

MMVR: Millimeter-wave Multi-View Radar Dataset and Benchmark for Indoor Perception

M. Mahbubur Rahman^{*,†}, Ryoma Yataka^{*,‡}, Sorachi Kato^{*,†}, Pu (Perry) Wang^{*,‡}, Peizhao Li[†], Adriano Cardace[†], Petros Boufounos[‡]

Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA 02139

Abstract. Compared with an extensive list of automotive radar datasets that support autonomous driving, indoor radar datasets are scarce at a smaller scale in the format of low-resolution radar point clouds and usually under an open-space single-room setting. In this paper, we scale up indoor radar data collection using multi-view high-resolution radar heatmap in a multi-day, multi-room, and multi-subject setting, with an emphasis on the diversity of environment and subjects. Referred to as the millimeter-wave multi-view radar (**MMVR**) dataset, it consists of 345K multi-view radar frames collected from 25 human subjects over 6 different rooms, 446K annotated bounding boxes/segmentation instances, and 7.59 million annotated keypoints to support three major perception tasks of object detection, pose estimation, and instance segmentation, respectively. For each task, we report performance benchmarks under two protocols: a single subject in an open space and multiple subjects in several cluttered rooms with two data splits: random split and cross-environment split over 395 1-min data segments. We anticipate that MMVR facilitates indoor radar perception development for indoor vehicle (robot/humanoid) navigation, building energy management, and elderly care for better efficiency, user experience, and safety. The MMVR dataset is available at <https://doi.org/10.5281/zenodo.12611978>.

Keywords: Indoor perception · radar · benchmark · open dataset

1 Introduction

Compared with popular automotive radar perception efforts and datasets [3–5, 12, 14, 15, 17–21, 23, 25, 28, 32], indoor radar perception receives less attention but it is essential in applications such as indoor vehicle (robot/humanoid) navigation, building energy management, and elderly care under low light and emergence (smoke, fire, and dust) situations with benefits of low device cost and less privacy concerns. As a result, its development lags and the open datasets are limited in terms of size, annotation, tasks, and benchmarks.

*: Equal contribution. †: The work of M. Rahman (Univ. of Alabama, USA), S. Kato (Osaka Univ., Japan), P. Li (Brandeis Univ., USA), and A. Cardace (Univ. of Bologna, Italy) was done during their internship at MERL. ‡: The work was done as a visiting scientist from Mitsubishi Electric Corporation, Japan. †: Project Lead.

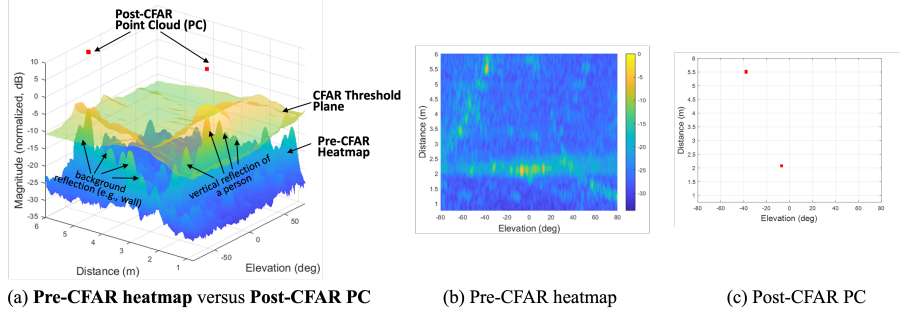


Fig. 1: Radar signal representations: pre-CFAR heatmap versus post-CFAR point cloud (PC) corresponding to a scene in our MMVR dataset. The heatmap ((a) for a 3D view and (b) for the top-down view) shows an extended vertical profile of a subject in terms of multiple clustered reflections over the elevation (height) domain. In contrast, the CFAR operation compares the heatmap with a threshold plane (the semi-transparent surface in (a)) to declare a few detection points (red squares in (a) and (c)), greatly suppressing weaker reflections from the subject.

A few lines of effort have explored radar signals and processed results for indoor perception. The earliest work is RF-Pose [33] using a T-shape antenna array [1] (16 horizontal and 4 vertical antennas) in the 5 – 7 GHz band. It takes two orthogonal radar heatmaps generated by the T-shape array as the input, uses a convolution-based autoencoder network to fuse features from the two radar views, and regresses multi-person keypoints for 2D image-plane pose estimation. It is noted that RF-Pose is not publicly accessible. More recently, HuPR has made a similar effort collecting two orthogonal LR radar heatmaps for action classification in an empty room (see Fig. 2 (a)) and pose estimation and publicly releases the dataset [13]. Due to the fact that their angular (azimuth and elevation) resolution is limited to 15° [13,33], we refer to RF-Pose and HuPR as **multi-view low-resolution (LR) radar heatmap** datasets in Table 1.

Table 1: Indoor Radar Perception Datasets.

Datasets	Year	Sensor	Views	Data	Tasks	Size	Public
RadHAR [26]	2019	LR (low-resolution)	Single PC (Point Cloud)	Action	Action	167K	✓
mm-Pose [24]	2020	LR	Multi	PC	Action, Pose	40K	✗
mmMesh [30]	2021	LR	Single	PC	3D Mesh	480K	✗
mRI [2]	2022	LR	Single	PC	Action, Pose	160K	✓
MM-Fi [31]	2023	LR	Single	PC	Action, Pose	320K	✓
RF-Pose [33]	2018	LR	Multi	Heatmap	Pose	-	✗
HuPR [13]	2023	LR	Multi	Heatmap	Pose	141K	✓
HIBER [29]	2023	HR (high-resolution)	Multi	Heatmap	Box, Pose, Seg.	179K	Partial
MMVR(ours)	2024	HR	Multi	Heatmap	Box, Pose, Seg.	345K	✓

LR: low-resolution radar with an angular resolution of 15° .

HR: high-resolution radar with an angular resolution of 1.3° .



Fig. 2: Snapshots of indoor radar **heatmap** datasets in Table 2. HuPR [13] and HIBER [29] were collected at a single location and, respectively, multiple locations inside a single open-space/open-foreground room. Our MMVR was collected in multiple locations over multiple open-foreground (d1-d4) and cluttered (d5-d9) rooms.

Following RF-Pose but with commercial radar sensors, other earlier efforts use TI’s single-chip millimeter-wave (mmWave) radar consisting of 3 transmitters and 4 receivers in the 60 – 64 GHz and 76 – 81 GHz bands. These single-chip radar sensors achieve a comparable angular resolution of around 15° as the RF-Pose due to a lower wavelength. Other than the multi-view radar heatmaps, indoor radar datasets such as RadHAR [26], mm-Pose [24], mRI [2], and more recently, MM-Fi [31] collect **single-view LR point cloud (PC)**, usually along the horizontal orientation, as shown in the first block of Table 1.

The difference between PC and heatmap may have significant consequences on the downstream tasks such as the image-plane pose estimation and segmentation. This is illustrated in Fig. 1, where radar PC is obtained by applying the constant false alarm rate (CFAR) operation [22] to the radar heatmap for a scene with a person moving at a distance of 2 m in our MMVR dataset. The heatmap ((a) for a 3D view and (b) for the top-down view) shows an extended vertical profile of a subject in terms of multiple clustered reflections over the elevation (height) domain. In contrast, simple CFAR operations compare the heatmap with an adaptively determined threshold plane (the semi-transparent surface in (a)) to declare a few detection points (red squares in (a) and (c)), greatly suppressing weaker reflections from the subject and missing fine-grained features for challenging tasks. As shown in Table 1, these single-view LR point cloud datasets are mainly used for tasks such as action classification [2, 24, 26, 31].

