Radar Point Cloud-Based Continuous Human Activity Classification Using Rényi Entropy Segmentation Methods

Nicolas C. Kruse, *Member, IEEE*, Alec Daalman, Francesco Fioranelli, *Senior Member, IEEE*, and Alexander Yarovoy, *Fellow, IEEE*

***Abstract*—Classification of human activities performed sequentially and with unconstrained durations using radar sensors has been studied in this work. A novel processing pipeline comprising a sequence segmentation stage, a segment processing stage, and a classification stage has been proposed to address this challenge. Specifically, the segmentation stage has been implemented by monitoring Rényi entropy for fluctuations in the radar data, with the entropy, derived from micro-Doppler spectrograms, functioning as a descriptive quantity of the activity being performed. The method has been experimentally verified on a challenging, publicly available dataset collected with a network of five simultaneously operating pulsed ultrawideband radars. Classification performance has been compared to reference works in the literature on the same dataset, and a test accuracy and macro *F1*-score of 89.3% and 82.0% have been, respectively, demonstrated.**

***Index Terms*—Human activity classification, radar network, Rényi entropy.**

I. Introduction

AUTOMATED human monitoring is a beneficial capability for healthcare professionals. Systems that offer these capabilities can monitor, e.g., vital signs [1], [2], [3], detect harmful events such as falls [4], [5], and can ensure a timely response of relevant personnel to assist vulnerable people or alert family and carers. Wearable sensors, such as inertial measurement units (IMUs), have been utilized for these tasks [6], [7] but are not always feasible, as subjects may forget or object to wear these devices. Camera-based observation is highly dependent on lighting conditions, with darkness and glaring reflections potentially impeding proper operation. In addition, video-based monitoring comes with elevated privacy concerns. In this context, radar sensing presents a remote monitoring solution that can overcome the limitations of the aforementioned sensor modalities and offers a promising, versatile platform for human monitoring tasks.

Received 20 December 2024; revised 30 April 2025 and 2 July 2025; accepted 9 July 2025. Date of publication 17 July 2025; date of current version 30 July 2025. This work was supported in part by the NWO KLEIN *RAD-ART* Project. *(Corresponding author: Nicolas C. Kruse.)*

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by TU Delft HREC under Application No. 1387.

The authors are with the Microwave Sensing, Signals and Systems (MS3) Group, Delft University of Technology, 2628 CD Delft, The Netherlands (e-mail: n.c.kruse@tudelft.nl; a.daalman@student.tudelft.nl; f.fioranelli@tudelft.nl; a.yarovoy@tudelft.nl).

Digital Object Identifier 10.1109/TRS.2025.3590161

Radar-based classification of activities of daily living (ADL) is, therefore, an active area of research, with a considerable amount of efforts being directed to the open challenge of continuous classification. Continuous classification here refers to the classification of extended sequences of human activities, with each activity of unknown duration, and with activities smoothly transitioning into each other. Three main approaches can be identified in the current literature on continuous activity classification: sliding window methods, recurrent neural networks (RNN) or similar models to process the entire data sequence, and segmentation-based methods. Sliding window approaches include, for instance, the work in [8], where a set of features is computed from a time-windowed amount of data, and used as input to classifiers such as, among others, support vector machines (SVMs). In [9], a coarse sliding window (i.e., 30 with 10-s overlap) is employed in conjunction with a convolutional neural network (CNN). Approaches where activity sequences are processed in their entirety include such classifiers as (Bidirectional) RNNs [(Bi)RNNs] [10], [11], hybrid models consisting of CNNs and (Bidirectional) long short-term memory [(Bi)LSTM] networks [12], [13], gated recurrent unit networks [14], and transformer-based models [15].

The approaches to continuous classification that focuses on segmentation of activity sequences generally do so by discriminating periods of motion from those without motion. In [16], the beginning and the end of an activity are identified based on fluctuations in Wi-Fi channel state information (CSI) variance. A similar approach is taken in [17], where the number of detections from the radar data is used as an indicator of an activity starting and stopping. The short-term average/long-term average (STA/LTA) change detection algorithm is employed in [18] and [19] to segment the original sequence into motion-detected intervals based on the spectrogram envelope. Motion-detected intervals can also be determined through machine learning methods, as demonstrated in [20]. The activity sequence is divided into fixed-size segments, and a CNN-LSTM network is utilized to determine the presence of an activity. A dynamic segment duration algorithm is proposed in [6], where an initial 1 s data segment is processed to yield information on the bandwidth. If the 1 s segment appears very dynamic, the segment length is extended.

The reviewed existing approaches in the literature for continuous classification of ADL have notable disadvantages. Sliding window methods with a nonzero overlap require input data to be processed multiple times, degrading the computational efficiency of these solutions. Furthermore, window sizes are primarily fixed, which may not necessarily be optimal for the classification of different activity types in realistic sequences. RNNs, transformers, and other neural networks that process complete sequences are generally less computationally lightweight than their counterparts that are designed to classify comparatively shorter data sequences with single activities and, in general, require a large amount of data for effective training. Among the segmentation methods, the prevalent means of identifying individual activities are the presence of a pause in motion between subsequent activities, revealed by the lack of returned power or detections in the data under test. This pause can be absent for activities that smoothly transition into the succeeding ones, often without a clear stop of the body movement from one to another.

To address the above issues, in this work, a continuous ADL classification method is proposed, which consists of three main components: a segmentation algorithm, a segment processing algorithm, and a classification network. Segmentation of the input activity sequences is based on a quantity derived from the micro-Doppler spectrograms, namely, the Rényi entropy [21]. This scalar quantity is constructed to be representative of the distribution of velocities at a given time [22], and monitoring this quantity for fluctuations gives an indication of the transition between activities. Classification of the individual segments extracted in this way is then achieved by first processing the data into a point cloud (PC) representation and then utilizing a point transformer (PT) network inspired by [23]. The proposed method, as well as two alternative segmentation methods, is validated on a publicly available experimental dataset, which consists of a variety of sequences of nine human activities. It is demonstrated that favorable performance metrics can be achieved on a challenging test set, with a test accuracy and macro *F1*-score of 89.3% and 82.0%, respectively. These metrics are an improvement on previous approaches validated on the same dataset, as detailed in Section IV. This research is an extension of our feasibility study in [24], where only segmentation in isolation has been investigated. The current work expands on this by completing an end-to-end classification approach and studying the suitability of the proposed segmentation method.

The contributions in this work can be summarized as follows.

1) A novel approach for classification of continuous sequences of ADL, based on segmentation with Rényi entropy, a quantity describing more complex fluctuations in the spectrogram data than simpler power-based indicators.
2) Experimental validation and performance evaluation of the proposed method with respect to reference methods from the literature, showing that the proposed method outperforms the reference approaches with a leave-one-person-out (L1PO) test accuracy and macro *F1*-score of, respectively, 89.3% and 82.0%.

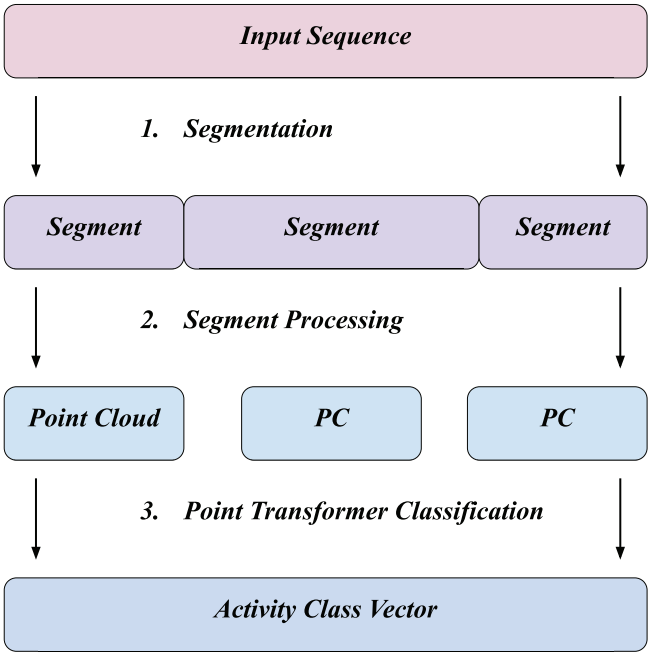

Fig. 1. Proposed three-stage classification pipeline. An input sequence of human activities in range-time format is segmented into single activities by a segmentation algorithm, the individual segments are processed to generate a PC format, and classification is, finally, achieved by means of a PT neural network.

