Intelligent Environments 2019
A. Muñoz et al. (Eds.)

doi:10.3233/AISE190067

A Framework for Long-Term Data Collection to Support Automatic Human Activity Recognition

Anna FERRARI [a], Daniela MICUCCI, [a,1], Marco MOBILIO [a] and Paolo NAPOLETANO [a]

[a] *University of Milano - Bicocca*

Abstract.
Human activity recognition (HAR) is a very active field of research and many techniques have been developed in recent years. Deep learning techniques applied to 3D and 4D signals such as images and videos have proven to be effective, achieving significant classification accuracy. Recently, these techniques are also being used for 1D signals and are exploited to recognize human Activities of Daily Living (ADLs). They process inertial signals such as those obtained from accelerometers and gyroscopes. However, to compute accurate and reliable deep learning models, a lot of sample data is required. Moreover, the creation of a dataset to be used with deep learning techniques is an onerous process that requires the involvement of a significant number of possibly heterogeneous subjects. The publicly available datasets are few and, with rare exceptions, contain few subjects. Furthermore, datasets are heterogeneous and therefore not directly usable all together. The goal of our work is the definition of a platform to support long-term data collection to be used in training HAR algorithms. The platform, termed Continuous Learning Platform (CLP), aims to integrate datasets of inertial signals in order to make available to the scientific community a large dataset of homogeneous signals and, when possible, enrich it with context information (e.g., characteristics of the subject, device position, and so on). Moreover, the platform has been designed to provide additional services such as the deployment of activity recognition models and online signal labelling services. The architecture has been defined and some of the main components have been developed in order to verify the soundness of the approach.

Keywords. dataset, deep learning, inertial data, integration platform, ADL, activity, machine learning

1. Introduction

Human Activity Recognition (HAR) is an active research field aimed at experimenting with new methods and techniques for the automatic recognition of Activities of Daily Living (ADLs) and, in some cases, also of falls [27,23,21]. Most of the proposed methods and techniques exploit sensors embedded in smartphones, smartwatches, fitness trackers, and ad-hoc wearable devices. The classification of sensor data with respect to the actions

[1]Corresponding Author: Daniela Micucci, University of Milano - Bicocca, Viale Sarca 336, 20126 Milan, Italy; E-mail: daniela.micucci@unimib.it.

performed by humans represents the main challenge of HAR. Indeed, in an ideal world, different activities present different signals so that the space of the signals is perfectly separated and easy to classify. In real world such a condition is not respected and signals often overlap each other.

The separation of data is a difficult task due to different aspects of the classification procedure, from the data acquisition to the definition of the activity. For instance, the *position* of the device influences the results of the performance [9]. Other aspects are related to the *intraclass variability* and the *interclass similarity* [8]. The former means that different people perform the same activity in different ways so a bijective association between signal and activity performed does not exist; the latter means that classes that are fundamentally different show very similar characteristics in the sensor data.

Preliminary applications of HAR exploited supervised machine learning techniques. These techniques present several challenges [18,8]. Firstly, the performances achieved in laboratory are difficult to be transferred to a real context without losing quality in the classification [3]. Secondly, it was noticed a strong dependency between the selection of the features and the performance of the classification, which corresponds to an inability of the algorithms to extract and organize discriminative information from the data [6]. Moreover, issues above mentioned, such as position, intraclass variability, and interclass similarity still strongly influence the classification performance [21,17,16].

In recent years, deep learning has been successfully applied in image analysis, natural language processing, sound recognition, and more recently it has been exploited also for 1D signals [19,28].

The widespread use of deep learning techniques is justified by their capability to overcome most of the presented issues, thanks to their properties of *local dependency* and *scale invariance* [36]. Furthermore, deep learning methodologies permit automated discovery of abstraction which overcomes the features extraction issue [6].

While deep learning techniques are powerful and achieve high performance, they rely on very complex models that strictly depend on the estimation of a large number of parameters, which requires a considerable amount of available data [7].

In recent years, some researchers have published their own dataset related to HAR. However, these datasets are heterogeneous and their standardization in a single unified dataset requires considerable effort. For example, signals are expressed in different units of measure, they may include gravity or not, and signals have different acquisition frequencies. Furthermore, labels are not aligned with a common dictionary and sometimes have different meanings among different datasets ('sitting' may refer to the state of being seated in a chair or the transition from standing to sitting). Thus, datasets are heterogeneous and cannot be used together without a significant effort to harmonize them.

