


HANDY: A Benchmark Dataset for Context-Awareness via Wrist-Worn Motion Sensors

Koray Açııcı^{1,*}, Çağatay Berke Erdaş¹, Tunç Aşuroğlu¹  and Hasan Oğul^{1,2}

¹ Department of Computer Engineering, Başkent University, Bağlıca Kampüsü, Fatih Sultan Mahallesi Eskişehir Yolu 18 Km, Ankara 06790, Turkey; berdas@baskent.edu.tr (Ç.B.E.); tuncasuroglu@baskent.edu.tr (T.A.); hogul@baskent.edu.tr (H.O.)

² Faculty of Computer Science, Østfold University College, P.O. Box 700, 1757 Halden, Norway

* Correspondence: korayacici@baskent.edu.tr; Tel.: +90-312-246-6661

Received: 23 May 2018; Accepted: 22 June 2018; Published: 24 June 2018



Abstract: Being aware of a personal context is a promising task for various applications, such as biometry, human-computer interactions, telemonitoring, remote care, mobile marketing and security. The task can be formally defined as the classification of a person being considered into one of predefined labels, which may correspond to his/her identity, gender, physical properties, the activity that he/she performs or any other attribute related to the environment being involved. Here, we offer a solution to the problem with a set of multiple motion sensors worn on the wrist. We first provide an annotated and publicly accessible benchmark set for context-awareness through wrist-worn sensors, namely, accelerometers, magnetometers and gyroscopes. Second, we present an evaluation of recent computational methods for two relevant tasks: activity recognition and person identification from hand movements. Finally, we show that fusion of two motion sensors (i.e., accelerometers and magnetometers), leads to higher accuracy for both tasks, compared with the individual use of each sensor type.

Dataset: available as the supplementary file or at www.baskent.edu.tr/~hogul/handy.

Dataset License: CC-BY.

Keywords: activity recognition; person identification; sensor data analysis; dataset; context-awareness; wearable computing

1. Introduction

While the ubiquitous use of mobile devices has been moderately affecting our daily life, several business needs have emerged to compete with the amount and value of information streamed from these devices. Nowadays, the spread of smart phones and other wearable machines throws up new challenges to identify the contexts of people who are using these devices. Context-awareness is the broad term to refer to this issue, which involves several tasks, such as person identification, activity recognition, and environment detection [1]. The term is also used to indicate any system that is capable of using contexts in the literature. Apart from this general definition, a system is accepted as context-aware if it uses contexts to provide relevant information and services to the user, where the relevance may refer to any environmental condition or action related to the user's task [2]. Applications that have context-awareness should provide information and services e.g., automatic execution of a service and tagging of context for later retrieval on account of users.

The solutions to context-awareness vary upon several design issues that should be addressed, such as which sensing devices to be used, where to place these devices and which task to be considered.

In this study, we target on wrist-worn motion sensors, which are usually embedded in smart watch phones or digital wristlets to detect the user context. Our argument is that many of the daily activities are performed by hand, and therefore, the motion signals obtained from hand would be an effective indicator for the context of the person who performs those activities. Furthermore, the wrist of a person has the highest wearability compared with the other parts of the body. Since many people already get used to wearing watches, bracelets or similar equipment on their wrists, having a small additional sensor kit will have the minimal discomfort. In fact, there will not be any need to have an additional kit, since this trendy equipment, such as watch phones, has already embedded these sensors by default.

We defined two sub-tasks related to context-awareness, which could be managed with wrist-worn motion sensors. The first task was the identification of the person, who wore that device based on his/her previous recordings. The second task was the recognition of the current hand activity that the user performed. To the best of our knowledge, there is no previous report that studied personal context-awareness from this point of view. Therefore, we first built an annotated benchmark dataset for the present work and made it publicly available for future studies (www.baskent.edu.tr/~hogul/handy). We introduced solutions to both tasks based on machine learning methods. To this end, we considered each task as a distinct classification problem and offered feature representation schemes for input sensor data to be fed into a variety of supervised learning models. We finally showed that simultaneous use of accelerometer and magnetometer sensors can improve the prediction accuracy for both tasks studied.

