time of roughly 50ms [58], the sequence duration should be set below 50ms.

3.3 Sensing Link Separation

To extract a unique sensing link for each target on the Rx side for RPD sake, the CSI component reflected by each metasurface needs

Table 1: Pre-coding sequences for four metasurfaces.

	bit1	bit2	bit3	bit4	bit5	**bit6**	bit7	**sync**
$\boldsymbol{c}_1$	π	0	0	0	0	π	π	π
$\boldsymbol{c}_2$	π	π	0	0	0	0	π	0
$\boldsymbol{c}_3$	π	π	π	0	0	0	0	0
$\boldsymbol{c}_4$	π	π	π	π	0	0	0	0

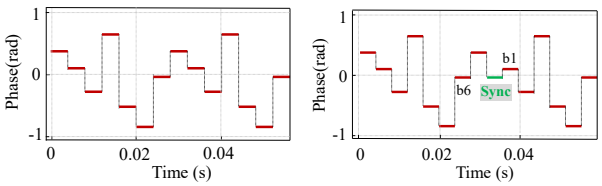

(a) W/o sync. (b) W/ sync.

Figure 8: With and without synchronization code.

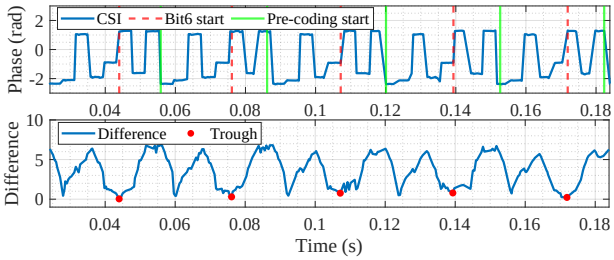

Figure 9: Estimation of the start time for pre-coding in the received signal.

to be separated from the composite CSIs. Although the problem is theoretically solvable given the properly designed pre-codings, the uncontrollable non-linearity in the complex coupling among multiple metasurface reflections largely incapacitates traditional statistical methods [4, 7]. To this end, we resort to a deep neural model to exploit the pre-coding sequences as prior knowledge and thus to enable an effective separation of the composite CSIs.

3.3.1 Determination of Pre-coding Start Time. To facilitate the effective separation of mixed signals on the Rx side using pre-coding information, it is essential to synchronize the start time between pre-coding sequences and the received CSI. To this end, our core approach involves identifying the synchronization bit pattern by computing the discrepancies between the CSI samples separated by one time slot. In particular, the difference at time slot t is obtained by calculating the absolute difference between all samples within time slot t and those within the time slot $t+2$. Since 'bit6' and the 'sync' share the same pattern, the difference exhibits a unique minimum value within each pre-coding period. To this end, a sliding window with the window length equal to the pre-coding period (e.g., 32ms) is moved to obtain the differences and in turn identify the trough within the window, as illustrated in Figure 9. To refine estimation, the detected minima and their corresponding timestamps are fitted using the linear least squares method, aiming to further improve detection accuracy. Given that each bit in the pre-coding sequence occupies the same time span (e.g., 4ms), the start time (i.e., 'bit1') of the next pre-coding sequence is just three time slots after that of 'bit6', as shown by Figure 9.

3.3.2 Network Architecture. After synchronizing the received CSI with the pre-coding sequences, we propose a model-based deep neural network (mDNN) to decompose the composite CSIs into individual metasurface-reflected components. The core principle of mDNN is to design a signal separation model based on prior pre-coding guide, which distinguishes it from conventional blind source separation approaches [7] that rely solely on statistical assumptions such as source independence. Specifically, our separation model employs an attention mechanism to learn the association between the unique signature of each pre-coding sequence and its corresponding signal component within the mixed CSI, thereby enabling the network to effectively handle the complex, non-linear couplings among reflections. Formally, our separation model implements a mapping $\mathcal{F}_\omega : \{\boldsymbol{H}^{\mathrm{En}}(t), \boldsymbol{c}_n\} \mapsto \{\boldsymbol{H}_n^{\mathrm{M}}(t)\}_{n=1}^{N}$, where $\mathcal{F}_\omega$ denotes the parameterized neural network with learnable weights ω. Here, $\boldsymbol{H}^{\mathrm{En}}(t) = [\boldsymbol{h}^{\mathrm{En,A1}}(t), \boldsymbol{h}^{\mathrm{En,A2}}(t)] \in \mathbb{C}^{T^{\mathrm{s}}\times 2}$ represents the CSI time-series over T^{s} samples from two receiving antennas, and

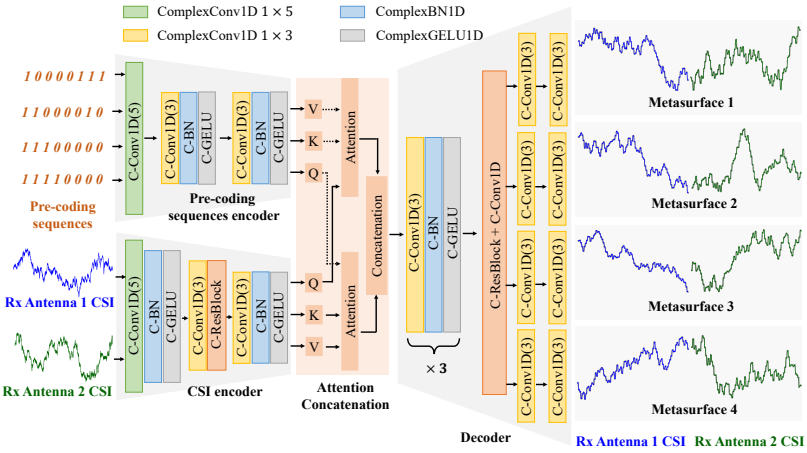

Figure 10: The mDNN architecture takes mixed CSIs and pre-coding sequences as input, processes them through a fusion module, and outputs the separated CSIs.

Module	Operations	Output dim
Input stream	Concat (CSI + Sequences)	$(B, 12, T)$
Sequence encoder	C-Conv1D (4 → 32)	$(B, 64, T)$
	C-Conv1D (32 → 64) + C-BN + C-GELU	$(B, 128, T)$
	C-Conv1D (64 → 128) + C-BN + C-GELU	$(B, 256, T)$
CSI encoder	C-Conv1D (2 → 32) + C-BN + C-GELU	$(B, 64, T)$
	C-Conv1D (32 → 64)	$(B, 128, T)$
	C-ResBlock(64 → 64)	$(B, 128, T)$
	C-Conv1D (64 → 128) + C-BN + C-GELU	$(B, 256, T)$
Concatenation	2× (*Parallel attention*)	
	C-Conv1D(128 → 128, $k = 1$)	Q,K,V: $(B, 256, T)$
	Dot-product attention	$(B, 256, T)$
	Concat	$(B, 512, T)$
	C-Conv1D(256 → 128, $k = 1$)	$(B, 256, T)$
Decoder	C-Conv1D (128 → 128) + C-BN + C-GELU	$(B, 256, T)$
	C-ResBlock + C-Conv1D	$(B, 256, T)$
	4× (*Parallel block*)	
	C-Conv1D (32 → 16)	$(B, 32, T)$
	C-Conv1D (16 → 2)	$(B, 4, T)$
Final Output		$(B, 16, T)$

Figure 11: Detailed configuration of the mDNN, where dimensions are denoted as (batch B, real + imaginary channels, time T).

$\{\boldsymbol{c}_n\}_{n=1}^{N}$ are the pre-coding sequences for N metasurfaces. The output $\boldsymbol{H}_n^{\mathrm{M}}(t) = [\boldsymbol{h}_n^{\mathrm{M,A1}}(t), \boldsymbol{h}_n^{\mathrm{M,A2}}(t)] \in \mathbb{C}^{T^s \times 2}$ represents the reconstructed CSI components reflected by the n-th metasurface for both receiving antennas. We adapt the classical encoder-decoder architecture [2, 56] and fine-tune it for our signal separation task. As illustrated in Figure 10, the separation model $\mathcal{F}_\omega$ consists of three components: i) an encoder extracts features from the complex CSI and pre-coding sequences, respectively, ii) a fusion module integrates the pre-coding features to selectively focus on different metasurface signals and extract their respective features, and iii) a decoder reconstructs the CSI for each metasurface individually by separating them at feature layer. In addition, to preserve phase information and minimize potential loss, we employ a dual-channel complex convolution layer [57] instead of simply splitting the CSI complex signals into real and imaginary channels. The detailed configuration of this module is presented in Figure 11.

