

MIRO: Multi-radar Identity and Ranging for Occupational Safety

Tirthankar Halder*
tirthankarhalder70@gmail.com
IIT Kharagpur, India

Argha Sen*
arghasen10@gmail.com
IIT Kharagpur, India

Swadhin Pradhan
swapradh@cisco.com
Cisco Systems, USA

Rijurekha Sen
riju@cse.iitd.ac.in
IIT Delhi, India

Sandip Chakraborty
sandipc@cse.iitkgp.ac.in
IIT Kharagpur, India

Abstract

Occupational exposure to airborne particulate matter (PM) poses a severe health risk in open industrial workspaces such as stone-cutting yards. Conventional monitoring solutions such as wearable PM sensors and camera-based tracking are impractical due to discomfort, maintenance issues, and privacy concerns. We present *MIRO*, a privacy-preserving framework that integrates continuous PM sensing with a multi-radar millimeter-wave (mmWave) re-identification (re-ID) backbone. A distributed network of PM sensors captures localized pollutant concentrations, while spatially overlapping mmWave radars track and re-associate workers across viewpoints without relying on visual cues. To ensure identity consistency across radars, we introduce a *GAN-based view adaptation network* that compensates for azimuthal distortions in range-Doppler (RD) signatures, combined with *correlation-based cross-radar matching*. In controlled laboratory experiments, our system achieves a re-ID F1-score of 90.4% and a mean Structural Similarity Index Measure (SSIM) of 0.70 for view adaptation accuracy. Field trials in rural stone-cutting yards further validate the system's robustness, demonstrating reliable worker-specific PM exposure estimation.

CCS Concepts

• **Human-centered computing** → *Ubiquitous computing*.

Keywords

mmWave Sensing, View Adaptation, Re-identification, Pollution Exposure

ACM Reference Format:

Tirthankar Halder, Argha Sen, Swadhin Pradhan, Rijurekha Sen, and Sandip Chakraborty. 2026. *MIRO: Multi-radar Identity and Ranging for Occupational Safety*. 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, 14 pages. <https://doi.org/10.1145/3774906.3802761>

*Both authors contributed equally in this work.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SenSys '26, Saint Malo, France

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-2309-4/26/05
<https://doi.org/10.1145/3774906.3802761>

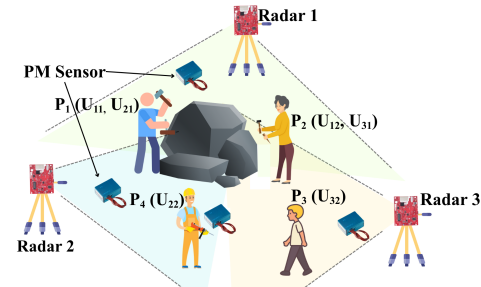


Figure 1: *MIRO*'s multi-mmwave radar re-identification (re-ID) in action. It resolves cross-view identity ambiguity due to overlapping regions, where each radar assigns distinct local user IDs (e.g., U_{11} , U_{21}) to the same worker (P_1), by performing activity-correlated re-ID, while co-located PM sensors enable personalized exposure estimation.

1 Introduction

Occupational safety issues arise from health risks due to the coupling between *what a worker is doing and where one is doing it*. In workspaces, such as construction, mining, stone cutting, and metal fabrication, environmental hazards are highly activity-dependent. If activities involve abrasive contact, like cutting, drilling, grinding, or welding, those emit intense bursts of particulate matter (PM) and fine respirable particles, which can cause chronic obstructive lung diseases, and cardiovascular complications [6, 7, 17]. Similarly, tasks such as torch cutting or solvent cleaning introduce transient chemical vapors, creating health problems due to exposure to problematic volatile organic compounds (VOCs) [48]. These exposures fluctuate quickly as workers alternate between operations, move across zones, and share confined air-spaces. So, accurate occupational-risk assessment, which is essential in improving the safety of the workspace, requires joint, time-aligned measurement of *both activity and pollutant concentration, not just coarse ambient averages*.

Despite advances in industrial IoT and safety analytics, such fine-grained exposure monitoring remains rare in outdoor and semi-outdoor workplaces, especially in developing regions [24, 57]. Cameras and wearable sensors, the two dominant approaches, each face fundamental limitations. Cameras provide rich behavioral context but raise severe privacy and compliance concerns, often leading to worker resistance [40], which is further validated by the survey done in our own experiments (see Figure 3). Wearable PM and motion sensors, while privacy-safe, are intrusive, require maintenance, and bias behavior under heavy workloads [20]. So, occupational-risk monitoring for ensuring safety needs a modality that is simultaneously continuous, privacy-preserving, and maintenance-free.

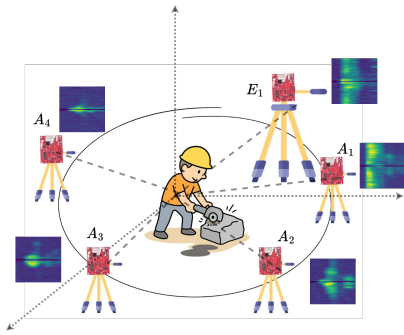


Figure 2: Illustration of azimuthal dependence of range-Doppler (RD) signatures. In stone-work, most motions (body/tool) occur within the azimuthal-plane, producing variations in azimuth angles (A_i). Elevation-plane motions (E_i) contribute less, yielding nearly identical RD responses.

Millimeter-wave (mmWave) radars fill this need for an infrastructure-embedded, device-free seamless sensing method [14, 46, 58]. Operating in the 30 – 300 GHz range with multi-gigahertz bandwidth, mmWave can capture centimeter-level range resolution and micro-Doppler signatures [30, 46, 51]. These signatures are fine-grained modulations caused by limb movements and gait cycles. Because of its high spatial resolution, mmWave is an ideal candidate for monitoring complex environments. Furthermore, the short wavelength (1-10 mm) is orders of magnitude larger than airborne dust particles. It makes propagation essentially immune to scattering and attenuation from particulates [4]. Moreover, mmWave’s small antenna sizes enable *compact, low-profile units* that can be integrated into poles or walls [27]. Beyond these resolution and robustness advantages, mmWave sensing also provides a distinct opportunity for *movement-pattern-based identification*. Unlike vision-based systems, where appearance cues such as clothing or background context can vary drastically across viewpoints, mmWave sensing depends primarily on coarse grained motion-induced micro-Doppler and range dynamics information. As a result, *identity signatures in mmWave remain largely invariant* to background and viewpoint, provided that the underlying movement pattern is consistent.

Yet scaling mmWave from controlled, single-sensor experiments to *large, real-world environments* introduces a fundamental systems challenge: **multi-radar re-ID**. A single radar’s coverage area is limited [2, 22]; monitoring large or irregularly shaped spaces requires deploying multiple radars with overlapping fields of view. As shown in Figure 1, such overlapping coverage introduces cross-view identity ambiguity, where the same worker may be assigned different local user IDs by separate radars observing from distinct viewpoints. This creates the need to consistently link detections of the same individual across sensors. Unlike camera-based re-ID, there are no visual appearance cues; unlike Wi-Fi or UWB [15, 33, 37, 44], there are no persistent MAC addresses or tags [26]. Without *robust identity persistence* across sensors, trajectories fragment where workers usually operate within semi-static task zones [41].

1.1 Closing the Multi-Radar re-ID Gap

Current mmWave re-ID methods typically operate under controlled conditions or rely on trajectory stitching [18, 32, 52]. In semi-static

industrial environments, workers often remain within fixed task areas, yet these approaches still struggle to maintain reliable identification. The micro-Doppler signatures serve as cues, but they remain sensitive to *changes in viewpoints* and *variability in motion*.

To address this, we design a *Pix2Pix-based micro-Doppler view adaptation network* tailored for multi-radar operation. Each radar produces time-synchronized 3D point clouds along with RD heatmaps for all detected users. To ensure robust per-radar user localization, we develop a clustering framework called *TDSCAN* (Temporal Doppler Spatial Clustering) that integrates Doppler-based filtering with temporal consistency matching (§4). Then, the Pix2Pix-based *view adaptation module* (§4) learns a supervised mapping between paired RD signatures captured from different radar viewpoints. This translation compensates for angle-dependent distortions in the *azimuth plane* of our target workspace.

In particular, the majority of motion in stone-cutting environments such as cutting, grinding, and chipping occurs predominantly within the azimuthal plane, while elevation-plane motion remains minimal, as shown in Figure 2. As a result, viewpoint variation primarily manifests as azimuthal shifts in the RD signature rather than significant elevation distortions. By modeling these azimuth-dependent transformations, the network captures the essential geometry of motion without overfitting to environment- or activity-specific nuances. Consequently, when the adaptation focuses solely on motion patterns intrinsic to the activity rather than on background clutter, the learned representation generalizes across unseen workspaces and task types.

1.2 Real-World Validation in Harsh Settings

To comprehensively assess the generalizability of *MIRO*, we conducted a two-phase evaluation in both controlled laboratory experiments and real-world field deployments in diverse stone-processing environments. This multi-stage validation enables us to isolate the performance of the view-adaptation and re-ID modules under ideal conditions and then assess their resilience to environmental variability, such as, dust exposure, and multipath effects.

In the controlled laboratory setup, we recruited 3 participants and instructed them to emulate representative stone-working activities, including chipping, grinding, and polishing gestures. Two mmWave radar nodes were positioned at different azimuthal viewpoints to capture paired RD data simultaneously for each activity. The participants’ ground-truth spatial trajectories were recorded using a Vicon Tracker 4.0 motion capture system with five infrared cameras operating at 100 Hz, providing sub-millimeter accuracy. This controlled setup allowed precise evaluation of the view-adaptation and cross-radar identity consistency modules. The paired multi-view data collected in this phase also used in training our Pix2Pix-based view transformation model.

To evaluate performance under real-world conditions, we deployed *MIRO* across three distinct stone-processing environments, each presenting unique spatial and operational characteristics. The first site, a large open-air marble processing facility, involved multiple workers simultaneously performing cutting, chipping, grinding, and polishing operations on heavy stone blocks. This site featured continuous dust emission, overlapping worker trajectories, and wide radar separation, testing both dust immunity and cross-node

identity persistence. The second site, a semi-mechanized stone-cutting factory, was dominated by high-speed cutting machines generating dense airborne particulates and strong mechanical Doppler signatures. This setting tested the system’s ability to maintain accurate worker identification in the presence of intense background motion. The third deployment took place in an indoor construction workspace where marble slabs were being shaped and installed.

