

Classifying Daily and Sports Activities Invariantly to the Positioning of Wearable Motion Sensor Units

Billur Barshan^{ID} and Aras Yurtman^{ID}

Abstract—We propose techniques that achieve invariance to the positioning of wearable motion sensor units on the body for the recognition of daily and sports activities. Using two sequence sets based on the sensory data allows each unit to be placed at any position on a given rigid body part. As the unit is shifted from its ideal position with larger displacements, the activity recognition accuracy of the system that uses these sequence sets degrades slowly, whereas that of the reference system (which is not designed to achieve position invariance) drops very fast. Thus, we observe a tradeoff between the flexibility in sensor unit positioning and the classification accuracy. The reduction in the accuracy is at acceptable levels, considering the convenience and flexibility provided to the user in the placement of the units. We compare the proposed approach with an existing technique to achieve position invariance and combine the former with our earlier methodology to achieve orientation invariance. We evaluate our proposed methodology on a publicly available data set of daily and sports activities acquired by wearable motion sensor units. The proposed representations can be integrated into the preprocessing stage of existing wearable systems without significant effort.

Index Terms—Accelerometer, activity recognition and monitoring, gyroscope, inertial sensors, Internet of Things (IoT), machine learning classifiers, magnetometer, position-invariant sensing, wearable motion sensors, wearable sensing.

I. INTRODUCTION

WITH the emergence of Internet of Things (IoT), products and practices are being transformed by communicating sensors and computing intelligence across many industries. Smart environments are continuously being developed and motion sensors such as low-cost inertial sensors are being embedded in many objects in the physical world that the users need to interact with in their daily lives (e.g., computer mouse, smartphone, tools, biomedical devices, kitchen and sports equipment). Bisio *et al.* [1] survey and compare the accelerometer signal classification methods to enable IoT for activity and movement recognition. The reference platforms used as elements of IoT in that article are smartphones at four different positions. Morales and Akopian [2] provide a detailed review of the studies that use smartphones in human activity

Manuscript received August 8, 2019; revised December 20, 2019; accepted January 14, 2020. Date of publication January 28, 2020; date of current version June 12, 2020. (Corresponding author: Billur Barshan.)

The authors are with the Department of Electrical and Electronics Engineering, Bilkent University, 06800 Ankara, Turkey (e-mail: billur@ee.bilkent.edu.tr).

Digital Object Identifier 10.1109/JIOT.2020.2969840

recognition. While developing algorithms for the interaction between the different elements of IoT and for processing the acquired sensory data, it is important to achieve position and orientation invariance in the placement of sensor units on these devices, which are either part of a smart environment or in wearable form. If the algorithms are restricted to operate only with predetermined sensor positions and orientations, this would require training the system for each possible sensor configuration and considerably increase the required amount of training time and training data. This is obviously not desirable.

Within the above context, recognizing human activities has attracted considerable interest in areas such as healthcare, sports science, fitness monitoring, and augmented/virtual reality [3]. Activity recognition and monitoring are performed either by motion sensor units worn by the user or sensors embedded in the environment such as cameras, accelerometers, vibration, and pressure sensors. These could be in the form of smart furniture, smart upholstery and mats, or smart floors. The former approach has become more preferable and advantageous as a result of the reduced size, weight, and longer battery life of wearable sensors as well as their integration into commonly used accessories such as smartphones, watches, and bracelets [4]. The latter approach restricts the user's mobility and raises privacy concerns.

Commonly employed motion sensor types are inertial sensors (accelerometers and gyroscopes) and magnetometers. These devices are typically triaxial, acquiring data on three mutually perpendicular axes x , y , and z . Recorded measurements of motion sensors depend on the position and orientation of the device. Naturally, users tend to place the sensor units with some uncertainty on their body parts, each time, at slightly different positions and orientations compared to the ideal. When this is the case, activity recognition accuracy tends to degrade. Expecting the users to place these devices every time in the same, predetermined way at their ideal positions and orientations is not only restrictive but also difficult to realize. This kind of restriction is impractical especially for elderly, disabled, or injured users who may need to put these devices on by themselves for fall detection, health monitoring, or physical therapy applications [5]–[7]. Even if the wearable devices are placed correctly at first, their positions and orientations may inevitably shift over time because of vibrations, movement, impacts on the body, etc. [8]. If the units are attached to clothing rather than directly on a body part, the problem is exacerbated. Sensor placement issue is also

present in prevalent smartphone applications because these devices are carried at different positions and orientations on the body such as in different pockets. Allowing the users to place the sensor units with some possible offset in position and orientation would bring them additional flexibility and make wearable activity recognition systems particularly advantageous and desirable compared to other approaches.

Uncertainty in the sensor unit placement is neglected in most applications of wearable sensing and it is unrealistically assumed that the users place each wearable exactly at the correct position and orientation on their body. Furthermore, most existing studies on activity recognition rely on methods that are sensitive to sensor placement [9]. Our earlier work on activity recognition addresses the orientations at which sensor units are worn on the body. We have proposed multiple techniques to transform sensory data such that they become invariant to the orientation of the sensor unit [10], [11], allowing the users to place the wearable units at any orientation at predetermined positions. To complement these studies, in this article, we tackle the issue of achieving invariance to the positioning of a sensor unit on a given rigid body part.

The main contribution of this article is to investigate how the activity recognition accuracy is affected by the random displacements of the sensor unit about its ideal spot while still being positioned on the same rigid body part. To reduce the degradation in accuracy, we propose the use of position-invariant sequences that can be extracted from short segments of recorded data independently. This allows the user to place each sensor unit anywhere on a given rigid body part on which it is supposed to be worn. We also provide a comparison based on the classification accuracy and the run times of the proposed approaches and an existing one.

The remainder of this article is organized as follows. Section II summarizes the related work on position invariance. We provide the methodology to achieve position invariance on a given rigid body part in Section III. In Section IV, we describe the data set and the activity recognition scheme that we employ in this article. We present the results in Section V. In Section VI, we provide and compare the run times of the proposed approaches and the classifiers considered in this article. We summarize our contributions, draw conclusions, and provide directions for future research in Section VII.

II. RELATED WORK

Existing methods to achieve robustness to the positioning of wearable motion sensor units can be grouped into four categories as described below [9], [12], [13], with their main attributes summarized in Table I.

A. Extracting Position-Invariant Information From Sensory Data

Some studies propose to heuristically transform the sensor data or extract heuristic features to achieve robustness to the positioning of the sensor units. Kunze and Lukowicz [8] argue that the acceleration caused by rotational movements depends

on the sensor position whereas the acceleration caused by linear movements is the same for all possible sensor positions on a given rigid body part. Based on this fact, the study neglects the acceleration data when there is a significant amount of rotational movement. This is decided to be the case if the difference between the magnitudes of the detected acceleration vector and the Earth's gravity (which approximately corresponds to the magnitude of motion-originated acceleration) is small compared to the magnitude of the angular acceleration derived based on the gyroscope output. Low-pass filtering the acceleration sequences brings out the gravitational component of acceleration which is considered to be invariant to the positioning of the unit on a given rigid body part [8], [14].

In addition to the activities of daily living, Hur *et al.* [15] consider two commuting activities during which vibrations are experienced by the whole body. Thus, the smartphone (whose motion sensors are used) is allowed to be placed at any position and orientation on the body. Classification is performed based on heuristic features extracted from the acceleration magnitude, discrete Fourier transform (DFT) of the vertical acceleration, and the speed measured by the global positioning system (GPS), which are obtained by using the built-in features of the Android mobile operating system.

B. Training Classifiers With Different Sensor Unit Positions

Another method to handle the varying positioning of the sensor units is to train an activity classifier in a generalized way to capture all (possible or considered) sensor unit positions. Some studies rely on such generalized classifiers mainly because data are inevitably acquired with multiple sensor configurations. This kind of natural variation in the data makes the activity recognition process inherently invariant to the positioning of the sensor units even though no specific techniques are developed or used for this purpose [16]–[20].

The data sets in [21]–[27] contain data from multiple sensor units and the data segments obtained from each unit are considered as separate training and test instances for generalized classification. In this scheme, the classifiers are trained with multiple unit positions and tested by using each position individually to demonstrate that the use of a single unit is sufficient for the recognition of activities. In [22]–[25], generalized classifiers trained with multiple sensor unit positions achieve accuracies slightly lower than position-specific classifiers. In [24], the accuracy further decreases when the leave-one-position-out method is used, where, for each position, a classifier trained with the data of the remaining positions is used. The studies [22], [23], [25]–[27] consider no more than several possible sensor unit positions and several activities, and the accuracy can drop abruptly if the numbers are increased. Förster *et al.* [24], on the other hand, classify aerobic movements with all the sensor units placed on the left leg and basic hand gestures with all the units worn on the right arm.

Förster *et al.* [24], Doppler *et al.* [28], Henpraserttae *et al.* [29], and Thiemjarus *et al.* [30] analyze the case where training and test data originate from different sensor unit positions and provide the accuracy for each position. When this is

TABLE I

ATTRIBUTES OF EXISTING WORKS ON POSITION INVARIANCE. DATA SETS ARE SEPARATED BY &. ACC: ACCELEROMETER, GYRO: GYROSCOPE, MAGN: MAGNETOMETER, GPS: GLOBAL POSITIONING SYSTEM, ADL: ACTIVITIES OF DAILY LIVING, Y: YES, N: NO. ACC AND GYRO ARE TRIAXIAL UNLESS STATED OTHERWISE

reference	sensors used	movement types	stationary activities	orientation invariance
[8] Kunze 2008	ACC, GYRO	forearm gestures & ADL, sports activities & gym exercises	N	Y (based on initial calibration, only for gesture recognition)
[9] Kunze 2014	ACC, GYRO	various data sets and methods analyzed	various data sets analyzed	Y (different methods analyzed separately)
[14] Jiang 2011	ACC	ADL, sports activities	Y	Y (based on initial calibration)
[15] Hur 2017	ACC, GYRO (for orientation invariance only), GPS	ADL (incl. transportation)	Y	Y
[16] Ravi 2016	ACC, GYRO	unknown & ADL & gait for Parkinson's disease & high-level activities	Y	Y
[17] Dembach 2012	ACC	ADL	Y	partial
[18] Wang 2009	ACC	ADL	Y (a single stationary class)	Y
[19] Khan 2013	ACC	ADL	Y (a single stationary activity)	N
[20] Khan 2014	ACC, air pressure sensor, microphone	ADL, sports activities	Y (a single stationary class)	N
[21] Khan 2010	ACC	ADL	Y (combined into a single stationary class)	N
[22] Lester 2006	ACC, microphone, barometer	ADL (incl. high-level activities)	Y	N
[23] Sun 2010	ACC	ADL	Y (a single stationary class)	partial
[24] Förster 2009	ACC	hand gestures (guided by geometric structures) & aerobic movements	N	N
[25] Reddy 2010	ACC, GPS	ADL (incl. transportation)	Y (combined into a single stationary class)	Y
[26] Anjum 2013	ACC, GYRO, GPS	ADL	Y (a single stationary class)	Y
[27] Siirtola 2013	ACC	ADL	Y (a single stationary class)	partial
[28] Doppler 2009	unknown (Euler angles used)	ADL (only three)	N	Y (based on initial calibration)
[29] Henpraserttae 2011	ACC	ADL	Y	partial
[30] Thiemjarus 2013	ACC	ADL	Y	Y
[31] Chavarriaga 2011	ACC	hand gestures (guided by geometric structures) & aerobic movements	N	partial (separate from position invariance)
[32] Förster 2009	ACC	synthetic data & hand gestures (guided by geometric structures) & aerobic movements	N	N
[33] Kunze 2005	ACC	ADL (incl. high-level activities)	Y	Y
[34] Xu 2012	ACC, biaxial GYRO	ADL (incl. transition activities)	N	N
[35] Lu 2010	ACC, microphone, GPS	ADL (incl. transportation)	Y (a single stationary class)	Y
[36] Martin 2013	ACC, GYRO, MAGN	ADL	Y	N
[37] Szttyler 2016	ACC	ADL	Y	partial
[39] Banos 2014	ACC	sports activities	N	partial
[40] Zhong 2014	ACC, GYRO	gait & gait	N	Y

the case, the accuracy significantly degrades because the data acquired only from a single unit position are not sufficient for training a generalized classifier. An acceptable accuracy level can be obtained if the training data include multiple positions, in particular, the position at which the test data are acquired.

