

# Dataset for Real-World Human Action Detection Using FMCW mmWave Radar

Dylan Jayabahu<sup>1</sup>, Parthipan Siva<sup>1,2</sup>  
<sup>1</sup>Chirp Inc.

<sup>2</sup>Vision and Image Processing Group, Systems Design Engineering, University of Waterloo  
dylanjayabahu@gmail.com, parthipan.siva@mychirp.com

## Abstract

Millimeter-wave (mmWave) radar is emerging as a privacy-preserving technology for in-home monitoring, particularly for older adults who prefer aging in place. This study addresses the need for accurate human action recognition (HAR) in real-world settings by introducing a dataset of natural activities collected from 28 homes. The dataset captures sit-down and stand-up actions—key indicators of mobility—using the Chirp Smart Home Sensor, which combines sparse 3D point cloud data from an mmWave radar with low-resolution thermal imaging for annotation.

The dataset comprises over 900 annotated actions across diverse residential environments, supplemented with augmentation techniques to mitigate class imbalance. A baseline convolutional neural network (CNN) model is evaluated using Doppler-time (DT) and positional-time data (XT, YT, ZT) features. Results reveal significant challenges, including high variability between training and testing environments, low precision due to diverse non-action classes, and limited spatial coverage of action locations.

## 1 Introduction

Millimeter-wave (mmWave) radar is rapidly emerging as a privacy-preserving sensing technology for home monitoring, with applications ranging from home automation to healthcare. As our population ages and more older adults prefer to age in place—driven by both personal choice and limited availability in care facilities—technology becomes crucial to scale in-home care. For mmWave sensors to effectively support in-home monitoring, they must accurately track individuals and recognize key actions like falls and transfers.




Figure 1: Chirp Smart Home Sensor (left) is installed in the homes at 122 cm (48") height (right).

While prior studies have explored human action recognition (HAR) using mmWave radar, they are limited to controlled environments with actors performing a set of predefined actions [1, 2, 3]. In this paper, we present a dataset collected in real homes, capturing natural activities across 28 residences; among this data we focus on transfer actions in and out of chairs. We evaluate the performance of an existing baseline CNN model for action detection in our real-world dataset to illustrate limitations and challenges.

## 2 Data Collection

In this section, we detail the data collection methods used to build a robust, real-world dataset for human action detection.

### 2.1 Sensor

Data was collected using the Chirp Smart Home Sensor [4], which incorporates the TI IWR6843AOP FMCW mmWave sensor (Figure 1). To reduce data storage and bandwidth issues, the sensor captures sparse 3D point cloud data at 10 frames per second (FPS) through a radar processing chain illustrated in Figure 2. Each point in the 3D point cloud includes coordinates  $(x, y, z)$ , ve-




Figure 2: Radar signal processing to obtain 3D point cloud.

locity, and signal-to-noise ratio (SNR). Using this radar point cloud data, Doppler-time, range-time, azimuth-time, elevation-time, and time-based images along each axis ( $x$ ,  $y$ ,  $z$ ) can be generated as per the approach presented in [1].

To support radar data annotation while maintaining privacy, we supplemented the mmWave sensor with a low-resolution  $32 \times 32$  pixel thermopile sensor, capturing thermal video at 3 FPS. Figure 3 illustrates sample actions as viewed by the thermal sensor. While the sparsity of the radar data alone complicates action annotation, the thermal sensor offers enough visual cues to accurately identify actions while ensuring individual privacy.




Figure 3: Low resolution thermopile data used for annotation. Here the sit-down action is illustrated from Frame 1 showing person standing to Frame 16 showing the person sitting.

## 2.2 Collection Sites

Data was collected from 28 unique homes, each occupied by a single older adult participant. Collection occurred over a continuous 24-hour period, allowing participants to follow their usual routines. Due to visitors, actions present within a single home may involve multiple in-




Figure 4: Distribution of sit-down and stand-up action locations in the dataset. Each arc is one meter apart.

dividuals. However, because participants were independently recruited, same individual will not appear in multiple homes.

Participants were given two to three devices and instructed to install them in common living areas, such as kitchens and living rooms. In total, data was collected across the following room types:

- Kitchens/Dining Rooms: 20 rooms
- Living Rooms: 24 rooms
- Multipurpose Rooms: 2 rooms
- Bathrooms: 2 rooms

Devices were generally installed at a height of 120 cm from the ground, roughly the standard height of a light switch, and placed to face areas where occupants commonly interacted (Figure 1). Installation height varied based on wall space availability and participant preferences.

## 2.3 Data annotation

Action annotation was achieved by manually reviewing the thermal video, with a focus on two key actions: sitting down and standing up. These actions correspond to transfers into and out of various types of furniture, such as chairs, stools, sofas, and toilets. Sit-down and stand-up actions were selected for their relevance to mobility assessment and their frequent use in clinical evaluations, such as the Timed Up and Go (TUG) test and the 30-Second Sit-to-Stand test.