2. Related Work

A common approach is to exploit audiovisual signals acquired by microphones and cameras that are located either on the body of individuals or at a convenient place around them for the task of recognizing person [3–8]. Although audiovisual-based solutions are often most accurate, they do not have widespread use since they require additional installation efforts. Furthermore, the processing time of these signals is often too costly. Another approach is to use device sensors that can detect changes of the environmental conditions using the data generated by home equipment and devices [9]. For example, microphone-based small sensors were attached to water pipes to monitor water usage, thereby to recognize the activity [10]. Electrical noise in residential power lines was monitored in another study to detect the use of electrical devices, which in turn could give idea about the environmental context about the residence [11]. Similarly, gas activity was sensed to monitor home infrastructure [12]. While these attempts have been successful to some degree, they have practical limitations. The same conditions, on which these methods have been developed, have to be realized in the environment where the methods will be applied. This condition is not easy to satisfy in many cases.

The motion sensors, such as accelerometers, which are attached with wearable devices, are promising means of having personal context-awareness. Owing to ubiquitous availability of these devices, the motion-sensor data analysis received a considerable attention of researchers studying context-awareness [13–17]. Since it is not often so feasible in daily life to wear many sensors on different parts of the body, the efforts on practical implementations of accelerometer-based context-awareness have mainly focused on recognition of human activities through mobile communications devices, which might be either carried on the body or attached with the clothes [18–22]. Whilst the majority of accelerometer analysis efforts have been focused on activity recognition, to the best of our knowledge, there is only one study that considers the person identification task via accelerometers for context-awareness [23]. They used a mobile phone attached on the chest to detect both the person and the activity.

The only study with wrist-worn devices proposed to use magnetic sensors to recognize other electric devices on the same environment for activity recognition [24]. Here, we offer to use motion sensors instead, which has the advantage that they do not need any interaction with the environment

or other devices around, but rather use only the signals generated from the action that the wearing individual is performing.

Several similar databases are publicly available for activity recognition tasks only. The dataset provided by Casale et al. [23] was collected from 15 people and labeled with 7 activities. In another dataset, provided by Bruno et al., [25] the data for 14 activities of daily living, which were collected from 16 volunteers, exist. Ugulino et al. [26] created a dataset with 5 classes collected on 8-h activities of 4 healthy subjects. Anguita et al. [27] created a human activity recognition database built from the recordings of 30 subjects, performing activities of daily living while carrying a waist-mounted smartphone with embedded inertial sensors. The PAMAP2 Physical Activity Monitoring dataset, which was collected by Reiss and Stricker [28], contains data of 18 different physical activities, performed by 9 subjects wearing 3 inertial measurement units and a heart rate monitor. In this study, we offer two novelties with a handy dataset: (1) the identities of participants are also included to make it useful for benchmarking person identifications tasks, and (2) extended numbers of hand activity labels are included.

3. Methods

3.1. Data Collection

The data were collected with a wearable sensor kit comprising an accelerometer, a gyroscope and a magnetometer that could synchronize in time. The streaming data from a wearable device were acquired through a Bluetooth connection with a personal computer. The sampling rate for each sensor was 52 Hz.

Since we defined the context-awareness over two main tasks, i.e., person identification and hand activity recognition, we first determined the persons and activity labels. Thirty people participated in the experiments, where they were asked to perform given activities using the equipment provided in the same environmental conditions. They all wore the sensor kit on their wrists with the same positioning during the activity. Each participant was asked to perform following activities with given durations:

