Human Pose and Seat Occupancy Classification with Commercial MMWave WiFi

Yu, Jianyuan; Wang, Pu; Koike-Akino, Toshiaki; Wang, Ye; Orlik, Philip V.

Abstract

Our previous studies introduced a mid-grained intermediate-level channel measurement — spatial beam signalto-noise ratios (SNRs) that are inherently available and defined in the 60-GHz IEEE 802.11ad/ay standards — for the fingerprinting-based indoor localization. In this paper, we take one step further to use the mid-grained channel measurement for human monitoring applications including human pose and seat occupancy classifications. The effectiveness of the mid-grained channel measurement is validated by an in-house experimental dataset that includes 5 separate data collection sessions using classical classification methods and modern deep neural networks. Our preliminary result shows that mmWave beam SNRs are capable of delivering high classification accuracy above 90%.

IEEE Global Communications Conference (GLOBECOM)

Human Pose and Seat Occupancy Classification with Commercial MMWave WiFi

Jianyuan Yu[1,2], Pu Wang[2], Toshiaki Koike-Akino[2], Ye Wang[2], Philip V. Orlik[2], and Haijian Sun[3]

[1]*Bradley Department of Electrical and Computer Engineering, Virginia Tech, Blacksburg, VA 24061*
[2]*Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA 02139, USA*
[3]*University of Wisconsin-Whitewater, Whitewater, WI 53190, USA*

Abstract—Our previous studies introduced a mid-grained intermediate-level channel measurement — spatial beam signal-to-noise ratios (SNRs) that are inherently available and defined in the 60-GHz IEEE 802.11ad/ay standards — for the fingerprinting-based indoor localization. In this paper, we take one step further to use the mid-grained channel measurement for human monitoring applications including human pose and seat occupancy classifications. The effectiveness of the mid-grained channel measurement is validated by an in-house experimental dataset that includes 5 separate data collection sessions using classical classification methods and modern deep neural networks. Our preliminary result shows that mmWave beam SNRs are capable of delivering high classification accuracy above 90%.

Index Terms—Millimeter-wave, WiFi sensing, beam training, beam SNR, CSI, human monitoring, deep learning.

I. INTRODUCTION

Radio-frequency (RF)-based human activity monitoring has gained attention over the past decade due to the decreasing cost and less privacy concerns (compared with camera-based approaches). WiFi-band (e.g., sub 6-GHz) and mmWave frequency-modulated continuous-wave (FMCW) signals from dedicated sensors have been used to track people behind the wall, determine personal identity, estimate poses/gestures, and track two-/three-dimensional (2D/3D) skeleton movements [1]–[6]. Almost all these studies built on their machine learning algorithms via *direct* measurements such as the angle-range-Doppler spectrum or its subsequent 2D/3D detection points at the output of these RF sensors to learn classification-related or regression-related features.

On the other hand, commercial WiFi signals at low frequency (2.4 and 5 GHz) bands are often *indirect* measurements of human activity using either coarse-grained received signal strength indicator (RSSI) or fine-grained channel state information (CSI). The conventional RSSI measurement suffers from the measurement instability and coarse granularity of the channel information, leading to limited accuracy for sensing applications. The CSI measurements for the IEEE 802.11n/ac standards are more fine-grained as defined as a group of complex amplitudes over sub-carriers. Instead of learning features from the direct measurement as in the case of dedicated FMCW sensors, the CSI measurements are trained via supervised learning or cross-modal deep learning for human sensing tasks such as device-free localization, activity recognition,

fall detection, personal identification, emotion sensing, and skeleton tracking [7]–[18]. Residual neural networks (ResNet) were applied in [16] to simultaneously perform gesture recognition and localization, and also predict the skeleton as a regression problem with the help of generative adversarial networks (GAN). More recently, the cross-modal learning approach showed the potential of commercial WiFi signals for human activity sensing applications. For example, [15] used annotations from camera images to train fine-grained CSI measurements over 30 subcarriers and 5 frames from three pairs of transmitting and receiving antennas. [19] explored the 3D velocity profile and distinguished posture-specific features from the static objects via recurrent neural networks (RNN).

