# Inferring Personal Attributes with a Mmwave Radar


1st Cinthya Celina Tamayo Gonzalez
*IMT School for Advanced Studies*
Lucca, Italy
cinthya.tgonzalez@imtlucca.it

2nd Simone Soderi
*IMT School for Advanced Studies*
Lucca, Italy
simone.soderi@imtlucca.it

3rd Julian Todt
*Karlsruhe Institute of Technology*
Karlsruhe, Germany
julian.todt@kit.edu

4th Thorsten Strufe
*Karlsruhe Institute of Technology*
Karlsruhe, Germany
thorsten.strufe@kit.edu

5th Mauro Conti
*University of Padua*
Padova, Italy
mauro.conti@unipd.it



***Abstract*—Various devices can be utilised to sense people in smart environments, and in this task, millimetre-wave (mmWave) radars have recently gained attention. This paper presents a novel investigation of the risk posed by mmWave radars that could be used in Internet of Things (IoT) scenarios, focusing on the potential for an attacker to access personal information without the user's consent. We introduce INFORMER: INFerring persOnal attRibutes with a MmwavE Radar, a new perspective on collecting information from personal features using gait. INFORMER utilises point clouds to capture the gait of volunteers. The mean information from the point clouds is then analysed to infer information from five attributes: height, weight, gender, waist, and arm span. Three variants of INFORMER scrutinised these sequences: one Deep Learning (DL) model and two Machine Learning (ML) models. The DL model showed accuracy up to 82%, demonstrating superior height and weight inference performance, while gender inference showed comparatively lower accuracy. Additionally, the first model ML, Singular Value Decomposition (SVD), tends to perform better in some attributes, such as waist, but worse attributes than the other models in features, such as weight. Lastly, the Principal Component Analysis (PCA) presents comparable results. Specifically, it shows the best performance in gender inference.**

***Index Terms*—Mmwave radar, point clouds, radar security.**


## I. Introduction

Gait analysis has significant potential for extracting important insights into an individual's characteristics. It is a versatile tool applied across various fields, including health monitoring, security checks, and human-computer interaction. Notably, millimetre wave (mmWave) radars have gained attention as an efficient approach for obtaining data for sensing people. Compared to camera-based solutions, mmWave radars display benefits such as remaining effective in non-line-of-sight scenarios [1]. These type of radars transmit electromagnetic waves with wavelengths in the millimetre range. The utilisation of these wavelengths confers numerous advantages, including high resolution [2]. Moreover, radars are compact and can operate in darkness, dust and high temperatures [3], making them suitable for Internet of Things (IoT) applications such as fall detection for older adults [4], sleep monitoring [5], human vital signs monitoring [6] and energy management [7].

However, although radars do not directly capture information about people's appearance or personal attributes, it is still possible to gather information about individuals within the system by analysing the signals they use. Moreover, it has been demonstrated that in mmWave systems, it is feasible to detect and analyse signals emitted by a mmWave radar from an adjacent room [8]. This can lead to the inference of soft biometric data, such as gender, which poses an important privacy risk [9]. In this paper, we present how it is possible to jeopardise the privacy of users in mmWave radar systems, which represents a significant security concern in smart environments.

In this investigation, we present INFORMER with three variants: a Deep Learning (DL) and two Machine Learning (ML) models. INFORMER aims to categorise individuals according to five basic personal attributes: gender, weight, height, arm span, and waist circumference. Our methodology entails utilising point clouds originating from the mmWave radar AWR1443 BOOST, eliminating the necessity for auxiliary devices by directly processing the radar's information. The radar data points are sparse, necessitating a method to remove noise; we used the clustering algorithm Density-Based Spatial Clustering of Applications with Noise (DBSCAN) [10]. The point clouds were then normalised to maintain consistency of point density across the dataset. In the DL implemented, the first part is based on the widely known PointNet [11] architecture, which is used to extract key features. The information's temporal dynamics underwent analysis with a causal dilated temporal convolutional layer. Similarly, the other two variants use the ML models, Singular Value Decomposition (SVD) and Principal Component Analysis (PCA), to classify volunteers based on their personal attributes. These models extract meaningful information from the point clouds, which are then classified using a Support Vector Machine (SVM). This approach successfully allows one to infer personal information from a mmWave system and enables evaluating the performance of a complex model and two other simpler models.