1.3 Contributions

In summary, this paper makes the following key contributions:

1. MIRO Framework. We present *MIRO*, a multi-radar re-ID framework built on top of robust multi-radar localization enabled by our *TDSCAN* clustering module (§4). *MIRO* introduces an *azimuth-aware Pix2Pix view-adaptation network* that translates RD signatures across radar viewpoints. Our model learns viewpoint transformations in *azimuthal motion geometry*, and adapts across unseen workspaces and activity types.

2. Cross-Radar Association and Identity Persistence. Beyond view adaptation, *MIRO* incorporates a *cross-radar association engine* that computes *spatio-temporal correlation scores* between view-adapted and observed RD maps, enabling reliable identity propagation across sensing zones and through occlusions. The resulting association graph ensures *global identity consistency* across all radar nodes. In lab tests, *MIRO* achieves a re-ID F1-score of up to 90.4% and a mean SSIM of 0.70.

3. In-the-wild Evaluation. *MIRO* was deployed in: (i) an open-air marble facility testing dust immunity and identity persistence, (ii) a semi-mechanized factory dominated by mechanical Doppler interference, and (iii) an indoor workspace with reflective surfaces assessing multi-path resilience. Together, these experiments validate *MIRO*’s *robustness* from controlled conditions to complex, real-world industrial settings.

2 Related Work

2.1 Occupational Exposure Monitoring

As occupational environments become increasingly instrumented through the Industrial Internet of Things (IIoT), numerous embedded sensing systems have been introduced to enable affordable, real-time monitoring of dust, gases, and volatile organic compounds (VOCs) [24, 57]. Recent systems combine low-cost particulate sensors with edge computing (e.g., LoRa/LoRaWAN) to produce spatial maps of pollutant concentrations in urban, manufacturing, and construction environments [8, 10, 25]. However, most deployments treat the workspace as a uniform environment, averaging readings over time or space. Such coarse spatiotemporal sampling fails to capture the short-term exposure spikes caused by transient, task-specific operations, which are key determinants of respiratory and cardiovascular outcomes [6, 7].

2.2 Activity and Worker Behavior Monitoring

To contextualize exposure, prior research has explored activity recognition using cameras, wearable sensors, and other modalities. Vision-based systems can infer detailed worker posture, task type, and safety compliance [3, 11, 56], offering rich behavioral information for risk analytics. However, the use of cameras in occupational settings raises privacy concerns, data governance, and regulatory

compliance [9, 21, 40, 45], often leading to worker resistance and legal restrictions. Wearable IMUs, physiological sensors, and smart PPE have also been employed for motion capture and workload estimation [19, 39, 55]. Despite privacy benefits, these devices cause discomfort, and behavioral bias, motivating infrastructure-based, *device-free* sensing that passively infers activity [12, 46].

2.3 mmWave and RF-Based Human Sensing

mmWave radars have demonstrated powerful capabilities for fine-grained human sensing. Operating in the 30–300 GHz spectrum with multi-gigahertz bandwidths, mmWave radars offer centimeter-level range resolution and the ability to capture subtle micro-Doppler modulations induced by limb movements [30, 46, 51]. These signatures have been utilized for gait recognition [51], gesture classification [14], vital signs [50, 59], and occupancy estimation [28, 44, 53]. Unlike optical sensors, mmWave operates reliably under poor illumination, dust, or smoke, and inherently preserves privacy [4]. However, existing studies often assume single-sensor, laboratory-controlled conditions with unobstructed line-of-sight targets. Extending these methods to multi-radar industrial environments requires addressing challenges of sensor coverage and identity continuity across viewpoints.

2.4 Multi-Sensor Cross-View Re-Identification

Cross-view re-ID has been extensively studied in computer vision, where visual appearance cues such as texture, color, and pose are exploited to link subjects across cameras [36, 47]. In contrast, RF-based sensing lacks explicit visual features and must instead rely on motion dynamics, spatio-temporal trajectories, or statistical similarity of doppler profiles. Early mmWave re-ID methods focused on trajectory stitching or clustering in range-time space [18, 32, 52]; however, these approaches are designed for single-radar setups and fail to maintain identity consistency across overlapping fields of view in multi-radar deployments. Wi-Fi [33, 44] and Ultra-Wideband (UWB) [26, 29, 49] systems provide identity persistence through MAC addresses or tag IDs, yet such mechanisms contradict the device-free design goals central to industrial sensing. Recent efforts have begun exploring quasi-biometric properties of micro-Doppler signatures [30, 51], but these features exhibit strong viewpoint dependence. As a result, maintaining consistent user identity across multiple mmWave radar viewpoints with overlapping fields of view remains an open challenge [41].

View adaptation and domain translation techniques have been highly successful in computer vision for cross-domain mapping and style transfer (e.g., Pix2Pix [23], CycleGAN [60]). However, the radar and mmWave re-ID literature has not yet converged on a standard approach for direct image-to-image translation of RD representations. Instead, most prior mmWave/RF methods rely on trajectory stitching, handcrafted or learned view-invariant descriptors, and supervised learning applied in single-radar settings [18, 32, 52]. These approaches work well when sensors operate independently or when coverage is non-overlapping, but they do not explicitly learn cross-view mappings of micro-Doppler structure and therefore struggle to preserve identity consistency when fields of view overlap or when viewpoint shifts substantially.

3 Background and Pilot Study

In this section, we first introduce the fundamentals of mmWave sensing to establish the radar processing concepts used throughout this work. We then shift focus to a specific real-world use case: monitoring occupational health risks in stone-cutting yards. To ground our study, we begin with a survey of workers' perceived challenges in their daily environment, followed by a pilot deployment of pollution sensors. Based on the limitations of pollution-only monitoring and the workers' concerns about privacy, we motivate the use of unobtrusive mmWave radar. Finally, we present our pilot experiments with radar-based activity monitoring.

3.1 Preliminaries: mmWave Radar Sensing

Commercial off-the-shelf mmWave radars typically operate using the FMCW principle [42]. In this scheme, the radar continuously transmits frequency-modulated chirps, while simultaneously mixing the TX signal with the RX echo from surrounding objects. This mixing process, often referred to as *dechirping*, produces an *Intermediate Frequency* (IF) signal.

3.1.1 Range estimation. The distance between the radar and a target is obtained from the IF signal offset [42]. This offset, known as the *beat frequency* (f_b), results from the round-trip delay τ . If the chirp of duration T_C sweeps over a bandwidth B , the chirp slope can be expressed as $S = \frac{B}{T_C} = \frac{f_b}{\tau}$. Since the delay is given by $\tau = \frac{2d}{c}$, where d is the object distance and c is the speed of light, the distance can be calculated as $d = \frac{c}{2} \cdot \frac{T_C}{B} \cdot f_b$. To estimate f_b , a Fast Fourier Transform (FFT), called the *range-FFT*, is applied to the IF signal. The resulting spectrum contains peaks at frequencies corresponding to object reflections, thereby providing the range estimate.

3.1.2 Velocity estimation. To extract velocity information, the radar transmits a sequence of N chirps, each separated by a duration of T_C . If a target moves at speed v , the motion introduces a phase shift between two consecutive received chirps. This phase shift, $\Delta\phi$, is given by $\Delta\phi = \frac{4\pi v T_C}{\lambda}$, where λ is the wavelength of the carrier. By performing a second FFT, referred to as the *Doppler-FFT*, on the phase variations across chirps, the velocity of the target can be estimated. This information is organized into a two-dimensional *range-Doppler* matrix $\mathbb{D}_{D \times R}$, where D and R denote the number of Doppler bins and range bins, respectively.

3.1.3 Point cloud estimation. The *point cloud* is extracted from the range-Doppler matrix using the Constant False Alarm Rate (CFAR) detector [38], which identifies prominent peaks corresponding to actual targets. The resulting point cloud encodes the spatial coordinates (x_i, y_i, z_i) , Doppler velocity (d_i), and received signal power (p_i) for each detection. Formally, for N detected objects, the point cloud set S can be expressed as, $S = \bigcup_{i=1}^N \{(x_i, y_i, z_i, d_i, p_i)\}$. Further details on point cloud generation are available in [43].

3.2 Survey of Stone-Cutting Workers

To understand the real challenges faced in stone workers' occupational environments, we conducted a survey among laborers in a

semi-urban marble processing factory. A total of 8 workers participated in the survey, with an average age of 43 years and more than 10 years of work experience in stone cutting.

Workers were then asked to rank the severity of common workplace risks, including (i) dust exposure (respiratory issues, silicosis risk), (ii) noise exposure, (iii) heat and harsh environmental conditions, (iv) repetitive strain and physical fatigue, and (v) risk of accidents from heavy machinery or tools. The results, summarized in Figure 3(a), clearly indicate that dust exposure is considered the most severe challenge, with an average severity score of 4.125 out of 5. Noise and heat were perceived as moderate risks, while repetitive strain and accident risk received relatively lower average scores. This confirms dust exposure as a dominant occupational hazard, emphasizing the need for localized pollution monitoring.

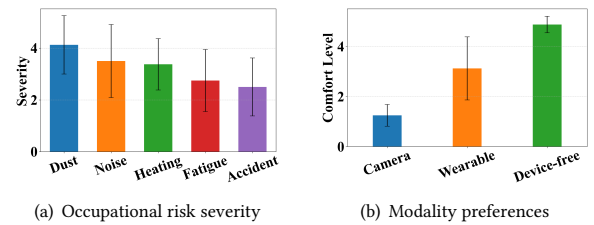


Figure 3: Survey results from stone-cutting workers.

3.3 Pilot Deployment with Pollution Sensors

Motivated by these findings, we first deployed *DALTON* PM sensors [25] to capture localized pollution levels in a large open-air marble processing facility. Each sensing node, shown in Figure 4(a), was capable of measuring $PM_{1.5}$, $PM_{2.5}$, and PM_{10} concentrations, along with other gases such as CO, CO₂, NO₂, and VOCs. The real-world deployment of these nodes in the yard is illustrated in Figure 4(b), while a schematic layout of the deployment is provided in Figure 4(c). In this pilot study we primarily focus on PM_x measurements, which have direct links to occupational hazards such as silicosis [6].

3.3.1 PM Exposure Analysis Across Activities. Stone-cutting tasks are significantly diverse in their dust-generation potential. We measured localized PM concentrations during three common activities: polishing, chipping, and grinding. Our results as shown in Figure 5(a) demonstrate a clear gradient in exposure levels. Grinding generated the highest pollution levels, often producing PM concentrations that exceed occupational safety thresholds by orders of magnitude [5]. The fine particles generated during grinding remained suspended longer and spread further. Polishing exhibited moderate PM levels, typically lower than grinding, but still significantly above background levels. Chipping produced the least amount of PM among the activities.