3) Comparison of the proposed segmentation method with two alternative segmentation approaches. This study reveals that a segmentation approach based on machine learning can be situationally more effective, achieving an L1PO macro *F1*-score of 86.0%.

The remainder of this work is organized as follows. The proposed method and the two reference segmentation methods are described in Section II. The experimental case study, designed to evaluate the performance of the proposed method with respect to reference works in the literature, is outlined in Section III. Results for the case study are presented and discussed in Sections IV and V, with conclusions following in Section VI.

II. Proposed Method

Ethics approval for this research was granted by the TU Delft HREC (ID 1387). The proposed classification method, shown in Fig. 1, consists of three main elements: a sequence segmentation algorithm, a processing algorithm for the individual segments, and a PT neural network as a classifier. In this section, these components will be described, as well as relevant preprocessing steps. In addition, two alternative segmentation methods are outlined in Section II-C.

A. Signal Model and Notation

Real-valued backscattering signal vectors are considered for a set $\mathcal{N}$ of distributed pulsed radar systems. The cardinality of the set $\mathcal{N}$ is denoted by N. The quadrature components of the N vectors are obtained through the application of a Hilbert transform, and the resulting complex-valued vectors

are reshaped into N complex-valued fast- and slow-time matrices. The fast- and slow-time dimensions correspond to range and time, respectively, and these range-time representations are denoted by $\mathcal{R}_{r,t}$.

Velocity–time (spectrogram) representations are computed from the complex $\mathcal{R}_{r,t}$. To this end, a fast Fourier transform is first applied in the range dimension, yielding an intermediate matrix. Subsequently, a short-time Fourier transform (STFT) is applied to the frequency bin corresponding to the center frequency of the system. This operation yields the spectrogram $\mathcal{V}_{v,t}$, which is the input to the subsequent step of segmentation.

B. Segmentation

Segmentation of an activity sequence is based on an earlier feasibility study in [24]. The aim of segmentation is to divide a continuous sequence of various activities into segments that contain only a single activity, simplifying the subsequent task of classification. To this end, a time-dependent quantity is extracted from the spectrogram representation that is indicative of the kinematics of the activity being performed. The quantity that is selected for this purpose is the Rényi entropy H_α [21], which is defined as

$$H_\alpha(\mathcal{P}) = \frac{1}{1-\alpha}\log\left(\sum_i p_i^\alpha\right) \tag{1}$$

for $0 < \alpha < \infty$ and $\alpha \neq 1$. In the above formula, $\mathcal{P}$ is a discrete probability distribution with p_i being a vector of probabilities of the outcomes making up the distribution. The parameter α weighs the contribution of individual elements in $\mathcal{P}$ on the overall entropy. For this research, the probability distribution is instead replaced by a velocity distribution. Specifically, at time t_i, a single time bin $\mathcal{V}_{v,t_i}$ of an input spectrogram is normalized by dividing the time bin vector by the sum of its constituent elements and subsequently utilized for the entropy calculation as

$$H_\alpha(t_i) \equiv H_\alpha\left(\mathcal{V}_{v,t_i}\right) = \frac{1}{1-\alpha}\log\left(\sum_v\left(\frac{\mathcal{V}_{v,t_i}}{\sum_v \mathcal{V}_{v,t_i}}\right)^\alpha\right). \tag{2}$$

Changes in the distribution of velocities of the human target result in corresponding changes in entropy extracted from the spectrogram. Sudden entropy changes are associated with changes in activity, and fluctuations are monitored to indicate these activity transition events. To this end, an inequality is established, which serves as an entropy difference threshold

$$\left|H_\alpha(t) - H_\alpha\left(t - T_{\text{lag}}\right)\right| \geq \beta\sigma_H \tag{3}$$

where T_{lag} is a parameter governing the time scale of fluctuations that will trigger an activity transition event. σ_H represents the standard deviation in entropy over a longer interval, for example, the full duration of the activity sequence. The parameter β is a constant that determines the required fluctuation magnitude to indicate an activity transition.

The utilization of the Rényi entropy over alternatives such as signal power or spectrogram envelope, such as in [19], is motivated by the invariance of the entropy under a set of key transformations of the distribution of velocities. Specifically, the value of the Rényi entropy will remain unchanged under both a scaling of the input velocity distribution and a translation. A scaling of the velocity distribution occurs when an otherwise identical motion is performed faster or slower, or when the motion is performed in different orientations, resulting in an altered projection of the target velocity profile onto the radar line-of-sight. In both these cases, it is desirable to have an unchanged entropy, as the nature of the motion remains the same. A translation of the velocity distribution corresponds to an offset in the bulk velocity of the target, implying that the same motion is performed while moving. When a human target is walking in various directions, the entropy, thus, remains constant.

Whenever the threshold value $\beta\sigma_H$ is exceeded, the precise time is recorded, yielding a vector of transition event time stamps. With the acquisition of the vector of transition time stamps, the input sequence can be segmented into a set of range-time matrices of varying durations. It is assumed that there is a minimum duration to human activities of interest, and a minimum segment duration is, thus, implemented as a parameter. Segments that are shorter than this parameter $T_{\min}$ are split evenly across the adjoining two segments.

C. Alternative Segmentation Methods

Two alternative reference segmentation methods will be employed to gauge the effectiveness of the proposed segmentation approach. They are described in this section.

1) STA/LTA: STA/LTA is a triggering algorithm based on the ratio between two moving averages of different window sizes. The STAs/LTAs of a generic signal $s(t)$ are given by

$$\text{STA}(s(t)) = \sum_{t'=t-T_s}^{t}\frac{s(t')}{T_s} \tag{4}$$

$$\text{LTA}(s(t)) = \sum_{t'=t-T_l}^{t}\frac{s(t')}{T_l} \tag{5}$$

where T_s and T_l indicate the short and long window durations, respectively. The initiation and termination of a segment are given by the following two sets of conditions, respectively:

$$\frac{\text{STA}(H_\alpha(t))}{\text{LTA}(H_\alpha(t))} > \sigma_2 \;\;\&\;\; \text{STA}(H_\alpha(t)) > \sigma_1 \tag{6}$$

$$\frac{\text{STA}(H_\alpha(t))}{\text{LTA}(H_\alpha(t))} < \sigma_2 \;\;\&\;\; \text{LTA}(H_\alpha(t)) < \sigma_3. \tag{7}$$

The entropy $H_\alpha(t)$ is here taken as the specific signal of interest, and $[\sigma_1, \sigma_2, \sigma_3]$ are method parameters. σ_1 and σ_3 are thresholds that govern the required entropy increase and decrease for detecting the onset of a segment. The required ratio between the short- and long-term averages to indicate the start of a segment is given by the remaining parameter σ_2. Together with the short and long window durations T_s and T_l, a total of five parameters are, thus, required to configure the STA/LTA algorithm.