This paper introduces a novel approach called INFORMER, which leverages point clouds to infer information from multiple attributes. The approach is specifically applied to gait analysis, whereby personal features are extracted. The results demonstrate that point clouds provide a rich data source capable of gathering personal attribute information, which raises potential privacy concerns in mmWave radar systems. It is noteworthy that this research represents the inaugural instance of utilising mmWave devices to extract physical traits, including arm span, weight, and waist measurement. Furthermore, we present the first single model capable of simultaneously gathering multiple attributes from users using mmWave data. The performance of the INFORMER system was evaluated on a cohort of 23 volunteers, who were asked to walk at two different speeds. The system successfully inferred many personal attributes, including gender, height, weight, waist measurement, and arm span.

The rest of the paper is organised as follows. Section II presents the data processing methodology. The experimental design is shown in section III, outlining the factors that informed our decisions about data collection. In section IV, we present the experimental findings and discuss possible limitations. Finally, section V presents our conclusions.

## II. Methodology

This section provides a detailed overview of the methodology illustrated in Figure 1. In subsection II-A, the details of the preprocessing stage are described, followed by the processing done with the three models in subsection II-B.

The first steps involve preparing the mmWave radar point clouds to improve information quality. After point cloud filtering, the information is normalised to ensure a uniform input data shape. The preprocessing steps are the same for all the variants of INFORMER; after preprocessing, each variant uses different methods to process the data.

The first model implemented is a DL architecture consisting of a PointNet inspired block and a convolutional temporal network. The architecture ends with a sigmoid function for binary classification. Secondly, we introduced two other variants of INFORMER with ML models: SVD and PCA. These models are used to compress dimensionality and extract the most meaningful information, which is then classified using an SVM for binary classification as well. Consistent parameters across models were used to derive information from five attributes: height, weight, gender, waist, and arm span.

### A. Preprocessing

To mitigate the impact of noise within the point cloud data, we employed the DBSCAN algorithm as a crucial preprocessing step. This algorithm facilitates the grouping of points into clusters, designating ungrouped points as noise and subsequently removing them. The selection of DBSCAN is motivated by the next reasons: a) it has been successfully used as a filtering algorithm for point clouds from mmWave radars in previous work [12]–[14], and b) the algorithm doesn't require the number of clusters beforehand, which allows more flexibility in sparse and noisy data.

After filtering, the number of points contained in the point cloud becomes variable, depending on the level of noise present in each frame. To maintain consistency across all frames, each point cloud undergoes an up-sampling procedure until it reaches a predetermined point limit. Initially, the required number of points ($N$) for each sample is calculated, after which N points are selected randomly and duplicated within the original set to ensure a uniform number of points. The radar used in our study can generate a maximum of $64$ points. To maintain consistency and streamline subsequent analyses, we set a limit of $70$ points, rounding up from the original number of point clouds prior to filtering.

To extract comprehensive insights from individuals walking at varying speeds, we conducted the analysis of gait in both normal walking and fast walking scenarios. Further elaboration on these two modalities is provided in the evaluation set-up section. Individuals demonstrate distinguishable walking speeds that are influenced by several factors, as discussed by the authors in [15]. Hence, variations in walking speeds among volunteers are expected in the two suggested scenarios. This, in turn, leads to a varying number of frames needed to capture individuals. All gathered data was analysed to determine the size of the collected samples. This step is essential for understanding the data set's variability and determining the parameter to normalise the data.

The collected samples varied between a minimum of $M_m$ frames and a maximum of $M_x$ frames. To evaluate how the number of frames affects the identification of personal characteristics, we used two normalisation methods. Firstly, we normalised the frames by the smallest sample size, resulting in all samples being analysed using only $M_m$ frames. The frames with more noise were eliminated until the $M_m$-frame limit was attained. Regarding the second normalisation method, $M_l$ frames were examined, which amounted to half the gap between the smallest and largest sample sizes, i.e., $M_x - M_m$. Therefore, some samples necessitated extra frames. To overcome this hurdle, we employed a practical approach by attaching Gaussian noise to the last frames of the samples. The additional frames were integrated into the original sequence, thereby creating a standardised sequence of $M_l$ frames.