Thus pollutants' exposure is not uniform across the workspace but is strongly tied to the nature of the activities being performed. Thus continuous task or activity recognition is crucial in this direction to have activity-specific insights.

3.3.2 Impact of Sensor Proximity to Dust Source. Another key observation from our pilot deployment is the strong effect of distance

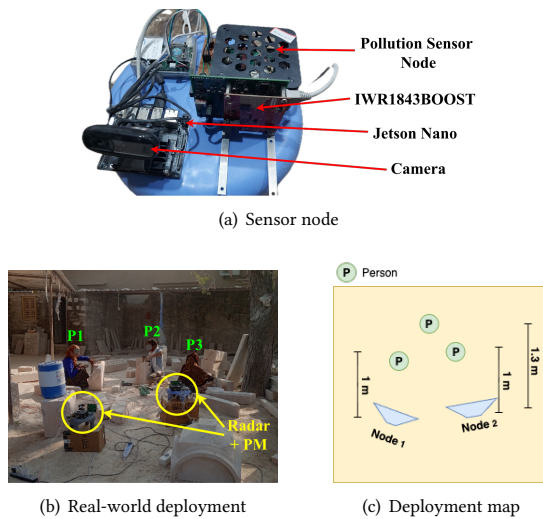


Figure 4: Detailed experimental setup.

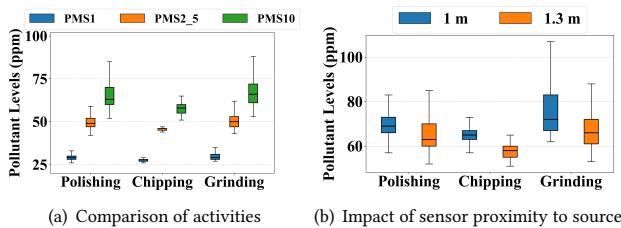


Figure 5: PM exposure analysis.

between the PM sensor and the pollutant source. As shown in Figure 4(b) we have stone cutting activities happening at different distances with respect to the pollutant sensors (≈ 1 m and 1.3 m as shown in Figure 4(c)). Sensors placed closer to the grinding activity recorded peak concentrations several times higher than those located only a few meters away (see Figure 5(b)). This distance sensitivity implies that sparse placement of PM sensors can lead to biased exposure estimates if worker proximity is not tracked.

Thus, it becomes essential to link worker trajectories with spatially varying PM measurements to estimate individualized exposures. While PM sensors provide valuable measurements of localized pollutant concentrations, they suffer from several limitations that prevent their use as a standalone solution. First, in high-dust environments such as stone-cutting yards, PM sensors can clog or saturate, leading to unreliable readings. Second, their measurements are highly sensitive to placement and orientation: a sensor located only a few meters away from the activity can record drastically lower exposure than one placed nearby. In addition, pollution levels alone cannot tell which worker is generating dust or which activity is responsible, making worker-specific exposure insights impossible. Due to these limitations, alternative monitoring approaches are needed that link exposure levels to worker positions and activities. Among the common options are camera-based monitoring

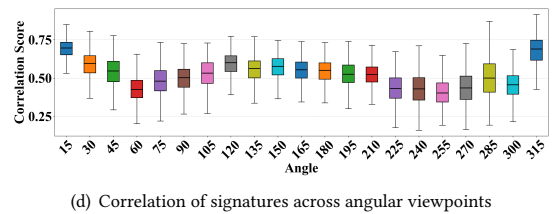
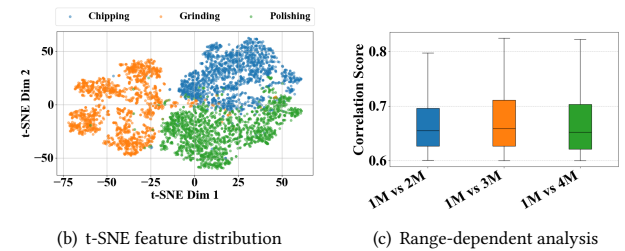
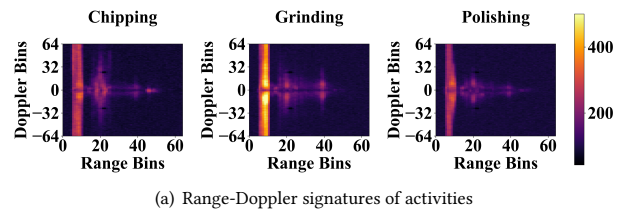


Figure 6: Radar-based activity analysis.

and wearable sensors, each with its own privacy, comfort, and reliability trade-offs. To better understand worker perceptions of these modalities, we conducted a survey on monitoring preferences.

3.4 Constraints on Monitoring Modalities

We asked workers about their comfort levels with different monitoring approaches. The survey results, summarized in Figure 3(b), indicate that *camera-based monitoring* was generally viewed as intrusive and raising strong privacy concerns, while *wearable-based monitoring* (e.g., smart bands) was reported as uncomfortable or impractical given the physical demands of stone-cutting. In contrast, workers expressed much higher acceptance for device-free monitoring solutions that do not rely on either cameras or wearables.

These findings suggest that mmWave radar can be a promising direction: it offers device-free, privacy-preserving robust sensing in dusty environment. However, a single mmWave radar cannot provide complete coverage due to limited FoV and frequent occlusions. This limitation motivates the exploration of multi-radar setups and also introduces new challenges in making activity recognition consistent across different viewpoints.

3.5 Activity Signature Analysis Using mmWave

Even though PM sensors measure localized concentrations, they cannot differentiate between workers or activities that generate pollution. To address this, we leverage mmWave radar's micro-Doppler signatures. As shown in Figure 6(a), different activities produce distinct range-Doppler spectrograms, enabling the possibility of activity classification and separation of workers' roles. To

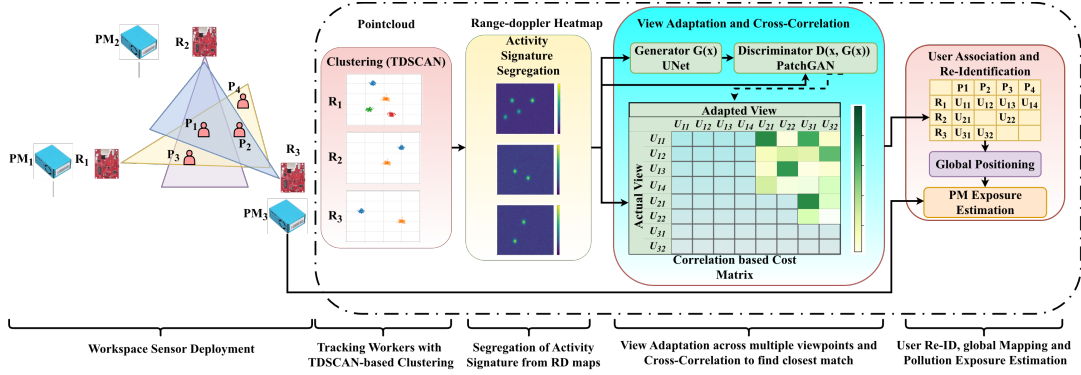


Figure 7: System overview diagram of MIRO.

highlight this, we apply t-distributed Stochastic Neighbor Embedding (t-SNE) [34] to project the features from raw range-Doppler data into a two-dimensional feature space. As shown in Figure 6(b), the embeddings for different activities form distinct clusters, demonstrating that micro-Doppler features can classify tasks.

However, this discriminability is not invariant to sensing geometry. When the same activity is observed from different angular viewpoints, the resulting spectrograms exhibit reduced correlation compared to a reference view taken at 0° along the radar boresight. For instance, grinding observed laterally (at 90°) produces a substantially different Doppler spread than when observed head-on, as shown in Figure 6(d). In contrast, changes in range have a weaker effect: signatures recorded at different distances but similar angles maintain higher correlation, as shown in Figure 6(c).

These findings highlight that although multiple sensors are required for full coverage, their differing viewpoints introduce distortions that can confuse recognition models. This motivates the design of a view-adaptation mechanism capable of aligning micro-Doppler features across arbitrary angles, thereby ensuring reliable activity discrimination in heterogeneous, multi-radar environments.

4 Methodology

We aim to achieve robust multi-radar user re-ID and personalized pollutant exposure estimation in the workplace. To accomplish this, we combine radar-based localization, activity signature extraction, multi-radar view adaptation, identity correlation, and pollution mapping with global positioning (shown in Figure 7).

4.1 User Localization via Pointcloud Clustering

A multi-radar environment makes accurate user localization difficult because each radar observes semi-static subjects from different viewpoints and under varied clutter and signal sparsity. Conventional clustering algorithms such as DBSCAN [13] are ineffective because radar detections are sparse, non-uniform, and temporally inconsistent across sensors. The *fixed density* and *neighborhood thresholds* used by DBSCAN make it sensitive to these variations, often causing fragmented clusters for weak detections or merged clusters when reflections from multiple users overlap. Such inconsistencies are further amplified in a multi-radar setup, where small clustering errors at each radar propagate across viewpoints, leading

to misaligned user representations and degraded performance in downstream view adaptation and cross-radar identity matching. To overcome these limitations, we propose **TDESCAN (Temporal Doppler Spatial Clustering)**, a modified clustering framework that integrates *Doppler-based filtering* with *temporal consistency matching* to ensure reliable localization across multiple radar nodes in cluttered industrial environments.

The key motivation behind TDESCAN is two-fold: (1) **Doppler-based filtering** removes static clutter points whose Doppler values are near zero, isolating reflections corresponding to slow human motions. Since workers in our scenario perform quasi-static activities (e.g., grinding and not walking/heavy exercises involving significant body movements), valid reflections exhibit small but non-zero Doppler velocities. Thus, we apply a lower and upper Doppler bound (τ_{\min}, τ_{\max}) to retain only those points satisfying $\tau_{\min} < |v_i| < \tau_{\max}$. In our implementation, thresholds of $\tau_{\min} = 0.05$ m/s and $\tau_{\max} = 1.0$ m/s effectively distinguish human workers from non-human static clutter and high-velocity noise. (2) **Temporal consistency matching** enforces continuity of detected clusters across consecutive frames. Rather than treating each frame independently, TDESCAN accumulates detections over a short temporal window (10 frames, or approximately 1 second at 10 FPS), then performs centroid-based track association using a distance-constrained assignment (Hungarian algorithm). This temporal linkage mitigates discontinuities in user localization and suppresses false positives from transient reflections. Formally, after Doppler-based filtering, we obtain a reduced set of points

$$\mathcal{P}' = \{p_i = (x_i, y_i, z_i, v_i) \in \mathcal{P} \mid \tau_{\min} < |v_i| < \tau_{\max}\}. \quad (1)$$

TDESCAN then operates on the Doppler-filtered point set \mathcal{P}' , aggregating spatially and temporally coherent points into user clusters. A point p_i is a *core point* if:

$$|\{p_j \in \mathcal{P}' \mid \|p_i - p_j\| \leq \epsilon\}| \geq \text{minPts}, \quad (2)$$

where ϵ is the neighborhood radius and *minPts* is the minimum number of neighboring points. Each cluster C_k represents a user track, with its centroid giving the instantaneous 2D position.

$$\mathbf{u}_k = \frac{1}{|C_k|} \sum_{p_i \in C_k} (x_i, y_i). \quad (3)$$

To maintain temporal coherence, cluster centroids from consecutive windows are matched based on spatial proximity. Tracks are

updated when a new centroid lies within a predefined distance from an existing track; otherwise, new IDs are initialized. Clusters that appear inconsistently or for fewer than a threshold number of frames are discarded, eliminating short-lived clutter responses. So *TDSCAN* integrates (i) Doppler-aware point selection, (ii) spatial density-based clustering, and (iii) temporal centroid tracking, producing robust, low-latency user localization.