For the purpose of this case study, the optimal parameters for the STA/LTA algorithm are determined using the built-in genetic algorithm (GA) optimizer in MATLAB. The objective function is a sum of two terms, both in the range $[0, 1]$. The first term expresses the suitability of a segment between two

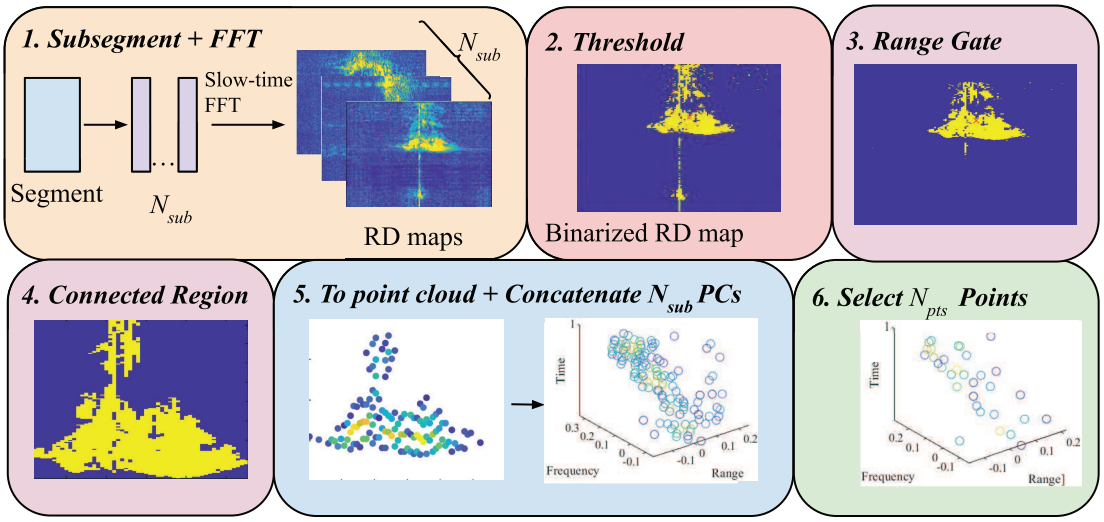

Fig. 2. Diagram of the proposed segment processing pipeline for the generation of PC samples suitable as inputs for PT networks.

detected transitions in terms of the most occurring activity label. The closer to 1, the better the segment captures a single activity. The second term penalizes the difference between the number of detected transitions and the number of transitions in the ground-truth (GT) target vector.

2) BiLSTM for Segmentation: A machine learning-based approach for segmentation is also investigated. Specifically, a BiLSTM network is trained to detect transition events from an entropy signal input $H_\alpha(t)$. Instead of using only a single value of the entropy as a feature for each time t, a feature vector of entropies for differing values of *alpha* is used as input to the BiLSTM. For this work, the feature vector $[H_{0.1}(t), H_{0.2}(t), \ldots, H_{0.9}(t)]$ is used, as it is empirically determined to balance computational complexity and salient information content.

The target sequence used to train this model is a binary vector of the same length as the entropy input. At every transition between two activities, the value of the target sequence is "1." It is "0" everywhere else. The target vector is relatively sparse, as the number of transitions in a sequence is generally orders of magnitude smaller than the duration of the sequence in time steps. This significant sparsity hinders the ability of the model to correctly train and predict transition events. To alleviate this problem, it is here proposed that the target vector is convolved with a rectangular function with a width of approximately 0.33 s. A physical interpretation of this convolution is the noninstantaneous nature of activity transitions and an inevitable degree of subjectivity in defining exactly when they happen. The value of 0.33 s is empirically found to provide good performance without making transition events unrealistically long. The sparsity of the target vector is a known problem in transition detection and is tackled in a similar manner in [25].

D. Segment Processing and Classification

Every range-time matrix containing a segment of activities is processed individually, yielding a PC representation. The PC representations are computationally more efficient to manipulate than using complete data matrices and allow for classification by the powerful PT family of neural networks [23]. This section describes the PC processing and classification.

1) Every segment is first evenly divided into N_{sub} subsegments, as shown in step (1) of Fig. 2. Each subsegment is, thus, a complex range time matrix $\mathcal{R}_{r,t}$, and an FFT along the time axis yields a set of N_{sub} range Doppler maps.
2) A threshold function is applied, yielding a binary map of where the signal power exceeds a fraction γ of the maximum signal power.
3) For the purpose of reduction of noise and clutter contributions, a range gate of 2 m is additionally employed, centered on the center of mass of each binary map, with the assumption that the largest detected region corresponds to the target of interest.
4) The three largest connected regions are then selected. Connected region in this case refers to a set of matrix elements that have nonzero neighbors. This step is performed since human anatomy and kinematics dictate a smoothly varying range-Doppler profile. Selecting the large connected regions, thus, suppresses speckle-type noise.
5) The N_{sub} binary maps are subsequently used to select the points of interest from the original N_{sub} range-Doppler maps. These points form a PC with a dimension range, Doppler, time, and signal power for each point.
6) For consistency between segments, a fixed number of points is required. Thus, the PC is upsampled or downsampled to N_{pts} points based on this requirement. Upsampling is achieved by duplication of existing points in the cloud, and downsampling is achieved by uniform subsampling of the cloud in range-Doppler space.

For classification of each segment, a PT neural network is utilized [23]. PT networks fall under the transformer family of deep learning architectures and are suited to tasks including object classification and scene segmentation. PCs serve as

inputs to these networks, and a modified attention mechanism [26] is the means of feature extraction. Details on the network can be found in the original paper [23]. For this work, based on previous research in [27] and inspired by the research in [28], the architecture proposed by [23] is selected. This choice is motivated by the demonstrated classification performance, respectively, to two alternative architectures that have been considered [29], [30]. Three parameters govern the implementation of the architecture into a specific neural network: the number of transformer blocks, the number of neighbors considered for each point, and the size of each PT layer. Based on the previous research in [27], they are set to 4, 16, and 128, respectively.

The PC corresponding to each individual segment is assigned an activity label based on the primary activity performed during the segment, which is determined through a majority ruling across the total number of time steps in the segment. An activity prediction is the output of the PT model, which is compared to the GT for that particular segment in order to determine the classification performance of the method.

III. Case Studies

To gauge the suitability of the proposed method for continuous human activity classification, as well as study the effects of alternative segmentation approaches, two case studies are performed. The studies are conducted using the publicly available dataset [31] that contains sequences of activities from 14 participants. Specifically, these case studies include the following.

1) A sample holdout validation of the classification performance of the proposed segmentation method and two alternative segmentation methods described in Section II-C.
2) An L1PO validation of the proposed method and the two alternative methods to compare to reference classification approaches in the literature. L1PO validation is further detailed in Section III-B and entails training on data from all but one participant and testing on the data of the omitted participant.

Comparing the L1PO and sample holdout results additionally grants insight into potential overfitting issues associated with any of the evaluated methods.

A summary of the parameters employed in the proposed method, based on Rényi entropy, is given in Table I. The selected values for these parameters are based on the segmentation feasibility study in [24] and the subsequent segment processing study in [27].