1. chopping: this activity involved chopping two 30-cm-thickness rods made by the dough using a small knife (Figure 1a). The activity took approximately 60 s.
2. cleaning table: this activity involved cleaning a table with a size of 150 cm × 110 cm using a small rag. The participant was asked to wipe the table, starting from the upper left corner and requested to cover the entire area twice without keeping his/her hand off (Figure 1b). The activity took about 30 s.
3. cleaning window: this activity involved cleaning an indoor window with a size of 140 cm × 55 cm using a small rag. The participant was asked to wipe the window, starting from the upper left corner and requested to cover the entire area twice without keeping his/her hand off (Figure 1c). The activity took about 30 s.
4. drinking water: this activity involved drinking water from a 33 cc porcelain cup (Figure 1d). The participant was asked to consume all content at a time before putting the cup back. The activity took about 30 s.
5. eating soup: this activity involved eating soup from a 33 cc porcelain cup using a spoon. We put water instead of real soup into the cup to simulate the activity, since the liquid content in the cup did not affect the motion signal (Figure 1e). The activity took about 50 s.
6. kneading dough: this activity involved kneading a piece of dough (about 30 g). The participant was not asked to perform a specific activity, but required to play freely with the dough at least for 60 s (Figure 1f).
7. using a tablet computer: this activity involved playing a game (blowing up the balloons that appear upward on the screen), using an application on a tablet computer (Figure 1g).

The participant was asked to finish the first level of the game with his/her best effort. The activity took exactly 60 s.

8. using a computer mouse: this activity involved moving 20 computer files in a folder on the left of the screen into another folder on the right of the screen by drag-and-drop using a mouse (Figure 1h). The files were located in 2×10 rectangle grid in the source folder, and they were expected to be placed as the same layout in the destination folder. The activity took about 60 s.
9. writing with a pen: this activity involved writing a given book text (359 characters) into an A4 paper using a standard pen (Figure 1i). The activity might take between 100 and 120 s depending on the participant's speed.
10. writing with a keyboard: this activity involved typing a given book text (359 characters) using a keyboard with the QWERTY layout (Figure 1j). The activity might take between 80 and 100 s depending on the participant's speed.

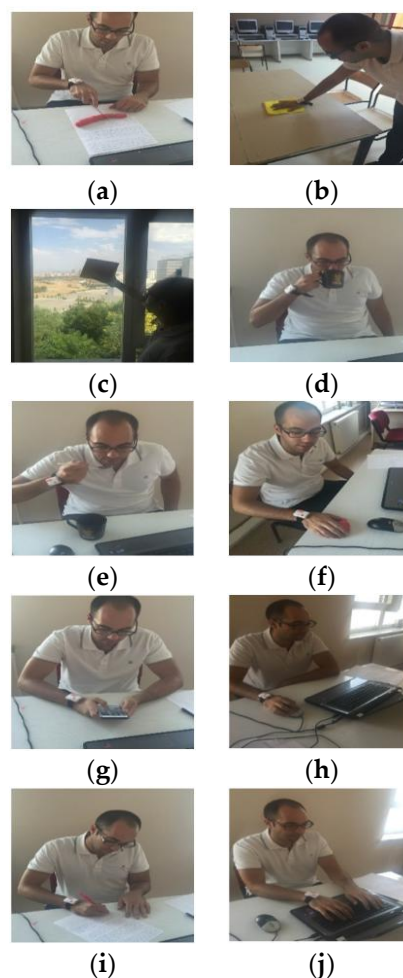


Figure 1. Images from hand activity recordings for (a) chopping; (b) cleaning table; (c) cleaning window; (d) drinking water; (e) eating soup; (f) kneading dough; (g) using a tablet computer; (h) using a computer mouse; (i) writing with a pen; and (j) writing with a keyboard.

Each entry in the dataset has been labeled by following attributes: (1) participant name (a fake name to identify the person); (2) participant age; (3) participant gender; (4) hand activity; (5) accelerometer signals on the x, y, and z axes; (6) gyroscope signals on the x, y, and z axes; and (7) magnetometer signals on the x, y, and z axes.

3.2. Framework for Context-Awareness

We adapted a conventional pattern classification framework to our tasks for context-awareness. We employed a typical discriminative model for supervised classification, which consisted of training and a prediction phase (Figure 2). Each phase involved feature extraction from input motion signals. The training phase built a classification model that mapped the extracted features into known context annotations, i.e., person identity and activity in our case. The prediction phase submitted the extracted features from sensed signals into a learned model and reported its output as the estimation of context.