The availability of a dataset containing a large number of samples, also obtained thanks to the integration of existing datasets, is a well known issue both in the field of ADLs recognition from inertial sensors and in other domains, such as that related to image processing [25]. Indeed, merging labels from different databases that resulted equals at semantic or syntactic level, may likely result in an inconsistent set of data, that does not improve the training of a classifier [12].

In the context of ADLs recognition from inertial signals, Bartlett *et al.* propose labels aggregation at semantic level [5]. Recently, Siirtola *et al.* propose a Matlab tool called Open HAR [31] that aggregates labels at a syntactic level but does not consider the semantics of the original signals. Obinikpo *et al.* propose a system for big data-d-health integration [24]. They split the integration in different layers, data acquisition, data processing, data analytics, and application. Nevertheless, the proposal is too general with respect to how homogeneize different sources of data, since its major concern is related to how handling missing values while integrating databases.

Since acquiring labeled time series is a costly procedure in terms of resource, time, and people involved, we think that the integration of existing datasets is the right direction despite the strong heterogeneity of the data.

In this article we propose a platform that semi-automatically integrates heterogeneous data and provides them in a homogenous form. The main contribution of this article is the definition of a new platform that:

- harmonizes heterogeneous data from inertial sensors, being them already existing datasets or coming from online acquisitions;
- distributes sets of labeled signals according to specific requests, such as, data related to a specific activity, data from subjects with specific physical characteristics, or, more generally, data from subjects that performed a set of specified activities;
- distributes activity recognition models;
- provides an online service by associating labels to series of signals in real-time.

The article is organized as follows: Section 2 provides an overview of our platform; Section 3 specifies how heterogeneous datasets can be acquired in order to be integrated in a unified dataset; Section 4 describes how the heterogeneous datasets can be unified so that they are organized in a uniform format; Section 5 provides insights about how the *Unified Dataset* can be exploited; finally, Section 6 sketches future directions.

2. *Continuous Learning Platform*

Human Activity Recognition (HAR) is a very active research field. Many techniques have been proposed, most of them based on the analysis of inertial signals from sensors embedded in smartphones, smartwatches, fitness trackers, and many others ad-hoc wearable devices. See the paper by Cornacchia *et al.* for a survey on activity detection using wearable devices [10].

The main aim of the *Continuous Learning Platform* (*CLP* in the sequel) is to make available (i) a large amount of labeled inertial signals related of ADLs and falls; (ii) a catalogue of downloadable activity recognition models, and (iii) a service that, given a set of raw data, identifies the corresponding ADL.

Labeled inertial signals can be used by researches to both validate new ADLs recognition techniques. Activity recognition models can be integrated into existing domain dependent applications that require ADLs recognition in order to provide application functionalities (e.g., estimation of energy expenditure [30,15], monitoring the development of the Parkinson's disease [29], and early detection of dementia [33]). Finally, if an application that requires to identify the ADLs performed by the user does not include an activity recognition model, it can rely on the service the platform provides each time a new series of data is acquired.

To fulfil the aim, *CLP collects* inertial signals from existing datasets or applications, *manages* the collected inertial signals, and *distributes* uniformed labeled inertial signals, trained activity recognition models, and labels assigned to series of inertial signals. Those functionalities are respectively reified by the *Data Collection*, *Data Management*, and *Data Distribution* components. Figure 1 sketches the overall architecture and its interaction with the external actors.

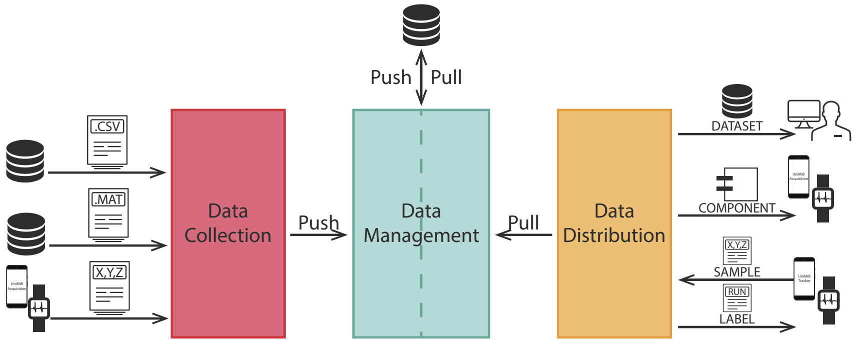

Figure 1. Overview of the platform.