Nevertheless, explicit utilization of channel measurements from mmWave WiFi devices for human activity monitoring was not yet reported in the literature, except our previous work on indoor localization [20]–[22]. Particularly, we introduced a mid-grained channel measurement — mmWave spatial beam signal-to-noise ratios (SNRs) that are inherently available and defined in the 60-GHz IEEE 802.11ad/ay standards — for this purpose. This mid-grained channel measurement provides the mmWave channel quality indicator as a function of probing beams. Compared with the fine-grained channel measurements at sub 6-GHz (e.g., CSI over sub-carriers), the beam SNR can be considered as a type of direct spatial-domain channel measurement which may provide more informative angle features; e.g., angle of arrival (AoA) and angle of departure (AoD). Moreover, the beam SNR measurement is easier to access with zero overhead as it is required to be reported for mmWave beam training under the protocol of IEEE 802.11ad/ay standards. In this paper, we take one step further to use the mid-grained channel measurement from the indoor localization application towards human sensing applications, particularly human pose and seat occupancy classification. The effectiveness of the channel measurement is validated by an in-house experimental dataset that includes 5 separate data collection sessions, using conventional classification methods and the deep neural network (DNN).

II. HUMAN ACTIVITY MONITORING USING MID-GRAINED CHANNEL MEASUREMENT

A. Beam SNR in IEEE 802.11ad/ay Standards

At mmWave frequency bands, WiFi signals in the 802.11ad/ay standards experience significantly larger path
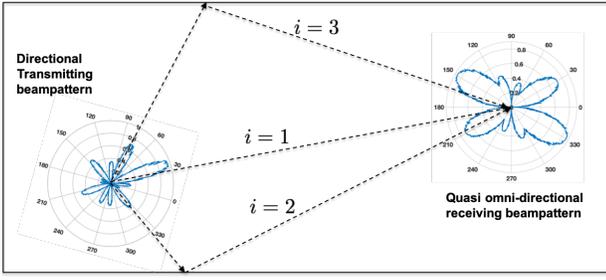
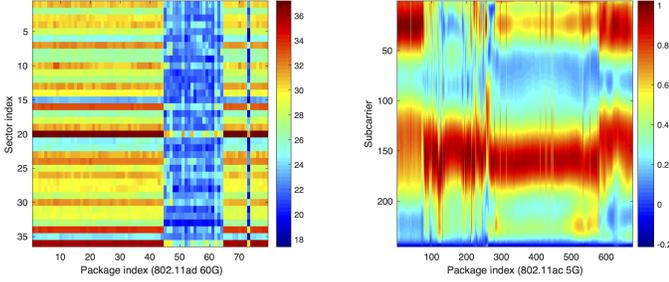Fig. 1: Illustration of beam SNR measurements as a function of transmitting and receiving beampatterns.

Fig. 2: The mid-grained beam SNR measurements at 60-GHz (left) and fine-grained CSI measurements at 5-GHz when a person walks through the monitoring area.

losses compared with those at sub 6-GHz frequency bands. As a result, beam training using a series of pre-defined directional beampatterns is required to identify propagation paths and establish the wireless link. For each directional beam, a beam SNR is collected and reported as a measure of spatial channel quality. For a given pair of transmitting and receiving beam patterns, the beam SNR can be defined as [22]

$$h_m = \frac{1}{\sigma^2} \sum_{i=1}^{I} \gamma_m(\theta_i)\zeta_m(\psi_i)P_i, \quad (1)$$

where m is the index of beampattern, I is the total number of paths, θ_i and ψ_i are the transmitting and receiving azimuth angles for the ith path, respectively. Here, P_i is the signal power at the ith path, $\gamma_m(\theta_i)$ and $\zeta_m(\psi_i)$ are the transmitting and receiving beampattern gains at the ith path for the mth beampattern, respectively, and σ^2 is the noise variance. In commercial-off-the-shelf (COTS) 802.11ad devices, the beam SNR is usually reported in dB with a certain quantization level.