### B. Processing

Once normalised, the data has undergone all the necessary pre-processing steps. Starting from the DL model, the core of the feature extraction stage is the implementation of the PointNet. This architecture is applied individually to each frame. It takes the $70$ points of the point cloud and processes them with two layers of a fully connected multi-layer perceptron. As each frame contains only one point cloud of the corresponding person walking, the result is a vector

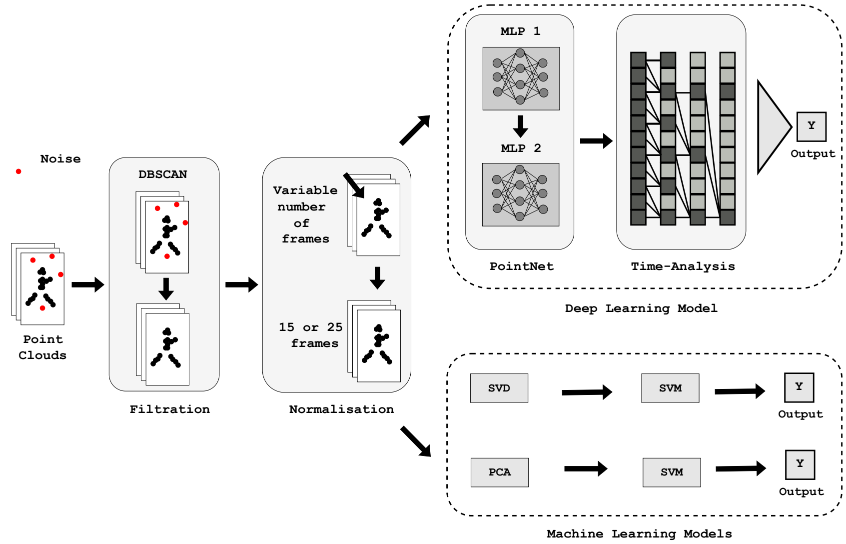


Fig. 1. The point clouds undergo two phases of preprocessing: filtration and normalisation. These processed inputs are then fed into three models. The DL model extracts feature vectors with PointNet, they pass through a temporal convolutional network, and a sigmoid function makes a binary classification. The second model uses SVD to extract meaningful information, which is then classified using SVM. Similarly, the third model uses PCA to extract mean features from the input, which are also classified using SVM.

of $128$ elements for each frame. These vectors are further analysed within a dilated temporal convolutional block. In this way, it is possible to extract the relevant information from each point cloud and to consider the temporal order of the sequence of point clouds as a factor that can contribute more information. We employ $3$ temporal convolution layers with filters of dimension $2$ and dilation rates of $1$, $2$, and $4$, respectively. After performing a temporal analysis, the model has a final layer to perform classification. This is a sigmoid function for binary classification, separating individuals into two groups according to different characteristics.

Taking a different approach to the methodology, we explore the implementation of ML models. We adopt a strategy where we combine the point clouds from multiple frames into a single entity rather than examining individual frames over time. The consolidated group is analysed by ML models to extract the most significant information. By training on this meaningful and amalgamated data, we aim to improve the generalisation capabilities of our models and compare their performance. The initial model used is SVD, which is applied to reconstruct the input data in lower dimensions by leveraging the largest singular values. After this transformation, we use the SVM with an RBF kernel for classification. The second ML model follows a similar path. It takes the combined set of points and applies PCA to extract the principal components of the input data. These components encapsulate the most crucial information from the point clouds, and SVM is then used for classification.

## III. Evaluation Setup

The experiments were conducted in a designated room, as shown in Figure 2. The line delineated by the crosses $A$ and $B$ indicates the farthest allowable distance within the designated area, guaranteeing that volunteers can participate in the necessary activities comfortably.

The volunteers were monitored using the AWR1443 Boost radar from TI. It was configured for short-range operation, incorporating clutter removal to facilitate the analysis of reflections exclusively from volunteers within the designated area while disregarding reflections from static objects. It was configured to generate $10$ point clouds per second to address memory constraints, as higher rates caused memory issues. In this study, we employed the Texas Instruments software in conjunction with the Robot Operating System (ROS) [16] to facilitate data acquisition and subsequent processing. Furthermore, we developed a model for point cloud data processing in Python, leveraging the TensorFlow $1.15$ version.