The following sections provide detailed information about how the components have been architecturally defined and about our preliminary development results.

3. *Data Collection*

The *Data Collection* is the component that acquires existing inertial signals and uniforms their organization to populate the repository of labeled inertial signals called *Unified Dataset*. This is achieved by integrating i) existing datasets (e.g., UniMiB SHAR [22]); ii) labeled inertial signals coming from ad-hoc applications that have been designed to acquire new labeled inertial signals from volunteers (e.g., UniMiB AAL [13]); and iii) signals from volunteers performing ADLs that the platform enriches by assigning them a proper label and then integrates in the repository (e.g., Sensor Data Logger [1]). Figure 2 sketches the modules we identified for the *Data Collection* component.

One of the main issues in handling multiple datasets and exploiting them to train a single classifier, is the lack of consistency in terms of *how* the data is stored in the file system and *what* are the information provided. For example, in the UCI HAR dataset [2] data are stored in two separated directories (train and test) which contain .txt data files. In contrast, in MobiAct [35] data are subdivided in 20 directories, each of them representing an activity or a scenario. Each directory includes .csv files each belonging to a specific subject and to a specific trail. Thus, the different structures used influence how data is stored and the general organization of the information.

In *Continuous Learning Platform* we enforce a single storage technology for all data and a single structure for the data. In order to standardize the format of the information, *CLP* offers a well-defined format of the data: from the signals to the supporting information that enriches the data. Thus, in order to be integrated, a new dataset must also be coupled with a software component (called *driver*) that is able to interpret the data of the source dataset and convert it into the format defined in *CLP*. The development of the *driver* is responsibility of the researcher who whishes to load his/her dataset into the

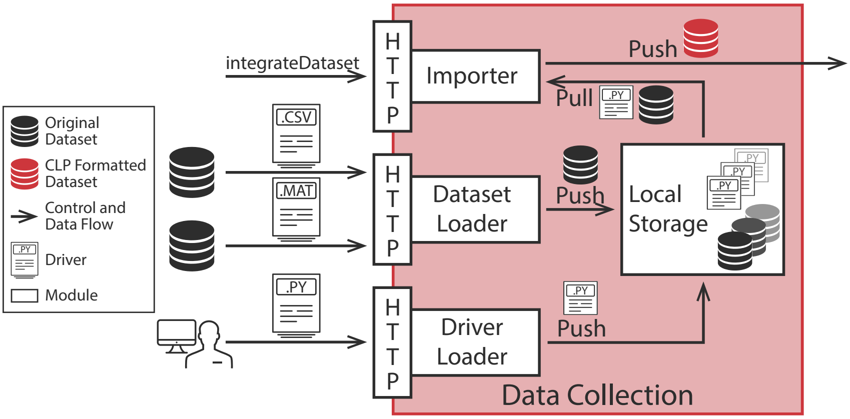

Figure 2. The *Data Collection* component.

Unified Dataset. Given the uniqueness of each datasets, it is unlikely that two datasets could share the same driver, thus each one will likely require an ad-hoc implementation. To ease the burden of writing such drivers, we provide a template interface for developing new drivers, which allows users to easily build new drivers that are compatible with *CLP*.

The *Data Collection* component includes the following modules: the *Driver Loader*, the *Dataset Loader*, and the *Importer* which respectively allow to load custom drivers developed to support specific datasets, to load datasets to be integrated in *Unified Dataset*, and to ask for the integration of the new datasets into the *Unified Dataset*. Separating the *Dataset Loader* from the *Driver Loader*, allows the dynamic on-boarding of the driver, which may require a reboot of the *Dataset Loader* service in order to be visible and exploitable from the service itself.

From an implementation point of view, all the modules are web services exposing the `loadDriver`, the `loadDataset`, and the `integrateDataset` functions.