C. Adapting Classifiers to New Sensor Unit Positions

Positioning the sensor units differently on the body causes variations in the features extracted from the acquired data. Chavarriaga *et al.* [31] and Förster *et al.* [32] assume that these variations only shift the class means in the feature space and calculate the amount of shifts in an unsupervised way (i.e., without the use of class labels) given new data obtained from a different sensor unit position. This assumption seems to hold for the position changes that occur on the same rigid body part (such as the torso or the left lower leg). However, both studies obtain unsatisfactory classification accuracies across different body parts. This is expected because the human body has 244 degrees of freedom where the body parts are connected at approximately 230 joints as a complex kinematic chain of links and joints. During an activity, in general, each body part exhibits different motion characteristics. A major drawback of

these adaptation-based methods is the difficulty of deciding when to start the adaptation process [31], [32].

D. Classifying Sensor Unit Positions

Some studies classify the sensor unit's position on the body during a predetermined set of activities assuming that there is a finite set of positions, which may not always be the case. Such position information can be used for context awareness or to select an activity classifier that is trained specifically for that position. Kunze *et al.* [33] distinguish the walking activity from the other activity types by training a generalized classifier for four predetermined sensor positions. Recordings of the walking activity of at least 1-min duration are used to classify the sensor unit's position. In this scheme, it is assumed that the sensor unit remains in the same position for at least a couple of minutes. Both classification techniques are invariant to the sensor unit orientations since the magnitudes of the acceleration vectors are used.

In [34], a sparse representation classifier is trained for each activity-sensor unit position pair. Then, the Bayesian fusion is used to recognize the activity type independently of the sensor unit position and to classify the position of the unit independently of the performed activity. Lu *et al.* [35] consider each

activity-position pair as a different class so that the activity and sensor unit position can be simultaneously classified. Another study [36] follows a two-stage approach by first classifying the sensor unit's position on the body and then recognizing the activity type using a classifier specifically trained for that position. By evaluating the accuracy through leave-one-subject-out cross validation (see Section IV) on the same data set, it shows that the two-stage approach performs considerably better than a single-stage generalized activity classifier trained by using all the sensor unit positions.

Sztyler and Stuckenschmidt [37] classify both the activity type and the sensor unit position following a three-stage approach which is more complicated. At each time segment, it first categorizes the activity into two (as static/dynamic) without the position information. Then, it classifies the sensor unit position by using the classifier specifically trained for the determined activity category. Finally, it recognizes the activity type by relying on the classifier trained for that particular sensor unit position. In all three steps, the classifiers are trained and tested separately for each subject. Hence, the method may not be generalizable to a new, unseen subject, considering that the activity recognition rate highly depends on the subject(s) from whom the training data are acquired [38].

E. Other Approaches

Banos *et al.* [39] rely on a machine-learning approach to fuse the decisions of multiple classifiers, each of which is trained specifically for one of the sensor units. (The conventional approach trains a single classifier by aggregating the features of all the units.) This method can tolerate incorrect positioning of a small subset of the sensor units by relying on the correctly placed ones during the classification process. Zhong and Deng [40] propose a transformation to handle differently oriented sensor units in gait classification, claiming that this method is invariant to the positioning of the units on laterally symmetric points on the body (e.g., left/right wrist).

In [41], a feature set which is independent of the orientation of the sensor unit and the movement speed is proposed and a two-stage signal processing algorithm is developed for activity/gesture recognition.

The publicly available data sets used in [42] contain activity data recorded by sensors at different positions on the body. To remove the effect of different sensor positions, activity data acquired from different positions are considered as sub data sets on which their proposed methodology is applied.

F. Discussion

Most of the existing techniques to achieve position invariance are not comparable with each other because of the differences in the sensor types and configuration, activity and movement types, classification and cross-validation techniques, and the way of evaluating the accuracy, as displayed in Table I. Moreover, the impact of the proposed position invariance methods on the accuracy is not always presented because of the lack of data acquired with correctly positioned sensor units.

A variety of activity or movement classes are considered in the previous studies (Table I), which highly affect the classification accuracy, as shown in [39]. Some studies consider only a single stationary activity (during which the subject is not even moving) [19], combine several activity types into a single class [18], [20], [21], [23], [25]–[27], [35], or do not include any [8], [24], [28], [31], [32], [34], [39], [40] (Table I). Activities that are often poorly classified or confused with each other are sometimes merged into a single class. For example, ascending and descending stairs are combined in [8], [17], and [33], which expectedly has a positive effect on the accuracy, given that these activities are classified with lower accuracy than the others in [9], [13], [19], [20], [22], [23], and [26]. In contrast, our data set includes a wide variety of stationary (static) and dynamic activities (see Section IV-A). We simultaneously classify a total of 19 daily and sports activities where activities similar to each other are not merged but considered as distinct classes. This is a more challenging problem than those addressed in many of the existing studies. Unlike most of the existing studies (other than [36]), we employ magnetometers in our proposed methodology and exploit magnetometer data to achieve invariance to the positioning of the sensor units in activity classification as well.

III. PROPOSED METHODOLOGY TO ACHIEVE POSITION INVARIANCE ON A RIGID BODY PART

Measurements acquired from motion sensor units are directly related to the linear and angular motion of the structure on which they are attached. We assume that the body part on which the sensor unit is placed (e.g., the lower arm) is rigid so that the relative position of any point with respect to another arbitrary point on the same body part remains constant in time during motion. In other words, the distance between any two arbitrary points is preserved. The motion of a rigid body at any time instant can be described by a translation and a rotation in 3-D space [43]. The points constituting the rigid body all have the same linear velocity and the same angular velocity. These velocities are represented by 3×1 column vectors \mathbf{v} and $\boldsymbol{\omega}$, respectively. The angular velocity (rate) vector $\boldsymbol{\omega}$ points along the instantaneous axis of rotation and its magnitude represents the rate of rotation. The direction of the rotation can be determined by using the right-hand rule. A triaxial gyroscope directly measures the angular rate vector $\boldsymbol{\omega}$ of the body part.

A magnetometer detects the vector sum of the Earth's magnetic field \mathbf{m} superposed with external magnetic sources, if any. The Earth's magnetic field vector points to the magnetic north and its magnitude and direction do not change significantly with the position of the sensor unit on the rigid body part, as well as throughout the human body. Hence, the three components of the magnetic field vector depend only on the orientation of the magnetometer but not on its position on a given body part.

Since both the gyroscope ($\boldsymbol{\omega}$) and the magnetometer (\mathbf{m}) sequences (as well as their magnitudes $|\boldsymbol{\omega}|$ and $|\mathbf{m}|$) are invariant to the positioning of the sensor unit on a given rigid body part, they can be directly used as position-invariant features

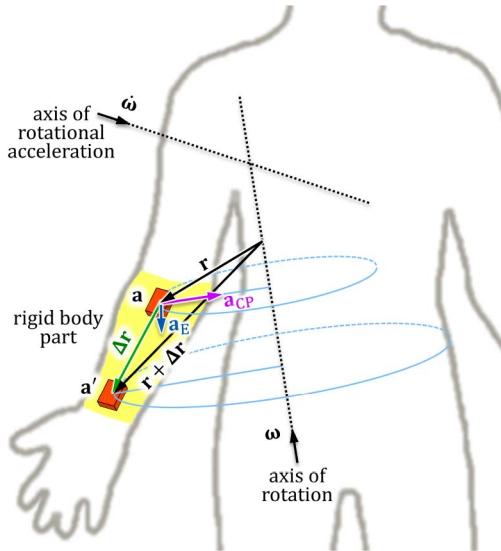


Fig. 1. Sensor unit positioning on a rigid body part (the lower arm). The displacement between two arbitrary positions as well as the centripetal and Euler components of the acquired acceleration vector are shown.

in the activity recognition process. On the other hand, the recorded acceleration sequences (\mathbf{a}) do depend on the position of the unit and the classification accuracy degrades when they are directly employed in the classification process, as we show later in Section V-B. Hence, we propose to select sequence sets that are invariant to the positioning of the sensor unit on a given body part and to use them in the classification process instead of the raw acceleration data.

According to the Coriolis theorem, an accelerometer detects the vector sum \mathbf{a} of multiple acceleration components which are the linear, centripetal, Euler, and Coriolis acceleration [43]

$$\mathbf{a} = \underbrace{\dot{\mathbf{v}} + \mathbf{g}}_{\mathbf{a}_L} + \underbrace{\omega \times (\omega \times \mathbf{r})}_{\mathbf{a}_{CP}} + \underbrace{\dot{\omega} \times \mathbf{r}}_{\mathbf{a}_E} + \underbrace{2\omega \times \dot{\mathbf{r}}}_{\mathbf{a}_C}. \quad (1)$$

Here, $\dot{\mathbf{v}}$ is the translational, \mathbf{g} is the gravitational, $\dot{\omega}$ is the angular acceleration, and \mathbf{r} is the position vector pointing from an arbitrary point on the axis of rotation to the center of the sensor unit, as illustrated in Fig. 1. The dot accent (') in (1) represents the first-order time derivative.

When the position of the sensor unit is shifted by $\Delta\mathbf{r}$ while still on the *same* rigid body part, the new position vector is $\mathbf{r}' = \mathbf{r} + \Delta\mathbf{r}$, where $\Delta\mathbf{r}$ is the sensor unit displacement vector (Fig. 1). The acceleration vector \mathbf{a}' of the displaced sensor unit can be expressed in terms of the acceleration vector \mathbf{a} at the original sensor unit position and the displacement $\Delta\mathbf{r}$ as follows:

$$\begin{aligned} \mathbf{a}' &= \dot{\mathbf{v}} + \mathbf{g} + \omega \times (\omega \times \mathbf{r}') + \dot{\omega} \times \mathbf{r}' + 2\omega \times \dot{\mathbf{r}}' \\ &= \dot{\mathbf{v}} + \mathbf{g} + \omega \times [\omega \times (\mathbf{r} + \Delta\mathbf{r})] + \dot{\omega} \times (\mathbf{r} + \Delta\mathbf{r}) \\ &\quad + 2\omega \times (\dot{\mathbf{r}} + \dot{\Delta\mathbf{r}}) \\ &= \mathbf{a} + \underbrace{\omega \times (\omega \times \Delta\mathbf{r})}_{\Delta\mathbf{a}_{CP}} + \underbrace{\dot{\omega} \times \Delta\mathbf{r}}_{\Delta\mathbf{a}_E} + \underbrace{2\omega \times \dot{\Delta\mathbf{r}}}_{\Delta\mathbf{a}_C}. \end{aligned} \quad (2)$$

We assume that once the user places the sensor unit on a certain body part, its position with respect to that body part remains fixed over time in the short term. (Here, short term

indicates the duration of a single data segment which is typically of the order of 1–10 s.) We represent this by keeping the sensor unit displacement $\Delta\mathbf{r}$ constant during each time segment of the data ($\dot{\Delta\mathbf{r}} = 0$). Thus, the Coriolis acceleration \mathbf{a}_C is not affected by the change in the sensor unit position on the same body part ($\Delta\mathbf{a}_C = 0$). This is also true for the linear acceleration component \mathbf{a}_L in (1) since both $\dot{\mathbf{v}}$ and \mathbf{g} are constant everywhere on the body part; the former, provided that $\Delta\mathbf{r} = \mathbf{0}$. Hence, shifting the position of the sensor unit while still on the same body part results in changes in the Euler and centripetal components ($\Delta\mathbf{a}_E$ and $\Delta\mathbf{a}_{CP}$) of the total acceleration vector.