3.3.3 Training Strategy. To train the mDNN as a generalized solver rather than a model overfitted to a specific physical space, we construct our training data from diverse Wi-Fi CSI datasets [19, 59] that contain various application scenarios. Specifically, we randomly select eight CSIs from two antennas and use the pre-coding sequences in Table 1 to generate reflected signals from n-th metasurface (i.e., ground truth), $\boldsymbol{H}_n^{\mathrm{M}}(t)$, and then according to Eqn. (7) construct a mixed signal, $\boldsymbol{H}^{\mathrm{En}}(t)$, as mDNN input. Based on this method, our experiment dataset comprises 55,000 synchronized CSI temporal segments partitioned into training (45,000 instances), validation (5,000 instances), and test (5,000 instances) sets. To ensure the model remains invariant to the time-varying hardware-induced phase offset (θ^{offset}), we perform data augmentation by injecting random Gaussian noise (mean is 0, standard deviation is 0.5) into the input signal phase. This encourages mDNN to learn stable relative phase relationships rather than relying on the absolute phase, which is corrupted by common-mode disturbances. Our strategy efficiently generates diverse training samples that provide a strong inductive bias, rather than aiming to reproduce real-world physics. While the synthetic test set serves as a preliminary validation (Figure 16), we further verify generalization in Section 5.2.

To preserve the phase and magnitude information of the complex signal, we design a complex loss function composed of two components: i) a mean squared error to minimize separation errors of real and imaginary, and ii) a smoothness loss to enforce temporal consistency in the reconstructed signals. This loss function is formulated as follows:

$$\begin{aligned}\mathcal{L} =& \lambda_1 \frac{1}{N} \sum_{n=1}^{N} \left(\boldsymbol{H}_n^{\mathrm{M}} - \hat{\boldsymbol{H}}_n^{\mathrm{M}}\right)^2 \\ &+ \lambda_2 \frac{1}{N} \sum_{n=1}^{N} \sum_{t=2}^{T^s} \left|\hat{\boldsymbol{H}}_n^{\mathrm{M}}(t) - \hat{\boldsymbol{H}}_n^{\mathrm{M}}(t-1)\right|,\end{aligned} \tag{10}$$

where $\hat{\boldsymbol{H}}_n^{\mathrm{M}}$ denotes the prediction result of n-th metasurface; The parameters λ_1 and λ_2 are weighting factors that balance the loss terms, set to 1 and 0.01, respectively. We train the model for 200 epochs with a batch size of 64 using the Adam optimizer ($\beta_1 = 0.9, \beta_2 = 0.999$). The initial learning rate is 1×10^{-3}, which is reduced by a factor of 10 at epoch 150. For model evaluation, we adopt the scale-invariant source-to-distortion ratio (SI-SDR) [32] as the primary metric, assessing performance on the validation set every 10 epochs.

4 PROTOTYPE AND EXPERIMENT SETUP

This section provides an overview of Mirror-Fi's implementation and the experiment setup used for evaluation.

4.1 Implementation

We start by presenting the design of the metasurface, including the meta-atom structure and the prototype implementation. This is followed by the overall construction of Mirror-Fi that leverages multiple metasurfaces.

4.1.1 Metasurface Prototype. Figure 12(a) presents the proposed 1-bit phase-shifting programmable meta-atom design. The meta-atom comprises three metal layers (i.e., top reflection layer, intermediate ground plane, and bottom control circuit layer) and three dielectric substrates (i.e., F4B, RO4450F, and FB-4) separating

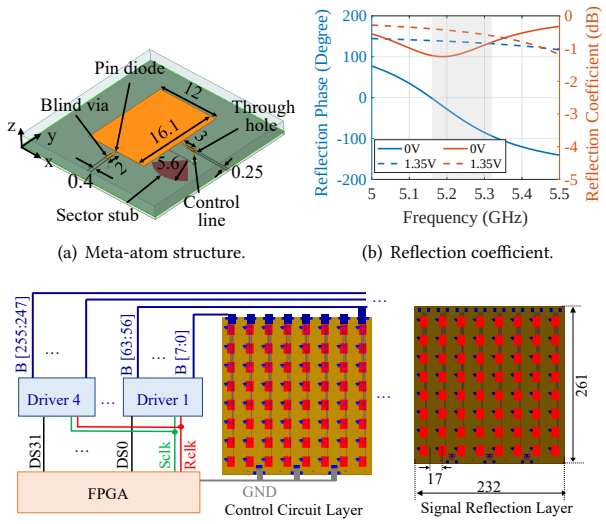

(a) Meta-atom structure. (b) Reflection coefficient.

(c) The prototype of metasurface.

Figure 12: Prototype of the metasurface: (a) structure of the meta-atom, (b) amplitude and phase of the reflection coefficient, and (c) control architecture of Mirror-Fi, with dimensions specified in millimeters (mm).

metallic layers. A surface-mounted PIN diode (MACOM MADP-000907-14020x) [24] is integrated into the structure, bridging the top reflection patch with the bottom control layer through a via hole. This setup allows for dynamic phase modulation by switching the diode's bias voltage. The performance of the meta-atom is characterized by its reflection coefficient, as shown in Figure 12(b). By toggling the bias voltage supplied to the PIN diode between 0 V (OFF state) and 1.35 V (ON state), the reflection phase of the meta-atom can be switched between two distinct states. Within our operational bandwidth centered around 5.2GHz (the shaded region), this switching induces a stable phase shift of nearly 180°. Furthermore, the reflection amplitude for both states remains high across this band, with a low reflection loss of less than 1.0 dB. The meta-atom's electrical design is inspired by a communication-oriented metasurface [71], adapted for our sensing application. As shown in Figure 12(c), considering the power requirements of all PIN diodes for the four metasurfaces, each metasurface is equipped with a driver bridging it with an FPGA control unit. To overcome input/output port limitations of the FPGA (e.g., ALINX AXU2CGB [23]), shift registers (e.g., SN74HC595 [25]) are integrated into the driver. These registers efficiently convert the FPGA's 1-bit-width serial output into an 8-bit parallel signal, allowing for synchronized control of metasurfaces. The power required by our metasurface system (i.e., diodes and FPGA) is about 9 W, which is less than half the power consumption of a router (e.g., ASUS RT-BE88U router [5]).

4.1.2 System Implementation. Mirror-Fi consists of a Wi-Fi AP, four metasurfaces controlled by an FPGA, and a receiver. The AP is an ASUS RT-BE88U router [5], while the receiver is an Acer TravelMate laptop equipped with an Intel AX210 Wi-Fi NIC [12], used for CSI collection via PicoScenes [27]. To maximize the Rx signal power, we equip each metasurface with a unique feed. Consequently, the signals received by a metasurface from non-front-facing feeds are significantly weakened and can thus be neglected [47]. As shown in Figure 13, we connect one antenna port of the Wi-Fi AP to a 1-to-4 power splitter so that each metasurface gets a feed (an antenna) placed in front of it, without affecting normal communications. The software components of Mirror-Fi are divided into three parts for implementation First, the low-level, real-time control logic is programmed on the FPGA using Verilog. This includes the hyper-metasurface configuration strategy (Section 3.1) and the phase pre-coding scheme (Section 3.2). The pre-coding sequence contains 8 bits spanning a 32ms period. Second, the high-level signal processing algorithms, including radio direction estimation (Section 3.1) and time synchronization (Section 3.3), are implemented on a laptop using MATLAB. Finally, the mDNN model (Section 3.3) is implemented in Python with PyTorch and trained on an NVIDIA RTX A5000 GPU.