To first validate the use of this mid-grained channel measurement for human activity monitoring, we provide a qualitative visualization of the mmWave beam SNRs in Fig. 2. The collected beam SNR measurements correspond to a simple experiment with a pair of COTS 802.11ad devices when one subject keeps still at the beginning for approximately 5 seconds, then moves one step forward in a few seconds, and stops at the end for another 5 seconds. As clearly shown in the left plot of Fig. 2, the beam SNRs over all beam sectors experienced a sudden drop, possibly due to the human blockage of mmWave WiFi propagation paths, while they are
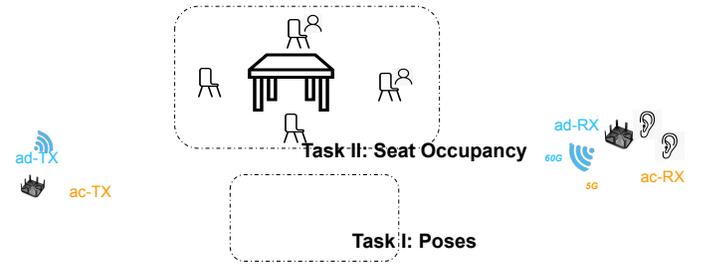

Fig. 3: Experimental setup for Task I of human pose recognition and Task II of seat occupancy sensing.

relatively stable at the beginning and end of the data collection period. This example suggests that, like the fine-grained CSI measurements at 5-GHz (see the right plot of Fig. 2), the mid-grained beam SNR measurement may be qualified as fingerprints for human activity sensing.

B. Human Activity Classification by mmWave Beam SNR

In the following, we utilize the beam SNRs to construct the offline training dataset for two scenarios: Task I of human pose recognition and Task II of seat occupancy sensing. We then apply conventional classification and deep learning methods to human activity recognition.

Fig. 3 shows the experimental setup for both tasks using a pair of COTS 60-GHz 802.11ad devices. For latter performance comparison, we also place a pair of 802.11ac routers communicating over 5-GHz links on the assigned channel. For Task I (human pose recognition), a human subject makes 8 distinct poses including gestures like 'stand', 'sit', and 'lift a hand' as shown in Fig. 4(a). The transmitter (TX) and receiver (RX) routers are placed on different stands with a height of 1.20 meters at a distance of approximately 2 meters. We repeat the data collection at 5 separate sessions such that we can train our classification model with one session and test on the other separate sessions.
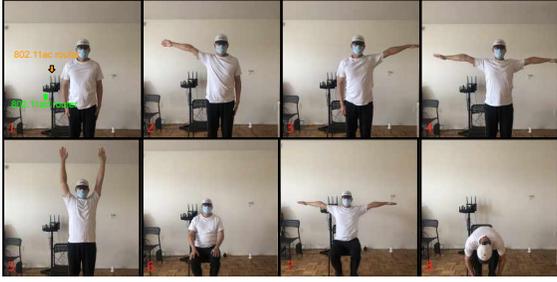
For Task II (seat occupancy sensing), considering the demand of social distancing monitoring in the post-COVID society, we are interested in multiple-seats, multiple-people sensing. The task is to figure out which seats are occupied. In this scenario, at most 2 subjects may be present in one of 4 chairs around the table in the middle. Particularly, we design 8 different combinations about how the 2 people sit at 4 chairs, as illustrated in Fig. 4(b). In this task, the TX and RX routers were 4 meters apart, and the distance between each occupancy varies from about 0.50 to 2.0 meters. The data collection is repeated for 5 sessions as well.
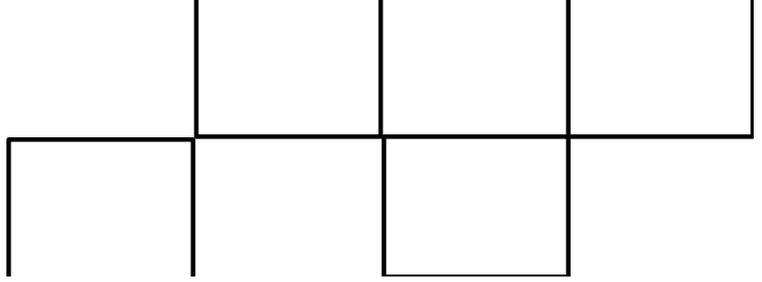
C. Offline Training Dataset