TABLE I

Summary of Parameters Used in the Proposed Method. Parameters Relating to Segmentation Are Found at the Top and Are Based on the Feasibility Study in [24], and Those Relating to the Subsequent Processing of the Segments Are Based on [27] and Are Found at the Bottom

Parameter	Value	Notes
Segmentation		
α	0.77	Influences dependence of entropy on velocity values with strong intensity
T_{lag}	3.1 s	Time scale of monitored entropy fluctuations
β	0.3	Required fluctuation magnitude to indicate a transition event
Processing		
N_{sub}	6	Number of subsegments extracted from segment
γ	0.8	Threshold value for detection (binarization) of range Doppler maps, as a fraction of maximum signal power
N_{pts}	1024	Number of points in the processed point cloud

TABLE II

Table Summarizing the Key Radar Parameters of Each PulsON P410 Pulsed UWB Sensor

Radar Parameter	Value	Unit
Center Frequency	4.3	GHz
Bandwidth	2.2	GHz
PRF	122	Hz
Channels	SISO	-

TABLE III

Nine Activity Classes Utilized in the Construction of the Activity Sequences

Activity		
Walking	Standing Up (Sitting)	Falling (Walking)
Stationary	Bending (Sitting)	Falling (Stationary)
Sitting Down	Bending (Standing)	Standing Up (Ground)

A. Dataset Description

1) Experimental Setup: The radar sensors used to capture the dataset are an ensemble of five Humatics PulsON P410 pulsed ultrawideband (UWB) single-input–single-output (SISO) sensors [32], operating as a network of distributed monostatic nodes. Each radar features a center frequency of 4.3 GHz and a bandwidth of 2.2 GHz, yielding a range resolution of approximately 6.8 cm. The PRF is set to 122 Hz, resulting in a maximum unambiguous velocity of $\pm$ 2.13 m/s. Each radar sensor is equipped with an antenna with a pattern that is symmetric in azimuth. Combined with the single-channel (SISO) nature of the sensors, this means that angle-of-arrival estimation cannot be performed. The key radar parameters are summarized in Table II.

The five sensors are arranged in a semicircle, as shown in Fig. 3. They are spaced at regular 45° intervals with a 6.38-m diameter of the semicircle. The activity area is a circle with a diameter of 4.38 m, concentric to the semicircle of the sensor network. The sensors are placed approximately 1 m above the floor.

2) Activity Sequences: For each of the 14 participants, 30 sequences of 2 min each are captured. The activities in the sequences are performed in random locations, facing random directions, and at random intervals. Nine different activities are used to construct the sequences and are listed in Table III. Four sample activities are displayed in Fig. 4, paired with their corresponding range-Doppler representations. It should be noted that the *Stationary* class represents mul-

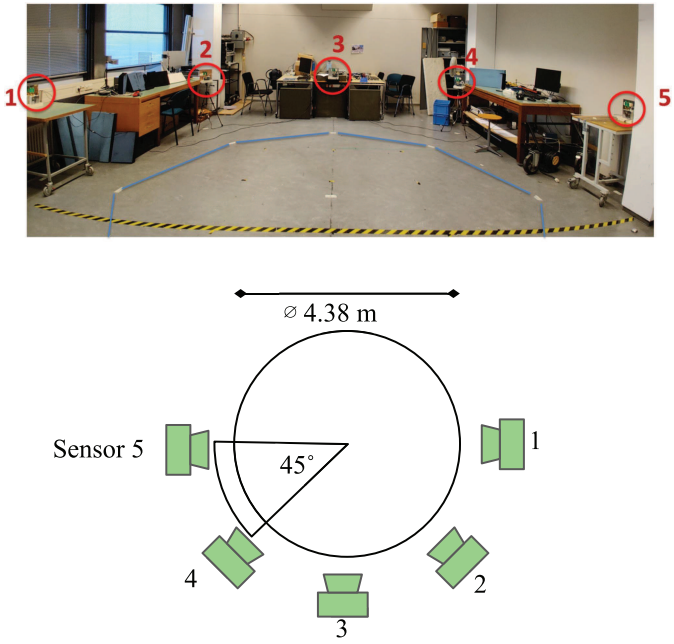

Fig. 3. Photograph (top) and diagram (bottom) of the experimental measurement area. Five sensors are arranged in a semicircle, spaced at regular 45 $^{\text{circ}}$ intervals with a 6.38-m diameter of the semicircle.

tiple static poses, such as standing still, sitting still, and lying still after a fall. Comprehensive descriptions of all 30 sequences are provided in the accompanying public dataset documentation [31].

B. Classification Approach

The approach for classification in the case studies conducted in this work is given as follows: the activity sequences are first segmented using the method under investigation. The found segments are processed in accordance with the approach described in Section II-D and subsequently used as input to the PT network for classification. A weighted cross-entropy loss function is used in the training of the PT, with the class weights inversely proportional to the sample frequency of the respective class. As an example, segments containing *Walking* are the most numerous in the dataset and consequently are assigned the lowest weight in the loss function. Comparing the output of the PT to the GT for the segment yields performance measures, in this research expressed in terms of test accuracy and macro *F1*-score. Test accuracy is selected due to its prevalence in classification tasks in the literature, and macro *F1*-score is reported as it can give a good insight into the performance on underrepresented classes.

The validation approaches taken for the case studies fall under two categories in terms of training/testing split.

1) An L1PO approach, where sequences from 13 of the participants are used to train the PT, and sequences from the remaining single participant are used to test the model. This process is repeated for each participant, and average performance figures are reported. This validation strategy yields the most comprehensive result, as the capability of the model to respond to unseen participants is evaluated.
2) A sample holdout method, where 80% of the segments are used to train the PT model and 20% are used for testing.

Comparing the results of the two evaluation methods gives insight into any overfitting issues that may present themselves for any of the methods under test.

TABLE IV
TEST ACCURACY AND MACRO F1-SCORE RESULTS FOR THE PROPOSED CLASSIFICATION METHOD AND TWO REFERENCE SEGMENTATION METHODS, DETAILED IN SECTION II-C. RESULTS ARE PRESENTED FOR A 80%/20% SAMPLE HOLDOUT SCHEME

Method	Test Accuracy	Macro F1-Score
Proposed	0.909±0.003	0.853±0.005
STA/LTA	0.903±0.006	0.829±0.006
BiLSTM	0.924±0.010	0.914±0.011

IV. RESULTS

This section reports the results of the case studies with related analysis and discussion.

A. Study on Segmentation Methods

Table IV contains the classification results of a comparison between the proposed method and two reference methods. As References, the STA/LTA and BiLSTM segmentation methods from Sections II-C1 and II-C2 are specifically shown. Each model is trained three times using a randomized 80%/20% sample holdout validation to analyze statistical fluctuations. The highest performance is achieved using the BiLSTM segmentation method, as this approach is able to learn more complex patterns in entropy rather than just strong fluctuations. The proposed segmentation method outperforms STA/LTA segmentation in terms of macro *F1*-score. Inspection of the segments created with the STA/LTA algorithm reveals that transitions are generally correctly found when transitioning from stationary to motion and vice versa, but that complex transitions between different types of motion are not detected as effectively.

B. L1PO Result

Fig. 5 shows the results of the L1PO validation of the proposed method, indicated with solid markers and lines. Additionally included is the performance for "perfect" segmentation where GT labels have been employed to yield segments that contain a single activity only, which is assumed to be "optimal" segmentation. These performance figures are shown with empty markers and dashed lines. It should be noted that this "optimal" result relies on information that is unavailable in a real scenario and is provided to indicate the effect of segmentation on the final classification effectiveness.

On average, the performance metrics for the segmentation based on GT information are higher than for segmentation following the proposed method based on the Rényi entropy. Average test accuracy for the former and latter are