To enable more fine-grained radar feature extraction, **multi-view high-resolution (HR) heatmaps** have been considered more recently for indoor radar perception such as HIBER [29]. However, HIBER was collected at multiple locations within the same room with open foreground (see Fig. 2 (b)). And there is no natural occlusion (table, chair, furniture) in front of subjects. To advance indoor radar perception in more complex environments (multiple rooms with cluttered space) and motivate challenging indoor downstream tasks using radar signals, we introduce a large-scale indoor radar perception dataset: Millimeter-wave Multi-View Radar (**MMVR**) with the following main contributions:

- First, our dataset MMVR has **345K** data frames collected from **25** human subjects over **6** different rooms (e.g, open/cluttered offices and meeting

- rooms) spanning over **9** separate days. To the best of our knowledge, this is the largest open-source indoor radar dataset in a truly **multi-day**, **multi-room**, and **multi-subject** setting.
- Second, MMVR consists of 2 parts: 1) 107.9K data frames of **P1: Open Foreground** in a single open-foreground space with a single subject; see the first row of Fig. 2 (c) for snapshots; and 2) 237.9K data frames of **P2: Cluttered Space** in 5 cluttered rooms with multiple subjects; see the second and third rows of Fig. 2 (c). **P1** is used to establish the best possible radar perception benchmarks, while **P2** is designed for more challenging scenarios and for cross-environment and cross-subject generalization.
 - Third, MMVR has annotated about **446K** bounding boxes, **7.59 million** keypoints, and **446K** segmentation instances (see Table 2 for detailed dataset statistics). We leverage state-of-art RGB-based pipelines, i.e., Mask2Former [9] and HRNets [27], to generate high-confidence image-plane annotation labels (bounding boxes, keypoints, and segmentation pixels) with human curation involved to support three perception tasks: 1) object detection, 2) pose estimation, and 3) instance segmentation, respectively.
 - Finally, we re-implement or modify state-of-art radar perception baseline methods: RF-Pose [33] and RFMask [29], for all three tasks and establish benchmarks under both **P1** and **P2**. Ablation studies show the impact of factors such as the number of input frames, the number of radar views, and different data splits on all considered perception tasks.

2 Related Work: Indoor Radar Perception Datasets

As shown in Table 1, the last few years have witnessed the release of several indoor radar perception datasets. Most datasets focus on exploring single-view LR point clouds for tasks such as action classification and pose estimation. Rad-HAR [26] collected a dataset consisting of point clouds using mmWave radar to classify only 5 human activities from 2 subjects. mm-Pose [24] used the CNN as the backbone network to estimate and track 25 body joints when the subject performed one of the 4 actions: walking, left-arm swing, right-arm swing, and both-arm swing. The dataset is rather limited with 39.7K frames in total. mmMesh [30] used single-view LR radar point clouds to directly estimate realistic-looking human meshes with encoded human body models to overcome

Table 2: Comparison of multi-view heatmap-based indoor radar datasets.

Dataset	Resolution			Annotations			Statistics					
	Range	Azi.	Ele.	BBox	KP	Seg	Rooms	Subjects	Actions	Sequences	Frames	Room Setting
RF-Pose [†]	10.0 cm	15°	15°	×	✓	×	50	100	free-form	-	-	-
HuPR	4.8 cm	15°	15°	141K	1.97M	×	1	6	static, walking	235	141K	open space
HIBER	12.2 cm	1.3°	1.3°	231K	3.23M	231K	1	10	free-form	152	179K	open foreground
MMVR(ours)	11.5 cm	1.3°	1.3°	446K	7.59M	446K	6	25	free-form	395	345K	open & cluttered

[†]RF-Pose is not publicly available.

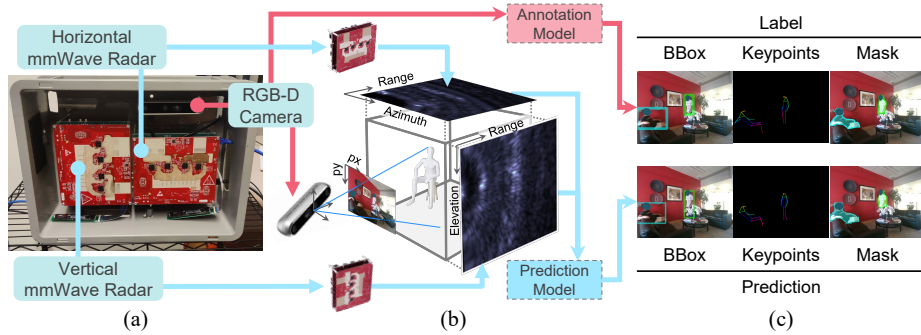


Fig. 3: MMVR sensor setup (a) with two mmWave radars and one RGB-D camera; (b) each frame includes two orthogonal high-resolution radar heatmaps and the synchronized RGB image; (c): the two radar heatmaps are inputs to radar perception models for the three perception tasks with the supervision from RGB-based labels.

the sparsity of point clouds. It collected radar data on 8 daily activities from 20 subjects within a confined open-space area, yielding 480K radar frames in total. Both mm-Pose and mmMesh datasets are not publicly accessible. More recently, mRI [2] and MM-Fi [31] targeted indoor rehabilitation activities and related pose estimation (14 daily activities and 13 rehabilitation exercises) from multi-modal sensors with single-view mmWave radar point cloud included.

RF-Pose [33] and HuPR [13] are the two datasets leveraging two orthogonal radar heatmaps with an angular (azimuth and elevation) resolution of 15° . While RF-Pose was comprehensive in terms of data size and the diversity of subjects (100) and rooms (50), it is not publicly accessible. On the other hand, HuPR released 141K data frames with 2D bounding boxes and keypoint annotations. However, as shown in Table 2, the diversity of rooms, subjects, and actions is rather limited and no segmentation annotations are included. In 2023, HIBER partially released 179K data frames collected from 10 subjects in a single room with an open foreground, using two orthogonal high-resolution radar heatmaps. To bridge the gap between open-source HIBER and inaccessible RF-Pose, our dataset MMVR emphasizes the diversity of rooms, subjects and tasks, setting a stage for cross-environment and cross-subject performance evaluation and advancing the generalization capability of radar perception models. Moreover, our dataset has almost double the size of annotation labels than that of HIBER.

3 MMVR Sensor Setup

For MMVR, we develop a portable sensor testbed with two high-resolution radar sensors and a paired RGB-D camera in Fig. 3 (a). Following the red line, we use the Intel RealSense D455 sensor to obtain images of the scene in Fig. 3 (b) and generate image-based annotation labels in Fig. 3 (c) via pretrained annotation models. Separately, following the light blue line, we use two TI AWR2243 mmWave cascade radars to generate two radar views in the range-azimuth and

range-elevation domains in Fig. 3 (b), and both views are fed to baseline prediction models to output prediction results for three perception tasks in Fig. 3 (c).

3.1 RGB-D Camera

Thanks to the small form factor of the RealSense camera, we place it on the upper right side of our testbed box, just right above the horizontal radar sensor. RealSense camera D455 provides up to 1280×800 resolution for RGB images and up to 1280×720 resolution for stereo depth. The depth operating range is about 6 m (varies with lighting conditions). In our data collection, we choose the image resolution of 480×640 with a frame rate of 15 fps. As shown in Fig. 3 (b), the RGB image can be projected from the 3D camera coordinate into the image plane using a pinhole camera model. The timestamp of RGB-D images is synchronized to the connected desktop with a sample timestamp accuracy of 50 microseconds.

3.2 Multi-View High-Resolution Radar

To achieve high resolution in both azimuth and elevation domains, we place two AWR2243 cascade radar sensors along the horizontal and vertical orientations in Fig. 3 (a). Coherently combining 4 single-chip FMCW radar sensors, the cascade radar forms a virtual array of 86 half-wavelength-spaced elements and offers an angular resolution of 1.3° at the boresight direction¹, significantly better than the angular resolution of 15° offered by the single-chip radar chips, e.g., IWR1443 and IWR1843, used by RadHAR [26], mRI [2], MM-Fi [31], and HuPR [13]. With the perpendicular radar configuration, the two radars achieve simultaneously high-resolution radar views in both the range-azimuth and range-elevation planes, as illustrated in Fig. 3 (b). Fig. 4 compares the range-azimuth heatmap of a corner reflector using the high-resolution (AWR2243) and low-resolution (IWR1843) radar sensors.




Fig. 4: Comparison between (a) low-resolution and (b) high-resolution radar heatmaps of a corner reflector.