3.1. *Preliminary Validation of the* Data Collection *Component*

In order to start validating the *Data Collection* component, we developed the *drivers* for the following datasets: Motion Sense [20], MobiAct [35], Real Word HAR [32], UmaFall [9], and UniMiB SHAR [22]. We selected these datasets for the following reasons. First, we were focused on datasets recorded by smartphone and smartwatch, because those kind of acquiring devices are not invasive devices and are widespread among the population. Second, we considered only datasets that have been acquired for HAR purposes. Indeed, such datasets may share the set of activities recorded. Third, we considered only datasets that are enriched with additional information related to the subjects' characteristics, such as sex, age, height, weight. This allows the *Data Distribution* component to provide sets of signals related to subjects with specified characteristics or to make available trained classifiers only with subsets of signals acquired from subjects with characteristics similar to the those required. The application of personalization, in fact, seems to provide better results in terms of accuracy [17,11]. Fourth, we selected data sets collected from 2016 until today for having comparable technology accuracy.

We have also developed a *driver* to import the datasets acquired by means of the UniMiB AAL suite [13]. UniMiB AAL includes two Andorid apps that ease both the acquisition of signals from sensors in a controlled environment and the labeling tasks required when building a dataset.

All the drivers have been implemented in the Python language, while the web service has been developed relying on Laravel[2] and exposing the services through RESTFul APIs.

4. *Data Management*

Once the signals have been standardized in terms of structure, they can be included in the *Unified Dataset* to be distributed. For this reason, the responsibility of the *Data Management* component is twofold: it integrates the new labelled signals into the *Unified Dataset* (*Data Storage*) and makes available sets of labelled signals to those who need them (*Data Access*).

Before being inserted into the *Unified Dataset*, signals require another elaboration to make them homogeneous both in terms of *representation* and *label*.

Sensors record data at a given sampling frequency, with a given range of intensity values, and so on. Each manufacturer design their own sensor with operating specifications that may be different from the typical ones. For instance, we may deal with accelerometer sensors that work at very different sampling frequency ranging from few to hundreds of Hertz. Machine learning methods require input data in a given format (e.g., number of samples per second and intensity range) that is consistent over time [4]. For this reason, raw data acquired by sensors need to be pre-processed before being processed by machine learning methods. Inevitably this produces an overhead of data to be handled and stored. Storage space management is carried out using cloud storage techniques that reduce space consumption by using capacity optimization, data deduplication and data compression tools [14].

Signals from different datasets may have assigned different (but semantical equal) labels for the same ADL (e.g., 'walk' vs 'walking'), same label for different ADL (e.g., 'sitting' may refer to the state of being seated in a chair or the transition from standing to sitting), and different labels for the same activity (e.g., 'running' vs 'jogging').

Thus, the aim of the *Data Management* component is to harmonize signals and to make them available for exploitation. The organization of the component is sketched in Figure 3.

The *Data Aligner* module is in charge of pre-processing the data from the *Data Collection* component in order to make them usable by any machine learning method. For example, an activity that is in charge of the *Data Aligner* module is the conversion to a same measurement unit.

The *Label Consolidator* module is in charge of uniforming the labels of the dataset to include to a common unified set. For example, if a dataset uses the label 'sitting' to label signals related to the transition (from standing to sitting down) and in the *Unified Dataset* is used 'sit down' to label signals related to the transition (from standing to sitting down), then the label will be changed to be consistent with the *Unified Dataset*. In view

[2] https://laravel.com/

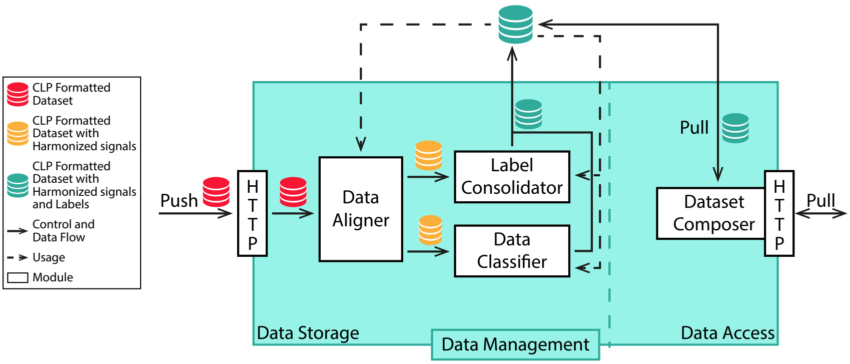

Figure 3. The *Data Management* component.

of the delicate nature of this procedure, this module is intended to be semi-automatic: it provides suggestions on the assignment of labels, but ultimately it is down to the end user to decide whether or not to accept the suggestions.

The *Data Classifier* module is in charge of assigning labels to inertial signals that have not been labeled. These types of signals may result from applications that acquire signals only without providing any classification. This module can exploit an activity recognition model already trained.