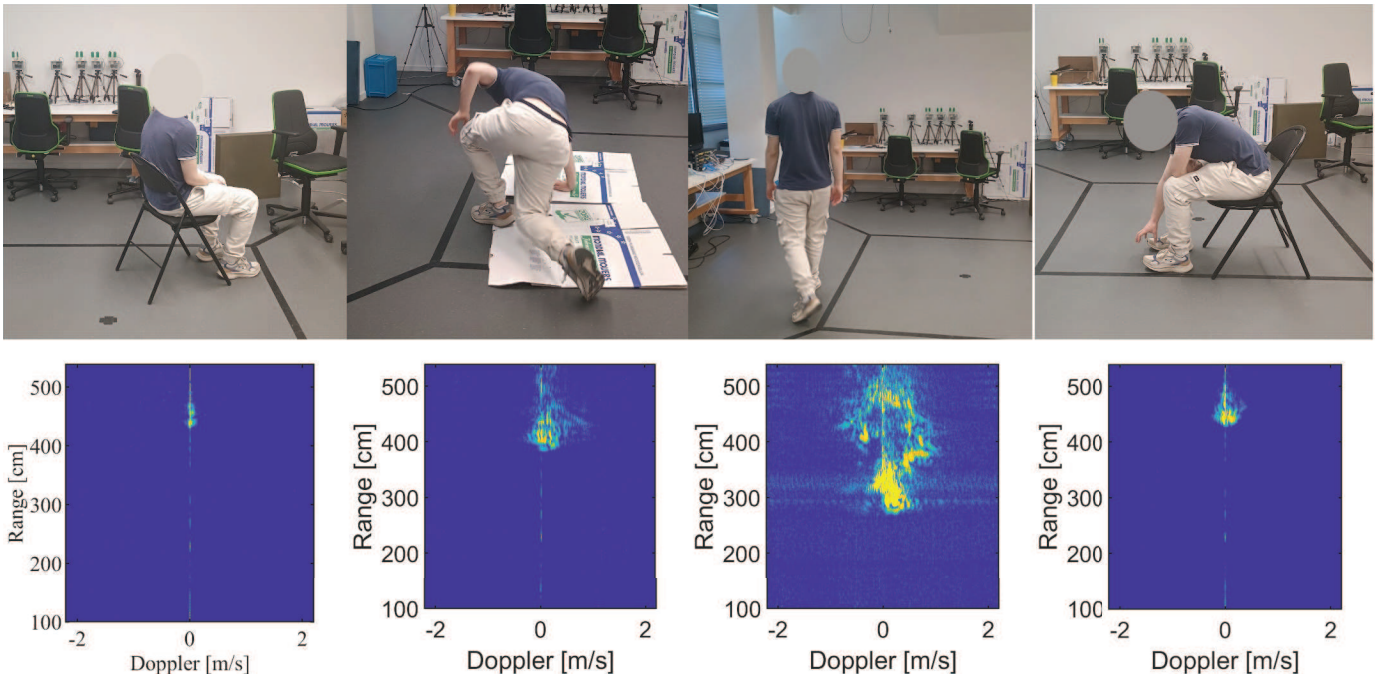

Fig. 4. Sample photographs of activities in the dataset, accompanied by corresponding range-Doppler representations. From left to right: stationary, standing up (from the ground), walking, and bending (from sitting).

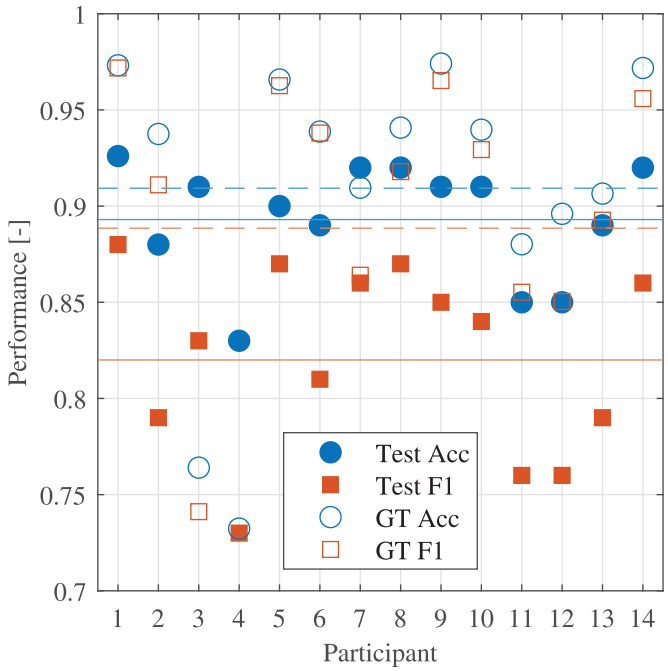

Fig. 5. Results for the proposed classification method, evaluated using an L1PO testing scheme. Test accuracy and macro *F1*-score results are shown for each of the 14 participants. Averages across all participants are indicated with horizontal lines. *GT acc* and *GT F1* refer to classification performance achievable when GT information is available about the moments where transitions occur, i.e., segments containing only a single activity. The averages for this "optimal" GT segmentation are indicated with dashed lines.

90.9% and 89.3%, respectively. Notable outliers, however, are participants #3 and #4. An important conclusion that can be drawn is that segmentation into human-interpretable, "perfect," boundaries is not necessarily the best strategy for classification. Single activity GT segments may align well with human understanding but do not necessarily facilitate automatic model-based classification. Breaking down complex or long segments into kinematically more homogeneous parts can enhance classification performance, as seen in the cases of participants #3 and #4.

TABLE V

TEST ACCURACY AND MACRO F1 -SCORE RESULTS FOR THE PROPOSED CLASSIFICATION METHOD AND FIVE ALTERNATIVE METHODS. GT SEGMENTS REFERS TO THE UTILIZATION OF THE PT NETWORK BUT WITH "PERFECT" SEGMENTATION USING THE GT DATA TO LOCATE TRANSITIONS. THE TOP FOUR ROWS CORRESPOND TO DIFFERENT METHODS DISCUSSED IN THIS WORK, AND THE LOWER TWO ROWS CORRESPOND TO REFERENCE WORKS FROM THE LITERATURE. ALL RESULTS ARE BASED ON THE SAME DATASET [31] AND THE SAME L1PO VALIDATION SCHEME. THE RESULTS IN THIS TABLE ARE FOR THE FULL NINE-CLASS CLASSIFICATION PROBLEM. *: [27]; §: [12]

Method	Test Accuracy	Macro F1-Score
Proposed	0.893	0.820
Proposed (GT Segments)	0.909	0.889
BiLSTM Segmentation	0.878	0.860
STA/LTA Segmentation	0.879	0.790
PT (Fixed Segments)*	0.869	0.787
CNN-BiGRU§	0.851	-

C. Comparison to Reference Methods in the Literature

Table V presents the L1PO test accuracy and macro *F1*-score for the proposed method, the BiLSTM segmentation

TABLE VI

GENERALIZATION GAPS FOR THE THREE METHODS EVALUATED IN THIS WORK. THE GENERALIZATION GAP CONSTITUTES THE DIFFERENCE IN PERFORMANCE BETWEEN SAMPLE HOLDOUT AND L1PO VALIDATION AND IS INDICATIVE OF OVERFITTING ISSUES

Method	Test Accuracy	Macro F1-Score
Proposed	1.6%	3.1%
BiLSTM Segmentation	4.6%	5.7%
STA/LTA Segmentation	2.4%	4.3%

TABLE VII

MERGING SCHEME FOR THE CONSOLIDATION OF THE FULL NINE ACTIVITY CLASSES INTO A SET OF FIVE ACTIVITY CLASSES. CLASSES ARE GROUPED BASED ON SIMILARITY

Constituent Classes	Merged Class
Walking	Walking
Stationary	Stationary
Sitting Down, Standing up (from sitting), Bending (from sitting), Bending (from standing)	In Situ
Falling (from walking), Falling (from stationary)	Falling
Standing up (from ground)	Standing up

TABLE VIII

TEST ACCURACY AND MACRO F1-SCORE RESULTS FOR THE PROPOSED CLASSIFICATION METHOD AND SEVERAL ALTERNATIVE METHODS. GT SEGMENTS REFERS TO THE UTILIZATION OF THE PT NETWORK BUT WITH "PERFECT" SEGMENTATION USING THE GT DATA TO LOCATE TRANSITIONS. THE TOP FOUR ROWS CORRESPOND TO METHODS DISCUSSED IN THIS WORK, AND THE LOWER SIX ROWS CORRESPOND TO REFERENCE WORKS FROM THE LITERATURE. ALL RESULTS ARE BASED ON THE SAME DATASET [31] AND THE SAME L1PO VALIDATION SCHEME. THE RESULTS IN THIS TABLE ARE FOR A FIVE-CLASS CLASSIFICATION PROBLEM. *: [27]; Y: SIGNAL FUSION [11]; AND Z: FEATURE FUSION [11]