3.3 Calibration

The camera-radar coordinate calibration is performed using measured 3D positions of a corner reflector at multiple locations in the camera and, respectively, radar coordinate systems. Denoting the measured camera-coordinate and radar-coordinate positions as $\mathbf{B}_{\text{camera}}, \mathbf{A}_{\text{radar}} \in \mathbb{R}^{3 \times N}$, where N is the number of calibration positions, the radar-camera calibration is done by finding a rotation

¹ The boresight direction is the direction of peak gain of the antenna array.

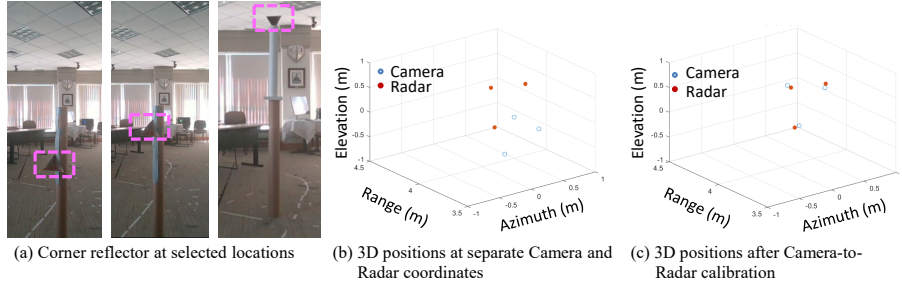


Fig. 5: Camera-Radar coordinate calibration.

matrix \mathbf{R} and a translation vector \mathbf{t} such at

$$\min_{\mathbf{R}, \mathbf{t}} \sum_{i=1}^N \|\mathbf{R}\mathbf{B}_{\text{camera}}[:, i] + \mathbf{t} - \mathbf{B}_{\text{radar}}[:, i]\|_2^2 \quad (1)$$

The above calibration can be solved using the singular value decomposition (SVD) to align two sets of points in the 3D space. Refer to Appendix D for the detailed calibration steps. Fig. 5 shows the corner reflector at multiple calibration locations (a) and its positions in the camera and radar coordinate systems before (b) and after (c) the camera-radar coordinate calibration.

4 Dataset

MMVR supports a range of perception tasks: object detection, pose estimation, and instance segmentation under a diverse setting of 6 rooms and 25 subjects.

4.1 Data Collection

MMVR consists of 345K data frames collected in 2 protocols:

- **Protocol 1 (P1: Open Foreground)**: 107.9K data frames in a single open-foreground room with a single subject; see Fig. 6 (a) for the floorplan of **P1** and the first row of Fig. 2 (c) for a snapshot. These data were collected over 4 separate days (d1-d4) with one or two sessions per day. The subject walking and jumping in the space remains unobstructed to both radar and RGB camera observations.

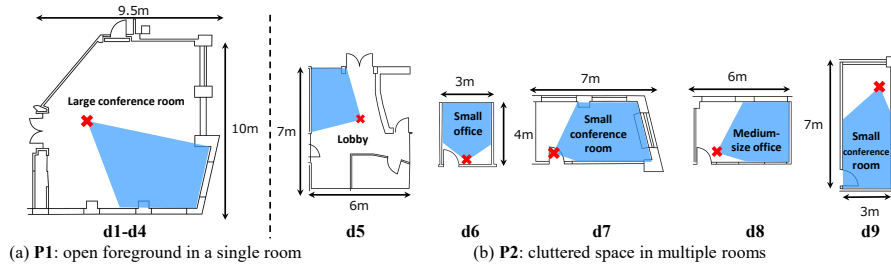


Fig. 6: Floorplans of 6 diverse room settings for **Protocol 1** and **Protocol 2**.

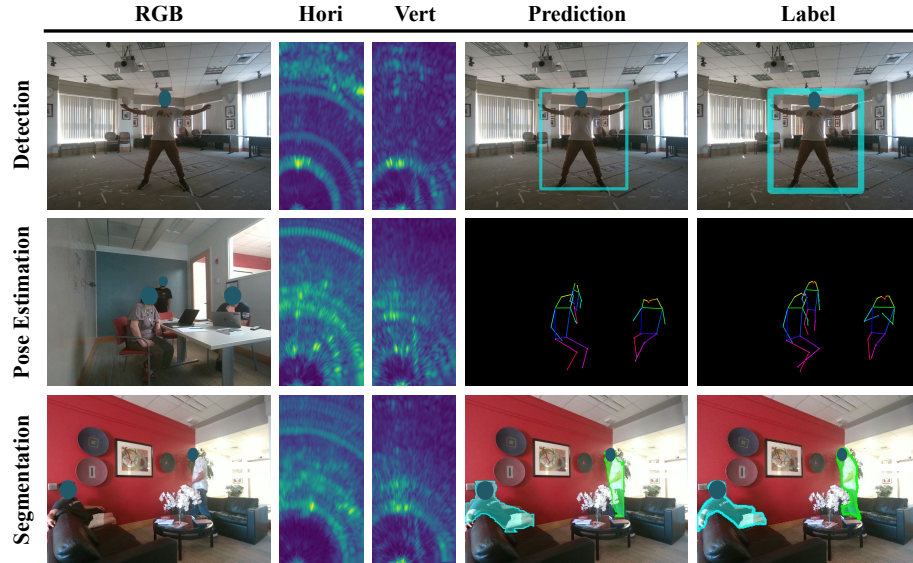


Fig. 7: Visualization of data collection (**P1: open foreground** in the first row and **P2: cluttered space** in the second and third rows), RGB (first column), corresponding multi-view (horizontal and vertical) radar heatmaps (second and third columns), baseline prediction results (fourth columns), and annotation labels (fifth columns).

- **Protocol 2 (P2: Cluttered Space):** 237.8K data frames in 5 cluttered rooms with single and multiple subjects; see Fig. 6 (b) for the floorplans of **P2** and the second and third rows of Fig. 2 (c) for snapshots. Starting from Day 5 (d5), 6 data sessions were collected in one room. During the data collection, the subjects were doing diverse activities such as walking, sitting, stretching, reading, writing on the board, and having conversations. Additional snapshots for all data sessions are provided in Appendix E.

We split all data frames into non-overlapping 1-min data segments, each having about 900 data frames, given the frame rate of 15 fps. Each data frame includes one RGB frame, one horizontal radar heatmap frame, one vertical radar heatmap frame, bounding box labels, segmentation labels, and keypoint labels. In total, we have 122 data segments for **P1** and, respectively, 273 data segments for **P2**. More details about the data collection protocol, number of segments in each data session, occlusion, and activities can be found in Table 8 of Appendix F.

4.2 Annotation

We utilize pretrained image-based deep learning models: Mask2Former [9] and HRNet [27], to generate annotation labels.

Object Detection We take our RGB image through a pretrained backbone, i.e., the Swin Transformer, to get a list of low-resolution feature maps. Then




Fig. 8: Annotation curation: remove spurious labels (magenta boxes) due to the human presence behind the glass, highly reflective whiteboards, and a hanging suit.

Mask2Former [8, 9] enhances these features using a pixel decoder module to get high-resolution features. The Mask2Former decoder takes in a set of input queries, cross-attends the queries to feature maps from the encoder, and transforms them into a set of bounding boxes and classes. For our dataset, we only extract the bounding boxes with the “person” class with a shape of $(n, 5)$ for n objects and $[x_{\min}, y_{\min}, x_{\max}, y_{\max}, \text{confidence score}]$ for each object. One example is shown in the last column (Label) of the first row of Fig. 7, where $n = 1$.

Pose Estimation HRNet [10, 27] maintains high-resolution representations by connecting high-to-low resolution convolutions in parallel, where there are repeated multiscale fusions across parallel convolutions. Such strong high-resolution representations can support multi-person semantic keypoint extraction. For each RGB frame, we applied a HRNet pretrained with the COCO Keypoints 2017 dataset to extract 17 COCO keypoints for each person. The shape of keypoint labels is given as $(n, 17, 3)$ for n objects and each keypoint represented by 3 elements: $(\text{pixel}_{\text{height}}, \text{pixel}_{\text{width}}, \text{confidence score})$. One example is shown in the last column (Label) of the second row of Fig. 7, where $n = 3$.