For a given position displacement vector $\Delta\mathbf{r}$, the components $\Delta\mathbf{a}_E$ and $\Delta\mathbf{a}_{CP}$ are perpendicular to $\dot{\omega}$ and ω , respectively. Their magnitudes are calculated as follows:

$$\begin{aligned} \|\Delta\mathbf{a}_E\| &= \|\dot{\omega} \times \Delta\mathbf{r}\| = \|\dot{\omega}\| \|\Delta\mathbf{r}\| \sin(\angle(\dot{\omega}, \Delta\mathbf{r})) \\ \|\Delta\mathbf{a}_{CP}\| &= \|\omega \times (\omega \times \Delta\mathbf{r})\| = \|\omega\|^2 \|\Delta\mathbf{r}\| \sin(\angle(\omega, \Delta\mathbf{r})). \end{aligned} \quad (3)$$

To determine which of the two components is dominant in general, we define the ratio

$$\rho \triangleq \frac{\|\Delta\mathbf{a}_E\|}{\|\Delta\mathbf{a}_{CP}\|} = \frac{\|\dot{\omega}\|}{\|\omega\|^2} \frac{\sin(\angle(\dot{\omega}, \Delta\mathbf{r}))}{\sin(\angle(\omega, \Delta\mathbf{r}))}. \quad (4)$$

If $\rho \gg 1$, then, we may neglect $\Delta\mathbf{a}_{CP}$. (Later in Section V-C, we show that this is indeed the case for our large data set.) Then, we can claim that the projection

$$p \triangleq \mathbf{a} \cdot \frac{\dot{\omega}}{\|\dot{\omega}\|} \quad (5)$$

of the total acceleration onto the direction of $\dot{\omega}$ is independent of the sensor unit displacement $\Delta\mathbf{r}$ because the dominant component $\Delta\mathbf{a}_E$ that originates from the shifted position is orthogonal to $\dot{\omega}$. Hence, we consider the component of the total acceleration \mathbf{a} along the direction of $\dot{\omega}$ as a position-invariant feature which is approximately invariant to the sensor unit position on the same body part. In the rest of this article, we will denote this projection by p , which is a scalar quantity that can change with time.

The orientation of the sensor unit with respect to the Earth frame can also be employed as a feature which is independent of position. To estimate the orientation of the sensor unit at each time sample based on the accelerometer, gyroscope, and magnetometer data, we use the novel orientation estimation method we proposed in [44] and represent the 3-D orientation at each time sample efficiently by a 4×1 quaternion vector \mathbf{q} , as a feature that is invariant to the position of the unit on a certain body part.

We propose to investigate different combinations of the position-invariant sequences $|\omega|$, $|\mathbf{m}|$, ω , \mathbf{m} , p , and \mathbf{q} for activity classification. Our purpose is to assess the performance of different sequence combinations and identify the best performing one(s) when the sensor units are incorrectly positioned on the same rigid body part.

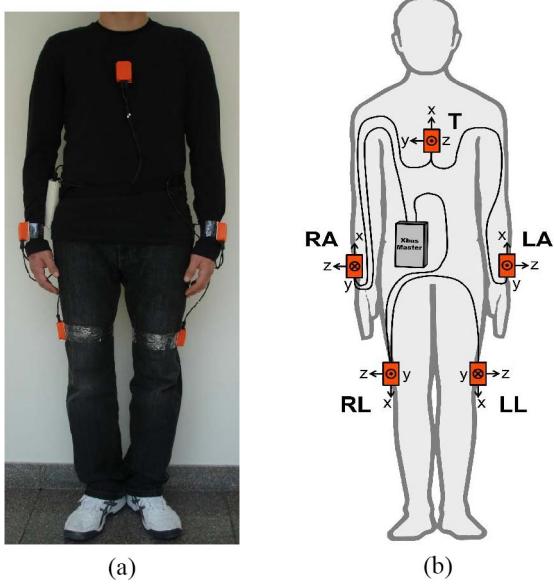


Fig. 2. (a) Configuration of the motion sensor units on the body. (b) Connection diagram of the units. [The body sketch in part (b) is from <http://www.clerk.com/clipart-male-figure-outline.html>; the cables, Xbus Master, and the motion sensor units were added by the authors.]

IV. DATA SET AND THE ACTIVITY RECOGNITION METHODOLOGY

A. Data Set

We use the publicly available daily and sports activities data set acquired by our research group [45], [46] by using five Xsens MTx sensor units [47] that were placed on the chest, on both wrists, and on the outer sides of both knees, as shown in Fig. 2. Each wearable unit contains three triaxial sensors, namely, an accelerometer, a gyroscope, and a magnetometer whose outputs are sampled at 25 Hz. Eight subjects performed the following 19 types of daily and sports activities:

Sitting; standing; lying on back; lying on right side; ascending stairs; descending stairs; standing still in an elevator; moving around in an elevator; walking in a parking lot; walking on a flat treadmill at a speed of 4 km/h; walking on a 15°-inclined treadmill at a speed of 4 km/h; running on a flat treadmill at a speed of 8 km/h; exercising on a stepper; exercising on a cross trainer; cycling on an exercise bike in horizontal position; cycling on an exercise bike in vertical position; rowing; jumping; and playing basketball.

The data set comprises 5-min recordings that consist of 7500 time samples each. For each activity performed by each subject, 45 (= 5 sensor units \times 3 sensor types \times 3 axes) time-domain sequences are recorded since each of the five motion sensor units contains three triaxial sensors.

B. Activity Recognition Scheme

To classify the activities, we follow the commonly used activity recognition scheme with the basic stages of data segmentation, feature extraction/normalization/reduction, and classification of the (possibly transformed) data [10], [11], [48]. We first divide the data into nonoverlapping segments of 5-s duration each. During the preprocessing stage, we either use the segmented data directly (the reference approach) or apply one of the two

transformations, described in Section III, to achieve robustness to sensor unit positioning.

We extract the following statistical features from each time-domain sequence of each segment: minimum, maximum, mean, variance, skewness, kurtosis, ten coefficients of the auto-correlation sequence for the lag values of 5, 10, ..., 45, 50 samples, and the five largest DFT peaks with the corresponding frequencies where the separation between any two peaks is taken to be at least 11 samples. There are 26 features for each time-domain sequence in each segment. For the reference approach that uses the **ωma** sequence set and does not involve any kind of position-invariant elements in the preprocessing stage, 1170 features (= 5 sensor units \times 9 axes \times 26 features) are concatenated to form a 1170-element feature vector for each segment. In general, the number of features depends on the number of vector elements and scalars comprising a given sequence set. For example, in the combination $|\mathbf{m}| \omega \mathbf{mpq}$, the total number of features per feature vector is 1560 (= 5 sensor units \times 12 elements \times 26 features). The features are normalized to the interval [0, 1] for each subject and the number of features is reduced to 30 through principal component analysis [49] which is a linear and orthogonal transformation where the transformed features are sorted to have variances in descending order.

We classify activities using seven state-of-the-art machine learning classifiers [11] described as follows.

Support Vector Machines (SVMs): The feature space is nonlinearly mapped to a higher-dimensional space through the use of a kernel function and divided into regions by hyperplanes. In this article, we select the kernel as a Gaussian radial basis function $f_{RBF}(\mathbf{x}, \mathbf{y}) = e^{-\gamma \|\mathbf{x} - \mathbf{y}\|^2}$ for any two (reduced) feature vectors \mathbf{x} and \mathbf{y} . The penalty parameter C (see [50, eq. (1)]) and the kernel parameter γ are jointly optimized by performing a two-level grid search and the optimal values of $C = 5$ and $\gamma = 0.1$ are used in SVM throughout this article. A binary SVM classifier is trained for each class pair and the decision of the classifier with the highest confidence level is taken [51]. The SVM classifier is implemented by using the MATLAB toolbox LibSVM [52].

Artificial Neural Networks (ANNs): We design a three-layer network of neurons where we select the input-output relationship of each neuron as a sigmoid function [53]. The number of neurons in the first (input) and the third (output) layers are equal to the reduced number of features (30) and the number of classes (K), respectively. When a test feature vector is provided at the input, the class decision is made by selecting the class corresponding to the neuron with the largest output (a scalar quantity). We select the number of neurons in the second (or hidden) layer as the integer nearest to the average of $[\log(2K)]/[\log 2]$ and $2K - 1$, with the former expression corresponding to the optimistic case where the hyperplanes intersect at different positions and the latter corresponding to the pessimistic one where they are parallel to each other. The weights of the linear combination calculated by each neuron are initialized randomly in the interval [0, 0.2]. During training, the weights are updated by the backpropagation algorithm [54] with a learning rate of 0.3. The algorithm terminates when the reduction in the error

(if any) compared to the average of the last ten epochs is less than 0.01.

Bayesian Decision Making (BDM): During the training phase, a multi-dimensional Gaussian distribution with an arbitrary covariance matrix is fitted to the training feature vectors of each class. Based on maximum-likelihood estimation, the mean vector is estimated as the arithmetic mean of the feature vectors and the covariance matrix is estimated as the sample covariance matrix for each class. In the test phase, the test vector's conditional probabilities given that it is associated with a particular class are calculated for each class. According to the maximum *a posteriori* decision rule, the class with the maximum conditional probability is selected [49], [53].

Linear Discriminant Classifier (LDC): The only difference of LDC from BDM is that the average of the covariance matrices, individually calculated for each class, is used overall. In this case, the Gaussians modeling the classes have identical covariance matrices but different mean vectors, causing them to be centered at different points in the feature space. Thus, the decision boundaries in the feature space correspond to hyperplanes, allowing the classes to be linearly separable [53].

k-Nearest Neighbor (k-NN): In the training phase, training vectors are stored with their class labels. In the classification phase, the class that the majority of the k training vectors that have the smallest Euclidean distance to the test vector belong to is selected [53]. Values of the k parameter between 1 and 30 are tested and $k = 7$ is employed.

Random Forest (RF): An RF classifier is a combination of multiple decision trees [55] where each tree is trained by randomly and independently sampling the training data. The splitting criterion at each node is the normalized information gain. The class decision is reached by majority voting over the tree decisions. We have used 100 decision trees and have observed that using more does not significantly improve the accuracy while increasing the computational cost considerably.

Orthogonal Matching Pursuit (OMP): The training phase consists of only storing the training vectors with their class labels as in *k*-NN. In the classification phase, each test vector is represented as a linear combination of a very small fraction of the training vectors with a bounded error. The vectors in this *sparse representation* are selected iteratively by using the OMP algorithm [56] where an additional training vector is selected at each iteration. The algorithm terminates when the desired representation error level (10^{-3}) is reached. Then, a residual for each class is calculated as the representation error when the test vector is represented as a linear combination of the training vectors of only that class, and the class with the minimum residual error is selected.

Cross-Validation Techniques: We have used two different cross-validation techniques to assess the accuracies of the classifiers: P -fold and leave-one-subject-out (L1O). In the former, the data set is randomly divided into $P = 10$ equal partitions. The feature vectors in each partition are classified using a classifier trained by the feature vectors in the remaining partitions and their accuracies are averaged out. The main difference of L1O from P -fold is that data are partitioned subjectwise so that each partition contains the data acquired from only one of the

eight subjects [53]. Thus, in L1O, there are eight partitions and the feature vectors of a given subject are left out while training the classifier with the remaining subjects' feature vectors. The left out subject's feature vectors are then used for testing (classification). This process is repeated for each subject. L1O is more challenging and highly affected by the variation in the data across the subjects, because the training and test sets are associated with different subjects, usually with larger variation between them [57]. It is usually employed to assess the generalizability of the system to an unseen subject and preferred over P -fold in scenarios where training data are not collected from the subjects who will use the system.