4.2 Segregation of Activity Signatures

Beyond localization, we require discriminative features for re-ID. For each of the localized clusters, we extract the RD heatmaps centered at the user’s range bin. Specifically, if the user is localized at range bin r_0 , we extract RD patches spanning $r \in [r_0 - 10, r_0 + 10]$ bins over the range axis. This yields a heatmap $\mathbf{H}_k(r, v)$ encoding the RD signature of user k . Let $\mathbf{R}(r, v)$ denote the full RD heatmap, then the user-specific activity signature is:

$$\mathbf{H}_k = \mathbf{R}(r, v)|_{r \in [r_0 - 10, r_0 + 10]}. \quad (4)$$

This selective extraction focuses on isolating motion-induced Doppler components associated with the user’s activity while excluding static reflections and background clutter present in other range bins. Since environmental reflections, machinery, and stationary objects contribute primarily to static or near-zero Doppler components, cropping around the user’s localized range bin effectively suppresses these environment-dependent features. The resulting \mathbf{H}_k thus captures only the intrinsic motion dynamics, such as limb oscillations or tool-hand interactions, that define the worker’s activity. By filtering out environmental influence at the feature level, these activity signatures become inherently *environment-agnostic*, and form the input for the view-adaptation module.

4.3 View-Adaptation via Generative Conversion

Micro-Doppler signatures are highly viewpoint-dependent – the same activity (e.g., grinding or polishing) produces significantly different RD spectrograms when observed from different radar angles. Such angular variations arise from differences in radial velocity components and relative motion geometry, which distort Doppler spread and intensity patterns. Consequently, direct comparison of RD signatures across radars becomes unreliable.

To mitigate this, we employ a generative view-adaptation module that learns to translate RD signatures between radar viewpoints, thereby normalizing viewpoint-specific distortions while preserving motion-specific structures. Intuitively, this module performs a *learned transformation of the radar “view”*, mapping the micro-Doppler representation of an activity as seen from radar m to how it would appear from radar n . This allows downstream identity embeddings to be compared in a common, view-normalized representation space, enabling robust multi-radar re-ID.

4.3.1 Pix2Pix-based View Transformation Network. We adopt a conditional GAN architecture based on the *Pix2Pix* framework [23], which allows paired, supervised training of cross-view mappings. Unlike unpaired image translation methods such as *CycleGAN* [60], *Pix2Pix* directly leverages aligned RD data collected from overlapping radar fields of view. The model is trained on a controlled indoor dataset where participants emulate typical stone-working activities such as cutting, chipping, grinding, and polishing while

being simultaneously observed by two radars placed at different azimuthal separations. This setup provides paired RD spectrograms for the same motion instance viewed from multiple azimuth angles. The resulting training corpus spans a full 360° coverage at 15° increments, capturing the continuous transformation of micro-Doppler signatures with changing azimuth viewpoint. Let $\mathbf{H}_k^{(m)}$ denote the RD signature of user k observed by radar m , and $\mathbf{H}_k^{(n)}$ the same activity observed by radar n . The generator network $G_{m \rightarrow n}$ learns to transform $\mathbf{H}_k^{(m)}$ into the target view, $\hat{\mathbf{H}}_k^{(n)} = G_{m \rightarrow n}(\mathbf{H}_k^{(m)})$.

4.3.2 Network Architecture. The generator G follows a *U-Net* encoder–decoder architecture with skip connections, enabling the preservation of low-level spectral structures while learning global viewpoint transformations. The discriminator D is implemented as a *PatchGAN* [23], which classifies local 70×70 patches rather than full images, encouraging high frequency consistency and reducing over smoothing in RD reconstructions. This architecture is well suited for RD spectrograms, where fine grained local textures encode micro-motion patterns.

4.3.3 Loss Function. We employ the Least-Squares GAN (LSGAN) objective [35] for improved training stability and gradient behavior compared to the original binary cross-entropy loss. The adversarial loss for the view transformation is given by:

$$\mathcal{L}_{LSGAN}(G, D) = \mathbb{E}_{\mathbf{H}^{(n)}} [(D(\mathbf{H}^{(n)}) - 1)^2] + \mathbb{E}_{\mathbf{H}^{(m)}} [D(G(\mathbf{H}^{(m)}))^2]. \quad (5)$$

To enforce content consistency between the synthesized and target RD maps, we add an L_1 reconstruction term:

$$\mathcal{L}_{L1}(G) = \mathbb{E}_{\mathbf{H}^{(m)}, \mathbf{H}^{(n)}} [\|\mathbf{H}^{(n)} - G(\mathbf{H}^{(m)})\|_1]. \quad (6)$$

The total training objective is thus:

$$\mathcal{L}_{total} = \mathcal{L}_{LSGAN}(G, D) + \lambda_{L1} \mathcal{L}_{L1}(G), \quad (7)$$

where λ_{L1} controls the trade-off between perceptual realism and structural fidelity. In our implementation, $\lambda_{L1} = 100$ provided optimal convergence across all activity types.

4.3.4 Effect of Azimuthal Dependence and Activity-Agnostic Learning. The angular dependence of micro-Doppler patterns primarily manifests in the *azimuth plane*, as most worker motions, such as arm swings, tool movements, and torso rotations, are planar and oriented parallel to the ground. In contrast, elevation-plane motion contributes minimally to the observed Doppler spread in typical stone-cutting and grinding postures. By learning transformations parameterized predominantly by azimuthal geometry, the network captures the underlying kinematic structure of motion rather than overfitting to activity semantics or environmental context.

Consequently, the learned mapping becomes both *activity-agnostic* and *environment-agnostic*. Since the model is trained on diverse activity instances that share similar geometric dependencies across azimuth viewpoints, it generalizes to unseen or composite activities that follow comparable motion trajectories. The learned transformation thus reflects a fundamental property of radar scattering geometry rather than any specific action, allowing reliable view adaptation across different workplaces, tasks, and deployment geometries. We demonstrate this through field testing in different workshop environments, *without retraining the pix-to-pix re-ID model* trained once in the lab setting (Section 7).

4.3.5 Effect of View Normalization. Through this learned mapping, the generator produces view-normalized RD representations that preserve intrinsic motion signatures (e.g., chipping vs. grinding) while suppressing angle-induced Doppler distortions. As a result, embeddings derived from these normalized spectrograms become directly comparable across radar viewpoints, enabling robust correlation-based cross-radar re-ID (§4.4) independent of activity type or environment.

4.4 User Association and Re-Identification

To achieve consistent user identities across radars, we perform a two-stage association process combining pairwise matching and identity consolidation. For each frame, we compute the normalized correlation between activity signatures from radar m and n . For user a at radar m and user b at radar n , the correlation coeff. is:

$$\rho(a_m, b_n) = \frac{\langle \hat{\mathbf{H}}_a^{(n)}, \mathbf{H}_b^{(n)} \rangle}{\|\hat{\mathbf{H}}_a^{(n)}\| \|\mathbf{H}_b^{(n)}\|}, \quad (8)$$

where $\hat{\mathbf{H}}_a^{(n)}$ is the view-adapted signature of user a in radar n 's perspective. Pairs with $\rho(a_m, b_n) \geq \tau$ (empirically, $\tau = 0.6$) are considered candidate matches.

We construct a correlation matrix for all detections between the two radars, convert it to a cost matrix by subtracting from its maximum, and solve a one-to-one assignment using the Hungarian algorithm. Formally, for radar pair (m, n) with N_m and N_n detections respectively, we define a similarity matrix $\mathbf{R} \in \mathbb{R}^{N_m \times N_n}$ where each entry corresponds to $\rho(a_m, b_n)$. Correlations below the threshold τ are suppressed prior to assignment. The matrix is converted into a cost matrix \mathbf{C} such that $C_{a,b} = \max(\mathbf{R}) - R_{a,b}$, ensuring that higher similarity results in lower assignment cost. Entries corresponding to suppressed correlations are assigned a large penalty to prevent forced matches. We apply a correlation threshold and optionally enforce a mutual-best check, retaining only the strongest bidirectional correspondences. The Hungarian algorithm is applied to cost matrix \mathbf{C} to derive a globally optimal one-to-one assignment that minimizes the total cost objective. Subsequently, assignments with a correlation coefficient below the threshold τ are pruned. In high-density scenarios, an optional mutual-best constraint is enforced to mitigate ambiguous associations. The surviving pairs constitute unique matches for that frame. Across multiple frames and/or additional radars, we aggregate all accepted pairwise matches into an undirected graph, where nodes represent local detection IDs (e.g., "U11", "U21") and edges represent confirmed matches. Connected components of this graph define transitive equivalence classes of detections belonging to the same user. Each component is assigned a deterministic global user ID (P1, P2, ...), and all member detections are recorded. To ensure identity persistence over time, this association graph is maintained across consecutive frames. Nodes are time-stamped, and short-term missed detections are tolerated within a predefined temporal window. If a user reappears within this window and within a spatial proximity threshold of its last known global position, the previous identity is reactivated to prevent identity fragmentation.

This pipeline guarantees strict one-to-one matches at the pairwise stage, robust multi-radar/multi-frame clustering via transitive closure, and stable global IDs suitable for downstream tracking.

In our case, the setup involves three radars participating in the pairwise association process. Sequentially, radar 1's transformed signatures are correlated with radars 2 and 3, followed by radar 2 correlated with radar 3. The resulting global associations are stored in a matrix \mathbf{A} , where entry $\mathbf{A}_{m,k} = g$ indicates that user k at radar m is assigned to global ID g .