Instance Segmentation Mask2Former [8, 9] extends the above object detection to instance segmentation by further regressing the set of input queries into a set of binary mask and class predictions, conditioned on the pixel decoder’s features. The format of the instance segmentation is a $(n, 480, 640)$ binary (Boolean) tensor, where 480 and 640 are the height and width of the RGB image, respectively. One example is shown in the last column (Label) of the third row of Fig. 7, where $n = 2$ for two subjects in the RGB image.

Annotation Curation Although the above annotation models facilitate automatic label generation at a high volume, spurious labels still exist. For instance, in the lobby setting of d5 in Fig. 8, RGB-based annotation pipelines generate high-confident annotation labels for passing-by people outside of the lobby (behind the glass), while radar sensors operating at 77 GHz have limited penetration capability through the glass. As a result, we remove these spurious annotation labels from frames when




Fig. 9: Quality of curated annotations.

there are people passing by. One can find other cases in offices where a highly reflective whiteboard is present (see d6 and d7 of Fig. 8) and where there is a hanging suit (see d8 of Fig. 8). After data curation, Fig. 9 shows the histograms of confidence scores of the annotated bounding box (BBBox) and keypoints. Please refer to Appendix G for other statistics of annotation in our dataset.

4.3 Preprocessing for Multi-View Radar Heatmaps

Fig. 10 shows our multi-view radar heatmap preprocessing chart from the two radar sensors forming two orthogonal virtual arrays of 86 half-wavelength-spaced elements while sending multiple pulses. By sampling the returned pulses, one can collect a 3D data cube along (horizontal/vertical) virtual array, ADC samples (intra-pulse or fast-time), and pulse (inter-pulse or slow-time) samples. By applying 3D fast Fourier transform (FFT) over the data-cube, one can obtain radar spectrum over the angle (azimuth for horizontal radar and elevation for vertical radar), range, and Doppler velocity domains. We further integrate the 3D radar spectrum along the Doppler domain to enhance the SNR, generating two (range-azimuth and range-elevation) radar heatmaps in the radar polar coordinate. We further project the two radar heatmaps into the radar Cartesian coordinate system and, with the help of the radar-camera calibration of Sec. 3.3, convert them into range-azimuth and range-elevation views in the camera Cartesian coordinate system.




Fig. 10: Multi-view heatmap preprocessing.

4.4 Synchronization

For synchronization between radar and camera, we use two alignment steps: an initial alignment and a refined alignment. For the initial alignment, we place a metal corner reflector in front of our MMVR sensor board at the beginning of each session such that we can identify the starting frame for the horizontal radar sensor, vertical radar sensor, and RGB camera; see Appendix H for an illustration. Even with the initial alignment, the clocks of the three sensors may drift over time. To this end, we perform a refined alignment to adjust the frames of all three sensors by synchronizing keyframes where notable motions, e.g., sit, stand up, are aligned at all three sensors.

5 Evaluation and Benchmarks

In this section, we introduce the benchmark setup, performance metrics, baseline methods, and main results.

5.1 Benchmark Setup

To evaluate the model, we provide two data split strategies.

- **Data Split 1 (S1: Random Split)** randomly splits the non-overlapping 1-min data segments (122 for **P1** and 273 for **P2**) into train, validation and test sets at a ratio of 80 : 10 : 10.
- **Data Split 2 (S2: Cross-Session and Unseen Split)** first splits all data segments in d5, d6, d7, and d9 into train, validation, and test sets. Then, we include all data in d8 in the test set such that one can assess the generalization performance of trained model for an unseen environment (d8).

The details about **S1** and **S2** can be found in Table 9 of Appendix I.

5.2 Metrics and Methods

Object Detection We adopt average precision on intersection over union (IoU) [11] as an evaluation metric. IoU is the ratio of the overlap to the union of a predicted BBox A and annotated BBox B as:

$$\text{IoU}(A, B) = \frac{|A \cap B|}{|A \cup B|}. \quad (2)$$

We present three variants of average precision: AP_{50} , AP_{75} , and AP , where the former two represent the loose and strict constraints of IoU, while AP is the averaged score over 10 different IoU thresholds in $[0.5, 0.95]$ with a stepsize of 0.05. We use the latest RFMask [29] as the baseline method but modify it such that we can directly calculate the BBox loss using the BBox annotation in the image plane rather than the BBox loss at the two radar-view planes in HIBER. More details are described in Appendix J.

Pose Estimation We adopt average precision on object keypoint similarity (OKS) [16] as an evaluation metric. OKS is scale-invariant and defined as


$$\text{OKS}_i = \exp\left(-\frac{d_i^2}{2s^2k_i^2}\right), \quad i = 1, \dots, 17, \quad (3)$$

for the i -th keypoint, where d_i is the Euclidean distance between the prediction and corresponding ground truth, k_i is a constant predefined uniquely for each joint, and s is the scale of the subject being targeted. Similarly, we calculate AP_{50} , AP_{75} , and AP for each keypoint and the average score over the 17 keypoints. We re-implement RF-Pose [33] from scratch for the baseline and further extend it to predict both keypoint-based heatmap and part affinity field (PAF) [6] for more robust multi-person association.

Instance Segmentation Instance masks are predicted by RFMask, and average precision over IoU is used as the evaluation metric. The segmentation IoU is computed by Eq. 2, where the area of interest is determined by pixels, instead of BBoxes, that are predicted or annotated to belong to a person [11].

Table 3: Baseline benchmarks for 3 tasks under 2 protocols and 2 data splits.

Protocol	Split	Object Detection			Pose Estimation			Instance Seg.
		AP	AP ₅₀	AP ₇₅	AP	AP ₅₀	AP ₇₅	IoU
P1	S1	25.53	67.30	15.86	46.24	62.88	47.45	71.98
	S2	24.46	66.82	11.22	29.82	43.03	30.29	67.03
P2	S1	31.37	61.50	27.48	32.13	44.22	32.58	65.30
	S2	6.03	22.77	0.88	7.11	11.98	6.76	56.07

**Fig. 11:** AP breakdown over keypoints under 2 protocols and 2 data splits.

5.3 Main Results

Object Detection Table 3 shows the baseline performance of object detection from RFMask by taking 4 consecutive radar frames as input. It achieves AP₅₀ scores at more than 65 under P1 for both data splits. These AP scores of our dataset MMVR are comparable to those numbers previously reported in the literature. For instance, RF-Pose [33] reported an AP score at 30.8 with 6 radar frames, which is similar to our benchmarks of AP at 25.53 and 24.46 under P1. Qualitatively, we observe more failure cases when the subject was about 1 m or closer to the radar sensors.

For P2, the performance under S1 is much better than that of S2 as there is a significant drop in all AP metrics, suggesting limited generalization capability of the baseline method, particularly towards unseen cluttered environments.

Pose Estimation Table 3 also consists of the baseline performance of pose estimation using RF-Pose with 4 consecutive radar frames. Under P1, it achieves AP₅₀ of about 63 and 43 for S1 and S2, respectively. On the other hand, under P2, it maintains reasonable performance for S1 with an AP₅₀ score at 44.22, while the performance drops significantly when the data split is S2, a similar effect we observed previously for object detection. Fig. 11 further breakdowns the AP score for multiple keypoints. It is seen that lower AP scores are more likely for keypoints such as limb and face parts.

Instance Segmentation In the last column of Table 3, we show relatively good and consistent performance of RFMask over the four combinations of 2 protocols and 2 data splits. On one hand, under the same protocol (P1 or P2), the performance drops slightly from S1 to S2, suggesting better generalization capability

Table 4: # of Frames: Using longer time horizon improves the performance.

# of frames	Detection			Pose Estimation			Segmentation
	AP	AP ₅₀	AP ₇₅	AP	AP ₅₀	AP ₇₅	IoU
4	25.53	67.30	15.86	46.24	62.88	47.45	71.98
12	33.12	71.17	26.84	46.86	63.22	48.14	71.81
18	32.78	72.03	25.44	46.70	63.07	47.95	73.28

over unseen environments. On the other hand, performance degradation appears to be slightly larger over protocols than over splits.