V. RESULTS

A. Random Position Displacement Model

To observe the effects of sensor unit positioning on the classification accuracy, we first consider the scenario where the units are randomly displaced from their ideal positions on the body parts at which they were originally placed while acquiring the training data. We generate a random displacement vector $\Delta\mathbf{r}$ independently for each sensor unit at the beginning of each time segment of the recorded data and then assume that it remains constant during that segment. We calculate the acceleration vector \mathbf{a}' for the displaced unit based on the originally acquired total acceleration vector \mathbf{a} using (2) (the last term being zero).

We assume that each sensor unit can be positioned on a disk with a given radius R centered at its ideal position. The random displacement $\Delta\mathbf{r}$ is confined to this disk which lies on the plane where the unit makes contact with the body part it is attached to. Note that all five sensor units make contact with the body on their x - y planes (see Fig. 2). Although we have considered three different plausible models in [12], here, we describe one of them where the points with displacement $\Delta\mathbf{r}$ from the origin are generated to have uniform distribution per unit area on the x - y plane. Two independent and identically (uniformly) distributed random variables Δr_x , Δr_y are generated such that Δr_x , $\Delta r_y \sim \mathcal{U}[-R, R]$, where $\Delta\mathbf{r} = [\Delta r_x, \Delta r_y, 0]^T$. The tip of the random vector $\Delta\mathbf{r}$ generated in this way falls uniformly onto a $2R \times 2R$ square region centered at the ideal position of the unit. If the tip remains outside the disk centered at the ideal position with radius R , the process is repeated as many times as needed until $\Delta\mathbf{r}$ resides inside the disk. (Note that the disk is inside the square and tangent to it on its sides.) Thus, the amount of displacement from the ideal spot is bounded by R in this model.

B. Random Position Displacement on a Rigid Body Part Without Attempting to Achieve Position Invariance

The standard activity recognition scheme is our reference case where the sensor units are fixed ideally at their correct positions and orientations. In this scheme, the originally recorded **omega** sequences are employed without attempting to achieve position invariance.

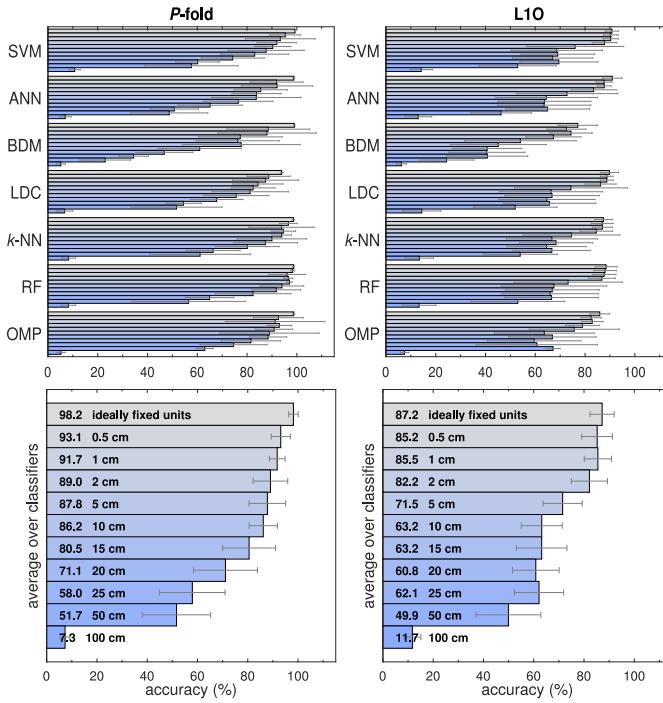


Fig. 3. Activity recognition accuracy for ideally fixed (reference case) and randomly shifted units. The lengths of the bars indicate the accuracy values for different R values. The thin sticks represent ± 1 standard deviation over the cross-validation iterations at the top, and over the classifiers at the bottom part of the figure.

To investigate the effect of randomly displacing the sensor units on the activity recognition accuracy, we shift the position of the sensor unit in the test data as described in the previous section, while keeping the training data associated with the correctly placed sensor units in their original form. Although it is more likely that users will put the units on their body with small displacements about their ideal positions (typically up to several centimeters), to determine the limitations of the standard activity recognition scheme, as well as the newly proposed schemes later, we consider R values between 0.5 and 100 cm.

In Fig. 3, we provide the classification accuracy of the reference system for each classifier separately at the top and by averaging over the seven classifiers at the bottom for the two cross-validation techniques. We observe that the activity recognition accuracy naturally degrades when the sensor units are attached to shifted positions about their ideal position on the body part they are supposed to be put on. Displacements up to a few centimeters can be tolerated by the standard activity recognition scheme whereas the accuracy significantly degrades for $R > 10$ cm. Such degradation in the accuracy is expected because the training data are associated with the correctly positioned sensor units while in the test data, positions of the sensor units are shifted randomly. The classifiers have not been trained and prepared for such displacement of the units. For $R = 100$ cm, the accuracy of L1O is higher than P -fold because the training data in L1O have wider variations (since each partition contains data acquired from one of the subjects) and the classifiers are trained to be more tolerant for possible variations in the test data.

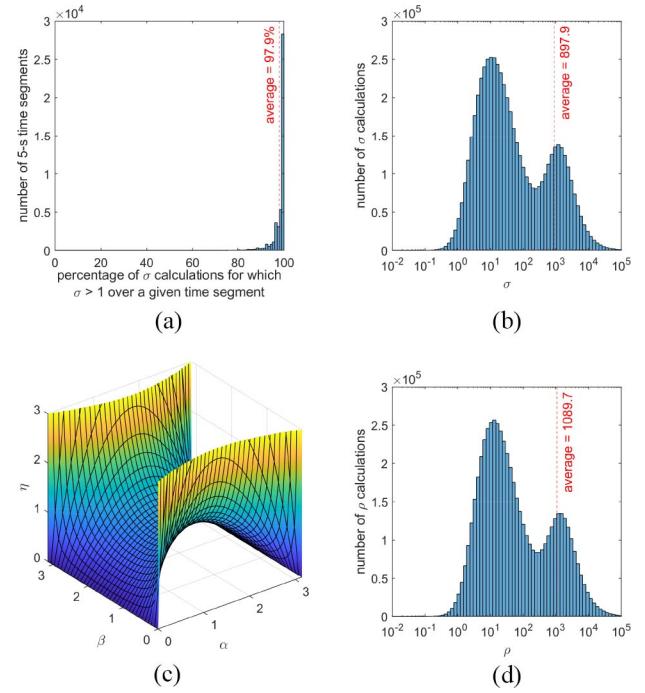


Fig. 4. Statistics of the quantities σ , η , and ρ that are related to \mathbf{a}_{CP} and \mathbf{a}_E . (a) Histogram of the percentage of σ calculations for which $\sigma > 1$ over one time segment. (b) Histogram of calculated σ values. (c) Surface plot for ρ on the α - β plane. (d) Histogram of calculated ρ values.

Note that up to this point, we have considered position shifts for the reference activity recognition system only, which uses the sequence set ω_{ma} , without any attempt to achieve position invariance on a given rigid body part.

C. Random Position Displacement on a Rigid Body Part With Position Invariance

First, we need to verify that the scalar quantity ρ , defined in (5) is indeed position invariant. By defining the following:

$$\sigma \triangleq \frac{\|\dot{\omega}\|}{\|\omega\|^2} \quad \alpha \triangleq \angle(\dot{\omega}, \Delta\mathbf{r}) \quad \eta \triangleq \frac{\sin \alpha}{\sin \beta} \quad \beta \triangleq \angle(\omega, \Delta\mathbf{r}) \quad (6)$$

the ratio in (4) may be expressed as $\rho = \sigma \eta$. We statistically analyze the quantities σ , η , and ρ in our data set as follows.

- Among all the 5-s time segments in the data set, $\sigma > 1$ in at least 68.8% of the time samples in each time segment. The histogram for the percentage of σ calculations for which $\sigma > 1$ per time segment is shown in Fig. 4(a).
- The average value of σ , over all the 5 700 000 values of σ calculated based on the data set, is $\bar{\sigma} = 897.9$. The histogram for σ is depicted in Fig. 4(b) where 97.3% of the σ values are greater than one.
- The ratio η is plotted as a function of the angles α and β in Fig. 4(c). The angles depend on the direction of the random displacement vector $\Delta\mathbf{r}$. The ratio η decreases as α approaches to 0 or π rad and increases as β approaches to 0 or π rad.
- Since the direction of $\Delta\mathbf{r}$ is uniformly distributed, the histogram for the distribution of ρ can be determined empirically. It is illustrated in Fig. 4(d) where we

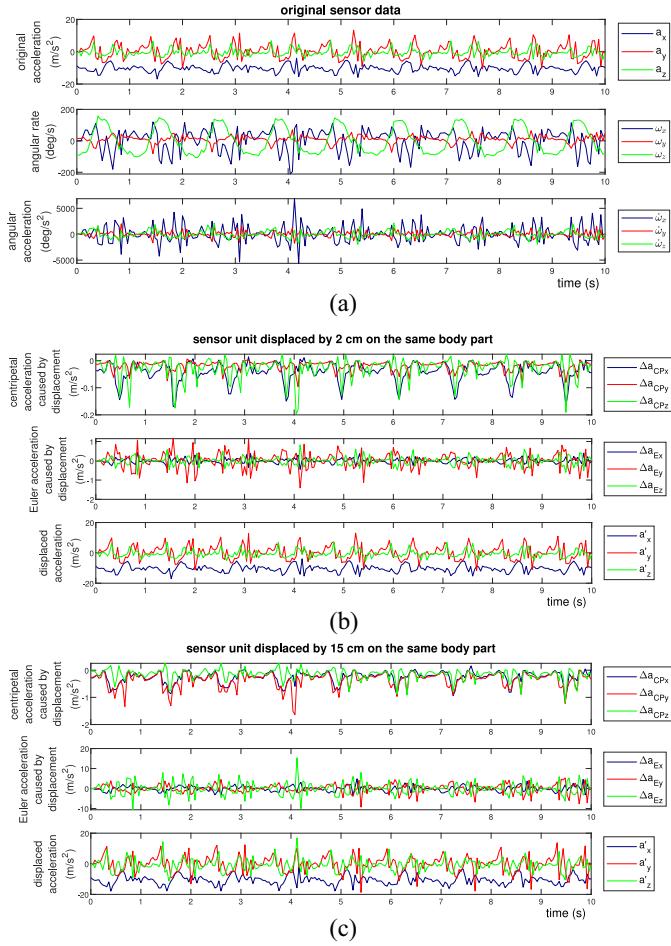


Fig. 5. Original and shifted acceleration data. (a) Acceleration, angular rate, and angular acceleration sequences acquired from the sensor unit at the original position. (b), (c) Centripetal, Euler, and shifted acceleration sequences calculated for the sensor unit when $R = 2$ cm and $R = 15$ cm.

observe that 97.8% of the calculated ρ values are greater than one.

These statistics indicate that $\rho \gg 1$; that is, $\|\Delta\mathbf{a}_E\| \gg \|\Delta\mathbf{a}_{CP}\|$ for the great majority of the time samples in the data set. Hence, we can rely on this fact and neglect the component $\Delta\mathbf{a}_{CP}$ to use p as a position-invariant feature on a given rigid body part.