4.5 Use-case Design: Global User Positioning and PM Exposure Estimation

Using the globally consistent user identities obtained from cross-radar association, we project all radar-localized coordinates into a unified global coordinate frame. We select one radar node, denoted as radar r_0 , as the reference and define all transformations relative to it. Let \mathbf{T}_m represent the rigid-body transformation (rotation and translation) from radar m 's local coordinates to the reference radar's coordinate system. The global position of user k observed by radar m is computed as $\mathbf{u}_k^{global} = \mathbf{T}_m \cdot \mathbf{u}_k$.

Once all users' trajectories are expressed in this common frame, we estimate their personalized exposure to PM. A sparse network of PM sensors (*DALTON* [25]) distributed across the workspace provides localized pollutant concentration readings $\{PM_j(s_j)\}$ at sensor locations s_j . Conceptually, the continuous PM concentration field can be modeled via inverse distance weighting (IDW):

$$P\hat{M}(\mathbf{x}) = \frac{\sum_j w_j(\mathbf{x}) PM_j}{\sum_j w_j(\mathbf{x})}, \quad w_j(\mathbf{x}) = \frac{1}{\|\mathbf{x} - s_j\|^p}, \quad (9)$$

where p controls the decay rate with distance and $P\hat{M}(\mathbf{x})$ denotes the interpolated pollution value at position \mathbf{x} .

In practice, since the number of PM sensors is limited, we approximate $P\hat{M}(\mathbf{x})$ using a discrete spatial heatmap representation rather than a fully continuous interpolation, as followed in several existing literature [31, 54]. Each PM sensor is associated with a spatial zone in the workspace corresponding to its physical placement, and the measured pollution levels are projected onto these zones. Zone-level pollution values are then smoothed to form a coarse pollution field, providing an interpretable spatial distribution of PM concentration over the monitored area. This implementation captures the relative concentration gradients necessary for per-worker exposure estimation while avoiding instability from sparse-sensor interpolation. For each user trajectory $\{\mathbf{u}_k^{global}(t)\}$, the personalized exposure is estimated as:

$$E_k = \frac{1}{T} \int_0^T P\hat{M}(\mathbf{u}_k^{global}(t)) dt, \quad (10)$$

where T is the observation duration. The resulting E_k values provide individualized, time-averaged exposure estimates by combining radar-tracked trajectories in the global frame with the discretized PM field. Radar localization and re-identification operate at 10 Hz, while PM sensors sample at 1 Hz. For synchronization all streams are timestamp-aligned, and exposure is computed over a 5-second aggregation window. Within each window, we compute the median re-ID outputs and corresponding PM readings before estimating E_k . This temporal smoothing reduces the impact of transient mismatches and sensor noise.

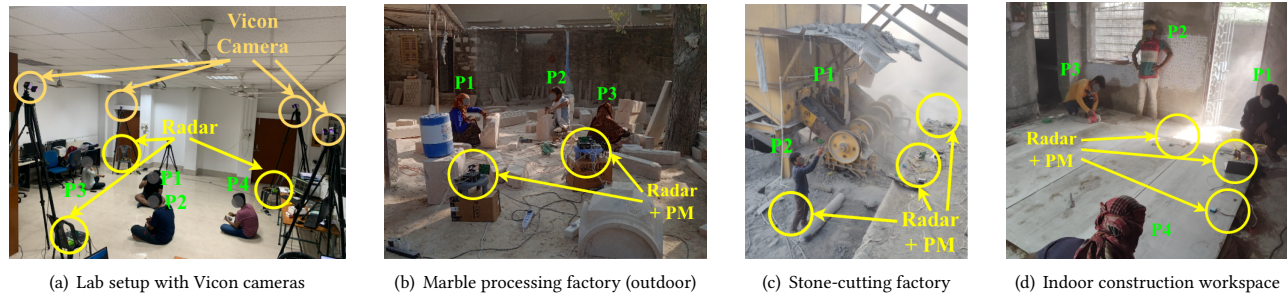


Figure 8: Multiple deployment scenarios of MIRO.

5 Implementation

The implementation of *MIRO* integrates a multi-radar network, micro-controller-based sensor interfaces, and a software pipeline consisting of data acquisition, pre-processing, and view adaptation.

5.1 Hardware Setup

Each sensing node is built with Texas Instruments IWR1843BOOST mmWave radar [2], which provides FMCW based range and Doppler signatures. The raw ADC data were captured using DCA1000EVM [1] data captured card and streamed over Gigabit Ethernet to a local embedded processor. Alongside the radar, we deployed an ESP32 NodeMCU microcontroller for handling serial and digital interfaces with the pollution sensors. We employed an NVIDIA Jetson Nano at each node as the central compute unit, receiving radar streams over ethernet and sensor data over Wi-Fi. Figure 4(a) shows one deployed radar-sensor node with its ESP32 and Jetson Nano stack. The radars were configured with TI mmWave Studio software. We used a chirp profile sweeping 4 GHz bandwidth (77–81 GHz) with a chirp duration of 72 μ s and a sampling rate of 4400 Msps. Each frame contained 182 chirps at a frame rate of 10 Hz with approximately 4 cm range resolution and 182 Doppler resolution to capture micro-motion signatures.

5.2 Software Framework

The collected raw ADC data from the Ethernet streams are first converted to RD heatmaps and pointclouds using custom Python scripts. The ESP32 NodeMCU is programmed using the Arduino framework, enabling serial data collection from PM and gas sensors, packetization, and transmission over Wi-Fi. On the receiving end, the Jetson Nano runs a lightweight server that synchronizes incoming Wi-Fi sensor streams with Ethernet radar streams. TI mmWave Studio is used for radar configuration. All collected data are stored locally on the Nano and mirrored to a central server for offline processing and model training. The view-adaptation pipeline is trained on a Tesla V100 GPU (two 16 GB GPUs) in a rack server with 256 GB RAM and an Intel Xeon(R) Gold 6152 CPU with 88 cores. The train-test split used is 80:20.

5.3 Deployment Scenarios

We evaluated *MIRO* in two distinct phases: (1) laboratory experiments with precise ground truth for validation of view-adaptation and re-ID pipeline, and (2) real-world field deployments.

5.3.1 Laboratory Validation Setup. We recruited 15 participants and instructed them to emulate typical stone-cutting and polishing motions, including chipping, grinding, and polishing gestures. The ground-truth spatial coordinates of each participant were captured using a Vicon motion capture system with five infrared cameras operating at 100 Hz, providing sub-millimeter accuracy. This setup enabled us to evaluate the end-to-end accuracy of the radar-based localization and the consistency of our view-adaptation and re-ID pipeline against high-precision user positioning. Furthermore, to train the view-adaptation network described in Section 4.3, we explicitly collected paired radar data from two radar nodes positioned at controlled azimuthal separations. One radar was fixed at 0° reference orientation, while the second radar was sequentially placed at 15°, 30°, 45°, and so on up to 360°. At each angular configuration, participants performed the same set of activities to capture paired RD representations of identical motions from multiple viewpoints. Each activity session lasted approximately 3–5 minutes per angular configuration at 10 FPS, for an effective cumulative recording duration of approximately 6 hours and yielding over 126,000 RD frames used for supervised training. This systematic multi-angle data collection enabled supervised training of the Pix2Pix-based view transformation model, enabling it to learn azimuthal dependencies and translate robustly across multiple radar perspectives. Following training in a lab-scale setup, we evaluated the pretrained model in real-world deployment scenarios, as discussed next.

5.3.2 Real-World Deployments. To assess robustness under actual working conditions, we deployed *MIRO* across three different stone-processing environments, as follows.

- **Site 1: Marble Processing Factory (Outdoor).** This large open-air marble factory hosts multiple work zones where 5 stone workers simultaneously perform cutting, chipping, grinding, and polishing operations on marble blocks. The workspace features continuous dust emission and overlapping worker trajectories, posing challenges for radar visibility.
- **Site 2: Stone-Cutting Factory.** This site represents a mechanized stone-cutting factory where dust generation originates primarily from high-speed stone-cutting machines. Despite only two workers, continuous machinery and heavy dust made it ideal for testing robustness of the re-ID pipeline.
- **Site 3: Indoor Construction Workspace.** The third deployment took place in a partially enclosed indoor construction area where marble slabs were being cut and installed, with four workers

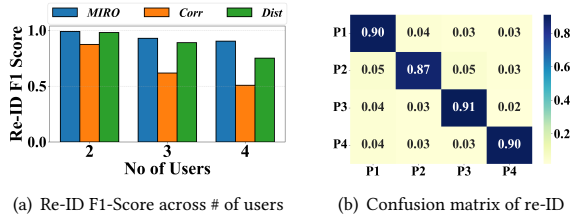


Figure 9: Overall re-ID performance.

present. Unlike the outdoor sites, this environment had strong multipath and limited fields of view.

Across all sites, three mmWave radar nodes were deployed in partially overlapping configurations, synchronized through Ethernet, and co-located with PM sensors. The system continuously captured multi-radar RD data streams and environmental PM readings. The overlapping radar placement ensures users remain observable despite *occlusions* or *FoV limits*, as at least one other node can maintain visibility. For example, P4 in Figure 8(d) is occluded in Radar 1’s FoV but remains visible to the other nodes.

Real-world data, totaling 20 hours and over 265,000 RD frames, were recorded across three sites to capture varying environmental conditions. The view-adaptation model, trained on laboratory data (80:20 split), was evaluated in these unseen field environments without fine-tuning to test generalization. To ensure rigorous re-identification assessment, a participant-independent protocol was used, keeping training and testing identities mutually exclusive.

6 Lab-Scale Evaluation

We first evaluate the overall performance of *MIRO* in the lab setup, followed by component-wise analysis to quantify their individual accuracy and robustness against baseline methods.

6.1 Overall Performance

To quantify the end-to-end performance of *MIRO*, we compare it against two baseline re-ID strategies: (i) a *distance-based* approach that associates detections across radars using nearest-neighbor matching in global coordinates, and (ii) a *correlation-based* approach that uses micro-Doppler activity correlation between radar nodes without any view normalization. Notably, to the best of our knowledge, our method is the first to be able to identify an object through multiple radars at the same time. Thus, we compare these two classical methods for radar point-cloud to human-subject association with our approach under multi-radar setup.