Visualization We provide visualization results for all three tasks in the column of Prediction of Fig. 7. In these snapshots, the considered baseline methods are able to distinguish radar signals from the subject from those from the background. Additional visualization results of good and failure cases under all 6 rooms are included in Appendix K.

5.4 Ablation Study

We report on several ablation studies using **S1** under **P1**.

Number of Frames From Table 4, longer sequences generally lead to better performance for each task. In particular, the object detection performance is increased significantly when 12 radar frames are used. At a frame rate of 15 fps, 18 frames for a duration of 1.2 seconds lead to slightly worse performance than the use of 12 frames.

Single View versus Dual View Table 5 presents a comparison between using a single view (feeding the horizontal view only), and using dual views (feeding both the horizontal and vertical views). The results indicate that using the dual views improves performance for both RFMask and RF-Pose. This improvement implies that, while the horizontal view is informative enough to determine the location and relative scale of the subjects in the image plane, the vertical view may provide additional features about a subject in the elevation domain and add more fine-grained and distinguished characteristics for different body parts.

Comparison of Detection Backbones We also take the more advanced DETR architecture [7] as another simple baseline method without significant customization towards radar frames. We compare the performance of object detection and instance segmentation between RFMask and DETR in Table 6. It appears that exploiting backbone feature self-attention and query-feature cross-attention in DETR leads to consistently better performance for the two considered tasks.

Table 5: Single/Dual: Using the dual views works the best.

Method	View	AP	AP ₅₀	AP ₇₅	IoU
RFMask	Single	24.64	66.17	15.21	70.64
	Dual	25.53	67.30	15.86	71.98
RF-Pose	Single	22.22	30.99	22.64	-
	Dual	46.24	62.88	47.45	-

Table 6: Method: DETR can improve the performance.

Method	View	AP	AP ₅₀	AP ₇₅	IoU
RFMask	Single	24.64	66.17	15.21	70.64
DETR	Single	35.78	78.59	28.02	72.65

Table 7: Baseline benchmarks between HIBER [29] and MMVR.

Room Setting	Dataset	Object Detection			Instance Seg.
		AP	AP ₅₀	AP ₇₅	IoU
Open	HIBER (Walk)	17.77	52.46	6.78	48.47
Foreground	MMVR (P1)	24.46	66.82	11.22	67.03

Comparison over Different Radar Datasets An interesting ablation study involves training the radar perception model on one radar dataset, such as MMVR, and evaluating its performance on another dataset, like HIBER, or vice versa. However, this study requires *dataset-dependent* 3D radar-camera co-ordinate calibration and 3D-to-2D projection in the camera coordinate, due to varying relative geometries between the camera and radar and the use of different cameras in different datasets. Instead, we directly compare the same tasks across different datasets under similar environmental settings: a single subject walking in an open-foreground room. This is achieved using the “walk” data split in HIBER and “P1” in our MMVR. As shown in Table 7, the performance between the two datasets is comparable. Note that we used refined bounding boxes in the 2D image plane for the HIBER evaluation as the original bounding boxes were excessively large around the subjects.

6 Conclusion

In this paper, MMVR scales up indoor radar data collection using multi-view high-resolution heatmaps in a multi-day, multi-room, and multi-subject setting, with an emphasis on the diversity of the environment and subjects. It complies with an extensive collection of 345K radar heatmap frames from 25 subjects and 6 different rooms to benchmark three mainstream perception tasks such as object detection, pose estimation, and instance segmentation with 446K bounding boxes/segmentation instances and 7.59 million keypoints annotated. We hope that MMVR can stimulate the development of robust radar perception models that can operate effectively in varied real-world applications in smart home systems, security, elderly care, and navigation assistance for visually impaired individuals, contributing positively to society.

References

1. Adib, F., Hsu, C.Y., Mao, H., Katabi, D., Durand, F.: Capturing the human figure through a wall. *ACM Trans. Graph.* **34**(6) (2015)
2. An, S., Li, Y., Ogras, U.: mRI: Multi-modal 3D human pose estimation dataset using mmWave, RGB-D, and inertial sensors. In: *Advances in Neural Information Processing Systems (NeurIPS)*. vol. 35, pp. 27414–27426 (2022)
3. Barnes, D., Gadd, M., Murcutt, P., Newman, P., Posner, I.: The Oxford radar robotcar dataset: A radar extension to the oxford robotcar dataset. In: *IEEE International Conference on Robotics and Automation (ICRA)*. pp. 6433–6438 (2020)
4. Bilik, I., Longman, O., Villeval, S., Tabrikian, J.: The rise of radar for autonomous vehicles: Signal processing solutions and future research directions. *IEEE Signal Processing Magazine* **36**(5), 20–31 (2019)
5. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuScenes: A multimodal dataset for autonomous driving. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 11618–11628 (2020)
6. Cao, Z., Simon, T., Wei, S.E., Sheikh, Y.: Realtime multi-person 2D pose estimation using part affinity fields. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1302–1310 (2017)
7. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: *European Conference on Computer Vision (ECCV)*. pp. 213–229 (2020)
8. Chen, K., Wang, J., et al.: MMDetection: OpenMMLab detection toolbox and benchmark. *CoRR*, abs/1906.07155 (2019)
9. Cheng, B., Misra, I., Schwing, A.G., Kirillov, A., Girdhar, R.: Masked-attention mask transformer for universal image segmentation. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1280–1289 (2022)
10. Contributors, M.: OpenMMLab pose estimation toolbox and benchmark. <https://github.com/open-mmlab/mmpose> (2020)
11. Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The pascal visual object classes (VOC) challenge. *International Journal of Computer Vision* **88**(2), 303–338 (2010)
12. Kim, G., Park, Y.S., Cho, Y., Jeong, J., Kim, A.: MulRan: Multimodal range dataset for urban place recognition. In: *IEEE International Conference on Robotics and Automation (ICRA)*. pp. 6246–6253 (2020)
13. Lee, S.P., Kini, N.P., Peng, W.H., Ma, C.W., Hwang, J.N.: HuPR: A benchmark for human pose estimation using millimeter wave radar. In: *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. pp. 5715–5724 (2023)
14. Li, P., Wang, P., Berntorp, K., Liu, H.: Exploiting temporal relations on radar perception for autonomous driving. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 17050–17059 (2022)
15. Lim, T.Y., Markowitz, S.A., Do, M.N.: RaDICAL: A synchronized FMCW radar, depth, IMU and RGB camera data dataset with low-level FMCW radar signals. *IEEE Journal of Selected Topics in Signal Processing* **15**(4), 941–953 (2021)
16. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common objects in context. In: *European Conference on Computer Vision (ECCV)*. pp. 740–755 (2014)
17. Ouaknine, A., Newson, A., Rebut, J., Tupin, F., Pérez, P.: CARRADA Dataset: Camera and automotive radar with range- angle- Doppler annotations. In: *International Conference on Pattern Recognition (ICPR)*. pp. 5068–5075 (2021)