In terms of data distribution, the *Data Composer* module simply intercepts requests for labeled signals, processes them, and returns the set of corresponding labeled signals. For example, a request can be: "all signals labeled *running*".

From an implementation point of view, the *Data Management* component is a web service exposing the `storeDataset`, and the `getDataset` functions.

4.1. Preliminary Validation of the Data Management *Component*

To start validating the *Data Management* component, we implemented the *Data Aligner*, the *Label Consolidator*, and the *Data Composer* modules. All the modules have been developed in Matlab. The modules manage acceleration signals only.

Our implementation of the *Data Aligner* module unifies the measurement units to g, removes gravity form the signals, and resamples to 50 Hz the frequencies (since this is the frequency usually used for ADLs recognition [26]). Removing the gravitational acceleration is not an exact process, however it is common practice and considered in literature [34].

The implementation of the *Label Consolidator* includes an automatic syntactic analysis of the labels based on the Levenshtein distance. Then the module relies on a k-Nearest Neighbor classifier in order to compute a confusion matrix that helps the user in deciding which activities are similar and then can be merged. This confusion matrix and, when necessary, some visual plots of activities, are exploited by the user in order to determine the concrete associations of the various activities and labels.

Finally the *Data Composer* module allows to request specific sets of labeled signals. The requests are parametrized. For example, the request `getDataset('activities', [1 2 3], 'gender', ['F', 'age', 24)` returns all the labeled signals for activities 1, 2, and 3 performed by females 24 years old.

5. *Data Distribution*

The *Data Distribution* component role is to provide i) sets of labeled signals according to specific needs; ii) trained classifiers; and iii) labels corresponding to the activities performed given frames of signals.

The availability of sets of labeled signals helps researchers that need to validate their techniques with a wide and public data set. The huge amount of data also allows to work with deep learning techniques, and publicly available datasets allows to generalize the results and to compare techniques.

Figure 4 sketches the *Data Distribution* component architectural organization.

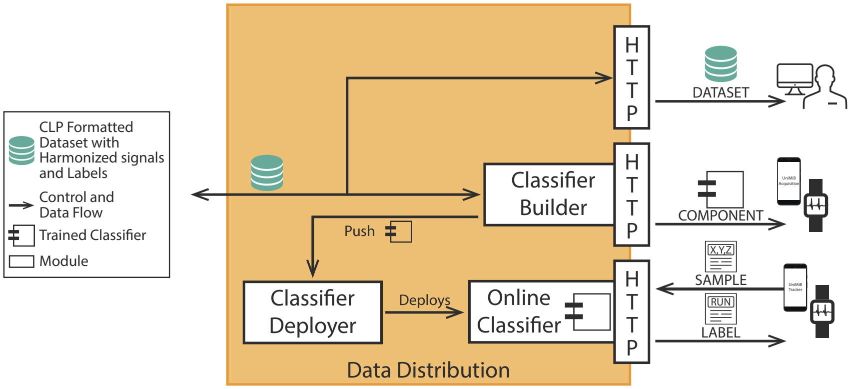

Figure 4. The *Data Distribution* component.

The *Classifier Builder* module is in charge of distributing activity recognition models that can be integrated in domain dependant applications. The module also relies on the *Classifier Deployer* to store the new trained activity recognition model to be used for online classification. Finally, the *Online Classifier* module provides online services related to classification: given a set of inertial signals, it provides information regarding the activity the subject is performing. Finally, a user asks for datasets by performing a HTTP request.

From an implementation point of view, the *Data Distribution* component is a web service exposing the `getDataset`, `getModule`, and `getLabel` functions.

5.1. *Preliminary Validation of the* Data Distribution *Component*

The component has been designed, but its implementation has been postponed in favour of the other components that are more challenging. Till now we have implemented the request for sets of signals as a RestFul endpoint.

6. Conclusions

Human activity recognition is a very challenging and active research field and has seen a rapid growth in the past few years.

The lack of large datasets penalizes the possibility of exploiting deep learning techniques that require a lot of data but providing a very good accuracy.

The goal of this work is to propose a platform which firstly integrates data from heterogenous sources and secondly provides several types of access to the data.

The framework has been partially implemented. We have prioritized the development of the most challenging modules: *Data Collection* and *Data Management* components. Till now, five datasets have been integrated.