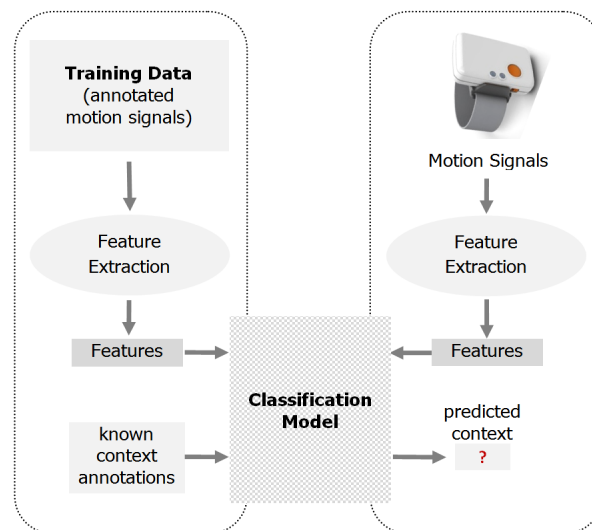


Figure 2. An overview of a classification framework for context-awareness.

3.3. Feature Extraction

Feeding motion signals directly to a classifier is not convenient for three reasons. First, a large majority of machine learning classifiers require that the input must be of equal length, which could not be satisfied for the motion signals in our case. Second, even if the signals are of equal length, they do not necessarily exist at the same time scale. Third, temporal observations may not reflect the behavior of motion in terms of the categories that the classifier is desired to assign. Therefore, a common practice in classification of such time-series signals is to extract a fixed number of representative features from raw signals and represent the input by a vector comprising these feature values instead of the direct measurements of the signal.

In this study, we extracted the several time-domain and frequency-domain features from any of the motion signals with three dimensions. Sixteen time-domain features were extracted: the mean, the median, the minimum value and its index, the maximum value and its index, the range, the root mean square (RMS), the interquartile range (IQR), the mean absolute deviation (MAD), the skewness, the kurtosis, the entropy, the energy, the power and the harmonic mean. In order to extract frequency-domain features, the fast Fourier transform (FFT) was applied to the signal. Then, seven frequency-domain features were extracted: the mean, the maximum value, the minimum value, the normalized value, the energy, the phase and the band power.

3.4. Classification

Here, we considered two classification problems, where a time-series motion signal was automatically assigned to one of pre-defined labels. In the person identification task, each class label corresponded to one of persons in our collection, while each of ten activities described referred to one class label in the hand activity recognition task. Having a fixed-length representation for any

input signal, the next step was to find a model that best separated the signals labeled with different classes. In this study, we compared several supervised classification methods, which have been shown to be successful in various applications for our particular tasks. Considering their common use and reported success in several domains [29,30], we chose four among these algorithms and evaluated their classification performance for our tasks: k-Nearest Neighbors (kNN), AdaBoost, Decision Tree (DT) and Random Forest (RF).

kNN is an instance-based classification algorithm. When the class of the query sample is determined, the algorithm decides by searching its nearest k neighbors according to predefined distance metrics, such as Euclidean, Manhattan and Minkowski distances.

When a single classifier does not provide satisfactory accuracy, many weak classifiers are constructed and combined into a powerful classifier instead of using one. To accomplish this task, Adaboost is one of the most widely used algorithms. The main idea of the Adaboost algorithm is to define a certain probability distribution of learning samples (a training set) based on previous results at each iteration. Firstly, a weight is assigned to each sample. If a sample is misclassified, its weight will be increased. In contrast, if a sample is correctly classified, its weight will be decreased, so that the new classifiers are focused on tough samples. At the end of the classification process, an ensemble classifier gives the final decision by using a majority voting.

The DT algorithm constructs decision trees from a set of training data using the concept of information entropy. The basic idea is to divide the data into a certain range based on, for example, the attribute values in the training set. The discrimination criterion is the normalized information gain. The node with the highest normalized knowledge gain is selected as the root node for decision-making. The DT algorithm then repeats on smaller sub-lists.

