Recognizing activities of daily living regardless of wearable device orientation

Mr. Aras Yurtman
Bilkent University, Electrical and Elecronics Engineering,
yurtman@ee.bilkent.edu.tr

Prof.Dr. Billur Barshan
Bilkent University, Electrical and Electronics Engineering,
billur@ee.bilkent.edu.tr

Keywords: wearable sensing, motion sensors, sensor placement, sensor orientation, orientation invariance, accelerometer, gyroscope, magnetometer

Human activity recognition has become an active research area associated with context-aware systems. The most commonly used approach is to utilize wearable motion sensors whose cost, size, weight, and wireless communication capabilities have improved in the last few decades. It is often assumed that each wearable sensor is placed at a pre-determined position and a pre-determined orientation on the body. Placing the sensor units on the body always at the same orientation is not practical because sensors are enclosed in small rectangular packages whose orientations are difficult to determine at a first glance. Moreover, many activity recognition systems utilize mobile phones containing built-in motion sensors to acquire data and to transfer and/or process them using the same device. The problem of incorrect sensor orientation is more prevalent for mobile phones because they can be naturally carried in a pocket at different orientations. Placing the sensors at correct orientations is a more challenging task for disabled, injured, or elderly people whose activities are monitored in healthcare applications. Hence, allowing the sensors to be fixed to the body at an arbitrary orientation is a valuable flexibility in human activity recognition and other wearable sensing applications.

Existing studies that focus on orientation-invariant sensing follow one of these approaches: (1) Orientation is classified in the long term among a pre-determined set of orientations during walking or a restricted set of activities. (2) Activity recognition is performed without taking care of the orientation information, relying on the ability of the classifiers to handle intra-class variations, where only a limited number of sensor orientations are considered. (3) The orientation of the sensor is estimated, assuming certain statistical properties of the body movements. (4) The magnitude (Euclidean norm) of tri-axial sensor data is used in classification. In the first three approaches, orientation-invariance can be achieved only for specific cases with impractical restrictions, whereas in the fourth approach, a significant portion of the data is lost and the accuracy decreases considerably.

In this study, however, we develop a transformation that removes the orientation information from the tri-axial sensor signals without losing a considerable amount of information. The result of the transformation is a multi-dimensional time-domain signal. Hence, to achieve orientation-invariance in various applications, the transformation can be applied to the sensor data in the pre-processing stage without extensively modifying the existing algorithm.

Changing the orientation of a tri-axial sensor corresponds to transforming the sensor data by a rotational transformation that is constant over time, considering that the orientation of the sensor with respect to the body remains the same throughout the recording session. Thus, the transformation must be invariant under rotational coordinate transformations. By satisfying this property, we transform the tri-axial sensor data to 15-axial data by using geometric properties of the data vectors in 3-D space. The transformed data contain the magnitudes of the acquired data vectors and the differences between the magnitudes of the data vectors that are up to four time samples apart, which are invariant under rotational transformations. The angle between any two data vectors at different time samples is also invariant under rotation. Therefore, we include the angles between the data vectors that are up to five time samples apart. Lastly, we calculate the rotation axes between the data vectors at consecutive time samples and include the angles between the rotation axes that are up to five time samples apart. Thus, we have 15 transformed elements as a function of time, where the first five of them are the magnitudes of difference vectors, the second five are the angles between the data vectors, and the remaining five are the angles between the axes of rotation.

We use five publicly available datasets containing diverse types of human activities acquired by different sensor types and configurations. One of the datasets also contains a number of sports activities. The number of subjects ranges from four to 40, the number of activities changes between five and 19, and the number of wearable sensor units ranges from one to five. The units contain accelerometers in all the datasets, and additionally gyroscopes and magnetometers in some datasets, where all the sensors are tri-axial. The time duration of the data is around 10 hours in each dataset, and the sampling rate changes between 8 and 100 Hz.

We apply the proposed transformation to the sensor data in the pre-processing stage of a generic human activity recognition scheme to demonstrate orientation invariance. Then, we divide the data into 5- or 2.5-sec segments and extract statistical features from each segment to construct a feature vector. The dimension of the feature space changes between 26 and 5850 depending on the transformation and the dataset. When the dimensionality is greater than 30, we reduce it to 30 through principal component analysis. We use four state-of-the-art classifiers: Bayesian decision making (BDM), k-nearest-neighbor (k-NN), support vector machines (SVM), and artificial neural networks (ANN). We optimize the parameters of k-NN and SVM and select them suitably for ANN. We evaluate the accuracy by P-fold and leave-one-subject-out (L1O) cross-validation techniques.

We observe that when the standard activity recognition system is used with randomly oriented sensors, the accuracy drops by 21.21% compared to the reference case with fixed sensor orientations, averaged over all the datasets, classifiers, and cross-validation techniques. Thus, the standard activity recognition system is highly vulnerable to changes in the orientation of the sensor units. Using only the magnitude of the sensor data, which is the most commonly used method in the literature, decreases the accuracy by 13.50% compared to the reference case. When the proposed method is used with the first five, the first ten, or all of the 15 axes, the average reduction in the accuracy is 14.14%, 8.40%, and 8.50%, respectively. Hence, the proposed method with ten axes performs the best, which is noticeably better than the magnitude method, especially in L1O cross validation. The last five axes corresponding to the angles between axes of rotation do not improve the accuracy, hence can be omitted. With the proposed transformation, an accuracy close to the fixed sensor orientation case can be obtained while allowing the users to place the sensors on their body at any orientation. The only required change in the system is to apply the transformation in the pre-processing stage, hence, this method can be applied to a wide range of existing wearable systems without much effort.