The x , y , and z components of the original acceleration \mathbf{a} , angular rate $\boldsymbol{\omega}$, and angular acceleration $\dot{\boldsymbol{\omega}}$ vectors are plotted as functions of time in Fig. 5(a) for the sensor unit on the right leg of a subject during the activity of walking on a treadmill in the flat position. The components of the vectors $\Delta\mathbf{a}_{CP}$ and $\Delta\mathbf{a}_E$ caused by the sensor unit displacement, as well as the acceleration \mathbf{a}' for the shifted sensor unit, are plotted as functions of time for $R = 2$ cm and $R = 15$ cm in Fig. 5(b) and (c), respectively. We observe that $\Delta\mathbf{a}_E$ has a magnitude greater than $\Delta\mathbf{a}_{CP}$ most of the time and thus has a stronger effect on the acceleration \mathbf{a}' measured by the displaced sensor unit. The acceleration component p and the elements of the orientation quaternion \mathbf{q} are plotted as functions of time in parts (a) and (b) of Fig. 6, respectively, for the same recording illustrated in Fig. 5.

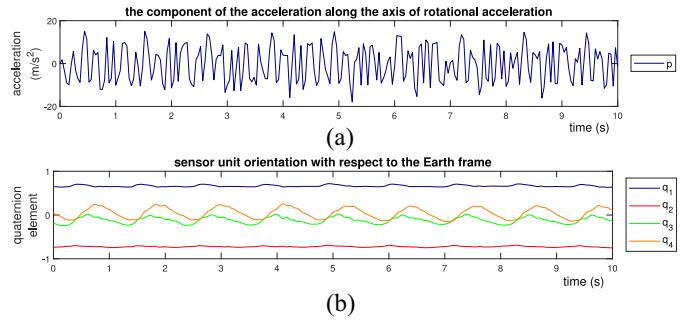


Fig. 6. Position-invariant sequences extracted from the sensor data. (a) Feature p . (b) Four elements of the quaternion \mathbf{q} .

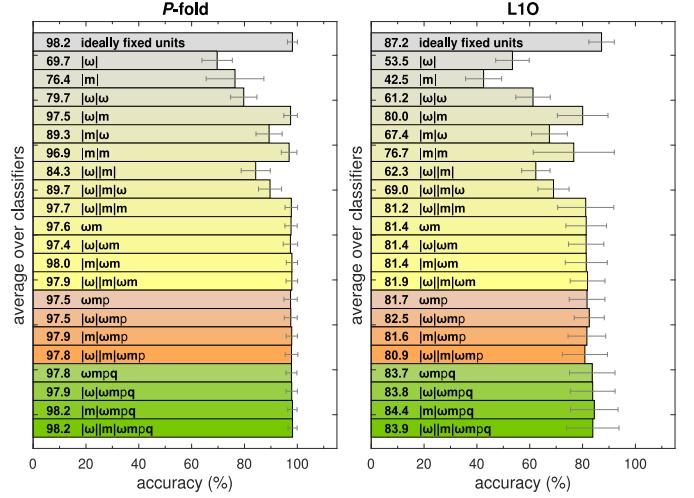


Fig. 7. Average activity recognition accuracy over the seven classifiers for the different sequence combinations considered to achieve position invariance on a given rigid body part.

We have considered a number of different combinations of the sequences $|\omega|$, $|\mathbf{m}|$, ω , \mathbf{m} , p , and \mathbf{q} to achieve position invariance on a given rigid body part. The results averaged over the classifiers are provided in Fig. 7, where we observe that the sequence combination $|\mathbf{m}|\omega\mathbf{m}\mathbf{p}\mathbf{q}$ results in the highest accuracy for P -fold. Since adding $|\omega|$ to this combination gives the same result for P -fold, results in a degradation for L10, and has a computational cost, we have not included $|\omega|$ in the selected sequence combinations.

In the following, we propose to use $|\mathbf{m}|\omega\mathbf{m}\mathbf{p}$ and $|\mathbf{m}|\omega\mathbf{m}\mathbf{p}\mathbf{q}$ to achieve invariance to sensor unit positioning when the units are shifted from their ideal spot on a given body part during the activity recognition scheme. We acquire $|\mathbf{m}|\omega\mathbf{m}\mathbf{p}$ (or $|\mathbf{m}|\omega\mathbf{m}\mathbf{p}\mathbf{q}$) based on the training data, whereas for the test data, we first randomly shift the positions of the sensor units and then acquire $|\mathbf{m}|\omega\mathbf{m}\mathbf{p}$ (or $|\mathbf{m}|\omega\mathbf{m}\mathbf{p}\mathbf{q}$).

Fig. 8 shows the activity recognition accuracy for the sequence set $|\mathbf{m}|\omega\mathbf{m}\mathbf{p}$ in combination with the random displacement model. The accuracy is not much affected by random sensor unit displacements up to $R = 50$ cm, whereas the maximum displacement of $R = 100$ cm causes a noticeable reduction in accuracy. On the other hand, the accuracy of the reference system starts degrading significantly after about $R = 10$ cm (compare Figs. 8 and 3). The results indicate that

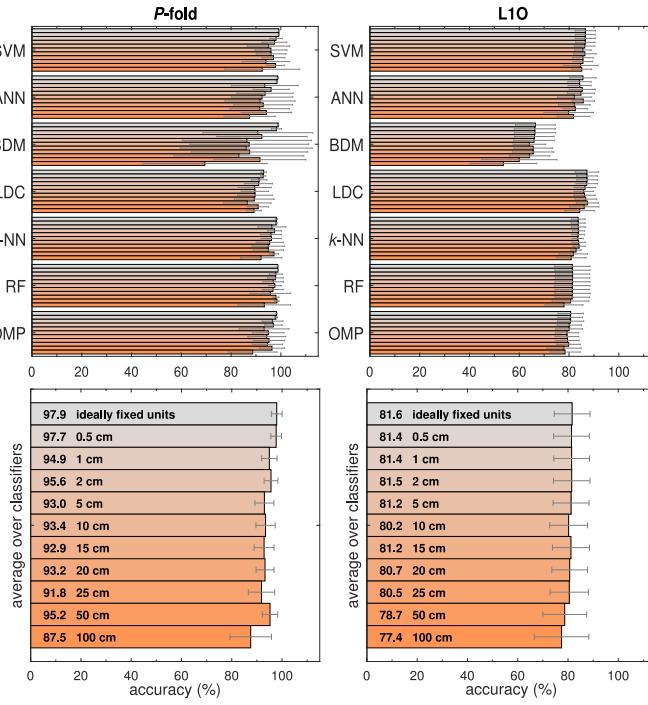


Fig. 8. Activity recognition accuracy for the $|m|omp$ sequence set for ideally fixed and randomly shifted units for different R values.

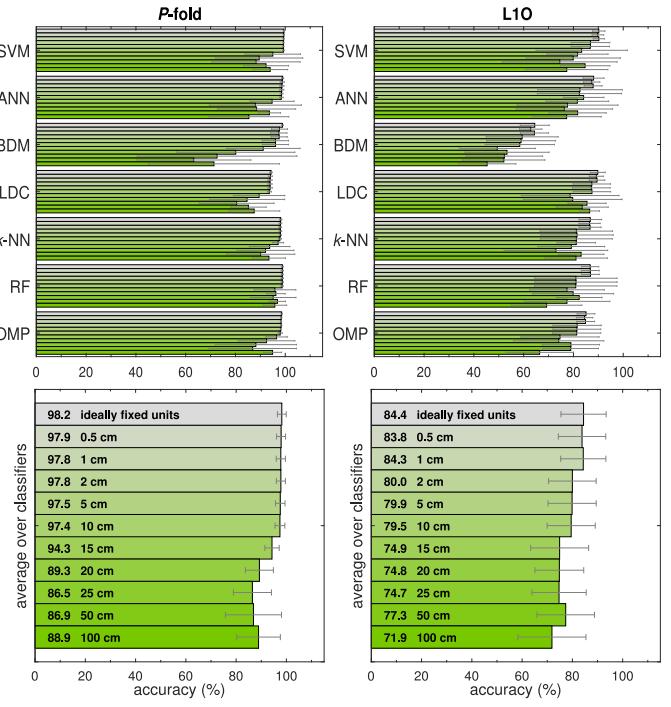


Fig. 9. Activity recognition accuracy for the $|m|ompq$ sequence set for ideally fixed and randomly shifted units for different R values.

the position-invariant features $|m|$ and p perform much better when used in place of the raw acceleration sequence a when the units are randomly displaced. Note also that using these scalar quantities is simpler compared to using the triaxial a .

The activity recognition accuracies for the sequence set $|m|ompq$ when combined with the random displacement model are provided in Fig. 9. Similar to $|m|omp$, using the sequence set $|m|ompq$ is robust to the displacement of the sensor units on a rigid body part up to about $R = 50$ cm. The accuracy of $|m|ompq$ is higher than $|m|omp$ on the average (compare Figs. 8 and 9).

D. Comparison of the Proposed Approach With an Existing One for Position Invariance on a Rigid Body Part

An existing approach that is applicable to our framework is to low-pass filter (LPF) the acceleration data [8], [14]. It is well known that the acceleration sequences recorded on Earth contain both gravitational and motion-originated components. Since the gravitational acceleration is independent of sensor unit placement, low-pass filtered acceleration sequences, dominated by gravity, are tolerant to the incorrect placement of the sensor units. We filter the acceleration sequences using a zero-phase Chebyshev type-II infinite impulse response LPF with a cut-off frequency of 10 Hz as proposed in [14] to extract the low-frequency components. In addition to the low-pass filtered acceleration sequence a_{LPF} , the gyroscope and magnetometer sequences, ω and m , are also used in the activity recognition process since they are already invariant to the positioning of the sensor unit on a given rigid body part.

Fig. 10 illustrates the activity recognition rates for the sequence set ωma_{LPF} [8], [14]. For random displacements

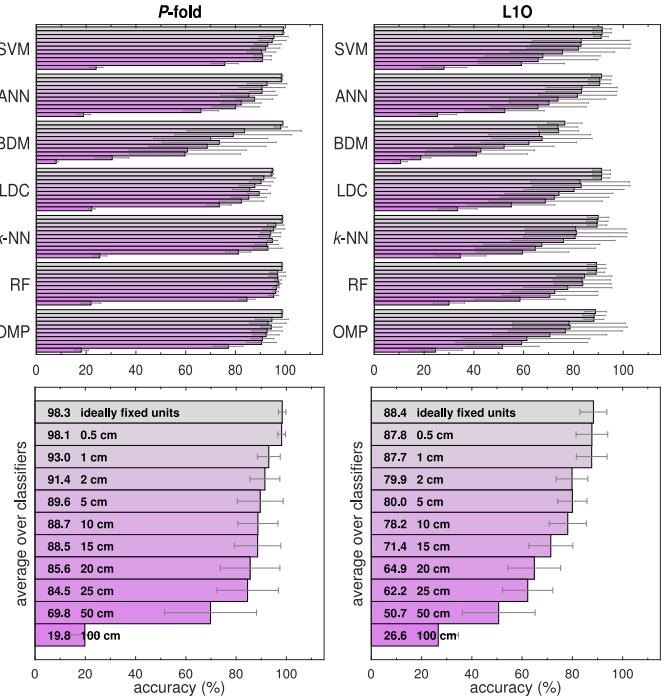


Fig. 10. Activity recognition accuracy for the ωma_{LPF} sequence set for ideally fixed and randomly shifted units for different R values.

bounded by a few centimeters, the accuracy achieved with this method is high but when the displacement exceeds several centimeters, it degrades at a much faster rate than those of the proposed $|m|omp$ and $|m|ompq$. This indicates that it is not as robust as the newly proposed sequence sets to the positioning of the sensor units. In particular, for the maximum sensor unit displacement of $R = 100$ cm, the existing approach

$\omega\mathbf{m}_{LPF}$ performs poorly, whereas the proposed $|\mathbf{m}|\omega\mathbf{m}$ and $|\mathbf{m}|\omega\mathbf{m}\mathbf{q}$ perform fairly well. This is mainly because the two sequence set combinations are selected in a way to be robust to position shifts on the same rigid body part.