This dataset thus captures common daily activities, enabling the development of mmWave-based HAR models capable of recognizing key actions that are critical to assessing and supporting mobility in home settings.

## 2.4 Dataset

Across 28 homes, a total of 458 sit-down and 454 stand-up actions<sup>1</sup> were annotated within the 48 rooms. The

<sup>1</sup>Since the day can end with person sitting down, the number of sit-down and stand-up differs.

spatial distribution of these actions is presented in Figure 4. For a robust evaluation process, the data was divided such that 3 homes were selected for testing, 3 for validation, and the remaining 22 were designated for training.

Non-action data was generated using a 5-second sliding window on each recording, ensuring a 0% overlap between windows and with annotated action locations. Any window not meeting the minimum required radar points was removed. The resulting data distribution is detailed in Table 1.

Table 1: Dataset split statistics

	Test	Valid	Train	Train + Aug
# of Homes	3	3	22	22
# of Rooms	6	6	36	36
sit-down	65	60	333	7659
stand-up	62	58	334	7682
other	N/A	6,666	29,803	29,803

An imbalance exists between non-action and action classes. To address this, action windows were over sampled with various augmentations applied to the 3D point cloud data. These random augmentations include:

- Scaling of signal-to-noise ratio (SNR)
- Scaling position along the radial axis
- Shifting along the x, y, and z axes
- Reflection along the x-axis
- Rotation around the z-axis
- Random drops of data points

This oversampling with augmentations produces a more balanced dataset, as shown in Table 1.

### 3 Experiment Setup

We use Doppler-time (DT), x-time (XT), y-time (YT), and z-time (ZT) data images as inputs to our CNN models. Each feature image is generated from the 3D point cloud data following the procedure outlined in [1].

Unlike existing methods that emphasize classification accuracy within a predefined set of actions [5, 6], our testing employs a sliding window detection approach. This approach better reflects real-world action detection scenarios, where “non-action” encompasses all other possible actions in a typical home environment.

We assess detection performance using standard metrics: Recall, Precision, and F1-Score for each class.

## 4 Detection Model

As a baseline model we use the 3-layer CNN architecture of [2, 1]. The architecture is illustrated in Figure 5. To evaluate baseline performance we study the effect of DT, DT+ZT, and DT+ZT+YT+XT as inputs to the model.

Adam optimizer and cross entropy loss was used to train the network from scratch for up to 25 epoch. Early stopping was used based on maximizing the F2-Score on validation dataset. F2 chosen to allow recall to be weighted higher due to the class imbalance problem.

## 5 Results & Discussion

The detection confusion matrix is presented in Table 2-4. Overall recall and precision are presented in Table 5

	Predicted		
	Other	Sit-Down	Stand-Up
Other	N/A	270	461
Sit-Down	33	32	0
Stand-Up	31	0	31

Table 2: DT Results

	Predicted		
	Other	Sit-Down	Stand-Up
Other	N/A	17	34
Sit-Down	35	30	0
Stand-Up	50	0	12

Table 3: DT+ZT Results

	Predicted		
	Other	Sit-Down	Stand-Up
Other	N/A	390	77
Sit-Down	12	53	0
Stand-Up	45	0	17

Table 4: DT+ZT+YT+XT Results

	Validate		Test	
	Recall	Precision	Recall	Precision
DT	0.729	0.079	0.496	0.074
DT+ZT	0.746	0.427	0.331	0.452
DT+ZT+YT+XT	0.873	0.106	0.551	0.130

Table 5: Overall recall and precision for detecting both actions.

While test results remain poor, validation recall results perform comparatively better, suggesting significant variability between the training, validation, and test sets. This variability arises not only from differences in individuals performing the actions but also from dis-