The experiments were divided into two segments: normal walking and fast walking. During the first phase, participants walked from point $A$ to point $B$, executed a turn when they arrived, and then returned to point $A$. This sequence was repeated five times at a typical pace, replicating a regular daily walk. In the following segment, participants reproduced

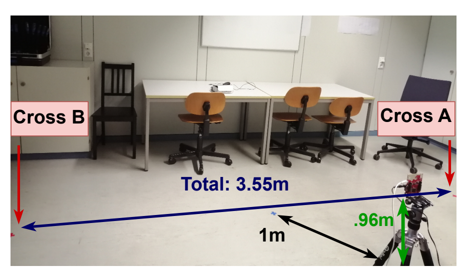


Fig. 2. Test room for the experiments.

the identical sequence from the first part; however, they were instructed to walk at a faster pace. This accelerated pace was repeated for an additional 5 iterations of the sequence.

After acquiring radar data from the participants, the procedure for collecting each anthropometric feature was as follows: The participants were asked to indicate their gender at birth for the gender attribute. Subsequently, the participants were weighed in kilograms using a calibrated scale, thus obtaining the weight attribute. To ascertain the height of the participants, they were instructed to stand in front of a wall while a measuring tape was used to determine the distance from the ground to the vertex of the head. To ascertain the arm span attribute, participants were instructed to extend their arms fully, after which the distance from the top of one hand's finger to the top of the other hand was measured. The waist attribute was determined by participants indicating the bottom of their ribs and the top of their hips; the waist circumference was then measured by placing a tape measure at the midpoint between these points and measuring around the waist.

### A. Dataset and testing

The study recruited 23 healthy volunteers with diverse personal attributes, comprising 13 males and 10 females. To ensure a uniform model for inferring information from the physical features of the volunteers and to work with a balanced dataset, the 23 volunteers were divided into two distinct groups (Group $A$ and Group $B$). These groups are defined by a division parameter that differs for each feature. The division parameter for each feature was set to ensure an even distribution, with one group consisting of 11 participants and the other group consisting of 12 participants. This approach enables using a uniform model for soft classification, in which the model predicts the corresponding group. The performance is evaluated using accuracy, which is defined as the ratio of correct predictions to the total number of predictions.

The radar was employed to record the walking activity of each participant, with the data collected from one cross to the other. This resulted in 10 recordings of individuals at normal walking speed and 10 recordings at high walking speed. Overall, 20 samples were collected per person. As stated in section II, data normalisation was conducted based on two criteria: the smallest $M_m$ and intermediate $M_l$ sample size whose size is between the smallest $M_m$ and the largest sample $M_x$. Afterwards, all collected samples were analysed to determine their sizes. The smallest sample obtained consisted of 15 frames, while the largest consisted of 35 frames. Therefore, the inference was performed for samples with sizes of $M_m = 15$ frames and $M_l = 25$ frames.

To circumvent overfitting, the data collected from 18 participants, designated as $D_{Train}$ set, was employed for the purpose of training the model. $D_{Train}$ encompasses both normal and fast walking scenarios from these 18 people. The remaining data from the remaining 5 participants, designated as $D_{Test}$ set, was reserved for testing purposes. This method ensures that the test data is gathered from participants who have not been used in the training phase, thus allowing us to evaluate the capability of our model for generalisation. The data split for training and validation was done randomly.

Owing to the limited amount of data accessible for the DL model, the arbitrary selection of subjects for training and testing data may cause irregularities in the model's performance. To address this issue, the inference process was executed five times, each time utilising a different randomly selected group. In each iteration, participants for $D_{Train}$ and $D_{Test}$ were selected at random, with the condition that no participant's data was included in both sets simultaneously.

## IV. Experimental Results

The results for the first variant of INFORMER, the DL model, are shown in Figure 3. As can be seen, height inference achieved 69.0% accuracy, weight inference achieved 73.0% accuracy, gender inference achieved 64.0% accuracy, waist inference achieved 62.0% accuracy, and arm span inference achieved 63.0% accuracy in the case of 15 frames. Standard deviations ranged from 2.1 to 2.45, indicating variability in the findings.

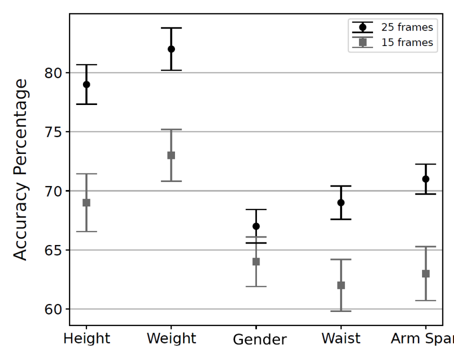


Fig. 3. DL model's inference accuracy with 15 and 25 frames.