Future directions include the development of the remaining modules and an intensive test of the overall platform. Once the tests are completed, we will make available the tool and the specifications for the design of the module for importing new datasets.

References

[1] Sensor data logger. `https://play.google.com/store/apps/details?id=net.steppschuh.sensordatalogger`, accessed on 22 March 2019.

[2] Davide Anguita, Alessandro Ghio, Luca Oneto, Xavier Parra, and Jorge Luis Reyes-Ortiz. A public domain dataset for human activity recognition using smartphones. In *Proceedings of the European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN13)*, 2013.

[3] Fabio Bagala, Clemens Becker, Angelo Cappello, Lorenzo Chiari, Kamiar Aminian, Jeffrey M Hausdorff, Wiebren Zijlstra, and Jochen Klenk. Evaluation of accelerometer-based fall detection algorithms on real-world falls. *PloS one*, 7(5):e37062, 2012.

[4] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. Multimodal machine learning: A survey and taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(2):423–443, 2019.

[5] James Bartlett, Vinay Prabhu, and John Whaley. Acctionnet: A dataset of human activity recognition using on-phone motion sensors. In *Proceedings of the International Conference on Machine Learning (ICML 17)*, 2017.

[6] Yoshua Bengio. Deep learning of representations: Looking forward. In *International Conference on Statistical Language and Speech Processing (SLSP13)*, 2013.

[7] Simone Bianco, Remi Cadene, Luigi Celona, and Paolo Napoletano. Benchmark analysis of representative deep neural network architectures. *IEEE Access*, 6:64270–64277, 2018.

[8] Andreas Bulling, Ulf Blanke, and Bernt Schiele. A tutorial on human activity recognition using body-worn inertial sensors. *ACM Computing Surveys (CSUR)*, 46(3):33, 2014.

[9] Eduardo Casilari, Jose A Santoyo-Ramón, and Jose M Cano-García. Umafall: A multisensor dataset for the research on automatic fall detection. *Procedia Computer Science*, 110:32–39, 2017.

[10] M. Cornacchia, K. Ozcan, Y. Zheng, and S. Velipasalar. A survey on activity detection and classification using wearable sensors. *IEEE Sensors Journal*, 17(2):386–403, 2017.

[11] Anna Ferrari, Daniela Micucci, Marco Mobilio, and Paolo Napoletano. On the personalization of classification models for human activity recognition. *IEEE Access*, submitted.

[12] Anna Ferrari, Daniela Micucci, Marco Mobilio, and Paolo Napoletano. On the homogenization of heterogeneous inertial-based databases for human activity recognition. In *Submitted to the IEEE World Congress on Services, Workshop on Big Data for Public Health Policy Making*, 2019.

[13] Davide Ginelli, Daniela Micucci, Marco Mobilio, and Paolo Napoletano. Unimib aal: An android sensor data acquisition and labeling suite. *Applied Sciences*, 8(8), 2018.

[14] Danny Harnik, Benny Pinkas, and Alexandra Shulman-Peleg. Side channels in cloud services: Deduplication in cloud storage. *IEEE Security & Privacy*, 8(6):40–47, 2010.

[15] Andrew P. Hills, Najat Mokhtar, and Nuala M. Byrne. Assessment of physical activity and energy expenditure: An overview of objective measures. *Frontiers in Nutrition*, 1(5), 2014.

[16] Christian Krupitzer, Timo Sztyler, Janick Edinger, Martin Breitbach, Heiner Stuckenschmidt, and Christian Becker. Hips do lie! a position-aware mobile fall detection system. In *IEEE International Conference on Pervasive Computing and Communications (PerCom18)*, 2018.

[17] Nicholas D Lane, Ye Xu, Hong Lu, Shaohan Hu, Tanzeem Choudhury, Andrew T Campbell, and Feng Zhao. Enabling large-scale human activity inference on smartphones using community similarity net-

works (csn). In *Proceedings of the 13th international conference on Ubiquitous computing (UbiComp11)*, 2011.

[18] Oscar D Lara, Miguel A Labrador, et al. A survey on human activity recognition using wearable sensors. *IEEE Communications Surveys and Tutorials*, 15(3):1192–1209, 2013.

[19] Yann LeCun and Yoshua Bengio. The handbook of brain theory and neural networks. pages 255–258, 1998.