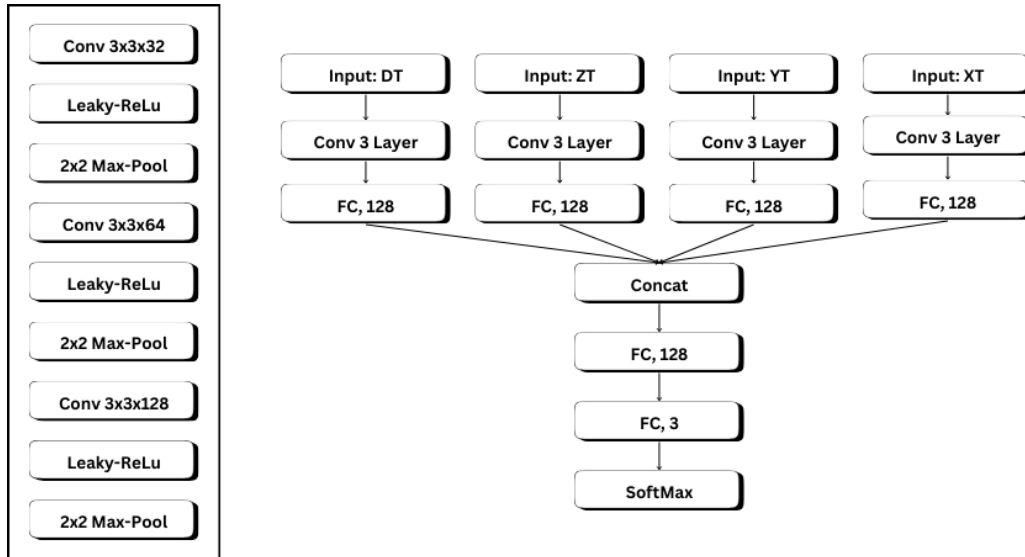


Figure 5: Baseline CNN model architecture from [1, 2].

tinct action locations within the environment. Although the dataset includes approximately 450 instances of sit-down and stand-up actions, there are only 48 distinct rooms, limiting the variety of spatial locations and orientations where these actions occur. This limitation is evident when examining the spatial distribution of action locations across the dataset splits (Figure 6); actions in the test set frequently occur in locations not represented in the training or validation sets.

Precision rates remain low, indicating high variation in the non-action class and prioritization of recall in early stopping. Despite using class weights and over-sampling, an optimal balance between recall and precision wasn't achieved.

Comparing DT (F1=0.137), DT+ZT (F1=0.382), and DT+ZT+YT+XT (F1=0.211), we observe that DT+ZT achieves a better balance of recall and precision. This is intuitive, as Z-axis variation for sit-down and stand-up actions remains consistent regardless of location within the room, while YT and XT vary by location. Additionally, ZT enhances the Doppler (speed) context provided by DT only.

## 6 Conclusions

We present a robust real-world dataset for action detection using mmWave radar in actual older adult's homes. Unlike existing datasets that simulate actions of interest, our dataset captures real-world activity in the homes of older adults and annotates actions within the captured data. This results in a more realistic dataset; however, it also leads to higher variance in both the environment and the actions performed. We demonstrate that, un-

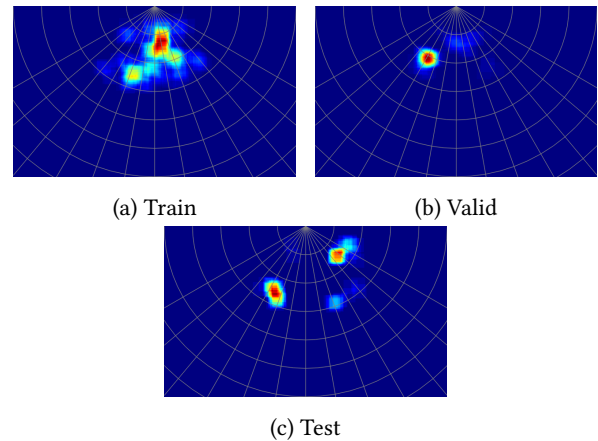


Figure 6: Distribution of action location by dataset splits.

der these challenging conditions, existing baseline approaches perform poorly.

## References

- [1] Z. Wu, Z. Cao, X. Yu, J. Zhu, C. Song, and Z. Xu, "A novel multiperson activity recognition algorithm based on point clouds measured by millimeter-wave mimo radar," *IEEE Sensors Journal*, vol. 23, no. 17, pp. 19 509–19 523, 2023.
- [2] F. Jin, R. Zhang, A. Sengupta, S. Cao, S. Hariri, N. K. Agarwal, and S. K. Agarwal, "Multiple patients behavior detection in real-time using mmwave radar and deep cnns," in *2019 IEEE Radar Conference (RadarConf)*. IEEE, 2019, pp. 1–6.
- [3] I. Ullmann, R. G. Guendel, N. C. Kruse, F. Fioranelli, and A. Yarovoy, "A survey on radar-based continu-

ous human activity recognition,” *IEEE Journal of Microwaves*, vol. 3, no. 3, pp. 938–950, 2023.

- [4] Chirp Inc, <https://mychirp.com/>, accessed: 2024-11-08.
- [5] P. Zhao, C. X. Lu, B. Wang, N. Trigoni, and A. Markham, “Cubelearn: End-to-end learning for human motion recognition from raw mmwave radar signals,” *IEEE Internet of Things Journal*, vol. 10, no. 12, pp. 10 236–10 249, 2023.
- [6] N. C. Kruse, F. Fioranelli, and A. Yarovoy, “Radar point cloud processing methods for human activity classification with point transformer networks,” *IEEE Transactions on Radar Systems*, 2023.