Conversely, increasing the number of frames improved accuracy across most attributes. The gender attribute was the exception, as its confidence intervals overlap, indicating no substantial statistical difference. In addition, the standard deviations of these attributes decreased between 1.26 and 1.79, suggesting that the longer sequence of frames enhances consistency and accuracy in the inference. Height inference accuracy increased to 79.0% increase, weight inference increased to 82.0%, gender inference rose to 67.0%, waist inference to 69.0%, and arm span inference reached 71.0%.

In the presented results, it can be seen that using 25 frames improves both accuracy and standard deviation for four attributes. This indicates that a longer sequence of frames not only enhances accuracy but also leads to a more dependable and consistent inference for these attributes. However, the DL model was the only one to show a noticeable improvement with a greater number of frames. The ML models used did not exhibit significant improvement with additional frames. Therefore, we abstain from presenting a comparative analysis

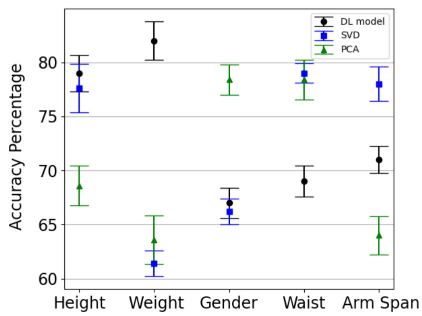


Fig. 4. Average inference accuracy.

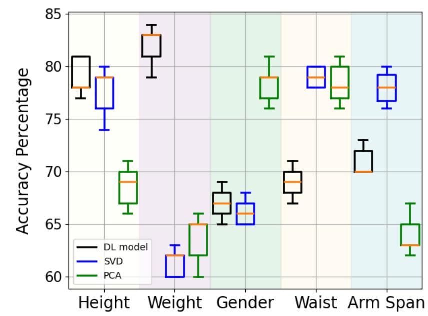


Fig. 5. Inference accuracy in 5 random test sets.

with a different number of frames, as was done with the DL model. Instead, we present the overall performance of all three models as presented in Figure 4.

Analysing the performance of the second variant of INFORMER with SVD, this model exhibited slightly lower accuracy in predicting gender, with an average of 66.2% and a lower standard deviation of 1.1%. Additionally, its performance in predicting height (average: 77.6%, standard deviation: 2.09) demonstrated increased variability compared to the DL model. The SVD model also exhibited poor performance in predicting weight (average: 61.4%, standard deviation: 1.9%). However, in comparison with the DL model, SVD performs better in the inferring waist and arm span, achieving accuracy levels of 79% and 78.2%, respectively.

The third variant of INFORMER, which uses PCA, demonstrated its best performance in predicting gender, with an average of 78.5% and a standard deviation of 1.87%. It also performed comparably well in inferring waist measurements (average performance: 78.1% and standard deviation: 2.24%), albeit slightly lower than SVD. However, it should be noted that PCA excels primarily in these two attributes. The performance of the other features is poor: height inference reaches 68.6% accuracy with a standard deviation of 1.8%, weight inference presents 63.6% of accuracy and 2.24% of standard deviation; finally, arm span inference shows a performance of 64% and standard deviation: 1.87%.

The results show that the DL model performed the best for two out of the five attributes: height and weight. In contrast, SVD performed the best on average in inferring waist and arm span. Although PCA performed the best for only one attribute, gender inference, it is worth noting that the other two models performed notably worse in this specific feature. In order to facilitate a more detailed analysis of the variability with regard to the impact of random selection on the training and validation data groups, as well as the distribution of outcomes, Figure 5 presents a box plot for each attribute across the three models. Starting from the results for the DL model, the box plots for the five attributes indicate a relatively consistent performance across datasets. However, one interquartile has a greater distance from the median than the others, specifically in the case of weight inference. The box plots for waist and gender accuracy exhibit similar levels of variability, as observed in the median and interquartiles.