E. Randomly Changing Both the Position and Orientation of the Sensor Unit on a Rigid Body Part Without Invariance

In this section, we combine the proposed methodology to achieve position invariance on a rigid body part with our earlier methodology to achieve orientation invariance. To represent sensor units whose positions and orientations are randomly changed while still on the same body part, we first shift the sensor units according to the model described in Section V-A. Then, we randomly rotate the sensor units in 3-D space about their shifted positions. For this purpose, we independently generate a random rotational transformation for each time-segmented window of data (of 5-s duration) simultaneously acquired from the nine axes of each sensor unit. The corresponding rotation matrix \mathbf{R} is calculated based on the three Euler angles ϕ , θ , and ψ (yaw, pitch, and roll) that are randomly and uniformly distributed in the interval $(-\pi, \pi]$ radians (where $s\phi \triangleq \sin \phi$, $c\phi \triangleq \cos \phi$, etc.).

$$\mathbf{R} = \begin{bmatrix} c\phi & -s\phi & 0 \\ s\phi & c\phi & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} c\theta & 0 & s\theta \\ 0 & 1 & 0 \\ -s\theta & 0 & c\theta \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & c\psi & -s\psi \\ 0 & s\psi & c\psi \end{bmatrix}. \quad (7)$$

We premultiply each of the 3×1 measurement vectors ω , \mathbf{m} , and \mathbf{a} with this rotation matrix to get the respective vectors $\mathbf{R}\omega$, $\mathbf{R}\mathbf{m}$, and $\mathbf{R}\mathbf{a}$ of the same size, that would have been obtained if the sensor unit were rotated in 3-D. Note that the measurement vectors of each of the three sensor types in the same unit are rotated in the same way throughout a given time segment. These transformations represent the case where each sensor unit is placed at a random position and orientation within a disk of radius R , whose center corresponds to the ideal position of the sensor unit on a given rigid body part.

The results of random displacements about the ideal position in combination with random rotation as described above, are provided in Fig. 11 for the reference system. Compared to ideally positioned and oriented units (topmost bars), keeping the position fixed while randomly rotating the units (second bars from the top) decreases the average accuracy abruptly, by more than 56%. When the units are displaced as well, the accuracy degradation is even more and keeps increasing with larger displacements from the ideal position.

Note that, as in Section V-B, we keep the training data in their original form. We randomly change both the position and the orientation of the sensor units only at the beginning of each time segment of the test data. This corresponds to the real-world scenario where the user puts the sensor units on with some error while using a wearable system previously trained when the units were ideally placed on his/her body.

Again, we note that in this section, we have considered random position and orientation shifts for the reference activity recognition system only, which uses the sequence triplet $\omega\mathbf{m}\mathbf{a}$, without attempting to simultaneously achieve position and orientation invariance on a given rigid body part yet.

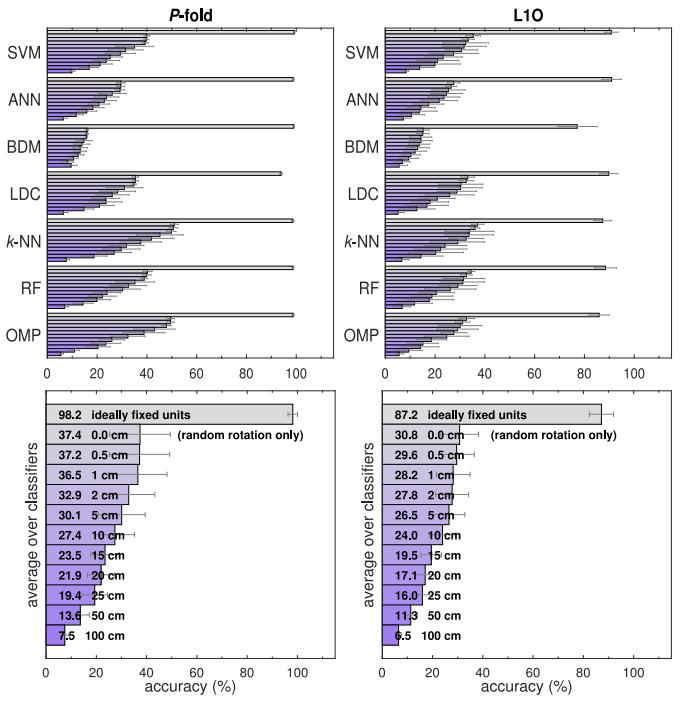


Fig. 11. Activity recognition accuracy for ideally fixed, randomly rotated, and both randomly shifted (for different R values) and randomly rotated units.

F. Methodology to Achieve Simultaneous Position and Orientation Invariance on a Rigid Body Part

We have considered achieving invariance to the orientation of wearable sensor units within the context of activity recognition in our earlier works [10], [11]. In [10], we have considered invariance to sensor unit orientation and developed two novel transformations (based on heuristics and singular-value decomposition) to remove the effect of absolute sensor orientation from the raw sensor data. The method proposed in [11] is based on transforming the recorded motion sensor sequences invariantly to sensor unit orientation by estimating the sensor unit orientation and representing the sensor data with respect to the Earth frame. We also determine the rotation of the sensor unit between consecutive time samples and represent it by quaternions with respect to the Earth frame.

To achieve orientation invariance in addition to position invariance on a given rigid body part, we replace the position-invariant sequences used in Section III with their counterparts that can also achieve orientation invariance. To do this, we first estimate the orientation of the sensor unit with respect to the Earth frame based on the sensor recordings by using the novel orientation estimation method we proposed in [44]. Based on the estimated orientation, we represent the position invariance ω and \mathbf{m} with respect to the Earth frame, denoting them with the superscript E. (Previously, these quantities were represented with respect to the sensor unit frame.) Note that \mathbf{p} is a scalar quantity independent of the reference frame.

In addition, we replace the orientation quaternion \mathbf{q} (with respect to the Earth frame) with the differential orientation quaternion $\delta\mathbf{q}$. To obtain $\delta\mathbf{q}$, we first calculate the differential rotation matrix \mathbf{D}_n that represents the rotation of the sensor unit frame between two consecutive time samples (n and $n+1$)

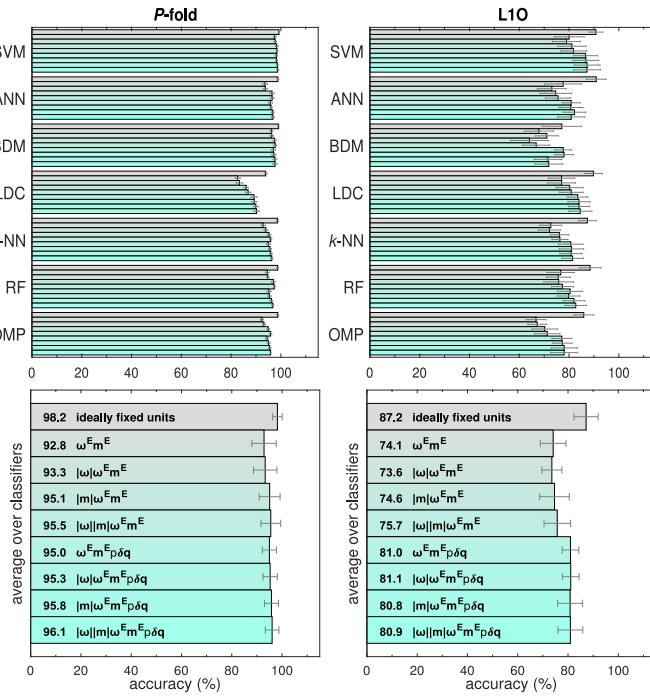


Fig. 12. Activity recognition accuracy for the different sequence combinations that are considered to achieve simultaneous position and orientation invariance.

with respect to the Earth frame [11]. Representing a rotational transformation with a 3×3 matrix (of nine elements) is inefficient because any 3-D rotation can be described by only three angles. Instead, we represent the differential rotation matrix \mathbf{D}_n compactly by a four-element differential quaternion $\delta\mathbf{q}_n$ (with respect to the Earth frame) as

$$\delta\mathbf{q}_n = \begin{bmatrix} \delta q_1 \\ \delta q_2 \\ \delta q_3 \\ \delta q_4 \end{bmatrix} = \begin{bmatrix} \frac{\sqrt{1+d_{11}+d_{22}+d_{33}}}{2} \\ \frac{d_{32}-d_{23}}{2} \\ \frac{4\sqrt{1+d_{11}+d_{22}+d_{33}}}{d_{13}-d_{31}} \\ \frac{4\sqrt{1+d_{11}+d_{22}+d_{33}}}{d_{21}-d_{12}} \\ \frac{4\sqrt{1+d_{11}+d_{22}+d_{33}}}{d_{31}-d_{13}} \end{bmatrix} \quad (8)$$

where d_{ij} ($i, j = 1, 2, 3$) are the elements of \mathbf{D}_n [58]. (Here, the dependence of the elements of $\delta\mathbf{q}_n$ and \mathbf{D}_n on n has been dropped from the notation for simplicity). We finally drop the subscript n from $\delta\mathbf{q}_n$ as well and simply denote it by $\delta\mathbf{q}$ in the following.

Since the sequences $|\omega|$, $|\mathbf{m}|$, ω^E , \mathbf{m}^E , p , and $\delta\mathbf{q}$ do not depend on the orientation at which the units are worn on the body, it is possible to achieve invariance to sensor unit orientation by employing combinations of these quantities.

We have considered a number of different combinations of the above-mentioned sequences to achieve simultaneous position and orientation invariance on a given rigid body part. The results are provided in Fig. 12 where we observe that the sequence combination $|\omega||\mathbf{m}|\omega^E\mathbf{m}^E p \delta\mathbf{q}$ gives the highest average accuracy for P -fold. On the other hand, the highest average accuracy for $L1O$ is obtained with $|\omega|\omega^E\mathbf{m}^E p \delta\mathbf{q}$. Replacing $|\omega|$ with $|\mathbf{m}|$ or the addition of $|\mathbf{m}|$ to the sequence set in $L1O$ degrades the average accuracy by 0.2%–0.3%. Compared to the reference system (topmost bars in the figure), using $|\omega||\mathbf{m}|\omega^E\mathbf{m}^E p \delta\mathbf{q}$ results in 2.1% lower accuracy for P -fold.

TABLE II
AVERAGE PROCESSING TIME TO TRANSFORM THE ORIGINAL SEQUENCE SET $\omega\mathbf{m}\mathbf{a}$ INTO A NEW REPRESENTATION PER 5-s TIME SEGMENT DURING THE PREPROCESSING STAGE

new representation	processing time per segment (ms)
$\omega\mathbf{m}\mathbf{a}_{LPF}$	6.43
$\omega\mathbf{m}\mathbf{p}$	2.57
$\omega\mathbf{m}\mathbf{p}\mathbf{q}$	11.19
$ \mathbf{m} \omega\mathbf{m}\mathbf{p}$	3.75
$ \mathbf{m} \omega\mathbf{m}\mathbf{p}\mathbf{q}$	11.58
$\omega^E\mathbf{m}^E p \delta\mathbf{q}$	32.86
$ \omega \omega^E\mathbf{m}^E p \delta\mathbf{q}$	35.96
$ \mathbf{m} \omega^E\mathbf{m}^E p \delta\mathbf{q}$	35.59
$ \omega \mathbf{m} \omega^E\mathbf{m}^E p \delta\mathbf{q}$	35.81

Using $|\omega|\omega^E\mathbf{m}^E p \delta\mathbf{q}$ in $L1O$ degrades the accuracy by 6.1% but achieves orientation invariance as well. All of the results given in Fig. 12 are for the ideal position and orientation of the sensor unit. We observe that the sequence sets considered here exhibit an acceptable drop in accuracy compared to the reference system. Since these sequence sets are selected to be both position and orientation invariant to begin with, when the sensor unit is randomly shifted in position and randomly rotated as well, we do not expect the accuracy to degrade as fast as that of the reference system (Fig. 11).