18. Paek, D.H., Kong, S.H., Wijaya, K.T.: K-Radar: 4D radar object detection for autonomous driving in various weather conditions. In: *Advances in Neural Information Processing Systems (NeurIPS)*. vol. 35, pp. 3819–3829 (2022)
19. Palffy, A., Pool, E., Baratam, S., Kooij, J.F.P., Gavrilu, D.M.: Multi-class road user detection with 3+1D radar in the view-of-delft dataset. *IEEE Robotics and Automation Letters* **7**(2), 4961–4968 (2022)
20. Pandharipande, A., Cheng, C.H., Dauwels, J., Gurbuz, S.Z., Ibanez-Guzman, J., Li, G., Piazzoni, A., Wang, P., Santra, A.: Sensing and machine learning for automotive perception: A review. *IEEE Sensors Journal* **23**(11), 11097–11115 (2023)
21. Rebut, J., Ouaknine, A., Malik, W., Pérez, P.: Raw high-definition radar for multi-task learning. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 17000–17009 (2022)
22. Scharf, L.L., Demeure, C.: *Statistical signal processing : detection, estimation, and time series analysis*. Addison-Wesley series in electrical and computer engineering, Addison-Wesley (1991)
23. Schumann, O., Hahn, M., Scheiner, N., Weishaupt, F., Tilly, J.F., Dickmann, J., Wöhler, C.: RadarScenes: A real-world radar point cloud data set for automotive applications. *CoRR*, abs/2104.02493 (2021)
24. Sengupta, A., Jin, F., Zhang, R., Cao, S.: mm-Pose: Real-time human skeletal posture estimation using mmWave radars and CNNs. *IEEE Sensors Journal* **20**(17), 10032–10044 (2020)
25. Sheeny, M., De Pellegrin, E., Mukherjee, S., Ahrabian, A., Wang, S., Wallace, A.: RADIATE: A radar dataset for automotive perception in bad weather. In: *IEEE International Conference on Robotics and Automation (ICRA)*. pp. 1–7 (2021)
26. Singh, A.D., Sandha, S.S., Garcia, L., Srivastava, M.: RadHAR: Human activity recognition from point clouds generated through a millimeter-wave radar. In: *ACM Workshop on Millimeter-Wave Networks and Sensing Systems*. pp. 51–56 (2019)
27. Sun, K., Xiao, B., Liu, D., Wang, J.: Deep high-resolution representation learning for human pose estimation. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 5686–5696 (2019)
28. Wang, Y., Jiang, Z., Li, Y., Hwang, J.N., Xing, G., Liu, H.: RODNet: A real-time radar object detection network cross-supervised by camera-radar fused object 3D localization. *IEEE Journal of Selected Topics in Signal Processing* **15**(4), 954–967 (2021)
29. Wu, Z., Zhang, D., Xie, C., Yu, C., Chen, J., Hu, Y., Chen, Y.: RFMask: A simple baseline for human silhouette segmentation with radio signals. *IEEE Transactions on Multimedia* **25**, 4730–4741 (2023)
30. Xue, H., Ju, Y., Miao, C., Wang, Y., Wang, S., Zhang, A., Su, L.: mmMesh: towards 3D real-time dynamic human mesh construction using millimeter-wave. In: *Annual International Conference on Mobile Systems, Applications, and Services*. pp. 269–282 (2021)
31. Yang, J., Huang, H., Zhou, Y., Chen, X., Xu, Y., Yuan, S., Zou, H., Lu, C.X., Xie, L.: MM-Fi: Multi-modal non-intrusive 4D human dataset for versatile wireless sensing. In: *Advances in Neural Information Processing Systems (NeurIPS)*. vol. 36, pp. 18756–18768 (2023)
32. Yataka, R., Wang, P., Boufounos, P., Takahashi, R.: SIRA: Scalable inter-frame relation and association for radar perception. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 15024–15034 (2024)
33. Zhao, M., Li, T., Alsheikh, M.A., Tian, Y., Zhao, H., Torralba, A., Katabi, D.: Through-wall human pose estimation using radio signals. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 7356–7365 (2018)

A Contributions

M. Rahman, P. Wang, and P. Boufounos contributed to the development of the MMVR testbed. M. Rahman and P. Wang coordinated and participated in data collection. P. Li and P. Wang established the annotation pipeline for labels used in object detection, pose estimation, and segmentation with P. Wang contributing to label curation. S. Kato and P. Wang worked on the pose estimation baseline with S. Kato evaluating benchmarks under various data splits and protocols and visualizing the results. R. Yataka, A. Cardace, and P. Wang contributed to the object detection and segmentation baselines with R. Yataka evaluating benchmarks under various data splits and protocols benchmarks and visualizing the results. M. Rahman and P. Wang worked on the initial alignment between the camera and radar frames, the camera-radar calibration, and the signal processing pipeline to generate radar heatmaps. R. Yataka contributed to the refined alignment between camera and radar frames. S. Kato, R. Yataka, and P. Wang contributed to data segmentation, sequence splits, and the organization of the final data structure. P. Wang initiated the MMVR effort for indoor perception applications at MERL and led the project from its inception. P. Wang and P. Boufounos supervised the project.

B Responsibility to Human Subjects

Our data collection for MMVR was approved by our institutional review board (IRB). We initiated a call-for-volunteers in our institution and recruited 25 participants. At the beginning of each data session, we informed participants about our experiment, its research goal, the procedure, potential exposure to high-frequency radio frequency waves, and the use of the camera. We notified the participants that de-identified data would be made publicly accessible for research purposes. A consent form was signed by all participants.

C Limitations and Potential Harms

Our dataset MMVR has limitations regarding annotations and benchmarks. Our annotation pipeline is based on RGB images. In cases of natural occlusion (chair, table, sofa) in a cluttered room of **P2**, the annotated BBox, keypoints, and segmentation pixels are also occluded or have extremely low reliability. As a result, it is unclear if radar-based approaches can lead to better perception performance than camera-based approaches under these natural occlusions. The maximum number of subjects in a session is limited to 3. Radar perception in a dense crowd is less explored and the dataset is extremely limited. Although radar perception has fewer privacy concerns than the camera, MMVR can be potentially utilized to classify and estimate attributes of subjects such as gender, size, height, and gait. It may be also used to advance technologies for indoor surveillance without acknowledgment or permission.

D Detailed Steps for Camera-Radar Calibration

Given N pairs of measured camera-coordinate positions $\mathbf{B}_{\text{camera}} \in \mathbb{R}^{3 \times N}$ and radar-coordinate positions $\mathbf{A}_{\text{radar}} \in \mathbb{R}^{3 \times N}$, the rotation matrix \mathbf{R} and a translation vector \mathbf{t} can be found by minimizing the Euclidean distance error in Eq. 1. This can be numerically solved in the following steps:

- **Centering the Calibration Position Sets:** Compute the centroids of both position sets $\mathbf{A}_{\text{radar}}$ and $\mathbf{B}_{\text{camera}}$ as $\boldsymbol{\mu}_{\text{radar}} = \text{mean}(\mathbf{A}_{\text{radar}}, \text{axis} = 1) \in \mathbb{R}^{3 \times 1}$ and $\boldsymbol{\mu}_{\text{camera}} = \text{mean}(\mathbf{B}_{\text{camera}}, \text{axis} = 1) \in \mathbb{R}^{3 \times 1}$. Then translate both sets to have their centroids at the origin:

$$\begin{aligned}\tilde{\mathbf{A}}_{\text{radar}} &= \mathbf{A}_{\text{radar}} - \boldsymbol{\mu}_{\text{radar}}, \\ \tilde{\mathbf{B}}_{\text{camera}} &= \mathbf{B}_{\text{camera}} - \boldsymbol{\mu}_{\text{camera}}.\end{aligned}\tag{4}$$

This initial alignment separates the calculation of the rotation matrix from that of the translation.

- **Computing the Covariance Matrix:** Calculate the covariance matrix between the centered calibration position sets $\tilde{\mathbf{A}}_{\text{radar}}$ and $\tilde{\mathbf{B}}_{\text{camera}}$:

$$\mathbf{H} = \tilde{\mathbf{A}}_{\text{radar}} \tilde{\mathbf{B}}_{\text{camera}}^T,\tag{5}$$

where $(\cdot)^T$ denotes the matrix transpose.

- **Singular Value Decomposition (SVD):** Perform the singular value decomposition on the covariance matrix:

$$[\mathbf{U}, \mathbf{S}, \mathbf{V}] = \text{SVD}(\mathbf{H}),\tag{6}$$

where \mathbf{U} and \mathbf{V} are unitary matrices, and \mathbf{S} is a diagonal matrix.

- **Determining the Rotation Matrix \mathbf{R} :** The rotation matrix is found as

$$\mathbf{R} = \mathbf{V}\mathbf{U}^T.\tag{7}$$

Since \mathbf{U} and \mathbf{V} are unitary matrices, \mathbf{R} is guaranteed to be a rotation matrix with a determinant of 1.