As shown in Figure 9(a), the *distance-based* baseline suffers with increasing user count (from 0.91 for two to 0.73 for four users), as it relies on one-to-one Hungarian matching in global coordinates and becomes unreliable under higher densities due to localization noise and spatial overlap. While the *correlation-based* baseline fails under dense settings as unaligned viewpoints distort temporal and angular features. Lacking view adaptation, pure correlation leads to frequent identity confusion. In contrast, *MIRO* sustains near-constant F1 scores across densities by employing a *Pix2Pix-LSGAN-based view adaptation* that learns azimuth-aware geometric transformations, producing view-consistent, activity-invariant RD representations. Consequently, *MIRO* achieves superior scalability and robustness

in dense multi-radar deployments, as evidenced by the per-user consistency in Figure 9(b). *MIRO* targets small-scale industrial work zones with a modest number of concurrent workers ($\approx 3-6$). While Fig. 9(a) evaluates up to four users, performance may degrade at higher densities due to radar resolution limits and overlapping range-Doppler signatures.

6.2 Microbenchmark Evaluation

We evaluated the performance of clustering, view adaptation, and cross-radar re-ID components of *MIRO*.

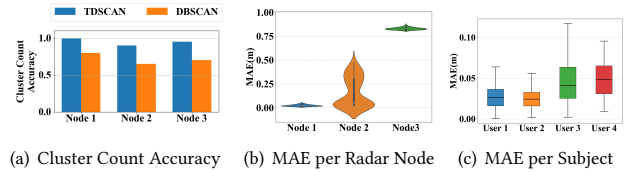


Figure 10: Comparison of radar-based position estimates against Vicon measurements (for the setup of Figure 8(a)).

6.2.1 Performance of Clustering. We first examine the accuracy of clustering pointclouds into user instances. Clustering performance is measured as $Acc = 1 - \frac{|C_{detected} - C_{actual}|}{C_{actual}}$, where $C_{detected}$ is the number of clusters formed and C_{actual} the ground-truth count of workers. As shown in Fig. 10(a), the proposed TDSCAN method consistently achieves high accuracy compared to standard DBSCAN based approach. As a result of TDSCAN’s design, which employs Doppler-based thresholding to suppress static clutter, and then stacks pointcloud frames in a timely manner, it creates a cleaner and more coherent representation of each worker. Beyond cluster detection, we assess how accurately TDSCAN estimates the spatial positions of the clusters. The mean absolute error (MAE) between each cluster centroid and its corresponding Vicon-tracked ground truth is summarized in Fig. 10(b) and Fig. 10(c).

Node-wise, we observe that Node3 can see four cluster IDs, resulting in a higher accumulated error compared to the other nodes. Node2 and Node1 only capture two cluster IDs each, producing similar but comparatively lower accumulated errors. Across all users, the average MAE remains below 10 cm, with slightly higher errors observed for users positioned near the radar’s field-of-view (FoV) boundaries due to sparser point returns. Increased user radar distance also contributes to higher errors, as it raises the likelihood of intermediate interference from other users’ activities entering the FoV and degrading the pointcloud density of the farthest user. In our setup, *user 4* was located at the largest distance from the radar (6 m), whereas *users 1-3* were positioned within 3 m.

For a single radar view, lateral position estimates rely on azimuthal angle measurements; thus, even a small angular error $\Delta\theta$ translates into a lateral displacement of approximately $R\Delta\theta$, which increases linearly with range. Moreover, radar path loss at longer distances reduces return power, lowers SNR, and increases CFAR misses, yielding sparser and noisier point sets. Under such conditions, DBSCAN becomes more fragile: clusters may fragment if `min_samples` is set too high, or may merge incorrectly if `eps` is enlarged to preserve cluster continuity.

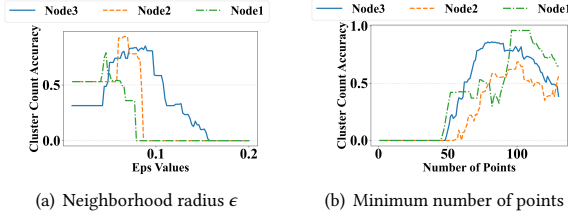


Figure 11: Clustering count accuracy of TDSCAN parameters.
Table 1: Performance of different view-adaptation methods.

Model	L1_mean	SSIM_mean	PSNR_mean
CycleGAN	12.8522	0.4967	19.48
Pix2Pix WGAN-GP	19.4237	0.2654	18.27
Pix2Pix LSGAN	6.1296	0.6988	28.02

After characterizing the performance of TDSCAN, we next tune the algorithm’s parameters by varying both the neighborhood radius ϵ and the minimum number of points (minPts) (see Figure 11). We observe a clear optimum at $\epsilon = .75$ m and minPts = 100 in Figure 11(a) and Figure 11(b) respectively, achieving an average cluster count accuracy of 94%. These parameters correspond well to the physical spacing between individuals in the worksite.

In this environment, each worker operates in close proximity to their stone block and tools, forming a compact point cloud roughly matching the size of a human body (around 0.5-1 m in spatial extent). If ϵ is too small, the algorithm splits a single worker into multiple fragments (for example, separating the body and nearby tools). Conversely, if ϵ is too large, neighboring workers may be merged into one cluster. The optimal value of $\epsilon = 0.75$ m effectively captures worker bodies and instruments, resulting in clusters that tightly enclose each individual’s workspace.

6.2.2 Performance of View-Adaptation. We compared three image-to-image translation frameworks for cross-view radar representation learning: *Pix2Pix with Least Squares GAN (Pix2Pix-LSGAN)* [23], *Pix2Pix with Wasserstein GAN and Gradient Penalty (Pix2Pix-WGAN-GP)* [16], and *CycleGAN* [60]. Each model was trained to translate RD heatmaps between radar viewpoints, enabling consistent micro-Doppler interpretation across different azimuth angles. To quantify view-adaptation quality, we evaluate both pixel-level reconstruction fidelity and structural preservation using L_1 loss, Structural Similarity Index (SSIM), and Peak Signal-to-Noise Ratio (PSNR). As shown in Figure 12(a)–(c) and Table 1, the *Pix2Pix-LSGAN* model achieves the best overall performance, with a mean L_1 loss of 6.12 ± 0.12 , SSIM exceeding 0.69, and PSNR above 28.0 dB across all activities.

The superior performance of *Pix2Pix-LSGAN* can be attributed to its stable adversarial formulation and balanced optimization objective. The least-squares loss penalizes samples that deviate from real RD heatmaps more smoothly than the binary cross-entropy in standard GANs. In contrast, the WGAN-GP variant introduces strong gradient constraints that can suppress high-frequency doppler details, while CycleGAN, often struggles to maintain pixel-wise consistency across views. Consequently, *Pix2Pix-LSGAN* produces sharper heatmaps with a better trade-off between adversarial realism and reconstruction accuracy across various radar views.

We further analyzed the per-activity adaptation behavior to interpret how motion patterns in each activity influence view-adaptation.

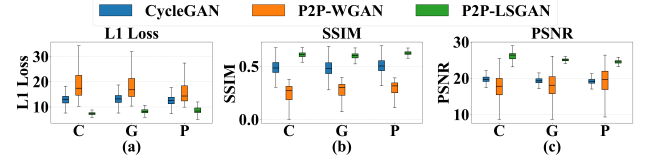


Figure 12: View adaptation performance: (a) L1-Loss, (b) SSIM, (c) PSNR (C: Chipping, G: Grinding, P: Polishing).

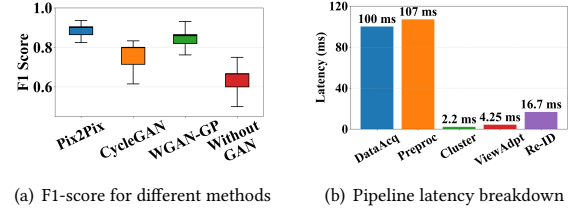


Figure 13: Performance evaluation of MIRO.

As shown in Figure 12 activities such as *grinding* yield more consistent adaptation than *polishing*, possibly as a result of their broader Doppler distributions. This indicates that GANs benefit from motion patterns with richer frequency content, as they provide stronger spatial-temporal cues for domain alignment. In contrast, *Pix2Pix*, which is primarily designed for paired image-to-image translation, captures the underlying structural patterns of RD signatures without a need for spatial correspondence. This is particularly important for RD data, where the global energy distribution over velocity and range conveys more discriminative information than exact pixel-level alignment. Overall, *Pix2Pix-LSGAN* demonstrates superior performance compared to both *CycleGAN* and *Pix2Pix-WGAN-GP*, which suffer from weaker structural consistency or over-regularization. By jointly enforcing adversarial realism and pixel-level reconstruction, *Pix2Pix-LSGAN* learns robust azimuth-invariant mappings that help for reliable cross-radar re-ID.

6.2.3 Performance of Re-Identification. We further assess the impact of different view-adaptation on multi-radar identity association. As summarized in Figure 13(a), the proposed fusion model employing *Pix2Pix-LSGAN* achieves a median re-ID F1-score of 90.4% across three radar nodes, outperforming both *Pix2Pix-WGAN-GP* (86.8%) and *CycleGAN* (80%). As a result, LSGAN-based view normalization enhances cross-radar matching by producing geometrically consistent and activity-invariant representations.

When GAN-based view adaptation is removed, performance drops by approximately 20%, highlighting the crucial role of viewpoint alignment in sustaining identity continuity. Overall, these results confirm that *MIRO*, when integrated with *Pix2Pix-LSGAN*-driven view adaptation, delivers the most stable and accurate re-ID across spatially distributed radar views.

6.2.4 Response Time Analysis. Figure 13(b) presents the end-to-end latency breakdown of the proposed system across its major processing stages. The overall response time is primarily dominated by the Data Acquisition (DataAcq) and Preprocessing (Preproc) stages, which account for 100 ms and 107 ms, respectively. The data acquisition latency is mainly constrained by hardware, as our

mmWave radar operates at 10 Hz. In contrast, the computational modules exhibit significantly lower latency. The Clustering module requires only 2.2 ms, while the View Adaptation (ViewAdapt) module completes within 4.25 ms. The final re-ID stage adds 16.7 ms. The complete system has a latency of 230 ms per prediction (excluding a 1-s initial buffer bootstrapping). Since stone-worker activity changes over minutes to hours, this latency is easily acceptable.

7 In-the-wild Evaluation

To assess the generalization of *MIRO* beyond the controlled lab setup, we deployed the pretrained view-adaptation model across unseen activities and three unseen environments (§5.3.2). We focus primarily on view-adaptation performance, as Vicon-based reference collection was infeasible in field conditions.

7.1 View-Adaptation under Unseen Activities

To further evaluate the activity invariance of the proposed view-adaptation model, we tested it on three previously *unseen actions* – *clapping*, *waving*, and *nodding*, captured from a single radar viewpoint. These activities were not part of the training dataset, which originally included only industrial gestures emulated in the lab environment (*chipping*, *grinding*, and *polishing*).