4.2 Experiment Setup

We begin with micro-benchmark studies involving 4 persons, followed by three case studies to evaluate Mirror-Fi under realistic sensing applications. The setups for the case studies share three key aspects: i) each person remains within the coverage of the metasurface without mutual occlusion, ii) all persons perform designated activities simultaneously to assess Mirror-Fi's capability in multi-target sensing, iii) all experiments take place in a typical classroom, and iv) four participants are invited to simultaneously and randomly perform one of the three sensing applications described below. Note that, for clarity, the results in Section 5.2 are presented separately for each application. Figure 13 illustrates the experiment setup for benchmark and case studies, showing the metasurfaces, persons, and Wi-Fi devices. For each trial, the persons are arranged to cover a random combination of azimuths within the range of

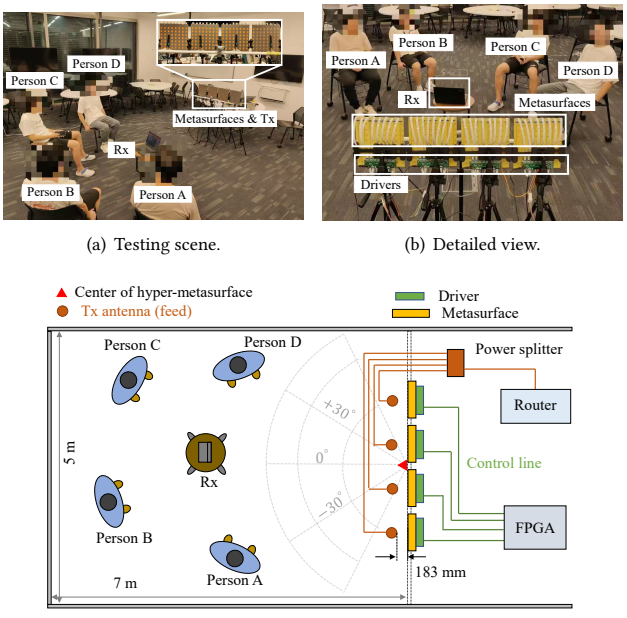

(a) Testing scene. (b) Detailed view.

(c) Layout.

Figure 13: (a) Experiment scene, (b) detailed view, and (c) layout of target arrangement for all three case studies.

±60°. Comparison baselines are obtained under the same configuration but without leveraging metasurfaces. Though RIScan [33] demonstrates the capability for multi-person indoor localization, its functionality is confined to this specific task. It cannot be directly extended to other fine-grained, multi-person sensing applications, as localization is an inherently coarse-grained task that provides only a target's spatial coordinates, which is fundamentally different from fine-grained sensing that requires analyzing subtle and dynamic features. Instead, our Mirror-Fi is designed to overcome the signal separation challenge for fine-grained sensing, so these two systems are not comparable in terms of sensing functionality.

Azimuth Estimation. The performance evaluation consists of 40 trials, each having different number of persons standing or sitting at various predefined positions, as illustrated in Figure 13(c). Each trial has 60 cycles, upon which average errors are derived. Each cycle involves the hyper-metasurface conducting full-angle scanning over the azimuthal range in 5° increments; each direction is scanned for 40 ms, resulting in a complete scanning cycle of one second.

Respiration Monitoring. We invite four persons to breathe simultaneously while using NeuLog chest belts [49] to collect ground truth. The total respiration recording duration is 80 minutes. After separating the metasurface-reflected CSIs using mDNN, we apply the method proposed by MoRe-Fi [76] to reconstruct their respiratory waveforms.

Gesture Detection. We evaluate six gestures, including push-pull (PP), up-down (UD), sweep (SW), draw a circle (DC), draw a zig-zag (DZ), and clap (CL). Each gesture is repeated 70 times per person in each trial, with a total of 20 trials. We leverage mDNN to separate the CSI for each person and employ the classifier proposed by Widar3.0 [77] for gesture recognition.

Activity Recognition. We also evaluate six activities, including bending (BD), jumping (JM), rotating (RT), sitting down (SD), standing up (SU), and walking (WL). Each activity is repeated 70 times per person in each trial, with a total of 20 trials. We employ mDNN to separate the CSI corresponding to each person and utilize RF-Net [13] to classify these activities.

Our evaluations aim to demonstrate Mirror-Fi's capability for multi-person sensing without requiring firm/hardware modifications to commodity Wi-Fi devices, rather than competing with existing Wi-Fi sensing proposals in specific tasks. All experiments strictly adhere to the IRB guidelines of our institute.

5 EVALUATIONS

In this section, we begin with two micro-benchmark studies, verifying the effectiveness of multi-target azimuth estimation and CSI separation. We then study three cases specified in Section 4.2, followed by discussions on several key design factors.

5.1 Micro-benchmark Studies

5.1.1 Performance of Azimuth Estimation. Although RIScan [33] has the capability of multi-person indoor localization, it does not provide openly accessible hardware for others to test. Therefore, we cannot directly compare Mirror-Fi with it even in terms of localization. Figure 14 shows that when the azimuth separation between targets exceeds 16°, the multi-target detection success rate

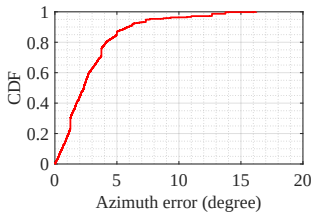

Figure 14: Estimation error.

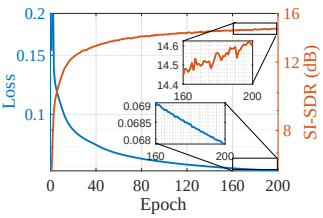

Figure 15: Loss & SI-SDR.

reaches 100%. Even with an azimuth separation error of 5°, Mirror-Fi achieves an azimuth estimation accuracy of 86.9%. These results evidently validate the reliability and effectiveness of Mirror-Fi's multi-target azimuth estimation capability enabled by its hyper-metasurface configuration.

5.1.2 Performance of Sensing Link Separation. To evaluate the effectiveness of mDNN in separating sensing links, we utilize the CSI time series dataset generated in Section 3.3.3. We assess mDNN using a hold-out validation method with a fixed, independent test set, an approach chosen for its statistical reliability and computational efficiency on our large dataset. The analysis is based on two key metrics: separation loss between the recovered and ground truth signals, as defined in Eqn. (10), and the SI-SDR metric mentioned there; it assesses the quality of signal separation on the test dataset by calculating the distortion between the reconstructed signal and the ground truth. Figure 15 demonstrates the changes in separation loss, which rapidly decreases during the initial training stages and eventually stabilizes at 0.068. Meanwhile, the SI-SDR of the reconstructed signals shows a steady increase, ultimately reaching up to 14.6 dB. The observed low loss and high SI-SDR value indicate that mDNN effectively separates signals reflected from different metasurfaces with minimal distortion.

To further investigate the performance of the separation model, Figure 16 provides examples of the separated CSI reflected by metasurfaces, including both magnitude and phase. i) *Amplitude*: The separated CSI magnitude aligns with the ground truth, accurately capturing the primary peaks and valleys. A notable observation is that the mDNN's output appears smoother than the ground truth, as it filters out some high-frequency transients. This behavior is expected due to the loss function (Eqn. (10)), which includes a smoothness term (λ_2) that penalizes rapid signal changes. Since the subsequent sensing stage inherently requires low-pass filtering, this smoothing effect does not compromise system performance. Moreover, although mDNN exhibits minor imperfections in recovering finer details, particularly at lower amplitude variations, these do not affect the primary sensing tasks. This slight deviation may result from the network's focus on minimizing overall loss rather than capturing subtle fluctuations. ii) *Phase*: While minor and nearly constant phase shifts can be observed in certain segments (e.g., Figures 16(a) and 16(b)), these are likely caused by slight residual interference that the network is unable to fully suppress. Importantly, in phase-based sensing applications such as respiration monitoring, it is the temporal variation in phase, rather than its absolute value, that carries meaningful information. As mDNN faithfully preserves the phase variation trend over time, these small static offsets have a negligible impact on the performance of downstream sensing tasks. These results evidently confirm mDNN's capability to effectively separate sensing links, ensuring high fidelity in both magnitude

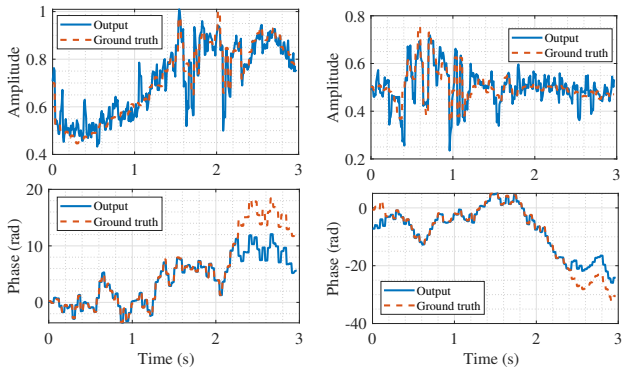

(a) The 1-st metasurface. (b) The 2-nd metasurface.