[20] Mohammad Malekzadeh, Richard G Clegg, Andrea Cavallaro, and Hamed Haddadi. Protecting sensory data against sensitive inferences. In *Proceedings of the 1st Workshop on Privacy by Design in Distributed Systems (P2DS18)*, 2018.

[21] Carlos Medrano, Raul Igual, Inmaculada Plaza, and Manuel Castro. Detecting falls as novelties in acceleration patterns acquired with smartphones. *PloS one*, 9(4):e94811, 2014.

[22] Daniela Micucci, Marco Mobilio, and Paolo Napoletano. Unimib shar: A dataset for human activity recognition using acceleration data from smartphones. *Applied Sciences*, 7(10), 2017.

[23] Daniela Micucci, Marco Mobilio, Paolo Napoletano, and Francesco Tisato. Falls as anomalies? an experimental evaluation using smartphone accelerometer data. *Journal of Ambient Intelligence and Humanized Computing*, 8(1):87–99, 2017.

[24] Alex Adim Obinikpo and Burak Kantarci. Big data aggregation in the case of heterogeneity: a feasibility study for digital health. *International Journal of Machine Learning and Cybernetics*, pages 1–13, 2019.

[25] Alexander Ratner, Stephen H Bach, Henry Ehrenberg, Jason Fries, Sen Wu, and Christopher Ré. Snorkel: Rapid training data creation with weak supervision. *Proceedings of the VLDB Endowment*, 11(3):269–282, 2017.

[26] Nishkam Ravi, Nikhil Dandekar, Preetham Mysore, and Michael L Littman. Activity recognition from accelerometer data. In *Proceedings of the 17th conference on Innovative applications of artificial intelligence (IAAI05)*, 2005.

[27] Jorge-L Reyes-Ortiz, Luca Oneto, Albert Samà, Xavier Parra, and Davide Anguita. Transition-aware human activity recognition using smartphones. *Neurocomputing*, 171:754–767, 2016.

[28] Charissa Ann Ronao and Sung-Bae Cho. Human activity recognition with smartphone sensors using deep learning neural networks. *Expert systems with applications*, 59:235–244, 2016.

[29] Erika Rovini, Carlo Maremmani, and Filippo Cavallo. How wearable sensors can support parkinson's disease diagnosis and treatment: A systematic review. *Frontiers in Neuroscience*, 11(555), 2017.

[30] Muhammad Shoaib, Stephan Bosch, Incel Ozlem Durmaz, Hans Scholten, and Paul J. M. Havinga. Fusion of smartphone motion sensors for physical activity recognition. *Sensors*, 14(6):10146–10176, 2014.

[31] Pekka Siirtola, Heli Koskimäki, and Juha Röning. Openhar: A matlab toolbox for easy access to publicly open human activity data sets. In *Proceedings of the 2018 ACM International Joint Conference and 2018 International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers (UbiComp18)*, 2018.

[32] Timo Sztyler and Heiner Stuckenschmidt. On-body localization of wearable devices: An investigation of position-aware activity recognition. In *IEEE International Conference on Pervasive Computing and Communications (PerCom16)*, 2016.

[33] Prabitha Urwyler, Reto Stucki, Luca Rampa, René Mü ri, Urs P. Mosimann, and Tobias Nef. Cognitive impairment categorized in community-dwelling older adults with and without dementia using in-home sensors that recognise activities of daily living. Scientific Technical Report 42084, PMC, 2017.

[34] Vincent T Van Hees, Lukas Gorzelniak, Emmanuel Carlos Dean Leon, Martin Eder, Marcelo Pias, Salman Taherian, Ulf Ekelund, Frida Renström, Paul W Franks, Alexander Horsch, et al. Separating movement and gravity components in an acceleration signal and implications for the assessment of human daily physical activity. *PloS one*, 8(4):e61691, 2013.

[35] George Vavoulas, Charikleia Chatzaki, Thodoris Malliotakis, Matthew Pediaditis, and Manolis Tsiknakis. The mobiact dataset: Recognition of activities of daily living using smartphones. In *Proceedings of the International Conference on Information and Communication Technologies for Ageing Well and e-Health (ICT4AgeingWell16)*, 2016.

[36] Ming Zeng, Le T Nguyen, Bo Yu, Ole J Mengshoel, Jiang Zhu, Pang Wu, and Joy Zhang. Convolutional neural networks for human activity recognition using mobile sensors. In *6th International Conference on Mobile Computing, Applications and Services (MobiCASE14)*, 2014.