- **Calculating the Translation Vector \mathbf{t} :** Once the rotation matrix \mathbf{R} is found in Eq. 7, the translation vector \mathbf{t} can be determined as

$$\mathbf{t} = \boldsymbol{\mu}_{\text{camera}} - \mathbf{R}\boldsymbol{\mu}_{\text{radar}}.\tag{8}$$

The above step first applies the rotation to the radar-coordinate calibration set centroid $\boldsymbol{\mu}_{\text{radar}}$ and then subtracts it from the camera-coordinate calibration set centroid $\boldsymbol{\mu}_{\text{camera}}$.

E Additional Snapshots of All Data Sessions

We provide additional snapshots of all data sessions in Fig. 12. Each row represents a day (d) while each column represents a session (s) with the session duration indicated.




Fig. 12: Gallery of additional snapshots.

F Details of Data Sessions

F.1 Content

Table 8 enumerates the details of all data sessions included in our dataset MMVR. Our dataset includes two protocols: **P1** and **P2**. Under each protocol, each session is split into non-overlapping one-minute data segments. The number of data segments in each session is listed in the column of **# of seg.**, along with the column of **# of frames**. Other statistics of each data session such as storage size, room type, obstacles for occlusion, number of subjects, and action types of the subjects are also listed.

Table 8: Details of data sessions.

Protocol	Session	# of seg.	# of frames	Room	Occlusion	# of Sbj.	Action
P1	d1s1	39	35,085	large conference	-	1	walking
	d1s2	20	17,618	large conference	-	1	walking
	d2s2	20	17,563	large conference	-	1	jumping
	d3s1	19	16,962	large conference	-	1	jumping
	d3s2	9	7,436	large conference	-	1	walking
	d4s1	15	13,238	large conference	-	1	walking
	Total d1-d4	122	107,902	-	-	-	-
P2	d5s1	9	7,808	lobby	chair, table, vase	1	walking, sitting, stretching, reading
	d5s2	9	7,407	lobby	chair, table, vase	1	walking, sitting, reading, writing, stretching
	d5s3	10	8,817	lobby	chair, table, vase	1	walking, sitting, reading, writing, stretching, eating, drinking
	d5s4	9	8,040	lobby	chair, table, vase	1	walking, sitting, reading, writing, stretching
	d5s5	6	5,382	lobby	chair, table, vase	2	walking, sitting, reading, drinking, stretching, talking
	d5s6	10	8,821	lobby	chair, table, vase	3	walking, sitting, reading, stretching, talking
	d6s1	10	8,592	small office	chair	1	walking, sitting, reading, writing, stretching
	d6s2	10	8,852	small office	chair	1	walking, sitting, writing, stretching
	d6s3	10	8,552	small office	chair	2	walking, sitting, reading, writing, stretching, talking
	d6s4	10	8,842	small office	chair	2	walking, sitting, reading, writing, stretching, talking
	d6s5	1	430	small office	chair	2	walking, sitting, reading, writing, stretching, talking
	d6s6	10	8,677	small office	chair	1	walking, sitting, reading, writing, stretching
	d7s1	10	8,757	small conference	chair, table	1	walking, sitting, reading, writing, stretching
	d7s2	10	8,392	small conference	chair, table	1	walking, sitting, reading, writing, stretching
	d7s3	10	8,855	small conference	chair, table	1	walking, sitting, reading, writing, stretching
	d7s4	10	8,844	small conference	chair, table	2	walking, sitting, reading, writing, stretching, talking
	d7s5	10	8,795	small conference	chair, table	2	walking, sitting, reading, writing, stretching, talking
	d8s1	10	8,623	medium office	chair, table	1	walking, sitting, reading, stretching, talking
	d8s2	10	8,177	medium office	chair, table	1	walking, sitting, reading, writing, stretching, playing
	d8s3	10	8,758	medium office	chair, table	1	walking, sitting, reading, writing, stretching
	d8s4	10	8,872	medium office	chair, table	1	walking, sitting, reading, writing, stretching
	d8s5	10	8,880	medium office	chair, table	2	walking, sitting, reading, writing, stretching, talking
	d8s6	10	8,805	medium office	chair, table	2	walking, sitting, reading, writing, stretching, talking, playing
	d9s1	9	7,964	small conference	chair, table	1	walking, sitting, reading, writing, stretching
	d9s2	10	8,795	small conference	chair, table	1	walking, sitting, reading, writing, stretching
	d9s3	10	8,770	small conference	chair, table	1	walking, sitting, reading, writing, stretching
	d9s4	10	8,677	small conference	chair, table	2	walking, sitting, reading, writing, stretching, talking
	d9s5	10	8,817	small conference	chair, table	2	walking, sitting, reading, writing, stretching, talking
	d9s6	10	8,797	small conference	chair, table	3	walking, sitting, reading, writing, stretching, talking
	Total d5-d9	273	237,798	-	-	-	-
Total		395	345,700	-	-	-	-

F.2 Folder Structure

Fig. 13 illustrates the folder structure of our dataset MMVR. Below the root folder, we have one data folder of $d_x s_y$ for each data session, where x and y are the indices for the day and session, respectively. Under each data folder or data session, we group the data frames into separate non-overlapping data segment folders named using a three-digit, zero-filled convention based on the chronological order of the data segment. For instance, the first one-minute data segment corresponds to the folder 000, while the second data segment folder is named after 001.

Under each data segment folder, there are three types of files: meta, radar, and annotation including bounding boxes, keypoints and segmentation masks for each radar frame, named using a five-digit, zero-filled convention. For instance, `d1s2/000/00001_meta.npz` contains the metadata for the second data frame 00001 in the first data segment 000 of the data session d1s2. Under the same data segment folder, `00001_radar.npz` contains the horizontal and vertical radar heatmaps. The annotation files are saved separately for each perception task to facilitate the data loading. For instance, `00001_bbox/pose/mask.npz` contain the bounding boxes, keypoints, and segmentation masks for the corresponding frame.



Fig. 13: The folder structure of our MMVR dataset.

G Statistics of Annotation Labels

Fig. 14 shows the average number of annotation labels for each data session between d5 and d9. Each color represents a session number (s1-s6) within a day. It is seen that earlier sessions, e.g., s1, s2, s3, are mostly single-subject sessions, while later sessions have multiple subjects. The distribution of single-person, two-person, and three-person annotations over data frames is shown in Fig. 15.

Fig. 16 further analyzes the distribution of annotated BBox sizes (height and width) in terms of pixel numbers. We categorize the BBoxes into small, medium, and large boxes in the height domain with a snapshot of each category shown in the picture above the distribution. This distribution indicates that the BBox annotations are well-balanced with a variety of distances between the camera and the subjects and sufficient diversity of subject postures.

Fig. 17 breaks down the keypoint annotation confidence scores over selected joints or body parts. Note that the displayed scores represent the average values for the left and right sides of all body parts except the nose. With human curation involved, Fig. 17 suggests that our annotation process yields high-confidence labels across the board. However, it's noteworthy that body parts located on the





Fig. 14: The number of labels per data session. **Fig. 15:** The distribution of the number of subjects over data frames.

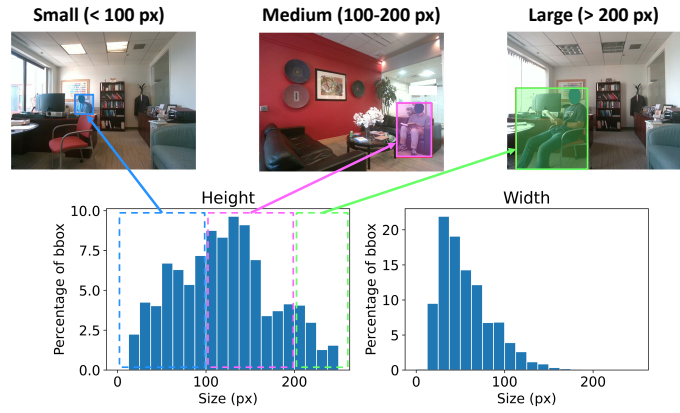


Fig. 16: The distribution of annotated bounding box sizes (height and width).

torso are annotated with significantly greater confidence compared to limbs and facial parts. Specifically, the nose, eyes, knees, and ankles exhibit a pronounced peak around zero confidence. This lower confidence level primarily arises from these parts being prone to occlusion, such as when a person is facing away from the camera or is seated behind a desk.