The RF has an ensemble learning strategy over many decision trees. Each tree is constructed independently by using of a bootstrap sample and randomly selected features. Bootstrap samples and a subset of features are selected from the input data. Finally, several decision trees are constructed. The RF decides the final prediction based on a majority voting rule which is formed by each decision tree.

An empirical comparison of these methods can be found in Caruana and Niculescu-Mizil [31]. For their detailed descriptions, the readers are referred to Bishop [32].

4. Results

4.1. Experimental Setup

To build an experimental evaluation environment for two classification tasks, we created a number of equal-length samples from sensor reads. Each sample had a length of 1 s (52 contiguous sensor reads) with a 50% overlap. This application included 992,976 samples for activity recognition and person identification tasks. Numbers of samples within each activity label are given in Table 1.

Table 1. Dataset content for activity recognition.

Activity Label	Number of Samples
Chopping	72,333
Cleaning table	45,117
Cleaning window	43,767
Drinking water	112,212
Eating soup	81,540
Kneading dough	96,174
Using a tablet computer	97,794
Using a computer mouse	106,299
Writing with a pen	169,152
Writing with a keyboard	168,588

A 10-fold cross-validation approach was used for evaluation. In this setup, the dataset was first divided into 10 balanced partitions, and at each fold, 9/10 of these partitions were used to train the model that would be used to predict the samples in remaining partitions. All samples were guaranteed to go through a prediction stage after repetition of the same experiments of 10 times with a different training set in each. The performance of systems was measured using overall accuracy of predictions, which was given by Equation (1):

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{FP} + \text{TN} + \text{FN}} \times 100 \quad (1)$$

where TP, TN, FP and FN refer to numbers of true positives, true negatives, false positives and false negatives, respectively.

4.2. Empirical Results

For activity recognition, we asked the system to predict the activity label of a given sample, regardless of the person who performed this activity. Table 2 shows the 10-fold cross-validation results of activity recognition with each motion sensor when different classification methods were applied. As shown, the highest accuracy was achieved when magnetometer data were fed into a RF classifier. The accuracy achieved by using an accelerometer sensor was competitive with that by magnetometer. The gyroscope could not perform well for the hand activity recognition task. The RF method can outperform consistently the other classifiers for all sensor types.

Table 2. Results of hand activity recognition (accuracy in %).

Classifier	Sensor		
	Accelerometer	Gyroscope	Magnetometer
kNN	44.7	45.5	60.6
AdaBoost	21.2	22.9	23.5
Decision Tree	54.1	43.8	60.5
Random Forest	72.2	61.4	98.9

For the person identification task, we first determined an activity and asked the system to predict the person who performed this activity by allowing the system to know its activity label. Considering their discriminative abilities, we selected three activities to realize person identification tests: writing with a pen, writing with a keyboard and using a tablet computer.

Tables 3–5 show the results of person identification experiments with different classifiers for each selected activity, respectively. Similar to activity recognition tests, a consistent result was that the RF method performed better than all other classifiers for the person identification task.

Table 3. Results of person identification from the activity of using a tablet computer (accuracy in %).

Classifier	Sensor		
	Accelerometer	Gyroscope	Magnetometer
kNN	26.4	11.1	33.7
AdaBoost	6.3	5.8	7.5
Decision Tree	59.9	14.8	57.6
Random Forest	78.4	29.9	78.2

Table 4. Results of person identification from the activity of writing with a pen (accuracy in %).


Classifier	Sensor		
	Accelerometer	Gyroscope	Magnetometer
kNN	22.1	22.5	35.7
AdaBoost	8.2	10.0	9.9
Decision Tree	82.7	37.5	75.4
Random Forest	82.9	58.4	88.1

Table 5. Results of person identification from the activity of writing with a keyboard (accuracy in %).

Classifier	Sensor		
	Accelerometer	Gyroscope	Magnetometer
kNN	20.3	12.7	30.3
AdaBoost	8.9	8.2	9.9
Decision Tree	55.1	32.2	68.3
Random Forest	67.8	32.2	80.4

Accelerometer and magnetometer signals were competitive in person identification from the activity of using a tablet computer. For other activities, higher accuracy could be achieved with magnetometer. The gyroscope was not helpful in identifying person as well.