To construct the offline training dataset, we stack all SNR measurements from all beam sectors as a training sample:

$$\mathbf{h} = [h_1, h_2, \ldots, h_M]^T, \quad (2)$$

where M is the number of beampatterns used for beam training and $[\cdot]^T$ denotes the transpose. In the case of multiple

(a) Task I: Human Pose Recognition

(b) Task II: Seat Occupancy Sensing

Fig. 4: Photo snapshots of (a) Task I: human pose recognition with 8 poses and (b) Task II: seat occupancy sensing for 8 patterns with 4 seats and 2 people.

access points (APs), we combine beam SNR measurements from all APs to form the training sample:

$$\tilde{\mathbf{h}} = [\mathbf{h}_1^T, \mathbf{h}_2^T, \ldots, \mathbf{h}_P^T]^T \in \mathbb{R}^{MP \times 1}, \quad (3)$$

where P is the number of APs. In our experiment, we use one AP with 36 beam SNRs which gives rise to $M = 36$ and $P = 1$.

For a given activity, R fingerprinting snapshots, $\tilde{\mathbf{h}}_1(l), \ldots, \tilde{\mathbf{h}}_R(l)$, are collected to construct the offline training dataset, where l is the label index for the activity. By collecting many realizations of beam SNR measurements over L activities, we have L sets of $MP \times R$ beam SNR measurements for the offline training dataset.

D. Online Human Activity Recognition

When new beam SNR measurements are available, our task is to identify the human activity among the trained activities. To this end, we apply both the conventional classification and deep learning methods for the human activity recognition.

1) Conventional Classification Methods: We compare several classic machine learning methods including 1) the k-nearest neighbor (kNN) that calculates the distance between the new measurement and all training samples to pick the most common label from the k nearest neighbors, 2) the support-vector machine (SVM), 3) linear discriminant analysis (LDA), 4) quadratic discriminant analysis (QDA), and 5) decision tree (DT) that predicts the label using a tree-like structure. These methods will serve as benchmarks for the following deep learning-based classification method.

2) Deep Learning Methods: We use a feed-forward DNN with 4 fully-connected hidden layers, each of width N_w. The input vector of beam SNRs $\tilde{\mathbf{h}}$ is mapped to the first hidden layer $\mathbf{y}_0$, via the transformation given by $\mathbf{y}_0 = \phi(\mathbf{W}_{\text{input}}\tilde{\mathbf{h}} + \mathbf{b}_{\text{input}})$, where the trainable parameters are the weight matrix $\mathbf{W}_{\text{input}} \in \mathbb{R}^{N_w \times MP}$ and bias vector $\mathbf{b}_{\text{input}} \in \mathbb{R}^{N_w \times 1}$, and ϕ denotes the element-wise application of the rectified linear unit (ReLU) as the non-linear activation function. Then, $\mathbf{y}_0$ is similarly processed to create a further $N_d = 3$ hidden layers, via

$$\mathbf{y}_l = \phi(\mathbf{W}_l\mathbf{y}_{l-1} + \mathbf{b}_l), \quad l = 1, \ldots, N_d, \quad (4)$$
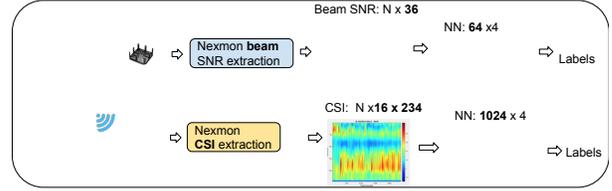

Fig. 5: Flow chart of collecting data and feeding into neural network for prediction.

where the weight matrices $\mathbf{W}_l \in \mathbb{R}^{N_w \times N_w}$ and bias vectors $\mathbf{b}_l \in \mathbb{R}^{N_w \times 1}$. Finally, an L-dimensional output score vector is generated via $\mathbf{u} = \mathbf{W}_{\text{output}}\mathbf{y}_{N_d} + \mathbf{b}_{\text{output}}$. For both human pose and seat occupancy recognition tasks, we have $L = 8$.