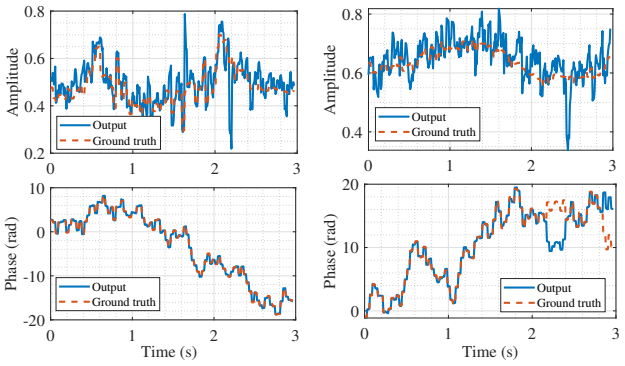

(c) The 3-rd metasurface. (d) The 4-th metasurface.

Figure 16: CSIs of each metasurface obtained by mDNN.

and phase reconstruction, and thus to achieve RPD. The inevitable "glitches" in both magnitude and phase virtually have no impact on the downstream sensing tasks.

5.1.3 Impact on Communication Performance. A critical requirement for any practical Wi-Fi sensing system is to avoid degrading the performance of the primary communication. This challenge involves two distinct aspects: i) ensuring the added sensing hardware (our metasurfaces) does not interfere with the link, and ii) managing the AP's resources for concurrent sensing and communication, which can be non-trivial (e.g., as explored in SenCom [17]). Our evaluation focuses on the first aspect, confirming the Mirror-Fi metasurface system has a negligible impact. We note that the second aspect, which concerns the AP's internal mechanisms, is orthogonal to our design. Mirror-Fi can be seamlessly integrated with SenCom-based APs to support concurrent sensing and communication if such compatibility is needed. This is achieved through two design principles: i) the metasurfaces operate entirely passively, reflecting signals from the existing Wi-Fi AP without emitting any radio frequency energy; and ii) each metasurface is controlled via an FPGA, ensuring control signaling remains fully isolated from data communication. We empirically validate the effectiveness of this non-interfering design by measuring throughput and latency. The experiment is conducted at five different user positions, comparing Mirror-Fi against a baseline setup operating without Mirror-Fi. We let the Wi-Fi AP transmit communication packets at a high frequency and reuse these packets for sensing. This enables passive CSI collection and allows throughput evaluation under full-load conditions. The results, presented in Figure 17, demonstrate that Mirror-Fi introduces negligible overhead to the communication link. As shown in Figure 17(a), Mirror-Fi consistently achieves throughput on par with the baseline at most positions. In terms of communication latency, Mirror-Fi shows no significant difference from the baseline. While its latency may be marginally higher at certain points (e.g., about 0.1 ms at position C), this minimal variation does not affect the communication function. Together, these findings confirm that Mirror-Fi can be deployed without adversely affecting primary communication channels while significantly enhancing multi-target awareness capabilities.

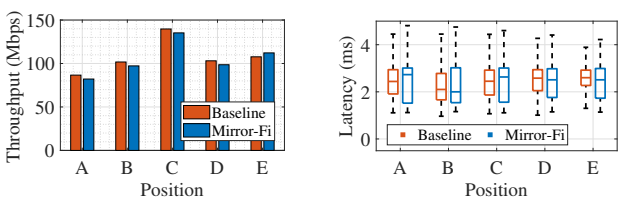

(a) Throughput. (b) Latency.

Figure 17: A comparison of communication performance between Mirror-Fi and a baseline, evaluating (a) throughput and (b) latency at five different positions.

5.2 Case Studies and Design Factors

5.2.1 Respiration Monitoring. We conduct experiments to monitor multi-person respiration using the setup described in Section 4.2. The objective is to evaluate the ability of Mirror-Fi to accurately extract respiration rates from multiple persons simultaneously, leveraging the separated CSI results for each metasurface. We first use Figure 18 to compare the respiration waveforms of a person obtained respectively by Mirror-Fi and the baseline method. Mirror-Fi effectively reconstructs respiration waveforms with high fidelity, evidently capturing inhalation and exhalation patterns. In contrast, the baseline method obtains only noisy mixed signals, hindering individual respiration detection. Furthermore, Figure 19(a) shows the accuracy of respiration rate estimation for both Mirror-Fi and the baseline. The results indicate that Mirror-Fi achieves multi-person respiration monitoring, with a mean respiration rate error below 1.2 bpm, as opposed to 6.59 bpm of the baseline, indicating the latter's ineffectiveness in multi-person scenarios.

To gain deeper insights into Mirror-Fi's capability for multi-person respiration sensing, we conduct a breath-holding experiment in which persons sequentially hold their breath for 15 seconds

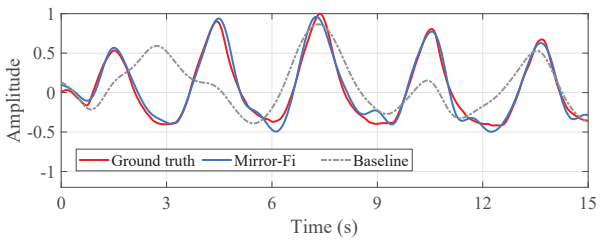

Figure 18: Comparison between Mirror-Fi and the baseline regarding respiration waveforms.

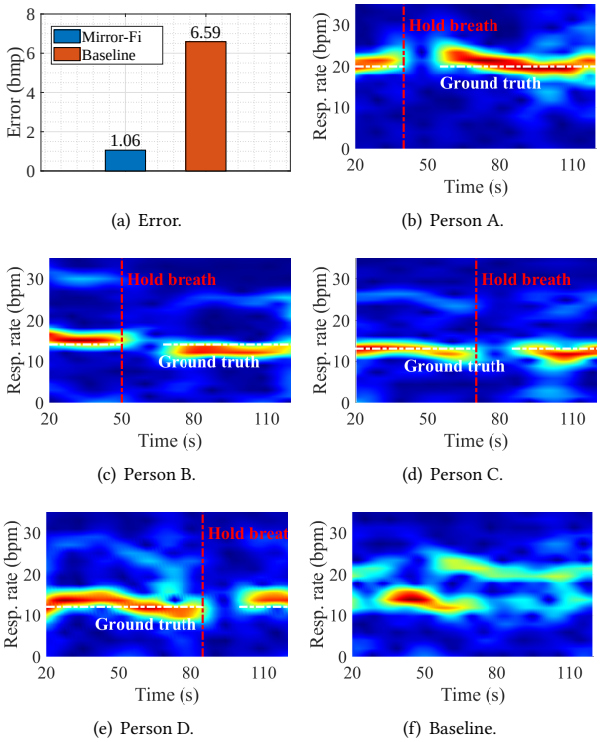

(a) Error. (b) Person A.

(c) Person B. (d) Person C.

(e) Person D. (f) Baseline.

Figure 19: Comparison and analysis on Mirror-Fi and the baseline in terms of the respiration.

each. The spectrograms of Mirror-Fi and the baseline are presented in Figure 19. The alignment between the breath-holding periods and signal interruptions (start time indicated by red line markers) on the spectrograms demonstrates that the respiration signals from different persons are well separated. As shown in Figures 19(b) to 19(e), the signal energy in the corresponding frequency bands decreases significantly during each person's breath-holding period. This indicates that Mirror-Fi can accurately separate and monitor breathing stops for different persons. In contrast, Figure 19(f) indicates that the spectrogram produced by the baseline is dominated by overlapping and noisy components, preventing the effective extraction of respiration rates. These results provide compelling evidence of the effectiveness of Mirror-Fi in multi-person respiration monitoring. By separating and reconstructing individual respiration signals, Mirror-Fi achieves high accuracy in respiration rate estimation while demonstrating robustness to interference and noise, making it a promising solution for device-free Wi-Fi sensing.

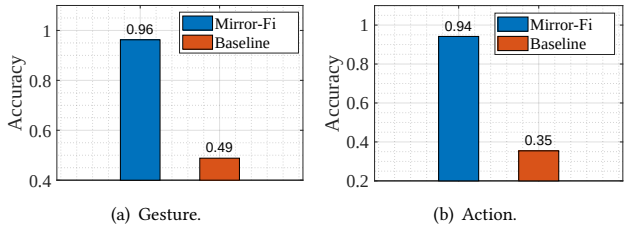

(a) Gesture. (b) Action.

Figure 20: Comparing gesture and action recognition accuracy between Mirror-Fi and the baseline.

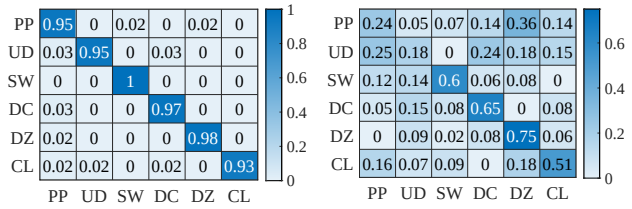

(a) Mirror-Fi. (b) Baseline.