An important step in the classification framework given in Figure 2 is the feature extraction from input motion signals. Which features to use is usually very crucial in prediction performance. Figure 3 compares the time-domain features, frequency-domain features and their fusion in terms of their contribution in the accuracy for several tasks that we considered here. In many of the classification tasks that we considered, time-domain features led to higher accuracy. As shown, their input-level fusion usually improved the prediction accuracy.

**Figure 3.** (a) Activity classification performance; (b) person identification performance from writing with a pen; (c) person identification performance from writing with a keyboard; and (d) person identification performance from using a tablet computer.

Fusion of streaming data is usually a challenging issue to improve the performance of the intelligent systems that utilize multiple sensors. Owing to their promising individual performances, we considered integrating accelerometer and magnetometer signals to increase the prediction accuracy of the tasks being studied. The fusion was performed at the input level; the features extracted from two distinct sensor signals were simply concatenated to make a final vector to be fed into the classifier used. Since the same features were extracted from normalized data of both sensor signals, no further action was needed to perform this data fusion. The experimental results revealed that this integration could significantly improve the hand activity classification accuracy by 4.5% and person identification accuracy by 6.6–13.0% upon the individual use of these sensors when the RF classifier was used (Table 6).

Table 6. Results of multisensor data fusion for both tasks.

Task	Sensor		
	Acc.	Mag.	Both
Hand activity recognition	72.2	78.9	84.4
Person identification from hand writing	78.7	88.1	94.7
Person identification from keyboard typing	67.8	80.4	88.9
Person identification from tablet using	78.4	78.2	91.4

5. Conclusions

Context-awareness is a vital issue for ubiquitous applications in several distinct domains. This paper approached the problem from a perspective of predicting contexts with a wrist-worn sensor. The contribution of the paper is threefold. First, a new modality of context-awareness was defined over multiple motions sensors placed on a single wrist-worn device. This wearable device mitigated a smart watch that potentially embedded similar motion sensors. In view of the increasing use of smart watch phones or specialized wearable sensors, we anticipate that this approach will find applications in several domains, such as personalized advertisement generation, telemonitoring children or dementia patients, and connected health. Second, a dataset was built using systematically collected human motion data with multiple labeling to define the context. The dataset is made publicly available for future studies. To the best of our knowledge, this is the first dataset released for context-awareness through multiple sensors worn on the wrist. Furthermore, a large set of classification models have been employed to create a baseline prediction performance to serve a benchmarking resource for future methods. Third, an evaluation of different motion sensors is reported for context-awareness. The results determine that the accelerometer and magnetometer are very useful in predicting both the person identity and the action performed, while the gyroscope is usually less effective. Another important finding is that the fusion of accelerometer and magnetometer signals can improve the prediction accuracy for both tasks.

The proposed model for context-awareness may find several applications in research and industry. A potential use is biometric identification. While it cannot offer a reliable identification method, such as fingerprints, the model can help mobile applications recognize their users. This may be extended to the customization of services based on the user identity and his/her activity. The public dataset is released as an attachment to the paper that can be used for benchmarking new methods of activity recognition and also generic algorithms developed for multi-modal and multivariate time-series data analysis.

A future direction is to analyze the effect of features extracted from motion signals on prediction accuracy. While we use a number of effective features recommended in similar studies, incorporating more descriptive features with domain-specific knowledge is an important challenge. On a basis of consideration into recent trends in machine learning field, feeding a deep learning network with raw sensor signals is another potential direction. Furthermore, we are keen to extend our dataset with additional participants and new activity types that will diversify the context.

Supplementary Materials: Dataset available online at <http://www.mdpi.com/2306-5729/3/3/24/s1>.

Author Contributions: H.O. contributed to the data collection and analysis and managed the participants during the experiments. Ç.B.E. contributed to the data collection and analysis. K.A. and T.A. participated in implementation, development and evaluation of the research. All authors significantly contributed to the writing and review of the paper.

Funding: This research received no external funding.