For each training input, the corresponding output of the last layer $\mathbf{u}$ is first normalized with the softmax operation to produce the likelihood vector $\mathbf{s} = [s_1, s_2, \ldots, s_L]$, as given by,

$$s_n = \exp(u_n) / \sum_{i=1}^{L} \exp(u_i), \quad n \in \{1, 2, \ldots, L\}, \quad (5)$$

Then, the cross-entropy loss function is computed with respect to the corresponding training label $c \in \{1, \ldots, L\}$ as

$$\ell_{\text{cross-entropy}} = -\log(s_c). \quad (6)$$

The average probability of successful classification (or accuracy) is calculated by the ratio between the number of correct estimations and total samples, i.e., $\Pr(\arg\max_i s_i = c)$ where $\Pr(\cdot)$ denotes the sample probability that the argument event is true.

III. PERFORMANCE EVALUATION

In this section, we will present the results for pose and occupancy classification tasks using the mid-grained beam SNR measurements. For comparison, we include the classification results using the fine-grained CSI measurements at 5-GHz.

A. Testbed and Data Collection

Fig. 5 shows the flow chart of collecting data and feeding into neural networks for prediction. The testbed consists of a

pair of 802.11ad routers (Talon AD 7200), a pair of 802.11ac routers (ASUS RT-AC86U), and a Ubuntu desktop as the host. Particularly, the Talon AD7200 routers use a Qualcomm QCA9500 60-GHz chipset that comes with a phased antenna array of 32 antenna elements and fully implements the IEEE 802.11ad standard. To extract the beam SNR measurements from Talon AD7200 routers, we used the *nexmon* firmware patching framework of [23] and followed the work in [24]–[26]. We applied medium access control (MAC) address or specific internet protocol (IP) filters to capture packages from expected transmitters.

The participants will conduct 5 sessions at different times with 8 poses for each session. The first session lasts 180 seconds for each pose with minor movements like leaning the body or moving the fingers. While for all other four sessions, each case lasts 18 seconds and participants try to remain fixed for each pose. As there is a 2 to 10 minutes delay between each session, we move out the data from the routers and check the data validity. We train the first session, while validating and testing on the other sessions. The final results are averaged over all testing sessions. Eventually, we obtained approximately 16k training samples and 4k testing samples.

B. Channel State Information (CSI)

For performance comparison, we also extract the fine-grained CSI measurements from the 802.11ac-compliant ASUS RT-AC86U routers by following the work in [27] with the *nexmon* firmware. At the sub 6-GHz band, the CSI can propagate through the wall and register the activities that are beyond the scene-of-interest. We list the comparison between beam SNR and CSI measurements in Table. I. For a fair comparison, we downsample the CSI by finding the nearest CSI measurement set, which is a batch of CSI with complete

TABLE I: Comparison between beam SNR and CSI

Type	beam SNR	CSI
frequency band (GHz)	60	2.4 / 5
bandwidth (Mbps)	1600	80
package rate (/second)	10	500
data dimension	36	16×234

TABLE II: Classification accuracy (in %) results of different methods and WiFi signals

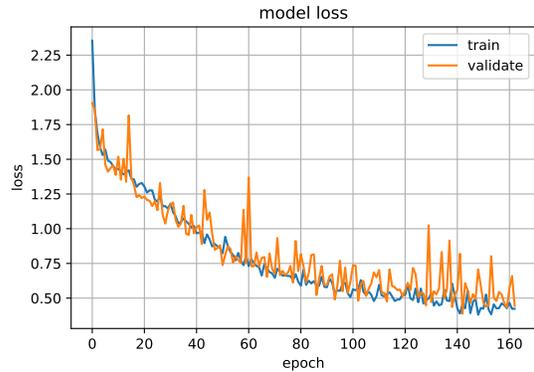| Methods | Task I | | Task II | |
	Beam SNR	CSI	Beam SNR	CSI
kNN	81.3	70.3	83.0	76.2
SVM	79.8	69.5	88.3	78.1
LDA	76.2	61.2	84.9	75.2
QDA	76.4	68.6	83.4	78.5
DT	56.0	55.2	46.4	36.2
DNN	**88.8**	79.7	**91.2**	80.9