Figure 21: Comparison between Mirror-Fi and the baseline in gesture confusion matrices.

5.2.2 Gesture Detection. We also conduct experiments to evaluate the performance of Mirror-Fi in detecting gestures, and summarize the results in Figure 20(a). The results indicate that Mirror-Fi achieves a mean accuracy exceeding 96%, significantly outperforming the baseline that obtains only 49%. This substantial accuracy gap highlights the effectiveness of Mirror-Fi in multi-person gesture detection scenarios. To gain deeper insights into this performance gap, we examine the confusion matrices of Mirror-Fi and baseline, as shown in Figure 21. The confusion matrices reveal that Mirror-Fi can correctly classify most gestures, while the baseline often confuses one gesture with others. Specifically, due to overlapping gestures from multiple persons, the PP and UD gestures perform the worst for the baseline with accuracy below 25%, as opposed to the 95% accuracy readily achieved by Mirror-Fi. The inferior performance of the baseline can be attributed to its inability to separate signals from interference caused by the simultaneous movement of the hands of multiple persons, resulting in an overall detection accuracy 47% lower than that of Mirror-Fi.

5.2.3 Activity Recognition. Similar experiment validations also extend to activity recognition, with comparative results depicted in Figure 20(b). Mirror-Fi achieves 94% average classification accuracy, indicating 2.69× improvement over 35% of the baseline. Confusion matrix analysis in Figure 22 reveals that Mirror-Fi achieves an accuracy exceeding 84% for all six activities, whereas the baseline remains below 50% in most cases, with the sole exception of SU reaching 71%. For instance, BD is frequently misclassified as JM at a rate of 32% or RT at 37%, while the recognition rate for RT itself is merely 13%. This poor performance stems from signal entanglement; without the RPD provided by Mirror-Fi, the baseline receives a single, convoluted signal where the concurrent activities of four individuals interfere, making them fundamentally difficult to separate. These results, along with previous case studies, collectively demonstrate the feasibility of multi-person Wi-Fi sensing

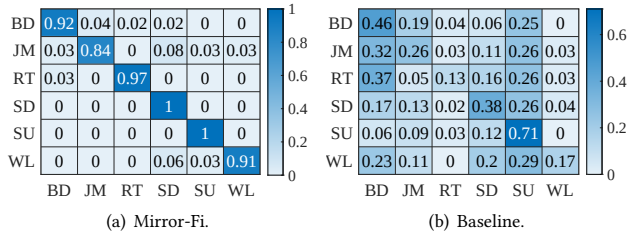

(a) Mirror-Fi. (b) Baseline.

Figure 22: Comparison between Mirror-Fi and the baseline in action confusion matrices.

in real-world scenarios, thanks to the significantly enhanced RPD offered by Mirror-Fi.

5.2.4 Deployment Flexibility and Practicality. Our work establishes a proof-of-concept for a new sensing paradigm, with future efforts directed toward cost reduction and miniaturization. Therefore, demonstrating deployment flexibility is key to addressing these practical concerns and showcasing a clear path toward real-world integration. To validate this, we move beyond the basic linear setup and evaluate a more practical configuration. As illustrated in Figure 23(a), four metasurfaces are arranged in an "L-shaped" corner configuration within a 5-meter-wide room. Four participants are distributed throughout the area and instructed to perform gestures simultaneously. We focus on gesture detection tasks due to their relevance to human-computer interaction (HCI) in extended reality (XR) applications [3, 30, 51]. Figure 23(b) shows that despite changes in metasurface deployment, multi-person gestures can still be effectively classified with an average accuracy of 97.4%, which closely aligns with the results in the linear deployment scenario shown in Figure 21(a). In summary, this flexibility confirms that the core principle of Mirror-Fi, enabling RPD for sensing, is maintained across different configurations.

5.2.5 Environmental Generalizability. To be viable in the real world, the system must generalize by overcoming limitations such as static occlusions and dynamic target mobility. In practice, occlusions between targets may prevent a metasurface from beamforming to a single target. As shown in Figure 23(a), the beamforming constructed by the metasurface 'MTS 3' covers two targets, with Person B occluded by Person A. To address this issue, we can effectively alleviate target occlusion by strategically re-arranging metasurfaces' locations and beamforming directions, thanks to the flexibility offered by Mirror-Fi. In addition, complex indoor environments, particularly those involving line-of-sight (LoS) occlusions, can substantially impair the accuracy of multi-target azimuth estimation and the robustness of subsequent perception links. A practical and effective mitigation strategy is a top-mounted deployment, in which the metasurface is installed on the ceiling to reduce occlusions from furniture and other indoor obstacles. However, due to installation constraints in our current laboratory setting, this work focuses on evaluating a horizontal deployment scheme under LoS conditions. While our current work focuses on quasi-static scenarios, the system can support mobile users with a "track-then-sense" approach: a quick initial scan finds targets, followed by rapid, targeted scans to track their movement while sensing is performed. The key challenge is distinguishing small actions (e.g., gestures) from the

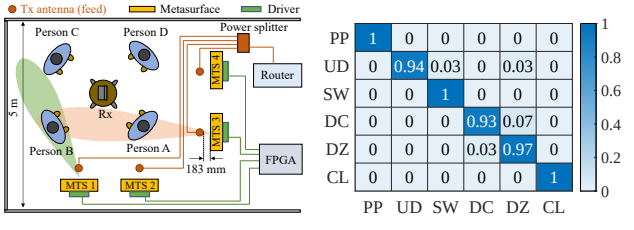

(a) Layout. (b) Gesture recognition.

Figure 23: Sensing accuracy of Mirror-Fi under nonlinear deployment conditions.

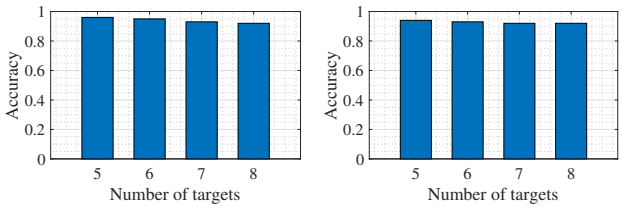

(a) Gesture. (b) Action.

Figure 24: Gesture and action recognition accuracy of Mirror-Fi when the number of targets exceeds the metasurfaces.

large-scale signal interference of whole-body movement (e.g., walking). Developing the advanced algorithms required to resolve this is a substantial research effort in itself, and it is beyond our scope.

5.2.6 Scalability for a Higher Number of Targets. The quantity matching is not an issue if target number is less, so we are only concerned with where the number of targets exceeds that of metasurfaces. To this end, we test an augmented metasurface control mechanism that applies multiple groups of pre-coding sequences to a metasurface. We combine code-division with time-division, assigning different pre-coding sequences to individuals at different time slices, ensuring that each metasurface constructs multiple independent sensing links. To validate this, we conduct experiments using four metasurfaces to perceive more than four individuals. By assigning unique time-division pre-coding sequences, the system successfully separated the signals and achieved robust performance in multi-person gesture and action recognition. As shown in Figure 24, Mirror-Fi maintains exceptionally high and stable recognition accuracy for both gestures (~93%) and actions (~92%), even as the number of people doubles the amount of metasurface. This result confirms that Mirror-Fi is not limited to one-to-one metasurface correspondence and has the potential to enable scalable multi-person sensing.

5.2.7 Potential Applications. With physical separability ensured by metasurface-created sensing links, Mirror-Fi is scalable to ubiquitous Wi-Fi infrastructures, enjoying an ever expanded scope of sensing applications. As demonstrated in Section 5.2, Mirror-Fi can simultaneously capture fine-gained body motions of multiple targets, including gestures and even breathing patterns. This capability extends to smart homes with the need for, for example, remote control and vital sign monitoring. Meanwhile, its gesture and activity recognition capabilities can enhance the convenience of XR interactions [51, 64] while potentially reducing the cost and complexity of HCI development. Of course, co-existing with other co-channel communication systems [35, 36, 69], as well as the incurred privacy infringement [20, 21] should be handled in the future.

6 RELATED WORK

In this section, we review existing works on Wi-Fi sensing and metasurfaces.