VI. RUNTIME ANALYSIS

The average processing times of the techniques considered in this article to transform the original sequence set $\omega\mathbf{m}\mathbf{a}$ into a new representation during the preprocessing stage are provided in Table II per 5-s time segment. This includes the existing approach $\omega\mathbf{m}\mathbf{a}_{LPF}$, the proposed sequence sets to achieve position invariance, and those to achieve simultaneous position and orientation invariance, as well as several other sequence combinations that we have considered. The processing was performed on a laptop computer containing a quad-core Intel Core i7-4720HQ processor with a clock speed of 2.6–3.6 GHz and 16 GB of RAM running 64-bit MATLAB R2018b.

Among the techniques that achieve only position invariance (the first five rows of the table), the proposed $|\mathbf{m}|\omega\mathbf{m}\mathbf{p}$ sequence set is computationally more efficient than the existing approach $\omega\mathbf{m}\mathbf{a}_{LPF}$, whereas the second proposed sequence combination $|\mathbf{m}|\omega\mathbf{m}\mathbf{p}\mathbf{q}$ takes longer to execute.

Note that some members of the sequence sets in the last four rows of the table are represented with respect to the Earth frame, requiring the estimation of sensor unit orientation. For this purpose, we employ the novel method we proposed in [44], which takes most of the processing time. Nevertheless, all of the run times in the table are much shorter than the duration of a single time segment (5 s), indicating that the new representations can be obtained in near real time.

Table III shows the average run times of the classifiers with their standard deviations over all the transformation techniques considered in this article (Figs. 7 and 12). In the second column of the table, we provide the average of the total run time

TABLE III

AVERAGE RUN TIME OVER ALL THE TRANSFORMATION TECHNIQUES
CONSIDERED WITH ONE STANDARD DEVIATION IN PARENTHESES

classifier	total run time (s)	training time (s)	classification time (ms)
SVM	11.91 (1.27)	9.94 (1.02)	0.48 (0.07)
ANN	6.99 (1.22)	6.98 (1.21)	0.01 (0.00)
BDM	1.65 (0.31)	0.01 (0.00)	1.60 (0.30)
LDC	2.05 (0.36)	0.45 (0.04)	0.04 (0.01)
<i>k</i> -NN	0.15 (0.05)	—	0.15 (0.04)
RF	34.92 (2.99)	30.83 (2.54)	1.15 (0.12)
OMP	245.63 (305.76)	—	236.70 (296.66)

(including the training phase, classification of all test feature vectors in the test phase, and programming overheads) per cross-validation iteration. We observe that the *k*-NN classifier has the shortest average total run time among the seven classifiers, whereas OMP has the longest.

Average training times of the classifiers per cross-validation iteration are given in the third column of Table III. Since the *k*-NN and OMP classifiers only store the training feature vectors, effectively, they have zero training time. On the other hand, the RF classifier takes the longest to train.

The average classification time per single test feature vector associated with a 5-s time segment is given in the fourth column of Table III. Although all of the classifiers can label a test feature vector in a duration much shorter than 5 s, the ANN and LDC classifiers perform this operation almost instantly, followed by *k*-NN, identifying the activity in no longer than 0.15 ms. The OMP classifier has the longest classification time because it executes an iterative algorithm independently for each test feature vector, but its run time is still much shorter than the segment duration, allowing near real-time implementation.

VII. DISCUSSION AND CONCLUSION

We have focused on the positioning of wearable sensor units and proposed methods that allow the user the flexibility to wear each sensor unit at shifted positions while keeping it on the same rigid body part. To achieve position invariance under these circumstances, we have proposed novel approaches based on the use of a set of position-invariant sequences. We have demonstrated their robustness to the positioning of a sensor unit on a rigid body part compared to an existing approach and the reference system. Since the reference system is not designed to achieve position invariance, it is highly vulnerable to position shifts. The proposed sequence sets cause small degradation in the activity recognition accuracy when the units are correctly placed and obtain much higher accuracies than the reference system when the position of the sensor unit is shifted while the unit remains on the same rigid body part. The main reason for the slower accuracy degradation is the appropriate choice of the position-invariant sequence sets in the proposed methods.

We have extended our methodology to achieve invariance to both the position and orientation of wearable motion sensor units at the same time. This scheme allows the user further

flexibility to place the wearable sensor units at any position and orientation on a given rigid body part, provided that data were acquired from that body part when the sensor unit was ideally configured and the system was trained with that data. More importantly, it substantially reduces the amount of time and data required to train the system since it is no longer necessary to train for each possible position and orientation of the sensor units.

Theoretically, our proposed methods are generalizable to position and orientation shifts *within* a rigid body part as well (rather than being restricted to the *surface* of the body part). However, in practice, sensor units are typically displaced *on* the surface of the body part unless they are implanted.

We have comparatively evaluated the proposed and existing approaches using our publicly available data set containing daily and sports activities which are larger in number and more complex than those considered in existing studies. We have employed seven state-of-the-art machine learning classifiers and two cross-validation techniques to demonstrate the robustness of our methodology.

In developing the techniques in this article, we have intentionally not used information on the activity types in the data set and the sensor unit positions because our purpose was to keep the proposed techniques sufficiently general to be applicable to a broad range of wearable systems and scenarios. The proposed transformations can be applied to each time segment of the acquired data independently. Hence, the impact of a shift or sudden change in the positions (and orientations) of the sensor units is limited to the time segment during which the change occurs. The classification accuracy in future time segments is not affected. The newly proposed techniques employ multi-dimensional time-domain sequences with a format similar to that of the raw data. As such, it is straightforward to integrate them into a wide variety of existing wearable systems by transforming the sensory data in the preprocessing stage without much effort. This way, the system becomes robust to variable sensor unit placement and its performance is not significantly affected by shifts in the positions and orientations of the sensor units on rigid body parts.

An interesting future research direction would be to investigate how the frequency content of the acquired acceleration signals changes with random shifts in the sensor unit position. This would depend on the body part the sensor unit is placed on, the type of activity being performed, and the rate of change of the angular velocity and acceleration during that activity on a given rigid body part. While $\Delta\mathbf{r}$ is assumed to be constant during each short time segment of data in this article, the case of $\Delta\mathbf{r}$ instantaneously changing with time (at each time sample) can be considered. We expect that this will also modify the frequency content of the Euler and centripetal acceleration components.

One may continue investigating additional features which are robust and invariant to the positioning of the sensor units as well as to the orientation, and the combination of both. Methods can be developed to achieve position and orientation invariance across different body parts and the lateral symmetry of the human body can be exploited. The dependence of the invariance techniques on the activity type or category can be

investigated. An important related problem is the investigation of position and orientation invariance of sensor units on the other elements of IoT which are not necessarily worn on the body, but are part of a smart environment.