Fig. 6: Loss function trajectory on the poses classification task using the mid-grained beam SNRs.

spatial indexing from 1 to 16, to the beam SNR time slot. Therefore, both CSI and beam SNR datasets have the same number of training samples.

For the CSI measurements, we apply standard preprocessing steps such as removing null subcarriers [23], shifting the sub-carrier amplitudes, and applying multi-variable regression [28] to smooth out the outlier.

C. Classification Performance with Different Methods

We applied classic classification methods, i.e., kNN with $k = 3$, SVM with a linear kernel, LDA, QDA and DT, to both the mid-grained mmWave beam SNRs and fine-grained CSI measurements. The deep learning method, we choose a validation rate of 0.20 out of the training dataset, set the early-stop with a patience of 20, and Adam as the optimizer. With 4 fully-connected hidden layers, we set the layer width to 64 for the beam SNRs and 1024 for the CSI measurement due to its large input dimension. Fig. 6 shows the training and validation losses over epochs which gradually converge with an early-stopping.

The result of pose recognition is summarized in Tab. II and Fig. 7 in terms of the confusion matrix $\mathbf{C}$:

$$\mathbf{C}(i,j) = \frac{1}{T_j} \sum_{t=1}^{T_j} \mathbb{1}[\hat{l}(\tilde{\mathbf{h}}_t(j)) = i], \qquad (7)$$

where i and j are indices, respectively, for the estimated and true labels (i.e., poses and occupancy patterns), and T_j is the number of samples in the test dataset for the index j. In addition, $\hat{l}(\tilde{\mathbf{h}}_t(j))$ is the pose/occupancy estimate by using the tth sample batch from the test data collected at jth label. It is seen that the use of the mid-grained beam SNRs offers better performance than the fine-grained CSI measurements in both tasks. The deep learning method gives the best performance among all considered classification methods.

D. Impact of Routers Location and Human Size

It is noted that the beam SNR measurements may be sensitive to the human body size and subject to geometry factors such as the height and orientation of the mmWave

Fig. 8: Confusion matrix of the pose recognition for a subject with a relatively small body size.
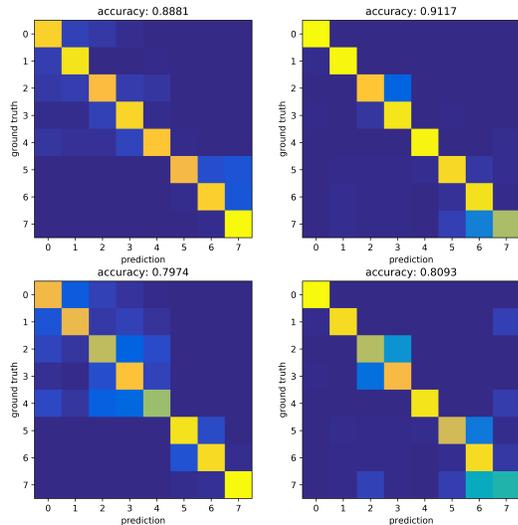
Fig. 7: Confusion matrices of Task I of human pose recognition with beam SNR (top-left) and CSI (bottom-left) and Task 2 of seat occupancy classification with beam SNR (top-right) and CSI (bottom-right). Beam SNR shows more robust performance over different sessions.

WiFi router, and the distance between the transmitter and receiver. Compared with the 802.11ac setup, low platform height and large distance between the transmitter and receiver can degrade the classification performance with the mmWave beam SNRs. For example, Fig. 8 shows degraded pose recognition performance for a subject with a smaller body size than the one in Fig. 7. Nevertheless, one can still distinguish different groups of activities, e.g., between standing (i.e., the categories $\{1, 2, \cdots, 5\}$) and sitting (i.e., the categories $\{6, 7, 8\}$) or between lifting one hand (the categories $\{2, 3\}$) and lifting both hands (the categories $\{4, 5\}$).