Method	Test Accuracy	Macro F1-Score
Proposed	0.928	0.880
Proposed (GT Segments)	0.947	0.943
BiLSTM Segmentation	0.913	0.900
STA/LTA Segmentation	0.925	0.870
PT (Fixed Segments)*	-	0.862
GRU†	0.909	0.778
LSTM†	0.910	0.769
bi-GRU†	0.933	0.844
BiLSTM†	0.931	0.836
BiLSTM‡	0.924	0.840

method, the STA/LTA segmentation method, and, finally, two reference methods from the literature using the same dataset [12], [27]. Additionally included is the result for classification of segments created using the GT labeling information (*GT segments*). Comparing the proposed method with the PT operating on fixed, 2-s windows [27] reveal an increase in test accuracy, as well as a $+3.3$ % increase in macro *F1*-score. This improvement highlights the benefit of using adaptive segments, mitigating the transition errors that occur when using fixed-duration segments. The CNN-BiGRU approach in [12] processes the activity sequences into a series of range-Doppler maps, where feature extraction is performed by a CNN on a per-map basis. The extracted features form a time series, which is then used as input to a bidirectional gated recurrent unit (BiGRU), which performs classification. To keep the size of this hybrid model computationally feasible, the range-Doppler maps are scaled down, possibly explaining the difference in performance with the proposed method.

Of note is the test accuracy of the BiLSTM segmentation method, which is lower than that of the proposed method and the STA/LTA-based method. This contrasts with the results achieved under the sample holdout validation scheme in Table IV, where the BiLSTM-based method outperforms the proposed method both in terms of test accuracy and macro *F1*-score. The gap in performance between the two validation schemes is shown in Table VI. This result highlights the problem of overfitting in machine learning approaches, where a neural network trained on a set of data can have issues with generalization capabilities outside of the training set. In this case, the test data of the unseen participant proves challenging to the BiLSTM segmentation algorithm. It should be noted that the implementation of the PT network is identical for all the investigated methods, and the generalization problem is, thus, associated with the BiLSTM network.

For the sake of completeness, the total amount of detected critical events is also reported, here defined as the fall activities in the dataset. For the BiLSTM-based method, 79.6% of a total of 988 fall events are correctly identified. The proposed method performs better, identifying 86.9% of the fall events. It is assumed that the likely cause for this improvement is the tendency of the proposed method to yield shorter segments, which diminishes the likelihood of completely absorbing a fall event into an adjacent activity.

Some of the reference works with methods benchmarked on the same dataset utilize an aggregated set of five activity classes. In this reduced set, displayed in Table VII, activities such as *falling from walking* and *falling while standing still* are, for example, joined into a singular *falling* class.

Table VIII shows the results for this five-class problem. The top four rows again correspond to methods discussed and proposed in this work; the lower rows show performance metrics reported in reference methods from the literature [11], [27]. The methods in [11] involve the computation of a spectrogram representation of the activity sequences, which is subsequently used as input to various RNN architectures. The performance of the proposed method, as well as the segmentation approaches utilizing a BiLSTM or STA/LTA, surpasses that of the reference method [11]. Notably, the combination of a BiLSTM with the PT model demonstrates superior results compared to the BiLSTM operating solely on spectrogram data. This improvement can be attributed to the more efficient utilization of the BiLSTM network in the combined approach. Specifically, the BiLSTM in the hybrid model processes only the Rényi entropy, rather than the entire spectrogram, and is tasked with predicting transitions rather than classifying specific activity types. Consequently, the combined model

TABLE IX

RMSE IN SECONDS OF THE ONSET AND OFFSET TIMES OF DETECTED FALL EVENT SEGMENTS. FOUR FALL TYPES ARE CONSIDERED

RMSE offset [s]	Walking	Standing
Slow	0.92±0.98	0.42±0.42
Fast	0.78±0.78	0.62±0.41

TABLE X

PMOL OF DETECTED FALL EVENT SEGMENTS. THE PMOL REPRESENTS THE FRACTION OF A SEGMENT DURATION THAT IS LABELED AS A FALL IN THE GT. FOUR FALL TYPES ARE CONSIDERED

PMOL [%]	Walking	Standing
Slow	71.4%±15.9%	88.7%±10.9%
Fast	84.7%±20.7%	78.0%±19.0%

remains more compact at 19 MB compared to the stand-alone BiLSTM of approximately 25 MB.

The L1PO experiments demonstrate promising results for the proposed method, with the BiLSTM-based segmentation delivering the highest performance in terms of macro *F1*-score. However, the proposed method based on Rényi entropy offers several key advantages in certain scenarios.

1) If minimizing model size is a significant priority, the proposed method is advantageous, as it relies solely on the PT model with minimal segmentation processing.
2) The segmentation approach used in the proposed method offers superior interpretability compared to the more complex BiLSTM-based segmentation.
3) The PT model takes inputs such as range, velocity, and time, which are consistent across experimental scenes and radar systems. In contrast, BiLSTM segmentation may need retraining across different scenarios, as the Rényi entropy patterns could vary. This argument is further strengthened by the disparity in performance of the BiLSTM-based method between the sample holdout experiment and the L1PO experiment, representing poor generalization capabilities.

D. Study on Transition Speed

To gauge the effectiveness of the proposed method for detecting transitions at different speeds, most notably for falls, a secondary dataset with 36 fall events is recorded. Experimental setup and method parameters are identical to those used with the main experimental dataset. Fall types are divided by initial conditions (from walking motion and from standing position) and speed (slow and fast). Each fall type is repeated four times by three participants.

For this study, only the segmentation component of the proposed method is evaluated. Two performance metrics are employed: the RMSE of the onset and offset times of each fall event and the quality of the detected fall segment in terms of the fraction that corresponds to the *Fall* label according to the recorded GT. The latter will be referred to as the percentage of the most occurring label (PMOL). When no segment exists with PMOL > 50 %, the fall is considered misdetected.

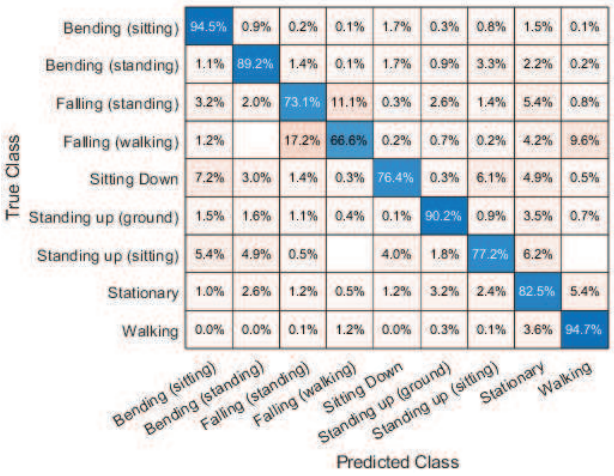

Fig. 6. Confusion matrix for the full L1PO experiment. Percentages are row-normalized.

Tables IX and X contain the performance metrics per fall type. Two fall events are misdetected: a fast fall from walking motion and a slow fall from a stationary position. From the tables, two statistically significant conclusions are drawn.

1) The slow fall from a walking motion features a higher error in activity onset and offset time compared to the slow fall from a stationary position.
2) The slow fall from a walking motion yields poorer segment quality in terms of PMOL compared to the slow fall from a stationary position and the fast fall from a walking motion.

Both conclusions indicate that the slow walking fall is most challenging to segment due to the slower fluctuation in entropy and the absence of a period of stationarity before and after the fall.

V. DISCUSSION

In this section, common classifier errors will be discussed, as well as the sensitivity of the method to different environments, together with computational considerations pertaining to the proposed method.