REFERENCES

- [1] I. Bisio, A. Delfino, F. Lavagetto, and A. Sciarrone, "Enabling IoT for in-home rehabilitation: Accelerometer signals classification methods for activity and movement recognition," *IEEE Internet Things J.*, vol. 4, no. 1, pp. 135–146, Feb. 2017.
- [2] J. Morales and D. Akopian, "Physical activity recognition by smartphones, a survey," *Biocybern. Biomed. Eng.*, vol. 37, no. 3, pp. 388–400, 2017.
- [3] Ö. Yürür, C. H. Liu, and W. Moreno, "A survey of context-aware middleware designs for human activity recognition," *IEEE Commun. Mag.*, vol. 52, no. 6, pp. 24–31, Jun. 2014.
- [4] O. D. Lara and M. A. Labrador, "A survey on human activity recognition using wearable sensors," *IEEE Commun. Surveys Tuts.*, vol. 15, no. 3, pp. 1192–1209, 3rd Quart., 2013.
- [5] J. Qi, P. Yang, M. Hanneghan, S. Tang, and B. Zhou, "A hybrid hierarchical framework for gym physical activity recognition and measurement using wearable sensors," *IEEE Internet Things J.*, vol. 6, no. 2, pp. 1384–1393, Apr. 2019.
- [6] A. T. Özdemir and B. Barshan, "Detecting falls with wearable sensors using machine learning techniques," *Sensors*, vol. 14, no. 6, pp. 10691–10708, Jun. 2014.
- [7] A. Yurtman and B. Barshan, "Automated evaluation of physical therapy exercises using multi-template dynamic time warping on wearable sensor signals," *Comput. Methods Programs Biomed.*, vol. 117, no. 2, pp. 189–207, Nov. 2014.
- [8] K. Kunze and P. Lukowicz, "Dealing with sensor displacement in motion-based onbody activity recognition systems," in *Proc. 10th Int. Conf. Ubiquit. Comput.*, Seoul, South Korea, Sep. 2008, pp. 20–29.
- [9] K. Kunze and P. Lukowicz, "Sensor placement variations in wearable activity recognition," *IEEE Pervasive Comput.*, vol. 13, no. 4, pp. 32–41, Oct.–Dec. 2014.
- [10] A. Yurtman and B. Barshan, "Activity recognition invariant to sensor orientation with wearable motion sensors," *Sensors*, vol. 17, no. 8, Aug. 2017, Art. no. 1838.
- [11] A. Yurtman, B. Barshan, and B. Fidan, "Activity recognition invariant to wearable sensor unit orientation using differential rotational transformations represented by quaternions," *Sensors*, vol. 18, no. 8, Aug. 2018, Art. no. 2725.
- [12] A. Yurtman, "Activity recognition invariant to position and orientation of wearable motion sensor units," Ph.D. dissertation, Dept. Elect. Electron. Eng., Bilkent Univ., Ankara, Turkey, Apr. 2019.
- [13] M. Shoaib, S. Bosch, Ö. D. Incel, H. Scholten, and P. J. M. Havinga, "A survey of online activity recognition using mobile phones," *Sensors*, vol. 15, no. 1, pp. 2059–2085, Jan. 2015.
- [14] M. Jiang, H. Shang, Z. Wang, H. Li, and Y. Wang, "A method to deal with installation errors of wearable accelerometers for human activity recognition," *Physiol. Meas.*, vol. 32, no. 3, pp. 347–358, Feb. 2011.
- [15] T. Hur, J. Bang, D. Kim, O. Banos, and S. Lee, "Smartphone location-independent physical activity recognition based on transportation natural vibration analysis," *Sensors*, vol. 17, no. 4, Apr. 2017, Art. no. 931.
- [16] D. Ravì, C. Wong, B. Lo, and G.-Z. Yang, "A deep learning approach to on-node sensor data analytics for mobile or wearable devices," *IEEE J. Biomed. Health Inform.*, vol. 21, no. 1, pp. 56–64, Jan. 2017.
- [17] S. Dernbach, B. Das, N. C. Krishnan, B. L. Thomas, and D. J. Cook, "Simple and complex activity recognition through smart phones," in *Proc. 8th Int. Conf. Intell. Environ.*, Guanajuato, Mexico, Jun. 2012, pp. 214–221.
- [18] Y. Wang *et al.*, "A framework of energy efficient mobile sensing for automatic user state recognition," in *Proc. 7th Int. Conf. Mobile Syst. Appl. Services (MobiSys)*, Kraków, Poland, Jun. 2009, pp. 179–192.
- [19] A. M. Khan, M. H. Siddiqi, and S.-W. Lee, "Exploratory data analysis of acceleration signals to select light-weight and accurate features for real-time activity recognition on smartphones," *Sensors*, vol. 13, no. 10, pp. 13099–13122, Sep. 2013.
- [20] A. M. Khan, A. Tufail, A. M. Khattak, and T. H. Laine, "Activity recognition on smartphones via sensor-fusion and KDA-based SVMs," *Int. J. Distrib. Sens. Netw.*, vol. 10, no. 5, May 2014, Art. no. 503291.
- [21] A. M. Khan, Y. K. Lee, and S. Y. Lee, "Accelerometer's position free human activity recognition using a hierarchical recognition model," in *Proc. 12th IEEE Int. Conf. e-Health Netw. Appl. Services*, Lyon, France, Jul. 2010, pp. 296–301.
- [22] J. Lester, T. Choudhury, and G. Borriello, "A practical approach to recognizing physical activities," in *Pervasive Computing (LNCS 3968)*, K. P. Fishkin, B. Schiele, P. Nixon, and A. Quigley, Eds. Berlin, Germany: Springer, 2006, pp. 1–16.
- [23] L. Sun, D. Zhang, B. Li, B. Guo, and S. Li, "Activity recognition on an accelerometer embedded mobile phone with varying positions and orientations," in *Ubiquitous Intelligence and Computing (LNCS 6406)*, Z. Yu, R. Liscanau, G. Chen, D. Zhang, and X. Zhou, Eds. Berlin, Germany: Springer, 2010, pp. 548–562.
- [24] K. Förster, P. Brem, D. Roggen, and G. Tröster, "Evolving discriminative features robust to sensor displacement for activity recognition in body area sensor networks," in *Proc. 5th Int. Conf. Intell. Sensors Sensor Netw. Inf. Process. (ISSNIP)*, Melbourne, VIC, Australia, Dec. 2009, pp. 43–48.
- [25] S. Reddy, M. Yun, J. Burke, D. Estrin, M. Hansen, and M. Srivastava, "Using mobile phones to determine transportation modes," *ACM Trans. Sensor Netw.*, vol. 6, no. 2, Feb. 2010, Art. no. 13.
- [26] A. Anjum and M. U. Ilyas, "Activity recognition using smartphone sensors," in *Proc. 10th IEEE Consum. Commun. Netw. Conf. (CCNC)*, Las Vegas, NV, USA, Jan. 2013, pp. 914–919.
- [27] P. Siirtola and J. Röning, "Ready-to-use activity recognition for smartphones," in *Proc. IEEE Symp. Comput. Intell. Data Min. (CIDM)*, Singapore, Apr. 2013, pp. 59–64.
- [28] J. Doppler *et al.*, "Variability in foot-worn sensor placement for activity recognition," in *Proc. Int. Symp. Wearable Comput. (ISWC)*, Linz, Austria, Sep. 2009, pp. 143–144.
- [29] A. Henpraserttae, S. Thiemjarus, and S. Marukat, "Accurate activity recognition using a mobile phone regardless of device orientation and location," in *Proc. Int. Conf. Body Sensor Netw. (BSN)*, Dallas, TX, USA, May 2011, pp. 41–46.
- [30] S. Thiemjarus, A. Henpraserttae, and S. Marukat, "A study on instance-based learning with reduced training prototypes for device-context-independent activity recognition on a mobile phone," in *Proc. IEEE Int. Conf. Body Sensor Netw. (BSN)*, Cambridge, MA, USA, May 2013.
- [31] R. Chavarriaga, H. Bayati, and J. del R. Millán, "Unsupervised adaptation for acceleration-based activity recognition: Robustness to sensor displacement and rotation," *Pers. Ubiquitous Comput.*, vol. 17, no. 3, pp. 479–490, Mar. 2013.
- [32] K. Förster, D. Roggen, and G. Tröster, "Unsupervised classifier self-calibration through repeated context occurrences: Is there robustness against sensor displacement to gain?" in *Proc. Int. Symp. Wearable Comput.*, Linz, Austria, Sep. 2009, pp. 77–84.
- [33] K. Kunze, P. Lukowicz, H. Junker, and G. Tröster, "Where am I: Recognizing on-body positions of wearable sensors," in *Location- and Context-Awareness (LNCS 3479)*, T. Strang and C. Linnhoff-Popien, Eds. Berlin, Germany: Springer, 2005, pp. 264–275.
- [34] W. Xu, M. Zhang, A. A. Sawchuk, and M. Sarrafzadeh, "Robust human activity and sensor location corecognition via sparse signal representation," *IEEE Trans. Biomed. Eng.*, vol. 59, no. 11, pp. 3169–3176, Nov. 2012.
- [35] H. Lu, J. Yang, Z. Liu, N. D. Lane, T. Choudhury, and A. T. Campbell, "The Jigsaw continuous sensing engine for mobile phone applications," in *Proc. 8th ACM Conf. Embedded Netw. Sensor Syst.*, Zürich, Switzerland, Nov. 2010, pp. 71–84.
- [36] H. Martín, A. M. Bernardo, J. Iglesias, and J. R. Casar, "Activity logging using lightweight classification techniques in mobile devices," *Pers. Ubiquitous Comput.*, vol. 17, no. 4, pp. 675–695, Apr. 2013.
- [37] T. Szttyler and H. Stuckenschmidt, "On-body localization of wearable devices: An investigation of position-aware activity recognition," in *Proc. IEEE Int. Conf. Pervasive Comput. Commun. (PerCom)*, Sydney, NSW, Australia, Mar. 2016.
- [38] B. Barshan and A. Yurtman, "Investigating inter-subject and interactivity variations in activity recognition using wearable motion sensors," *Comput. J.*, vol. 59, no. 9, pp. 1345–1362, Sep. 2016.
- [39] O. Banos, M. A. Toth, M. Damas, H. Pomares, and I. Rojas, "Dealing with the effects of sensor displacement in wearable activity recognition," *Sensors*, vol. 14, no. 6, pp. 9995–10023, Jun. 2014.
- [40] Y. Zhong and Y. Deng, "Sensor orientation invariant mobile gait biometrics," in *Proc. IEEE Int. Joint Conf. Biometrics*, Clearwater, FL, USA, Sep./Oct. 2014.

- [41] J. Wu and R. Jafari, "Orientation independent activity/gesture recognition using wearable motion sensors," *IEEE Internet Things J.*, vol. 6, no. 2, pp. 1427–1437, Apr. 2019.
- [42] W. Lu, F. Fan, J. Chu, P. Jing, and Y. Su, "Wearable computing for Internet of Things: A discriminant approach for human activity recognition," *IEEE Internet Things J.*, vol. 6, no. 2, pp. 2749–2759, Apr. 2019.
- [43] J. L. Meriam, L. G. Kraige, and J. N. Bolton, *Engineering Mechanics: Dynamics*, 9th ed. New York, NY, USA: Wiley, Mar. 2018.
- [44] A. Yurtman and B. Barshan, "Novel non-iterative orientation estimation for wearable motion sensor units acquiring accelerometer, gyroscope, and magnetometer measurements," *IEEE Trans. Instrum. Meas.*, early access. [Online]. Available: <https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=8768414>, doi: [10.1109/TIM.2019.2930437](https://doi.org/10.1109/TIM.2019.2930437).
- [45] K. Altun and B. Barshan, *Daily and Sports Activities Dataset*, IEEE Data Port, Feb. 2019, doi: [10.21227/at1v-6f84](https://doi.org/10.21227/at1v-6f84).
- [46] K. Altun and B. Barshan, *Daily and Sports Activities Dataset*, UCI Mach. Learn. Repository, School Inf. Comput. Sci., Univ. California at Irvine, Irvine, CA, USA, 2013. [Online]. Available: <http://archive.ics.uci.edu/ml/datasets/Daily+and+Sports+Activities>
- [47] *MTI, MTx, and XM-B User Manual and Technical Documentation*, Xsens Technol. B.V., Enschede, The Netherlands, 2020. [Online]. Available: <http://www.xsens.com>
- [48] K. Altun, B. Barshan, and O. Tunçel, "Comparative study on classifying human activities with miniature inertial and magnetic sensors," *Pattern Recognit.*, vol. 43, no. 10, pp. 3605–3620, Oct. 2010.
- [49] A. R. Webb, *Statistical Pattern Recognition*, 2nd ed. West Sussex, U.K.: Wiley, 2002.
- [50] C.-W. Hsu, C.-C. Chang, and C.-J. Lin, "A practical guide to support vector classification," Dept. Comput. Sci., Nat. Taiwan Univ., Taipei, Taiwan, Rep., 2003.
- [51] K.-B. Duan and S. S. Keerthi, "Which is the best multiclass SVM method? An empirical study," in *Multiple Classifier Systems* (LNCS 3541), N. C. Oza, R. Polikar, J. Kittler, and F. Roli, Eds. Berlin, Germany: Springer, 2005, pp. 278–285.
- [52] C.-C. Chang and C.-J. Lin, "LIBSVM: A library for support vector machines," *ACM Trans. Intell. Syst. Technol.*, vol. 2, no. 3, Apr. 2011, Art. no. 27. [Online]. Available: <http://www.csie.ntu.edu.tw/cjlin/libsvm/>
- [53] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, 2nd ed. New York, NY, USA: Wiley, 2001.
- [54] S. Haykin, *Neural Networks: A Comprehensive Foundation*, 3rd ed. Upper Saddle River, NJ, USA: Prentice-Hall, May 2007.
- [55] I. H. Witten, E. Frank, M. A. Hall, and C. J. Pal, *Data Mining: Practical Machine Learning Tools and Techniques*, 4th ed. Cambridge, MA, USA: Elsevier, 2016.
- [56] Y. C. Pati, R. Rezaifar, and P. S. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in *Proc. 27th Asilomar Conf. Signals Syst. Comput.*, Pacific Grove, CA, USA, Nov. 1993, pp. 40–44.
- [57] L. Wang, L. Cheng, and G. Zhao, *Machine Learning for Human Motion Analysis: Theory and Practice*. Hershey, PA, USA: IGI Glob., 2010.
- [58] J. Diebel, "Representing attitude: Euler angles, unit quaternions, and rotation vectors," Dept. Aeronaut. Astronaut., Stanford Univ., Stanford, CA, USA, Rep., Oct. 2006. [Online]. Available: <http://www.swarthmore.edu/NatSci/mzucker1/papers/diebel2006attitude.pdf>



Billur Barshan received the B.S. degrees in electrical engineering and in physics from Boğaziçi University, Istanbul, Turkey, in 1986, and the M.S. and Ph.D. degrees in electrical engineering from Yale University, New Haven, CT, USA, in 1988 and 1991, respectively.

From 1991 to 1993, she was a Postdoctoral Researcher with the Robotics Research Group, University of Oxford, Oxford, U.K. In 1993, she joined the faculty of Bilkent University, Ankara, Turkey, where she is currently a Professor with the Department of Electrical and Electronics Engineering. Her current research interests include wearable sensing, wearable robots and mechanisms, intelligent sensing, motion capture and analysis, detection and classification of falls, machine learning, pattern classification, and multisensor data fusion.

Dr. Barshan was a recipient of the TÜBİTAK Incentive Award in 1998, the METU Mustafa Parlar Foundation Research Award in 1999, and two best paper awards. She served on the Management Committee of the COST-IC0903 Action MOVE from 2010 to 2013.



Aras Yurtman received the B.S., M.S., and Ph.D. degrees in electrical and electronics engineering from Bilkent University, Ankara, Turkey, in 2010, 2012, and 2019, respectively.

He was a Research and Teaching Assistant and an Administrative Assistant with Bilkent University. He has authored and coauthored articles published in the journals *IEEE TRANSACTIONS ON INSTRUMENTATION AND MEASUREMENT*, *Computer Methods and Programs in Biomedicine*, *Applied Artificial Intelligence*, *Computer Journal*, and *Sensors*. His current research interests include wearable sensing, human motion analysis, sensor signal processing, machine learning, and pattern recognition.