IV. CONCLUSION AND FUTURE WORK

Our preliminary study shows the potential of using the mid-grained channel measurement, i.e., mmWave beam SNRs, for human activity monitoring, particularly, pose recognition and seat occupancy sensing using commercial-off-the-shelf mmWave WiFi routers. With the same amount of data from multiple separate sessions, the classification performance with the beam SNRs is better than the fine-grained CSI measurement. We also discuss possible limiting factors, e.g., human size and router location, to the classification performance.

REFERENCES

[1] F. Adib, C.-Y. Hsu, H. Mao, D. Katabi, and F. Durand, "Capturing the human figure through a wall," *ACM Trans. Graph.*, vol. 34, no. 6, Oct. 2015.

[2] C.-Y. Hsu, Y. Liu, Z. Kabelac, R. Hristov, D. Katabi, and C. Liu, "Extracting gait velocity and stride length from surrounding radio signals," in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 2017, pp. 2116–2126.
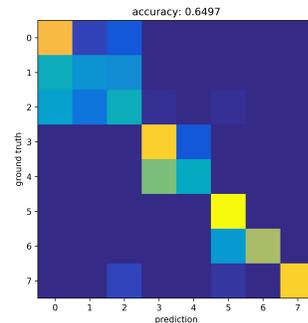
[3] M. Zhao, T. Li, M. A. Alsheikh, Y. Tian, H. Zhao, A. Torralba, and D. Katabi, "Through-wall human pose estimation using radio signals," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2018, pp. 7356–7365.

[4] M. Zhao, Y. Tian, H. Zhao, M. Alsheikh, T. Li, R. Hristov, Z. Kabelac, D. Katabi, and A. Torralba, "RF-based 3D skeletons," in *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, 2018, pp. 267–281.

[5] P. Zhao, C. X. Lu, J. Wang, C. Chen, W. Wang, N. Trigoni, and A. Markham, "mID: Tracking and identifying people with millimeter wave radar," in *2019 15th International Conference on Distributed Computing in Sensor Systems (DCOSS)*, May 2019, pp. 33–40.

[6] A. Singh, S. Sandha, L. Garcia, and M. Srivastava, "RadHAR: Human activity recognition from point clouds generated through a millimeter-wave radar," in *Proceedings of the 3rd ACM Workshop on Millimeter-Wave Networks and Sensing Systems*, 2019, pp. 51–56.

[7] X. Lu, H. Wen, H. Zou, H. Jiang, L. Xie, and N. Trigoni, "Robust occupancy inference with commodity WiFi," in *2016 IEEE 12th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*, Oct 2016, pp. 1–8.

[8] Y. Zeng, P. H. Pathak, and P. Mohapatra, "WiWho: WiFi-based person identification in smart spaces," in *2016 15th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, April 2016, pp. 1–12.

[9] D. Wu, D. Zhang, C. Xu, H. Wang, and X. Li, "Device-free WiFi human sensing: From pattern-based to model-based approaches," *IEEE Communications Magazine*, vol. 55, no. 10, pp. 91–97, Oct 2017.

[10] H. Zou, Y. Zhou, J. Yang, W. Gu, L. Xie, and C. Spanos, "WiFi-based human identification via convex tensor shapelet learning," in *Proceedings of the 2018 AAAI Conference on Artificial Intelligence*, 2018, pp. 1711–1718.

[11] J. Yang, H. Zou, H. Jiang, and L. Xie, "Device-free occupant activity sensing using WiFi-enabled IoT devices for smart homes," *IEEE Internet of Things Journal*, vol. 5, no. 5, pp. 3991–4002, Oct 2018.

[12] Y. Gu, T. Liu, J. Li, F. Ren, Z. Liu, X. Wang, and P. Li, "Emosense: Data-driven emotion sensing via off-the-shelf WiFi devices," in *2018 IEEE International Conference on Communications (ICC)*, May 2018, pp. 1–6.