As shown in Figure 14(a), the model consistently achieves high reconstruction fidelity across all three unseen activities, demonstrating strong generalization to motion patterns beyond the training distribution. The *SSIM* values remain in the range of 0.65 – 0.75, indicating that the synthesized RD representations preserve fine structural and spatial details. Correspondingly, the *PSNR* values range between 25 dB and 35 dB, reflecting low pixel-wise distortion. Overall, these results confirm that the proposed architecture generalizes effectively across *unseen activity categories*.

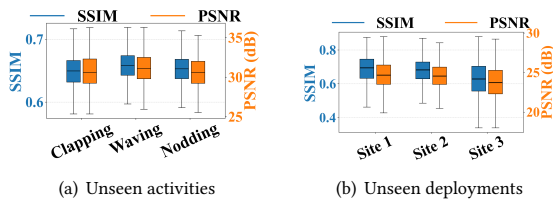


Figure 14: In-the-wild view adaptation performance.

7.2 View-Adaptation in Unseen Environments

To assess the robustness of the view-adaptation module, we evaluated its performance across three distinct deployment environments as introduced in Section 5.3.2. The same view-adaptation network, trained solely on in-lab industrial gestures, was directly applied to these unseen deployment environments without any fine-tuning. As shown in Figure 14(b), the model maintains stable reconstruction quality across all three sites, achieving *SSIM* values in the range of 0.6 – 0.8 and *PSNR* values between 22 – 30 dB. The outdoor marble processing factory (Site 1) achieves the highest *SSIM* and *PSNR* thanks to clear line-of-sight and well-separated workers. The stone-cutting factory (Site 2) shows moderate degradation due to high-speed cutting machinery, whose vibrations interfere with human motion patterns in the RD domain. The indoor construction

Table 2: Pollution exposure at various deployment locations.

Deployment	Person	Activity	PM ₁	PM _{2.5}	PM ₁₀
Marble Processing Factory	P1	Polishing	44	106	139
	P2	Chipping	161	360	439
	P3	Polishing	44	106	139
	P4	Grinding	245	603	703
	P5	Grinding	233	563	665
Stone-Cutting Factory	P1	Grinding	126	544	701
	P2	Standing	61	117	125
Indoor Construction Site	P1	Sitting	112	362	415
	P2	Standing	118	376	449
	P3	Grinding	128	633	898
	P4	Grinding	119	717	978

workspace (Site 3) experiences further challenges from strong multi-path and confined geometry, which introduce viewpoint-dependent distortions. Despite these factors, the view-adaptation model reliably preserves overall motion intensity across views.

7.3 Occupational Pollution Exposure Analysis

Table 2 and Figure 15 present the pollution exposure analysis across different activities, workers, and deployment environments. The results show distinct PM levels across activities and worksites. *Grinding* consistently produces the highest PM levels (PM_{2.5} exceeding 500 $\mu\text{g}/\text{m}^3$), often an order of magnitude higher than those from *polishing* (approximately 100 $\mu\text{g}/\text{m}^3$) or *chipping* (around 350 $\mu\text{g}/\text{m}^3$). Hence, high-friction, abrasive operations like grinding result in significantly higher dust exposure than polishing or standing idle. Figure 15 compares the interpolated PM_{2.5} fields across three representative deployment sites. Due to continuous cutting and high airflow, the outdoor yard shows PM hotspots near grinding areas, whereas the indoor site has more confined PM pockets.

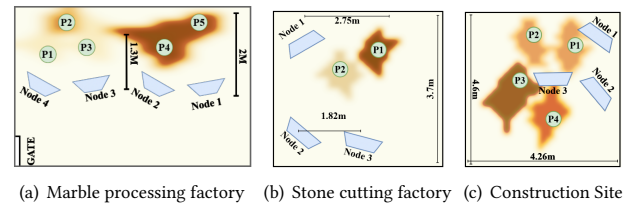


Figure 15: Interpolated PM_{2.5} exposure maps across locations.

Another critical insight from Table 2 and Figure 15 is that even passive activities can experience substantial exposure. For instance, at the indoor construction site, a worker *sitting* near a grinding area (P1) records a PM_{2.5} level of 362 $\mu\text{g}/\text{m}^3$, comparable to moderate grinding exposure. Similarly, *standing* workers near grinding or chipping tasks experience significantly elevated PM readings despite minimal personal activity. This proximity-driven exposure underscores the need to jointly consider *activity type*, *spatial position*, and *co-located operations* for occupational health risks. It is worth noting that the global PM map generation incurs an additional delay of approximately 5 s, as the PM sensors operate at 1 Hz and aggregate readings over a 5 s observation window, resulting in an overall latency of about 5.23 s (230 ms of latency in the re-ID pipeline §6.2.4).

8 Conclusion

We presented *MIRO*, a multi-radar privacy-preserving mmWave framework for personalized PM exposure monitoring in harsh industrial settings. Extensive field deployments across marble yards and construction sites demonstrate robust performance, achieving sub-10 cm localization error and 90.4% cross-radar re-ID F1-Score. The system further enables fine-grained, per-worker exposure estimation. Overall, *MIRO* demonstrates mmWave sensing as an effective means of monitoring occupational health and safety in environments where optical or wearable methods are not preferred.

The current *MIRO* framework assumes partial radar overlap and predominantly azimuthal human motion. While effective for the targeted stone-working scenarios, its generalizability may be limited by extreme sensor orientations, sparse node density, or restricted FoV overlap. Environmental factors, such as ghost reflections and dense multi-path interference, also present potential sources of association ambiguity. Resolving these geometric and interference-related challenges defines our primary direction for future research.

Ethical Considerations

All experimental procedures were approved by the Institute Ethical Committee under approval number IIT/SRIC/DEAN/2025/29, dated 25-09-2025 and conducted in accordance with its guidelines. Informed consent was obtained from all participants.

Acknowledgment

The authors acknowledge the support of the Department of Science and Technology (DST), Government of India, through the NGP Division under Grant No. NGP/GS-02/Sandip/IIT-K/WB/2023(C), dated 09-01-2024. This work has also been partially supported by the joint DST-US National Science Foundation (NSF) program under Grant No. DST/INT/USA/NSF-DST/NeTS_Sandip/P4/2024 (G), dated 10-01-2025. In addition, the authors acknowledge support from the Google Academic Research Award for Society-Centric AI for indoor pollution analysis and measurements (DALTON module).

References

- [1] 2018. DCA1000EVM. <https://www.ti.com/tool/DCA1000EVM>. [Accessed: June 27, 2026].
- [2] 2019. IWR1843BOOST. <https://www.ti.com/tool/IWR1843BOOST>. [Accessed: June 27, 2026].
- [3] Raiyaan Abdullah, Yogesh Singh Rawat, and Shruti Vyas. 2025. iSafetyBench: A video-language benchmark for safety in industrial environment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 1433–1442.
- [4] Esmail MM Abuhdima, Gurcan Comert, Pierluigi Pisu, Chin-Tser Huang, Ahmed Elqaoua, Chunheng Zhao, Shakendra Alston, Kirk Ambrose, and Jian Liu. 2021. The effect of Dust and Sand on the 5G Millimeter-Wave links. In *2021 IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE)*. IEEE, 60–65.
- [5] Team Airveda. 2017. Measuring Data - Understanding Particulate Matter and How It Impacts Our Health. <https://www.airveda.com/blog/Understanding-Particulate-Matter-and-Its-Associated-Health-Impact> [Accessed: June 27, 2026].
- [6] Farhang Akbar-Khanzadeh, Sheryl Milz, April Ames, Pamela P Susi, Michael Bisesi, Sadik A Khuder, and Mahboubeh Akbar-Khanzadeh. 2007. Crystalline silica dust and respirable particulate matter during indoor concrete grinding—wet grinding and ventilated grinding compared with uncontrolled conventional grinding. *Journal of Occupational and Environmental Hygiene* 4, 10 (2007), 770–779.
- [7] Farhang Akbar-Khanzadeh, Sheryl A Milz, Cynthia D Wagner, Michael S Bisesi, April L Ames, Sadik Khuder, Pam Susi, and Mahboubeh Akbar-Khanzadeh. 2010. Effectiveness of dust control methods for crystalline silica and respirable suspended particulate matter exposure during manual concrete surface grinding. *Journal of Occupational and Environmental Hygiene* 7, 12 (2010), 700–711.
- [8] Wael Alsafery, Omer Rana, and Charith Perera. 2023. Sensing within smart buildings: A survey. *Comput. Surveys* 55, 13s (2023), 1–35.
- [9] Kieran Browne, Ben Swift, and Terhi Nurmi-Koivu. 2020. Camera adversaria. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–9.
- [10] Xuxu Chen, Yu Zheng, Yubiao Chen, Qiwei Jin, Weiwei Sun, Eric Chang, and Wei-Ying Ma. 2014. Indoor air quality monitoring system for smart buildings. In *Proceedings of the 2014 ACM international joint conference on pervasive and ubiquitous computing*. 471–475.
- [11] Mejdi Dallel, Vincent Havaud, David Baudry, and Xavier Savatier. 2020. Inhard-industrial human action recognition dataset in the context of industrial collaborative robotics. In *2020 IEEE International Conference on Human-Machine Systems (ICHMS)*. IEEE, 1–6.
- [12] Han Ding, Chen Qian, Jinsong Han, Ge Wang, Zhiping Jiang, Jizhong Zhao, and Wei Xi. 2016. Device-free detection of approach and departure behaviors using backscatter communication. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. 167–177.
- [13] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. 1996. A density-based algorithm for discovering clusters in large spatial databases with noise. In *kdd*, Vol. 96. 226–231.
- [14] Chao Feng, Nan Wang, Yicheng Jiang, Xia Zheng, Kang Li, Zheng Wang, and Xiaojiang Chen. 2022. Wi-learner: Towards one-shot learning for cross-domain wi-fi based gesture recognition. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 3 (2022), 1–27.
- [15] Nakul Garg, Irtaza Shahid, Ramanujan K Sheshadri, Karthikeyan Sundaresan, and Nirupam Roy. 2025. Large Network {UWB} Localization: Algorithms and Implementation. In *22nd USENIX Symposium on Networked Systems Design and Implementation (NSDI 25)*. 1187–1203.
- [16] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. 2017. Improved training of wasserstein gans. *Advances in neural information processing systems* 30 (2017).
- [17] Johanne Ø Halvorsen, Pål Graff, Elin Lovise Folven Gjengedal, and Torunn K Ervik. 2025. Measurements of dust and respirable crystalline silica during indoor demolition and renovation. *Annals of Work Exposures and Health* 69, 1 (2025), 48–58.
- [18] Chong Han, Siyu Chen, Biyun Sheng, Jian Guo, and Lijuan Sun. 2025. mmReID: Person Re-Identification Based on Commodity Millimeter-Wave Radar. *IEEE Internet of Things Journal* (2025).
- [19] Jun Han, Shijia Pan, Manal Kumar Sinha, Hae Young Noh, Pei Zhang, and Patrick Tague. 2018. Smart home occupant identification via sensor fusion across on-object devices. *ACM Transactions on Sensor Networks (TOSN)* 14, 3-4 (2018), 1–22.
- [20] Katrin Hänsel, Romina Poguntke, Hamed Haddadi, Akram Alomainy, and Albrecht Schmidt. 2018. What to put on the user: Sensing technologies for studies and physiology aware systems. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [21] Qianyi Huang, Youjing Lu, Zhicheng Luo, Hao Wang, Fan Wu, Guihai Chen, and Qian Zhang. 2023. Rethinking privacy risks from wireless surveillance camera. *ACM Transactions on Sensor Networks* 19, 3 (2023), 1–21.
- [22] Cesar Iovescu and Sandeep Rao. 2020. The fundamentals of millimeter wave radar sensors. *Texas Instruments* (2020).
- [23] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1125–1134.
- [24] Lene N Johannessen, Hans Jørgen Grimstad, Jens Emil Skjetne, Ida Nord Myklebust, and Kristin V Hirsch Svendsen. 2020. Embedded systems and the Internet of Things: Can low-cost gas sensors be used in risk assessment of occupational exposure? *Journal of Occupational and Environmental Hygiene* 17, 10 (2020), 495–503.
- [25] Prasenjit Karmakar, Swadhin Pradhan, and Sandip Chakraborty. 2024. Exploring indoor air quality dynamics in developing nations: A perspective from india. *ACM Journal on Computing and Sustainable Societies* 2, 3 (2024), 1–40.
- [26] Yakup Kilic, Henk Wymeersch, Arjan Meijerink, Mark J Bentum, and William G Scanlon. 2013. Device-free person detection and ranging in UWB networks. *IEEE journal of selected topics in signal processing* 8, 1 (2013), 43–54.
- [27] Hao Kong, Cheng Huang, Jiadi Yu, and Xuemin Shen. 2024. A survey of mmwave radar-based sensing in autonomous vehicles, smart homes and industry. *IEEE Communications Surveys & Tutorials* 27, 1 (2024), 463–508.
- [28] Peizheng Li, Jagdeep Singh, Han Cui, and Carlo Alberto Boano. 2023. BmmW: A DNN-based Joint BLE and mmWave Radar System for Accurate 3D Localization. In *2023 19th International Conference on Distributed Computing in Smart Systems and the Internet of Things (DCOSS-IoT)*. IEEE, 47–54.
- [29] Yilong Li, Ramanujan K Sheshadri, Karthik Sundaresan, Eugene Chai, Yijing Zeng, Jayaram Raghuram, and Suman Banerjee. 2025. Medusa: Scalable Multi-View Biometric Sensing in the Wild with Distributed MIMO Radars. In *Proceedings of the 31st Annual International Conference on Mobile Computing and Networking*. ACM, 78–92.