In Fig. 6, the confusion matrix for the full L1PO experiment is displayed. The three most prevalent error types are confusions between the two fall types in the experimental dataset and confusions between the *walking* and *falling from walking* activity classes. The latter error likely stems from the partially subjective boundary between these two activity classes. The confusion between fall types can, in part, be explained by the lack of temporal context available to the classifier, as no information from previous segments is available to the PT network. Less frequent errors are between activities with distinct kinematic profiles in the vertical direction, such as sitting down on a chair versus bending down from the chair. These vertical motion components cannot be captured with the sensor network geometry that is used for the recording of the dataset.

The effect of different room layouts on the performance of the proposed method can be anticipated in part by the effect

of static clutter on the segmentation algorithm and the PT network. The effect of static clutter on the entropy is expected to be an approximately constant offset, with fluctuations due to human motions preserved. The parameter α in (1) governs the influence of single, high probability events on the entropy. This parameter can, thus, potentially be used to compensate for more or less cluttered environments. The PT network in the proposed method is operating on physical variables such as target range, velocity, and time. It is, therefore, expected that the effect of a different environment will be limited, and no complete retraining of the model will be required.

With regards to computational requirements, the processing of 120 s of data into a spectrogram representation combined with computing the entropy takes 1.4 s on a 3.40-GHz CPU. On the same CPU, computation of PC samples takes [separate-uncertainty = true]25.9(9.5)ms per sample. For this work, PT models are trained on an Nvidia RTX 3090 board. Training time per epoch is approximately 18 s per epoch, with an inference time of approximately 0.63 ms per PC sample. A typical model size is 17 MB.

VI. Conclusion

This study proposes a novel method for the classification of continuous sequences of human activities. The proposed method consists of three main elements: segmentation of the sequences, processing of the segments, and classification of the segments. As a key step of the proposed processing, segmentation is achieved through the monitoring of fluctuations in Rényi entropy, a scalar quantity computed from micro-Doppler spectrograms. The proposed method offers a solution to the problem of continuous activity classification that is more reliable in terms of classification performance compared to reference methods from the literature. In addition, it is also computationally efficient due to the effectiveness of the Rényi entropy as an indicator of activity changes.

The proposed method is experimentally validated on a publicly available dataset. An L1PO test accuracy of 89.3% and a macro *F1*-score of 82.0% are achieved on the dataset, which consists of a variety of sequences of nine human activities. Alternative segmentation methods are also investigated. These include the STA/LTA change detection algorithm and a BiLSTM network taking Rényi entropy as input. Between the proposed method and the BiLSTM segmentation method, the highest classification performance metrics are attained by the BiLSTM. The proposed method is, however, preferable in terms of computational efficiency and interpretability and outperforms reference methods from the literature on the same dataset.

An idealized form of segmentation is also studied, where GT labels are used to create segments that contain only a single activity. Classification performance in this case is generally higher than on segments produced by the various segmentation algorithms. In some cases, however, the opposite is true. It is, therefore, concluded that segmenting to the GT of an activity sequence is not necessarily optimal, if at all possible in a realistic scenario, and segmentation into classifier-interpretable, homogeneous segments might be preferable.

To further enhance the capabilities of the proposed method, the addition of information pertaining to temporal context may enhance the classification performance for classes such as *falling from walking*, where a preceding activity is a key indicator of the nature of the activity. A multitarget adaptation of the method could also be realized by including a detection stage prior to the segmentation step. Moreover, the inclusion of a multilabel classification method could be considered, such as that presented in [33], to further mitigate the problem of multiple activities in a segment. In addition, alternative segmentation methods can be investigated that strike a balance between interpretability, computational efficiency, and classification performance.

Acknowledgment

The authors are grateful to all the volunteers who participated in the data collection. They thank the anonymous reviewers for their constructive comments and suggestions in the review process.

References

[1] X. Dang, Z. Chen, and Z. Hao, "Emotion recognition method using millimetre wave radar based on deep learning," *IET Radar, Sonar Navigat.*, vol. 16, no. 11, pp. 1796–1808, Nov. 2022.

[2] K.-C. Peng, M.-C. Sung, F.-K. Wang, and T.-S. Horng, "Noncontact vital sign sensing under nonperiodic body movement using a novel frequency-locked-loop radar," *IEEE Trans. Microw. Theory Techn.*, vol. 69, no. 11, pp. 4762–4773, Nov. 2021.

[3] F. Fioranelli, J. Le Kernec, and S. A. Shah, "Radar for health care: Recognizing human activities and monitoring vital signs," *IEEE Potentials*, vol. 38, no. 4, pp. 16–23, Jul. 2019.

[4] Y. Yao et al., "Fall detection system using millimeter-wave radar based on neural network and information fusion," *IEEE Internet Things J.*, vol. 9, no. 21, pp. 21038–21050, Nov. 2022.

[5] S. A. Shah, A. Tahir, J. Le Kernec, A. Zoha, and F. Fioranelli, "Data portability for activities of daily living and fall detection in different environments using radar micro-Doppler," *Neural Comput. Appl.*, vol. 34, no. 10, pp. 7933–7953, May 2022.

[6] H. Nematallah and S. Rajan, "Adaptive hierarchical classification for human activity recognition using inertial measurement unit (IMU) time-series data," *IEEE Access*, vol. 12, pp. 52127–52149, 2024.

[7] J. J. Steckenrider, B. Crawford, and P. Zheng, "GPS and IMU fusion for human gait estimation," in *Proc. IEEE 24th Int. Conf. Inf. Fusion (FUSION)*, Nov. 2021, pp. 1–7.

[8] R. Yu, Y. Du, J. Li, A. Napolitano, and J. Le Kernec, "Radar-based human activity recognition using denoising techniques to enhance classification accuracy," *IET Radar, Sonar Navigat.*, vol. 18, no. 2, pp. 277–293, Feb. 2024.

[9] I. Ullmann, R. G. Guendel, N. C. Kruse, F. Fioranelli, and A. Yarovoy, "Radar-based continuous human activity recognition with multi-label classification," in *Proc. IEEE Sensors*, Oct. 2023, pp. 1–4.

[10] L. Werthen-Brabants, G. Bhavanasi, I. Couckuyt, T. Dhaene, and D. Deschrijver, "Split BiRNN for real-time activity recognition using radar and deep learning," *Sci. Rep.*, vol. 12, no. 1, pp. 1–12, May 2022.

[11] R. G. Guendel, F. Fioranelli, and A. Yarovoy, "Distributed radar fusion and recurrent networks for classification of continuous human activities," *IET Radar, Sonar Navigat.*, vol. 16, no. 7, pp. 1144–1161, Jul. 2022.

[12] S. Zhu, R. G. Guendel, A. Yarovoy, and F. Fioranelli, "Continuous human activity recognition with distributed radar sensor networks and CNN–RNN architectures," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, 2022, Art. no. 5115215.

[13] N. Kruse, R. Guendel, F. Fioranelli, and A. Yarovoy, "Distributed radar fusion for extended target location and velocity reconstruction," in *Proc. IEEE Radar Conf. (RadarConf)*, May 2024, pp. 1–6.

[14] Q. Jian, S. Guo, P. Chen, P. Wu, and G. Cui, "A robust real-time human activity recognition method based on attention-augmented GRU," in *Proc. IEEE Radar Conf. (RadarConf)*, May 2021, pp. 1–5.