Acknowledgments: This study was supported by the Scientific and Technological Research Council of Turkey (TUBITAK) under the Project 115E451.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Yürür, Ö.; Liu, C.H.; Moreno, W. A survey of context-aware middleware designs for human activity recognition. *IEEE Commun. Mag.* **2014**, *52*, 24–31. [[CrossRef](#)]
2. Dey, A.K. Understanding and using context. *Pers. Ubiquitous Comput.* **2001**, *5*, 4–7. [[CrossRef](#)]
3. Aggarwal, J.K.; Ryoo, M.S. Human activity analysis: A review. *ACM Comput. Surv. (CSUR)* **2011**, *43*, 16. [[CrossRef](#)]
4. Clarkson, B.; Sawhney, N.; Pentland, A. Auditory context awareness via wearable computing. In Proceedings of the 1998 Workshop on Perceptual User Interfaces, San Francisco, CA, USA, 5–6 November 1998.
5. Ke, S.R.; Thuc, H.L.U.; Lee, Y.J.; Hwang, J.N.; Yoo, J.H.; Choi, K.H. A review on video-based human activity recognition. *Computers* **2013**, *2*, 88–131. [[CrossRef](#)]
6. Moeslund, T.B.; Hilton, A.; Krüger, V. A survey of advances in vision-based human motion capture and analysis. *Comput. Vis. Image Underst.* **2006**, *104*, 90–126. [[CrossRef](#)]
7. Robertson, N.; Reid, I. A general method for human activity recognition in video. *Comput. Vis. Image Underst.* **2006**, *104*, 232–248. [[CrossRef](#)]
8. Vrigkas, M.; Nikou, C.; Kakadiaris, I.A. A review of human activity recognition methods. *Front. Robot. AI* **2015**, *2*, 28. [[CrossRef](#)]
9. Philipose, M.; Fishkin, K.; Perkowitz, M. Inferring activities from interactions with objects. *IEEE Pervasive Comput.* **2004**, *3*, 50–57. [[CrossRef](#)]
10. Fogarty, J.; Au, C.; Hudson, S.E. Sensing from the basement: A feasibility study of unobtrusive and low-cost home activity recognition. In Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology, Montreux, Switzerland, 15–18 October 2006; pp. 91–100.
11. Patel, S.N.; Robertson, T.; Kientz, J.A.; Reynolds, M.S.; Abowd, G. At the flick of a switch: Detecting and classifying unique electrical events on the residential power line. In Proceedings of the International Conference on Ubiquitous Computing, Innsbruck, Austria, 16–19 September 2007; pp. 271–288.
12. Cohn, G.; Gupta, S.; Froehlich, J.; Larson, E.; Patel, S.N. GasSense: Appliance-level, single-point sensing of gas activity in the home. In Proceedings of the International Conference on Pervasive Computing, Newcastle, UK, 17–20 May 2010; pp. 265–282.
13. Casale, P.; Pujol, O.; Radeva, P. Human activity recognition from accelerometer data using a wearable device. In Proceedings of the Iberian Conference on Pattern Recognition and Image Analysis, Las Palmas de Gran Canaria, Spain, 8–10 June 2011; pp. 289–296.
14. Chen, L.; Hoey, J.; Nugent, C.D.; Cook, D.J.; Yu, Z. Sensor-based activity recognition. *IEEE Trans. Syst. Man Cybern. Part C Appl. Rev.* **2012**, *42*, 790–808. [[CrossRef](#)]
15. Preece, S.J.; Goulermas, J.Y.; Kenney, L.P.; Howard, D. A comparison of feature extraction methods for the classification of dynamic activities from accelerometer data. *IEEE Trans. Biomed. Eng.* **2009**, *56*, 871–879. [[CrossRef](#)] [[PubMed](#)]
16. Ravi, N.; Dandekar, N.; Mysore, P.; Littman, M.L. Activity Recognition from Accelerometer Data. In Proceedings of the 17th Conference on Innovative Applications of Artificial Intelligence, Pittsburgh, PA, USA, 9–13 July 2005; pp. 1541–1546.
17. Yürür, Ö.; Liu, C.H.; Sheng, Z.; Leung, V.C.M.; Moreno, W.; Leung, K.K. Context-awareness for mobile sensing: A survey and future directions. *IEEE Commun. Surv. Tutor.* **2016**, *18*, 68–93. [[CrossRef](#)]
18. Kwapisz, J.R.; Weiss, G.M.; Moore, S.A. Activity recognition using cell phone accelerometers. *ACM SigKDD Explor. Newsl.* **2010**, *12*, 74–82. [[CrossRef](#)]