[13] Y. Gu, X. Zhang, Z. Liu, and F. Ren, "Besense: Leveraging WiFi channel data and computational intelligence for behavior analysis," *IEEE Computational Intelligence Magazine*, vol. 14, no. 4, pp. 31–41, Nov 2019.

[14] Z. Wang, K. Jiang, Y. Hou, W. Dou, C. Zhang, Z. Huang, and Y. Guo, "A survey on human behavior recognition using channel state information," *IEEE Access*, vol. 7, pp. 155 986–156 024, 2019.

[15] F. Wang, J. Feng, Y. Zhao, X. Zhang, S. Zhang, and J. Han, "Joint activity recognition and indoor localization with WiFi fingerprints," *IEEE Access*, vol. 7, pp. 80 058–80 068, 2019.

[16] F. Wang, S. Zhou, S. Panev, J. Han, and D. Huang, "Person-in-WiFi: Fine-grained person perception using WiFi," in *Proceedings of (ICCV) International Conference on Computer Vision*, July 2019.

[17] Y. Ma, G. Zhou, and S. Wang, "WiFi sensing with channel state

information: A survey," *ACM Computing Surveys (CSUR)*, vol. 52, no. 3, pp. 1–36, 2019.

[18] L. Zhang, X. Ruan, and J. Wang, "WiVi: A ubiquitous violence detection system with commercial WiFi devices," *IEEE Access*, vol. 8, pp. 6662–6672, 2020.

[19] W. Jiang, H. Xue, C. Miao, S. Wang, S. Lin, C. Tian, S. Murali, H. Hu, Z. Sun, and L. Su, "Towards 3D human pose construction using WiFi," in *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, 2020, pp. 1–14.

[20] P. Wang, M. Pajovic, T. Koike-Akino, H. Sun, and P. Orlik, "Fingerprinting-based indoor localization with commercial mmwave WiFi–Part II: Spatial beam SNRs," in *2019 IEEE Global Communications Conference (GLOBECOM)*, Dec 2019.

[21] M. Pajovic, P. Wang, T. Koike-Akino, H. Sun, and P. V. Orlik, "Fingerprinting-based indoor localization with commercial MMWave WiFi–Part I: RSS and beam indices," in *2019 IEEE Global Communications Conference (GLOBECOM)*, Dec. 2019.

[22] T. Koike-Akino, P. Wang, M. Pajovic, H. Sun, and P. V. Orlik, "Fingerprinting-based indoor localization with commercial MMWave WiFi: A deep learning approach," *IEEE Access*, vol. 8, pp. 84 879–84 892, 2020.

[23] M. Schulz, D. Wegemer, and M. Hollick, "Nexmon: The c-based firmware patching framework," 2017.

[24] D. Steinmetzer, D. Wegemer, and M. Hollick, "Talon tools: The framework for practical IEEE 802.11 ad research," 2017.

[25] G. Bielsa, J. Palacios, A. Loch, D. Steinmetzer, P. Casari, and J. Widmer, "Indoor localization using commercial off-the-shelf 60 GHz access points," in *IEEE INFOCOM 2018*, 2018, pp. 2384–2392.

[26] S. K. Saha, H. Assasa, A. Loch, N. M. Prakash, R. Shyamsunder, S. Aggarwal, D. Steinmetzer, D. Koutsonikolas, J. Widmer, and M. Hollick, "Fast and infuriating: Performance and pitfalls of 60 GHz WLANs based on consumer-grade hardware," in *2018 IEEE SECON*, 2018, pp. 1–9.

[27] F. Gringoli, M. Schulz, J. Link, and M. Hollick, "Free your CSI: A channel state information extraction platform for modern Wi-Fi chipsets," ser. WiNTECH 19, 2019, p. 2128.

[28] A. Sobehy, E. Renault, and P. Muhlethaler, "CSI based indoor localization using ensemble neural networks," in *International Conference on Machine Learning for Networking*. Springer, 2019, pp. 367–378.