H Synchronization

The initial alignment between the camera and two radar sensors is done with the use of a passive device, a corner reflector. A corner reflector is a passive device designed to reflect radar waves back toward the radar transceivers, regardless of the angle of incidence, with a geometric shape with three mutually perpendicular, flat surfaces or planes. As an ideal target for calibration and testing of radar sensors, we place it in front of our MMVR sensor testbed such that both camera and radar sensors can identify its presence at the beginning of each data collection session.




Fig. 17: The distribution of keypoint annotation confidence scores over selected joints, e.g., nose, shoulder, hip, etc..




Fig. 18: Initial alignment between camera and two radar sensors using a corner reflector.

One example is shown in Fig. 18 for the data session *d5s1*. The corner reflector is visualized in an RGB image frame which is identified as Frame 0 for the image stream. Meanwhile, we preprocess the radar raw waveform, obtain the range-azimuth heatmap for the horizontal radar, and identify the frame with strong reflection (bright yellow) at the close distance (less than 50 cm) as Frame 0 for the vertical radar. A similar procedure is done for the vertical radar in the range-elevation heatmap. Once identifying Frame 0 for the camera and two radar sensors, the consequent frames are automatically aligned due to the use of the same frame rate of 15 fps. For the refined alignment, please refer to the enclosed demo videos.

I Details about Data Splits S1 and S2

Table 9 presents the detailed list of data segments within the two data splits **S1** and **S2** under both protocols **P1** and **P2**. These non-overlapping data segments were randomly selected once and then fixed to training, validation, and test sets,

Table 10: Details about evaluation parameters.

	Name		Notation	Value
Data	P1	S1	# of training	- 86579
			# of validation	- 10538
			# of test	- 10785
		S2	# of training	- 70266
			# of validation	- 24398
			# of test	- 13238
	P2	S1	# of training	- 190441
			# of validation	- 23899
			# of test	- 23458
		S2	# of training	- 118280
			# of validation	- 33841
			# of test	- 85677
Input radar heatmap size		$H \times W$	256×128	
Original segmentation mask size		$H \times W$	480×640	
Input segmentation mask size		$H \times W$	240×320	
Fixed-height size (pixel)		H	36	
Heatmap/PAF size for pose estimation		$H \times W$	46×46	
Scale		-	log	
Training	batch size		-	32
	epoch		-	100
	early stopping	patience	-	5
	early stopping	check val every N epoch	-	2
	optimizer		-	Adam
	learning rate		-	1e-4
	sheduler Linear × 0.1		-	50
	weight decay		-	1e-4

the training part, we include hyper-parameters such as batch size, number of epochs, parameters related to early stopping, learning rate, etc.

J.2 Baseline Methods for Object Detection and Segmentation

We use RFMask [29] as one baseline method for BBox and segmentation tasks. Unlike the HIBER dataset with the BBox labels on the two radar views, our dataset MMVR annotates the BBox labels directly on the image plane. As a result, we modify RFMask in a way that the BBox loss is calculated directly on the image plane and backpropagates to trainable parameters in an end-to-end fashion.

Fig. 19 shows the modified RFMask architecture with the loss function computed directly on the image plane. Specifically, we add an **Img BBox Regression** module alongside a **Hori BBox Regression** module, enabling the conversion of BBox offsets to the image plane. By computing loss with respect to these offsets, we can directly learn BBoxes on the image plane. Additionally, the region pro-

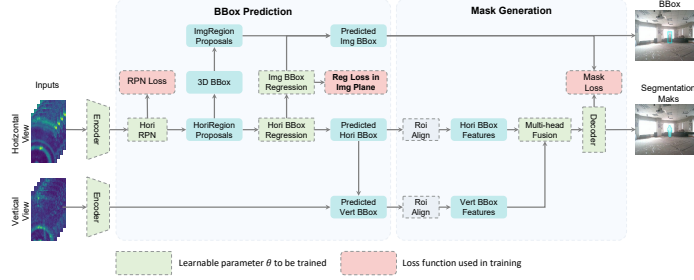


Fig. 19: Baseline method for object detection and segmentation: RFMask [29].

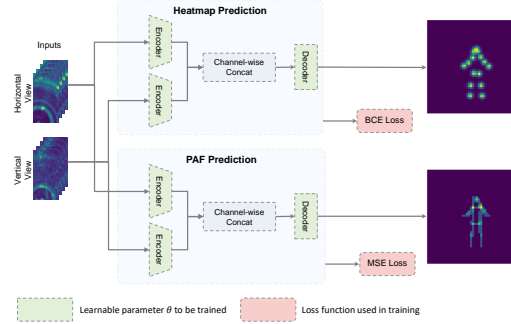


Fig. 20: Baseline method for pose estimation: RF-Pose [33].




Fig. 21: Additional baseline for Object Detection and segmentation: DETR [7].

posals estimated by the region proposal network (RPN) are transformed into 3D BBoxes based on the fixed-height size (see Table 10), the same as the original RFMask. These 3D BBoxes are then projected onto the image plane using the Camera-Radar calibration of Sec. 3.3 and a 3D-to-2D projection using intrinsic parameters obtained from the RealSense camera D455.

In Sec. 5.4, we also consider an additional baseline method that borrows the more advanced Detection Transformer (DETR) [7] for the radar perception tasks of object detection and instance segmentation. The DETR-based architecture is shown in Fig. 21, with a modification to convert BBoxes predicted on the horizontal plane to the image plane. The BBoxe loss is calculated on both the horizontal and image planes. For the segmentation task, UNet is added and the mask loss is computed. During the training, only the parameters of the Trans-

former Decoder and BBox head are initially updated, and then, the parameters of the segmentation head and UNet are updated while keeping the parameters of the Transformer Decoder and BBox head fixed.

J.3 Baseline Method for Pose Estimation

We use RF-Pose [33] as a baseline method for the evaluation of pose estimation performance on our dataset MMVR. The original RF-Pose aims at predicting only a confidence heatmap for each joint, which is then used for multi-person association. We extend its architecture by using two encoder-decoder sets to predict both heatmap and part affinity map (PAF), to follow the common practice in the bottom-up pose estimation in computer vision and provide better multi-person association [6]. Fig. 20 shows our implementation of RF-Pose with the exact the same encoder and decoder structures as RF-Pose [33, Sec. 4.3] (e.g., 3D convolution, stride and convolution kernel parameters) with the loss function combining both the keypoint heatmap and PAF.

K Additional Visualization Results

We provide additional qualitative inspection of results in Fig. 22 for object detection, Fig. 23 for pose estimation, and Fig. 24 for instance segmentation, respectively. For each perception task, we select one frame from d1 in **P1** and one frame from each of the 5 days (d5-d9) in **P2**. Meanwhile, we show failure cases for each perception task in Fig. 25.

K.1 Object Detection

As shown in Fig. 22, it is possible to localize the subject in the image plane using two radar views. The results for d1 (the first row) show that the radar features can support the BBox prediction that includes the spread arms and legs of a subject. On the other hand, the first two rows of Fig. 25 show missing detection when the subject is close to the radar and camera, and when the subject is stationary for a long duration of time.

K.2 Pose Estimation

Fig. 23 shows reasonably good performance of baseline pose estimation in the single-person session such as d1 and the performance degrades in multi-person sessions. While the location and relative scale of subjects are estimated with small errors, the end of limbs like wrists and ankles tends to have larger errors, especially in more complex environments in **P2**. For the failure cases shown in the third and fourth rows of Fig. 25, one can observe that certain connections between keypoints can be twisted or a subset of keypoints were missed or estimated with small confidence scores.

K.3 Instance Segmentation

Finally, for instance, segmentation results in Fig. 24, the baseline RFMask can capture the rough shape of subjects. On the other hand, as shown in the last two rows of Fig. 25, segmentation mask predictions may be missing in the multi-person scenarios, and mask pixels in between legs can be identified as part of the predicted instance mask.




Fig. 22: Visualization of baseline object detection results.




Fig. 23: Visualization of baseline pose estimation results.




Fig. 24: Visualization of baseline instance segmentation results.




Fig. 25: Visualization of failure cases.