SVD also presents no considerable variable results, where the interquartiles are close to the median in all attributes except for height. The variability in this specific feature is considerably higher than in the other performances. On the other hand, PCA presents more variable results. It is evident that not only are the standard deviations larger than the other two models, but also, in three out of the two models, there are interquartiles more distant from the respective median, suggesting more instability than the other two models.

Overall, the model showcases some variability across datasets, thereby highlighting the potential impact of dataset characteristics on the prediction. When evaluating the model's discriminative abilities, it is crucial to set a baseline performance. To establish a comparative benchmark, we have considered a basic distinguisher that is accurate up to 52.17% for parameters such as weight, height, waist, and arm span. For gender, the accuracy is 56.52%, which is attributable to an equal representation of male and female subjects in the database. Given this, the model outperforms the aforementioned baseline across all attributes, achieving accuracies above 60%. These results confirm the model's ability to capture information from the volunteer's personal attributes.

Our work revealed that an adversary could leverage the inherent broadcasting nature of these waves to gather data on individuals within a specific area. The analysis of these waves may yield information that could be used to infer soft biometric details. This information is of a sensitive nature, as it represents characteristics that are common to all human beings. Consequently, an attacker can construct a dataset comprising acquaintances, which can subsequently be employed in the identification of individuals not previously encountered by the model. In contrast with the approach outlined in the existing literature [17], [18], the model does not require the inclusion of previously known information about the target in the training data to infer information from a subject. Our methodology is not restricted to scenarios involving active spy mmWave radars. It is also applicable in situations where an attacker passively analyzes signals emitted by other devices. It is also noteworthy that the mmWave radar can operate at frequencies of approximately 80 GHz, enabling the detection of signals through obstacles or walls [19]. This capacity could facilitate scenarios in which an attacker is not physically inside the area under attack.

### A. Limitations

It is necessary to consider the attacker's perspective to address potential limitations. To illustrate, if the objective is to ascertain the victim's gender, the PCA-informer model may be employed. The DL-informer model is sufficient if the objective is to monitor the victim's weight over time to detect any increases. The three variants of the informant can be combined to obtain the maximum amount of information. Our findings illustrate the potential capabilities of an attacker in this domain. Further developments could analyse and infer additional information with enhanced accuracy.

To overcome the limitation of the dataset size, we have implemented rigorous data collection and preprocessing procedures to ensure the highest possible data quality. Notwithstanding the modest dataset size, our model can effectively learn and extract valuable information from the biometric data. We recognise the limitations of our approach and intend to address this by utilising larger datasets in future work.

## V. Conclusion

In this paper, we apply the INFORMER model with three variants to analyse gait and extract key personal attributes, including height, weight, gender, waist circumference, and arm span, for classification purposes. The approach utilises sparse point clouds generated from a mmWave radar, comprising a multi-step process that includes clustering-based filtering, normalisation, mean feature extraction, and dilated convolutions. The highest accuracy rates for the five attributes were achieved as follows: the DL model reached 79% accuracy for height and 82% for weight, the SVD model achieved 79% accuracy for waist and 78.2% for arm span, and the PCA model attained 78.5% accuracy for gender inference. This research highlights significant privacy vulnerabilities in mmWave radar systems, demonstrating that these systems can be exploited to covertly gather personal information without the user's consent. In contrast to previous methodologies that depend on intricate representations such as micro-Doppler signatures [20] or raw signals [17], our approach employs a streamlined representation comprising just 64 points and three coordinates. This suggests reduced processing demands and raises concerns about the simplicity of malicious exploitation. Further work should investigate additional scenarios and alternative data capture and analysis techniques.

## VI. Acknowledgments

This work was partially funded by the Topic Engineering Secure Systems of the Helmholtz Association (HGF) and supported by KASTEL Security Research Labs, Karlsruhe.

## References

[1] Z. Meng, S. Fu, J. Yan, H. Liang, A. Zhou, S. Zhu, H. Ma, J. Liu, and N. Yang, "Gait recognition for co-existing multiple people using millimeter wave sensing," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 01, 2020, pp. 849–856.

[2] Z. Shen, J. Nunez-Yanez, and N. Dahnoun, "Advanced millimeter-wave radar system for real-time multiple-human tracking and fall detection," *Sensors*, vol. 24, no. 11, p. 3660, 2024. [Online]. Available: https://www.mdpi.com/1424-8220/24/11/3660

[3] W. A. Ahmad, J. Wessel, H. J. Ng, and D. Kissinger, "Iot-ready millimeter-wave radar sensors," in *2020 IEEE Global Conference on Artificial Intelligence and Internet of Things (GCAIoT)*, 2020, pp. 1–5.