[15] Y.-S. Chen, K.-H. Cheng, and Y.-A. Xu, "Transformer-sequential-based learning for continuous HMR with high similarity using mmWave FMCW radar," in *Proc. IEEE Wireless Commun. Netw. Conf. (WCNC)*, Glasgow, U.K., Mar. 2023, pp. 1–6.
[16] M. J. Bocus, K. Chetty, and R. J. Piechocki, "UWB and WiFi systems as passive opportunistic activity sensing radars," in *Proc. IEEE Radar Conf. (RadarConf)*, May 2021, pp. 1–6.
[17] A. Gorji, H.-U.-R. Khalid, A. Bourdoux, and H. Sahli, "On the generalization and reliability of single radar-based human activity recognition," *IEEE Access*, vol. 9, pp. 85334–85349, 2021.
[18] S. Z. Gurbuz, E. Kurtoglu, M. M. Rahman, and D. Martelli, "Gait variability analysis using continuous RF data streams of human activity," *Smart Health*, vol. 26, Dec. 2022, Art. no. 100334.
[19] E. Kurtoglu, A. C. Gurbuz, E. A. Malaia, D. Griffin, C. Crawford, and S. Z. Gurbuz, "ASL trigger recognition in mixed activity/signing sequences for RF sensor-based user interfaces," *IEEE Trans. Human-Mach. Syst.*, vol. 52, no. 4, pp. 699–712, Aug. 2022.
[20] Z. Yang, H. Wang, P. Ni, P. Wang, Q. Cao, and L. Fang, "Real-time human activity classification from radar with CNN-LSTM network," in *Proc. IEEE 16th Conf. Ind. Electron. Appl. (ICIEA)*, Aug. 2021, pp. 50–55.
[21] A. Rényi, "On measures of information and entropy," in *Proc. 4th Berkeley Symp. Math., Statist. Probab.*, vol. 1, 1960, pp. 547–561.
[22] M. Liuni, A. Röbel, M. Romito, and X. Rodet, "Rényi information measures for spectral change detection," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2011, pp. 3824–3827.
[23] H. Zhao, L. Jiang, J. Jia, P. Torr, and V. Koltun, "Point transformer," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 16239–16248.
[24] N. C. Kruse, R. G. Guendel, F. Fioranelli, and A. G. Yarovoy, "Segmentation of micro-Doppler signatures of human sequential activities using Rényi entropy," in *Proc. Int. Conf. Radar Syst. (RADAR)*, 2022, pp. 435–440.
[25] R. Yin, H. Bredin, and C. Barras, "Speaker change detection in broadcast TV using bidirectional long short-term memory networks," in *Proc. Interspeech*, Aug. 2017, pp. 3827–3831.
[26] A. Vaswani et al., "Attention is all you need," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 6000–6010.
[27] N. C. Kruse, F. Fioranelli, and A. Yarovoy, "Radar point cloud processing methods for human activity classification with point transformer networks," *IEEE Trans. Radar Syst.*, vol. 2, pp. 1–12, 2024.
[28] Z. Guo, R. G. Guendel, A. Yarovoy, and F. Fioranelli, "Point transformer-based human activity recognition using high-dimensional radar point clouds," in *Proc. IEEE Radar Conf. (RadarConf)*, May 2023, pp. 1–6.
[29] N. Engel, V. Belagiannis, and K. Dietmayer, "Point transformer," *IEEE Access*, vol. 9, pp. 134826–134840, 2021.
[30] M.-H. Guo, J.-X. Cai, Z.-N. Liu, T.-J. Mu, R. R. Martin, and S.-M. Hu, "PCT: Point cloud transformer," *Comput. Vis. Media*, vol. 7, no. 2, pp. 187–199, Jun. 2021.
[31] R. G. Guendel, M. Unterhorst, N. Kruse, A. Yarovoy, and F. Fioranelli, "Dataset of continuous human activities performed in arbitrary directions collected with a distributed radar network of five nodes," Delft Univ. Technol., Delft, The Netherlands.
[32] Time Domain, *PulsON410*, 320-0284C, 2013.
[33] I. Ullmann, R. G. Guendel, N. Christian Kruse, F. Fioranelli, and A. Yarovoy, "Classification strategies for radar-based continuous human activity recognition with multiple inputs and multilabel output," *IEEE Sensors J.*, vol. 24, no. 24, pp. 40251–40261, Dec. 2024.

Nicolas C. Kruse (Member, IEEE) received the B.Sc. degree in applied physics from Delft University of Technology, Delft, The Netherlands, in 2017, the M.Sc. degree in physics from the University of Groningen, Groningen, The Netherlands, in 2020, and the Ph.D. degree from Delft University of Technology in 2025.

He joined the Microwave Sensing, Systems, and Signals Group, Faculty of Electrical Engineering, Delft University of Technology, in 2021. His research interests pertain to classification algorithms for continuous human activity sequences through micro-Doppler signatures.

Alec Daalman received the B.Sc. degree in electrical engineering from Hague University of Applied Sciences, The Hague, The Netherlands, in 2024. Currently, he is pursuing the M.Sc. degree in electrical engineering with Delft University of Technology, Delft, The Netherlands.

He joined the Microwave Sensing, Systems, and Signals Group, Faculty of Electrical Engineering, Delft University of Technology, for his thesis on Adaptive Segmentation for Radar Point Cloud-based Continuous Activity Classification.

Francesco Fioranelli (Senior Member, IEEE) received the Ph.D. degree from Durham University, Durham, U.K., in 2014.

He was an Assistant Professor with the University of Glasgow, Glasgow, U.K., from 2016 to 2019, and a Research Associate at University College London, London, U.K., from 2014 to 2016. He is currently an Associate Professor at TU Delft, Delft, The Netherlands. He has authored over 190 peer-reviewed publications and edited the books on *Micro-Doppler Radar and Its Applications* and *Radar Countermeasures for Unmanned Aerial Vehicles* (IET-Scitech, 2020). His research interests include the development of radar systems and automatic classification for human signatures analysis in healthcare and security, drones and UAVs' detection and classification, and automotive radar.

Dr. Fioranelli received four best paper awards and the IEEE AESS Fred Nathanson Memorial Radar Award 2024.

Alexander Yarovoy (Fellow, IEEE) received the Diploma degree (Hons.) in radiophysics and electronics, and the Candidate Phys. and Math.Sci. and Doctor Phys. and Math.Sci. degrees in radiophysics from Kharkiv State University, Kharkiv, Ukraine, in 1984, 1987, and 1994, respectively.

In 1987, he joined the Department of Radiophysics, Kharkiv State University, as a Researcher, where he became a Full Professor in 1997. From September 1994 to 1996, he was with the Technical University of Ilmenau, Ilmenau, Germany, as a Visiting Researcher. Since 1999, he has been with Delft University of Technology, Delft, The Netherlands, where he has been leading the Chair of microwave sensing, systems and signals since 2009. He has authored or co-authored more than 600 scientific or technical articles, 11 patents, and 14 book chapters. His main research interests are in high-resolution radar, microwave imaging, and applied electromagnetics (in particular, UWB antennas).

Dr. Yarovoy has been a member of numerous conference steering and technical program committees. He was a recipient of European Microwave Week Radar Award for the paper that best advances the state-of-the-art in radar technology in 2001 (together with L. P. Ligthart and P. van Genderen) and in 2012 (together with T. Savelyev). In 2023, together with Dr. I. Ullmann, N. Kruse, R. Gündel, and Dr. F. Fioranelli, he got the Best Paper Award at IEEE Sensor Conference. In 2010, together with D. Caratelli, he got the Best Paper Award of the Applied Computational Electromagnetic Society (ACES). He served as the General TPC Chair for the 2020 European Microwave Week (EuMW'20), the Chair and the TPC Chair for the 5th European Radar Conference (EuRAD'08), and the Secretary for the 1st European Radar Conference (EuRAD'04). He also served as the Co-Chair and the TPC Chair for the Xth International Conference on GPR (GPR2004). From 2008 to 2017, he served as the Director of European Microwave Association (EuMA). He serves on various editorial boards such as that of IEEE TRANSACTIONS ON RADAR SYSTEMS. From 2011 to 2018, he served as an Associate Editor for *International Journal of Microwave and Wireless Technologies*.