19. Mannini, A.; Sabatini, A.M. Machine learning methods for classifying human physical activity from on-body accelerometers. *Sensors* **2010**, *10*, 1154–1175. [[CrossRef](#)] [[PubMed](#)]
20. Mitchell, E.; Monaghan, D.; O'Connor, N.E. Classification of sporting activities using smartphone accelerometers. *Sensors* **2013**, *13*, 5317–5337. [[CrossRef](#)] [[PubMed](#)]
21. Shoaib, M.; Scholten, H.; Havinga, P.J.M. Towards physical activity recognition using smartphone sensors. In Proceedings of the IEEE 10th International Conference on Ubiquitous Intelligence and Computing and 10th International Conference on Autonomic and Trusted Computing, Vietri sul Mare, Italy, 18–21 December 2013; pp. 80–87.
22. Zhong, M.; Wen, J.; Hu, P.; Indulska, J. Advancing Android activity recognition service with Markov smoother: Practical solutions. *Pervasive Mob. Comput.* **2017**, *38*, 60–76. [[CrossRef](#)]
23. Casale, P.; Pujol, O.; Radeva, P. Personalization and user verification in wearable systems using biometric walking patterns. *Pers. Ubiquitous Comput.* **2012**, *16*, 563–580. [[CrossRef](#)]
24. Maekawa, T.; Kishino, Y.; Sakurai, Y.; Suyama, T. Activity recognition with hand-worn magnetic sensors. *Pers. Ubiquitous Comput.* **2013**, *17*, 1085–1094. [[CrossRef](#)]
25. Bruno, B.; Mastrogiovanni, F.; Sgorbissa, A.; Vernazza, T.; Zaccaria, R. Analysis of human behavior recognition algorithms based on acceleration data. In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Karlsruhe, Germany, 6–10 May 2013; pp. 1602–1607.
26. Ugulino, W.; Cardador, D.; Vega, K.; Velloso, E.; Milidiu, R.; Fuks, H. Wearable computing: Accelerometers' data classification of body postures and movements. In Proceedings of the 21st Brazilian Symposium on Artificial Intelligence, Curitiba, Brazil, 20–25 October 2012; pp. 52–61.
27. Anguita, D.; Ghio, A.; Oneto, L.; Parra, X.; Reyes-Ortiz, J.L. A public domain dataset for human activity recognition using smartphones. In Proceedings of the 21st European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN), Bruges, Belgium, 24–26 April 2013.
28. Reiss, A.; Stricker, D. Introducing a new benchmarked dataset for activity monitoring. In Proceedings of the 16th IEEE International Symposium on Wearable Computers (ISWC), Newcastle, UK, 18–22 June 2012.
29. Erdaş, Ç.B.; Atasoy, I.; Açıcı, K.; Oğul, H. Integrating features for accelerometer-based activity recognition. *Procedia Comput. Sci.* **2016**, *98*, 522–527. [[CrossRef](#)]
30. Wong, T.T. Performance evaluation of classification algorithms by k-fold and leave-one-out cross validation. *Pattern Recognit.* **2015**, *48*, 2839–2846. [[CrossRef](#)]
31. Caruana, R.; Niculescu-Mizil, A. An empirical comparison of supervised learning algorithms. In Proceedings of the 23rd International Conference on Machine Learning, Pittsburgh, PA, USA, 25–29 June 2006; pp. 161–168.
32. Bishop, C.M. *Pattern Recognition and Machine Learning*, 1st ed.; Springer: New York, NY, USA, 2006; ISBN 978-1-4939-3843-8.



© 2018 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).