Wi-Fi Sensing. The ability of Wi-Fi sensing to enable localization [31, 52, 68, 74] and human behaviors [10, 13, 17, 26, 61, 72, 76, 77] using wireless signals (i.e., CSI [16]) has gained increasing recognition for its advantage of leveraging commodity devices. Due to

page limit, we focus only on exploiting (physical) *diversity* for multi-person sensing (MPS) in the following. Widar2.0 [52] uses multiple antennas to enhance spatial diversity, providing potential for MPS. Subsequently, Karanam et al. [29] utilize magnitude measurements from three dual-antenna Rx to actually achieve MPS. Later proposals enhance antenna-induced spatial diversity via specifically crafted algorithms, typically including mD-Track [66] to leverage spatial search and expectation-maximization for separating multi-path signals, and MultiSense [73] to employ blind source separation for extracting individual signals. In fact, such approaches are still confined by diversity regardless of what ingenuity has been put into algorithm design: for example, MultiSense can recognize at most 4 persons with 5 antennas. Consequently, latest proposals directly target diversity enhancement, which includes, among others, MUSE-Fi [22] exploiting near-field dominance in link diversity while UWB-Fi [41] and μCeiver-Fi [40] significantly expanding frequency diversity via compressive channel sampling and multi-link receivers. Quite unfortunately, these proposals either require additional user devices (e.g., smartphones) for assistance or cannot be readily deployed due to the need for modifying Wi-Fi firm/hardware. In contrast, Mirror-Fi offers a device-free solution without modifying commodity Wi-Fi devices; it exploits link diversity (i.e., RPD) created via multiple external metasurfaces.

Metasurfaces. Extensive research has explored metasurfaces, classified into passive [45, 50, 53, 62] and active (programmable) [11, 28, 38, 39, 48, 75] types, to modulate electromagnetic waves for enhancing wireless environments. These approaches include RIScan [33] that applies metasurface-based beamforming for multi-person indoor localization as the only sensing application; the remaining contributions center on communication performance. RoS [50] pioneers in improving roadside-to-vehicle communication by embedding digital information into reflected signals using passive metasurfaces. Building on this, AutoMS [45] introduces an automated framework to optimize mmWave coverage by strategically placing metasurfaces in a more convenient manner. In contrast to the static nature of passive solutions, active metasurfaces enable real-time reconfiguration, offering adaptability in dynamic environments. RFlens [15], for example, explores the beamforming capabilities of a reconfigurable metasurface, lifting signal strength for both communication and sensing. Extending this concept, RF-Bouncer [39] introduces a dual-band programmable metasurface enabling simultaneous beamforming on two sub-6 ISM bands, which is later utilized to enhance signal coverage [42] or to attack wireless sensing systems [28]. Finally, a recent work STCM [8] claims to perform communication beamforming via coding, yet it appears to be very similar to that proposed by RIScan and is rather irrelevant to tackling the issue of signal separation required in our context. In summary, none of the existing methods effectively tackles the key challenges inherent in general-purpose, multi-target sensing. Mirror-Fi is, on the contrary, the first work to leverage metasurfaces for achieving RPD and thus supporting general multi-target sensing.

7 CONCLUSION

Taking an essential step towards ubiquitous human sensing, Mirror-Fi has pioneered in multi-target sensing by addressing the major challenge of physically separating multiple targets without modifying Wi-Fi firm/hardware. By leveraging metasurfaces to create a unique sensing link for each individual, Mirror-Fi has successfully achieved reflected path diversity (RPD) and thus demonstrated multi-person sensing for respiration monitoring, gesture detection, and activity recognition. Two key technical advancements have contributed to this success: i) a pre-coding scheme to enhance the distinguishability of metasurface-reflected signals, and ii) a model-based neural network employed to separate sensing links effectively. Our extensive evaluations have clearly demonstrated that Mirror-Fi can be a more readily deployable alternative to existing device-free Wi-Fi sensing systems. Looking ahead, we believe Mirror-Fi holds significant potential for diverse applications, including healthcare, smart homes, and security. Additionally, we plan to scale Mirror-Fi for deployment in larger environments to evaluate its performance in more diversified scenarios.

ACKNOWLEDGEMENT

This work is supported in part by the Natural Science Foundation of Jiangsu Province (Key Program) under Grant No. BK20243040, National Natural Science Foundation of China under Grant No. 92467202, National Natural Science Foundation of China under Grant No. 62272216, Collaborative Innovation Center of Novel Software Technology and Industrialization, and the National Research Foundation Singapore and DSO National Laboratories under the AI Singapore Programme (AISG Award No: AISG2-GC-2023-006).

REFERENCES

[1] Fadel Adib, Zachary Kabelac, and Dina Katabi. 2015. Multi-Person Localization via RF Body Reflections. In *Proc. of the 12th USENIX NSDI*. 279–292.
[2] Kyle Aitken, Vinay Ramasesh, Yuan Cao, and Niru Maheswaranathan. 2021. Understanding How Encoder-Decoder Architectures Attend. In *Proc. of the 35th NeurIPS*. 22184–22195.
[3] A'aeshah Alhakamy. 2024. Extended Reality (XR) Toward Building Immersive Solutions: The Key to Unlocking Industry 4.0. *Comput. Surveys* 56, 9 (2024), 1–38.
[4] Shun-ichi Amari, Andrzej Cichocki, and Howard Yang. 1995. A New Learning Algorithm for Blind Signal Separation. *Advances in neural information processing systems* 8 (1995).
[5] ASUSTeK Computer Inc. 2025. User Guide: RT-BE88U BE7200 Dual Band Wi-Fi Router. https://dlcdnets.asus.com/pub/ASUS/wireless/RT-BE88U/E23805_RT-BE88U_UM_WEB.pdf?model=RT-BE88U/. Online; accessed 1 March 2025.
[6] James R Bitner, Gideon Ehrlich, and Edward M Reingold. 1976. Efficient Generation of the Binary Reflected Gray Code and Its Applications. *Commun. ACM* 19, 9 (1976), 517–521.
[7] J-F Cardoso. 1998. Blind Signal Separation: Statistical Principles. *Proc. IEEE* 86, 10 (1998), 2009–2025.
[8] Xiao Qing Chen, Lei Zhang, Yi Ning Zheng, Shuo Liu, Zhuo Ran Huang, Jing Cheng Liang, Marco Di Renzo, Vincenzo Galdi, and Tie Jun Cui. 2025. Integrated Sensing and Communication based on Space-Time-Coding Metasurfaces. *Nature Communications* 16, 1 (2025), 1836.
[9] Zhe Chen, Guorong Zhu, Sulei Wang, Yuedong Xu, Jie Xiong, Jin Zhao, Jun Luo, and Xin Wang. 2021. M^3: Multipath Assisted Wi-Fi Localization with a Single Access Point. *IEEE Trans. on Mobile Computing* 20, 2 (2021), 588–602.
[10] Linsong Cheng and Jiliang Wang. 2019. Walls Have No Ears: A Non-Intrusive WiFi-Based User Identification System for Mobile Devices. *IEEE/ACM Transactions on Networking* 27, 1 (2019), 245–257.
[11] Kun Woo Cho, Mohammad H Mazaheri, Jeremy Gummeson, Omid Abari, and Kyle Jamieson. 2023. mmWall: A Steerable, Transflective Metamaterial Surface for NextG mmWave Networks. In *Proc. of the 20th USENIX NSDI*. 1647–1665.
[12] Intel Corporation. 2025. Intel® Wi-Fi 6E AX210. https://www.intel.com/content/www/us/en/products/sku/204836/intel-wifi-6e-ax210-gig/specifications.html. Online; accessed 1 March 2025.
[13] Shuya Ding, Zhe Chen, Tianyue Zheng, and Jun Luo. 2020. RF-Net: A Unified Meta-Learning Framework for RF-enabled One-Shot Human Activity Recognition. In *Proc. of the 18th ACM SenSys*. 517–530.
[14] Xiaoran Fan, Longfei Shangguan, Richard Howard, Yanyong Zhang, Yao Peng, Jie Xiong, Yunfei Ma, and Xiang-Yang Li. 2020. Towards Flexible Wireless Charging