[4] W. Li, D. Zhang, Y. Li, Z. Wu, J. Chen, D. Zhang, Y. Hu, Q. Sun, and Y. Chen, "Real-time fall detection using mmwave radar," in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 16–20.

[5] S. Yue, Y. Yang, H. Wang, H. Rahul, and D. Katabi, "Bodycompass: Monitoring sleep posture with wireless signals," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 4, no. 2, jun 2020. [Online]. Available: https://doi.org/10.1145/3397311

[6] Y. Wang, W. Wang, M. Zhou, A. Ren, and Z. Tian, "Remote monitoring of human vital signs based on 77-ghz mm-wave fmcw radar," *Sensors*, vol. 20, no. 10, p. 2999, 2020.

[7] E. Cardillo, C. Li, and A. Caddemi, "Heating, ventilation, and air conditioning control by range-doppler and micro-doppler radar sensor," in *2021 18th European Radar Conference (EuRAD)*. IEEE, 2022, pp. 21–24.

[8] J.-T. González-Partida, P. Almorox-Gonzalez, M. Burgos-García, B.-P. Dorta-Naranjo, and J. I. Alonso, "Through-the-wall surveillance with millimeter-wave lfmcw radars," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 47, no. 6, pp. 1796–1805, 2008.

[9] J. L. Kröger, P. Raschke, and T. R. Bhuiyan, "Privacy implications of accelerometer data: a review of possible inferences," in *Proceedings of the 3rd International Conference on Cryptography, Security and Privacy*, ser. ICCSP '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 81–87. [Online]. Available: https://doi.org/10.1145/3309074.3309076

[10] K. Khan, S. U. Rehman, K. Aziz, S. Fong, and S. Sarasvady, "Dbscan: Past, present and future," in *The fifth international conference on the applications of digital information and web technologies (ICADIWT 2014)*. IEEE, 2014, pp. 232–238.

[11] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 652–660.

[12] J. Wu, Z. Zhu, and H. Wang, "Human detection and action classification based on millimeter wave radar point cloud imaging technology," in *2021 Signal Processing Symposium (SPSympo)*. IEEE, 2021, pp. 294–299.

[13] H. Xie, P. Han, C. Li, Y. Chen, and S. Zeng, "Lightweight midrange arm-gesture recognition system from mmwave radar point clouds," *IEEE Sensors Journal*, vol. 23, no. 2, pp. 1261–1270, 2022.

[14] J. Pegoraro and M. Rossi, "Real-time people tracking and identification from sparse mm-wave radar point-clouds," *IEEE Access*, vol. 9, pp. 78 504–78 520, 2021.

[15] A. Barliya, L. Omlor, M. A. Giese, A. Berthoz, and T. Flash, "Expression of emotion in the kinematics of locomotion," *Experimental brain research*, vol. 225, pp. 159–176, 2013.

[16] W. R. Drumheller and D. C. Conner, "Documentation and modeling of ros systems," in *SoutheastCon 2021*. IEEE, 2021, pp. 1–7.

[17] Y. Chen, R. Ou, Y. Deng, and X. Yin, "Wiage: A gait-based age estimation system using wireless signals," in *2021 IEEE Global Communications Conference (GLOBECOM)*. IEEE, 2021, pp. 01–06.

[18] Z. Wang, Z. Meng, K. Saho, K. Uemura, N. Nojiri, and L. Meng, "Deep learning-based elderly gender classification using doppler radar," *Personal and Ubiquitous Computing*, pp. 1–13, 2022.

[19] J.-T. González-Partida, P. Almorox-González, M. Burgos-García, B.-P. Dorta-Naranjo, and J. I. Alonso, "Through-the-wall surveillance with millimeter-wave lfmcw radars," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 47, no. 6, pp. 1796–1805, 2008.

[20] G. Garreau, C. M. Andreou, A. G. Andreou, J. Georgiou, S. Dura-Bernal, T. Wennekers, and S. Denham, "Gait-based person and gender recognition using micro-doppler signatures," in *2011 IEEE Biomedical Circuits and Systems Conference (BioCAS)*. IEEE, 2011, pp. 444–447.