- [30] Hankai Liu, Xiulong Liu, Xin Xie, Xinyu Tong, and Keqiu Li. 2024. Pmtrack: Enabling personalized mmwave-based human tracking. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 7, 4 (2024), 1–30.
- [31] Liang Liu, Wu Liu, Yu Zheng, Huadong Ma, and Cheng Zhang. 2018. Third-eye: A mobilephone-enabled crowdsensing system for air quality monitoring. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 1 (2018), 1–26.
- [32] Ruofeng Liu, Tianshun Yao, Ruili Shi, Luoyu Mei, Shuai Wang, Zhimeng Yin, Wencho Jiang, and Shuai Wang. 2024. Mission: mmwave radar person identification with rgb cameras. In *Proceedings of the 22nd ACM Conference on Embedded Networked Sensor Systems*. 309–321.
- [33] Yiheng Liu, Wengang Zhou, Mao Xi, Sanjing Shen, and Houqiang Li. 2020. Vision meets wireless positioning: Effective person re-identification with recurrent context propagation. In *Proceedings of the 28th ACM International Conference on Multimedia*. 1103–1111.
- [34] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, Nov (2008), 2579–2605.
- [35] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. 2017. Least squares generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision*. 2794–2802.
- [36] Jingke Meng, Sheng Wu, and Wei-Shi Zheng. 2019. Weakly supervised person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 760–769.
- [37] Hessam Mohammadoradi, Shengrong Yin, and Omprakash Gnawali. 2017. Room occupancy estimation through WiFi, UWB, and light sensors mounted on doorways. In *Proceedings of the 2017 International Conference on Smart Digital Environment*. 27–34.
- [38] Ramon Nitzberg. 1972. Constant-false-alarm-rate signal processors for several types of interference. *IEEE TAES* (1972), 27–34.
- [39] Amashi Niwarthana, Pamuditha Somarathne, Pierre Qian, Ken-Tye Yong, and Anusha Withana. 2024. Efficient and Robust Heart Rate Estimation Approach for Noisy Wearable PPG Sensors Using Ideal Representation Learning. In *Proceedings of the 2024 ACM International Symposium on Wearable Computers*. 1–8.
- [40] G Daryl Nord, Tipton F McCubbins, and Jeretta Horn Nord. 2006. E-monitoring in the workplace: privacy, legislation, and surveillance software. *Commun. ACM* 49, 8 (2006), 72–77.
- [41] Savvas Papaioannou, Andrew Markham, and Niki Trigoni. 2016. Tracking people in highly dynamic industrial environments. *IEEE Transactions on mobile computing* 16, 8 (2016), 2351–2365.
- [42] Sandeep Rao. 2017. Introduction to mmWave sensing: FMCW radars. *Texas Instruments (TI) mmWave Training Series* (2017), 1–11.
- [43] S Rao. 2020. Introduction to mmwave radar sensing: Fmcw radars. *Texas Instruments* (2020), 1–70.
- [44] Yili Ren, Yichao Wang, Sheng Tan, Yingying Chen, and Jie Yang. 2022. Person re-identification using wifi signals. In *Proceedings of the 28th annual international conference on mobile computing and networking*. 829–831.
- [45] Mikko Rinta-Homi, Naser Hossein Motlagh, Agustin Zuniga, Huber Flores, and Petteri Nurmi. 2021. How low can you go? performance trade-offs in low-resolution thermal sensors for occupancy detection: A systematic evaluation. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 3 (2021), 1–22.
- [46] Argha Sen, Anirban Das, Swadhin Pradhan, and Sandip Chakraborty. 2024. Continuous multi-user activity tracking via room-scale mmwave sensing. In *2024 23rd ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE, 163–175.
- [47] Yang Shen, Weiyao Lin, Junchi Yan, Mingliang Xu, Jianxin Wu, and Jingdong Wang. 2015. Person re-identification with correspondence structure learning. In *Proceedings of the IEEE international conference on computer vision*. 3200–3208.
- [48] Feng-Chiao Su, Melissa C Friesen, Aleksandr B Stefaniak, Paul K Henneberger, Ryan F LeBouf, Marcia L Stanton, Xiaoming Liang, Michael Humann, and M Abbas Virji. 2018. Exposures to volatile organic compounds among healthcare workers: modeling the effects of cleaning tasks and product use. *Annals of Work Exposures and Health* 62, 7 (2018), 852–870.
- [49] Kai Sun, Shuai Wang, Renjie Zhao, Ruofeng Liu, Weiwei Chen, Zhimeng Yin, Wei Gong, and Shuai Wang. 2025. Wuloc: Achieving extremely long-range high-precision localization via wi-fi-uwband connection. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 9, 1 (2025), 1–24.
- [50] Mingyue Tang, Pranshu Teckchandani, Jizheng He, Hanbo Guo, and Elahe Soltanaghai. 2024. Bsense: In-vehicle child detection and vital sign monitoring with a single mmwave radar and synthetic reflectors. In *Proceedings of the 22nd ACM Conference on Embedded Networked Sensor Systems*. 478–492.
- [51] Dequan Wang, Xinran Zhang, Kai Wang, Lingyu Wang, Xiaoran Fan, and Yanyong Zhang. 2024. Rdgait: A mmwave based gait user recognition system for complex indoor environments using single-chip radar. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 8, 3 (2024), 1–31.
- [52] Yingqi Wang, Min Xu, Zhongqin Wang, Yongze Wang, and J Andrew Zhang. 2025. User Re-Identification Through mmWave Radio Imaging. *IEEE Internet of Things Journal* (2025).
- [53] Teng Wei and Xinyu Zhang. 2015. mtrack: High-precision passive tracking using millimeter wave radios. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*. 117–129.
- [54] Di Wu, Tao Xiao, Xuewen Liao, Jie Luo, Chao Wu, Shigeng Zhang, Yong Li, and Yike Guo. 2020. When sharing economy meets IoT: Towards fine-grained urban air quality monitoring through mobile crowdsensing on bike-share system. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 2 (2020), 1–26.
- [55] Qingxin Xia, Joseph Korpela, Yasuo Namioka, and Takuya Maekawa. 2020. Robust unsupervised factory activity recognition with body-worn accelerometer using temporal structure of multiple sensor data motifs. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 3 (2020), 1–30.
- [56] Naoya Yoshimura, Jaime Morales, Takuya Maekawa, and Takahiro Hara. 2024. Openpack: A large-scale dataset for recognizing packaging works in IoT-enabled logistic environments. In *2024 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, 90–97.
- [57] Chunshan Zhou, Shijie Li, and Shaojian Wang. 2018. Examining the impacts of urban form on air pollution in developing countries: A case study of China's megacities. *International journal of environmental research and public health* 15, 8 (2018), 1565.
- [58] Rui Zhou, Songlin Li, Hongwang Zhang, Chenxu Liu, and Jiajun Sun. 2025. mmMulti: Multi-person Action Recognition Based on Multi-task Learning Using Millimeter Waves. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 9, 2 (2025), 1–25.
- [59] Jiangyifei Zhu, Kuang Yuan, Akarsh Prabhakara, Yunzhi Li, Gongwei Wang, Kelly Michaelson, Justin Chan, and Swarun Kumar. 2025. Measuring multi-site pulse transit time with an AI-enabled mmWave radar. *arXiv preprint arXiv:2510.18141* (2025).
- [60] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*. 2223–2232.