for Medical Implants Using Distributed Antenna System. In *Proc. of the 26th ACM MobiCom*. 1–15.
[15] Chao Feng, Xinyi Li, Yangfan Zhang, Xiaojing Wang, Liqiong Chang, Fuwei Wang, Xinyu Zhang, and Xiaojiang Chen. 2021. RFlens: Metasurface-enabled Beamforming for IoT Communication and Sensing. In *Proc. of the 27th MobiCom*. 587–600.
[16] Daniel Halperin, Wenjun Hu, Anmol Sheth, and David Wetherall. 2011. Tool Release: Gathering 802.11n Traces with Channel State Information. *ACM SIGCOMM Comput. Commun. Rev.* 41, 1 (2011), 53.
[17] Yinghui He, Jianwei Liu, Mo Li, Guanding Yu, Jinsong Han, and Kui Ren. 2023. SenCom: Integrated Sensing and Communication with Practical WiFi. In *Proc. of the 29th ACM MobiCom*. 1–16.
[18] Yinghui He, Mingming Xu, Zhe Chen, Fu Xiao, and Jun Luo. 2025. Beam-Fi: Integrated Sensing and Communication via MU-MIMO upon Commodity Wi-Fi. *Proc. ACM IMWUT* 9, 3 (2025), 1–22.
[19] Yinghui He, Guanding Yu, Yunlong Cai, and Haiyan Luo. 2024. Integrated Sensing, Computation, and Communication: System Framework and Performance Optimization. *IEEE Trans. on Wireless Communications* 23, 2 (2024), 1114–1128.
[20] Jingzhi Hu, Xin Li, Jin Gan, and Jun Luo. 2025. Poison to Cure: Privacy-preserving Wi-Fi Multi-User Sensing via Data Poisoning. In *Proc. of the 31st ACM MobiCom*. 1–16.
[21] Jingyang Hu, Hongbo Wang, Tianyue Zheng, Jingzhi Hu, Zhe Chen, Hongbo Jiang, and Jun Luo. 2023. Password-Stealing Without Hacking: Wi-Fi Enabled Practical Keystroke Eavesdropping. In *Proc. of the 30th ACM CCS*. 239–252.
[22] Jingzhi Hu, Tianyue Zheng, Zhe Chen, Hongbo Wang, and Jun Luo. 2023. MUSE-Fi: Contactless MUti-person SEnsing Exploiting Near-field Wi-Fi Channel Variation. In *Proc. of the 29th ACM MobiCom*. 23–31.
[23] Advanced Micro Devices Inc. 2025. ALINX AXU2CGB. https://www.xilinx.com/products/boards-and-kits/1-1eu3go3.html. Online; accessed 1 February 2025.
[24] MACOM Technology Solutions Inc. 2025. MADP-000907-14020P. https://www.macom.com/products/product-detail/MADP-000907-14020P. Online; accessed 1 February 2025.
[25] Texas Instruments Inc. 2025. SN74HC595. https://www.ti.com/product/SN74HC595. Online; accessed 1 February 2025.
[26] Wenjun Jiang, Hongfei Xue, Chenglin Miao, Shiyang Wang, Sen Lin, Chong Tian, Srinivasan Murali, Haochen Hu, Zhi Sun, and Lu Su. 2020. Towards 3D Human Pose Construction Using WiFi. In *Proc. of the 26th ACM MobiCom*. 1–14.
[27] Zhiping Jiang, Tom H. Luan, Xincheng Ren, Dongtao Lv, Han Hao, Jing Wang, Kun Zhao, Wei Xi, Yueshen Xu, and Rui Li. 2021. Eliminating the Barriers: Demystifying Wi-Fi Baseband Design and Introducing the PicoScenes Wi-Fi Sensing Platform. *IEEE Internet of Things Journal* (2021), 1–21.
[28] Jiang, Chenghan and Yang, Jinjiang and Li, Xinyi and Li, Qi and Zhang, Xinyu and Ren, Ju. 2024. RISiren: Wireless Sensing System Attacks via Metasurface. In *Proc. of the 31st ACM CCS*. 3332–3345.
[29] Chitra R Karanam, Belal Korany, and Yasamin Mostofi. 2019. Tracking from One side: Multi-Person Passive Tracking with WiFi Magnitude Measurements. In *Proc. of the 18th ACM IPSN*. 181–192.
[30] Z. Jonny Kong, Qiang Xu, Jiayi Meng, and Y. Charlie Hu. 2023. AccuMO: Accuracy-Centric Multitask Offloading in Edge-Assisted Mobile Augmented Reality. In *Proc. of the 29th ACM MobiCom*. 1–16.
[31] Manikanta Kotaru, Kiran Joshi, Dinesh Bharadia, and Sachin Katti. 2015. SpotFi: Decimeter Level Localization Using WiFi. *ACM SIGCOMM Comput. Commun. Rev.* 45, 4 (2015), 269–282.
[32] Jonathan Le Roux, Scott Wisdom, Hakan Erdogan, and John R Hershey. 2019. SDR – Half-baked or Well Done?. In *Proc. of the 35th IEEE ICASSP*. 626–630.
[33] Chenggao Li, Qianyi Huang, Yuxuan Zhou, Yandao Huang, Qingyong Hu, Huangxun Chen, and Qian Zhang. 2023. RIScan: RIS-aided Multi-user Indoor Localization Using COTS Wi-Fi. In *Proc. of the 21st ACM SenSys*. 445–458.
[34] Changming Li, Mingjing Xu, Yicong Du, Limin Liu, Cong Shi, Yan Wang, Hongbo Liu, and Yingying Chen. 2024. Practical Adversarial Attack on WiFi Sensing Through Unnoticeable Communication Packet Perturbation. In *Proc. of the 30th ACM MobiCom*. 373–387.
[35] Feng Li, Jun Luo, Gaotao Shi, and Ying He. 2013. FAVOR: Frequency Allocation for Versatile Occupancy of Spectrum in Wireless Sensor Networks. In *Proc. of the 14th ACM MobiHoc*. 39–48.
[36] Feng Li, Jun Luo, Gaotao Shi, and Ying He. 2017. ART: Adaptive Frequency-Temporal Co-Existing of ZigBee and WiFi. *IEEE Trans. on Mobile Computing* 16, 3 (2017), 662–674.
[37] Tianxiang Li, Mohammad Hossein Mazaheri, and Omid Abari. 2024. Enabling On-Demand Low-Power mmWave Repeaters via Passive Beamforming. In *Proc. of the 30th ACM MobiCom*. 618–632.
[38] Xinyi Li, Chao Feng, Fengyi Song, Chenghan Jiang, Yangfan Zhang, Ke Li, Xinyu Zhang, and Xiaojiang Chen. 2022. Protego: Securing Wireless Communication via Programmable Metasurface. In *Proc. of the 28th Mobicom*. 55–68.
[39] Xinyi Li, Chao Feng, Xiaojing Wang, Yangfan Zhang, Yaxiong Xie, and Xiaojiang Chen. 2023. RF-Bouncer: A Programmable Dual-band Metasurface for Sub-6 Wireless Networks. In *Proc. of the 20th USENIX NSDI*. 389–404.
[40] Xin Li, Yinghui He, and Jun Luo. 2025. μCeiver-Fi: Exploiting Spectrum Resources of Multi-link Receiver for Fine-Granularity Wi-Fi Sensing. In *Proc. of the 31st ACM MobiCom*. 1–15.
[41] Xin Li, Hongbo Wang, Zhe Chen, Zhiping Jiang, and Jun Luo. 2024. UWB-Fi: Pushing Wi-Fi towards Ultra-wideband for Fine-Granularity Sensing. In *Proc. of the 22nd ACM MobiSys*. 42–55.
[42] Xinyi Li, Gaoteng Zhao, Ling Chen, Xinyu Zhang, and Ju Ren. 2024. RFMagus: Programming the Radio Environment With Networked Metasurfaces. In *Proc. of the 30th ACM MobiCom*. 16–30.
[43] Jian Liu, Yan Wang, Yingying Chen, Jie Yang, Xu Chen, and Jerry Cheng. 2015. Tracking Vital Signs During Sleep Leveraging Off-the-shelf WiFi. In *Proc. of the 16th ACM MobiHoc*. 267–276.
[44] Jun Luo, Hangcheng Cao, Hongbo Jiang, Yanbing Yang, and Zhe Chen. 2024. MIMOCrypt: Multi-user Privacy-Preserving Wi-Fi Sensing via MIMO Encryption. In *Proc. of the 45th IEEE S&P*. 2812–2830.
[45] Ruichun Ma, Shicheng Zheng, Hao Pan, Lili Qiu, Xingyu Chen, Liangyu Liu, Yihong Liu, Wenjun Hu, and Ju Ren. 2024. AutoMS: Automated Service for mmWave Coverage Optimization using Low-cost Metasurfaces. In *Proc. of the 30th ACM MobiCom*. 62–76.
[46] Si Yu Miao and Feng Han Lin. 2022. Light-Controlled Large-Scale Wirelessly Reconfigurable Microstrip Reflectarrays. *IEEE Trans. on Antennas and Propagation* 71, 2 (2022), 1613–1622.
[47] Gabriele Minatti, Enrica Martini, and Stefano Maci. 2017. Efficiency of Metasurface Antennas. *IEEE Trans. on Antennas and Propagation* 65, 4 (2017), 1532–1541.
[48] Kim Minseok, Ahn Namjo, and Min Kim Song. 2024. NR-Surface: NextG-ready µW-reconfigurable mmWave Metasurface. In *Proc. of the 21st USENIX NSDI*. 1641–1657.
[49] NeuLog. 2017. Respiration Monitor Belt Logger Sensor NUL-236. https://neulog.com/respiration-monitor-belt/. Online; accessed 1 February 2025.
[50] John Nolan, Kun Qian, and Xinyu Zhang. 2021. RoS: Passive Smart Surface for Roadside-to-Vehicle Communication. In *Proc. of the 35th ACM SIGCOMM*. 165–178.
[51] Yongtae Park, Sangki Yun, and Kyu-Han Kim. 2019. When IoT met Augmented Reality: Visualizing the Source of the Wireless Signal in AR View. In *Proc. of the 17th ACM MobiSys*. 117–129.
[52] Kun Qian, Chenshu Wu, Yi Zhang, Guidong Zhang, Zheng Yang, and Yunhao Liu. 2018. Widar2.0: Passive Human Tracking with a Single Wi-Fi Link. In *Proc. of the 16th ACM MobiSys*. 350–361.
[53] Kun Qian, Lulu Yao, Xinyu Zhang, and Tse Nga Ng. 2022. MilliMirror: 3D Printed Reflecting Surface for Millimeter-wave Coverage Expansion. In *Proc. of the 28th ACM MobiCom*. 15–28.
[54] Kun Qian, Lulu Yao, Kai Zheng, Xinyu Zhang, and Tse Nga Ng. 2023. UniScatter: a Metamaterial Backscatter Tag for Wideband Joint Communication and Radar Sensing. In *Proc. of the 29th ACM MobiCom*. 1–16.
[55] Ralph Schmidt. 1986. Multiple Emitter Location and Signal Parameter Estimation. *IEEE Trans. on Antennas and Propagation* 34, 3 (1986), 276–280.
[56] Ui-Hyeop Shin, Sangyoun Lee, Taehan Kim, and Hyung-Min Park. 2024. Separate and Reconstruct: Asymmetric Encoder-Decoder for Speech Separation. In *Proc. of the 38th NeurIPS*. 1–20.
[57] Chiheb Trabelsi, Olexa Bilaniuk, Ying Zhang, Dmitriy Serdyuk, Sandeep Subramanian, Joao Felipe Santos, Soroush Mehri, Negar Rostamzadeh, Yoshua Bengio, and Christopher J Pal. 2018. Deep Complex Networks. In *Proc. of the 6th ICLR*.
[58] Deepak Vasisht, Swarun Kumar, and Dina Katabi. 2016. Decimeter-Level Localization with a Single WiFi Access Point. In *Proc. of the 13th USENIX NSDI*. 165–178.
[59] Fei Wang, Yizhe Lv, Mengdie Zhu, Han Ding, and Jinsong Han. 2024. XRF55: A Radio Frequency Dataset for Human Indoor Action Analysis. *Proc. ACM IMWUT* 8, 1 (2024), 1–34.
[60] Hao Wang, Da-Gang Fang, and Y Leonard Chow. 2008. Grating Lobe Reduction in a Phased Array of Limited Scanning. *IEEE Trans. on Antennas and Propagation* 56, 6 (2008), 1581–1586.
[61] Wei Wang, Alex X. Liu, Muhammad Shahzad, Kang Ling, and Sanglu Lu. 2015. Understanding and Modeling of WiFi Signal Based Human Activity Recognition. In *Proc. of the 21st ACM MobiCom*. 65–76.
[62] Timothy Woodford, Kun Qian, and Xinyu Zhang. 2023. Metasight: High-Resolution NLoS Radar with Efficient Metasurface Encoding. In *Proc. of the 21st ACM SenSys*. 308–321.
[63] Dan Wu, Youwei Zeng, Ruiyang Gao, Shenjie Li, Yang Li, Rahul C. Shah, Hong Lu, and Daqing Zhang. 2023. WiTraj: Robust Indoor Motion Tracking With WiFi Signals. *IEEE Trans. on Mobile Computing* 22, 5 (2023), 3062–3078.
[64] Nan Wu, Kaiyan Liu, Ruizhi Cheng, Bo Han, and Puqi Zhou. 2024. Theia: Gaze-driven and Perception-aware Volumetric Content Delivery for Mixed Reality Headsets. In *Proc. of the 22nd ACM MobiSys*. 70–84.
[65] Yaxiong Xie, Zhenjiang Li, and Mo Li. 2015. Precise Power Delay Profiling with Commodity WiFi. In *Proc. of the 21st ACM MobiCom*. 53–64.
[66] Yaxiong Xie, Jie Xiong, Mo Li, and Kyle Jamieson. 2019. mD-Track: Leveraging Multi-Dimensionality for Passive Indoor Wi-Fi Tracking. In *Proc. of the 25th ACM MobiCom*. 1–16.

[67] Tong Xin, Bin Guo, Zhu Wang, Pei Wang, Jacqueline Chi Kei Lam, Victor Li, and Zhiwen Yu. 2018. FreeSense: A Robust Approach for Indoor Human Detection Using Wi-Fi Signals. *Proc. ACM IMWUT* 2, 3 (2018), 1–23.

[68] Jie Xiong, Karthikeyan Sundaresan, and Kyle Jamieson. 2015. ToneTrack: Leveraging Frequency-Agile Radios for Time-Based Indoor Wireless Localization. In *Proc. of the 21st ACM MobiCom*. 537–549.

[69] Ruitao Xu, Gaotao Shi, Jun Luo, Zenghua Zhao, and Yantai Shu. 2011. MuZi: Multi-Channel ZigBee Networks for Avoiding WiFi Interference. In *Proc. of the 4th IEEE/ACM CPSCOM*. 323–329.

[70] Huanhuan Yang, Fan Yang, Shenheng Xu, Yilin Mao, Maokun Li, Xiangyu Cao, and Jun Gao. 2016. A 1-Bit 10×10 Reconfigurable Reflectarray Antenna: Design, Optimization, and Experiment. *IEEE Trans. on Antennas and Propagation* 64, 6 (2016), 2246–2254.

[71] Huanhuan Yang, Fan Yang, Shenheng Xu, Yilin Mao, Maokun Li, Xiangyu Cao, and Jun Gao. 2016. A 1-Bit 10×10 Reconfigurable Reflectarray Antenna: Design, Optimization, and Experiment. *IEEE Trans. on Antennas and Propagation* 64, 6 (2016), 2246–2254.

[72] Youwei Zeng, Jinyi Liu, Jie Xiong, Zhaopeng Liu, Dan Wu, and Daqing Zhang. 2022. Exploring Multiple Antennas for Long-range WiFi Sensing. *Proc. ACM IMWUT* 5, 4 (2022), 190–219.

[73] Youwei Zeng, Dan Wu, Jie Xiong, Jinyi Liu, Zhaopeng Liu, and Daqing Zhang. 2020. MultiSense: Enabling Multi-person Respiration Sensing with Commodity WiFi. *Proc. ACM IMWUT*. 4, 3 (2020), 1–29.

[74] Chi Zhang, Feng Li, Jun Luo, and Ying He. 2014. iLocScan: Harnessing Multipath for Simultaneous Indoor Source Localization and Space Scanning. In *Proc. of the 12th ACM SenSys*. 91–104.

[75] Lei Zhang, Ming Zheng Chen, Wankai Tang, Jun Yan Dai, Long Miao, Xiao Yang Zhou, Shi Jin, Qiang Cheng, and Tie Jun Cui. 2020. A Wireless Communication Scheme based on Space-and Frequency-division Multiplexing Using Digital Metasurfaces. *Nature Electronics* 4, 3 (2020), 218–227.

[76] Tianyue Zheng, Zhe Chen, Shujie Zhang, Chao Cai, and Jun Luo. 2021. MoRe-Fi: Motion-robust and Fine-grained Respiration Monitoring via Deep-Learning UWB Radar. In *Proc. of the 19th ACM SenSys*. 111–124.

[77] Yue Zheng, Yi Zhang, Kun Qian, Guidong Zhang, Yunhao Liu, Chenshu Wu, and Zheng Yang. 2019. Zero-Effort Cross-Domain Gesture Recognition with Wi-Fi. In *Proc. of the 17th ACM MobiSys*. 313–325.