8-27-2016

Important Considerations for Human Activity Recognition Using Sensor Data

Matt Buckner
Rose-Hulman Institute of Technology, bucknemj@rose-hulman.edu

IMPORTANT CONSIDERATIONS FOR HUMAN ACTIVITY RECOGNITION USING SENSOR DATA

Matthew Buckner
Mathematics and Computer Science Double Major
Rose-Hulman Institute of Technology
Terre-Haute, Indiana, USA
bucknemj@rose-hulman.edu

Advisor: Yosi Shibberu
Mathematics
Rose-Hulman Institute of Technology
Terre-Haute, Indiana, USA

ABSTRACT

Automated human activity recognition has received much attention in recent years due to increasing focus on interconnected devices in The Internet of Things (IoT) and the miniaturization and proliferation of sensor systems with the adoption of smartphones. In this work, we focus on the current status of human activity recognition across multiple studies, including methodology, accuracy of results, and current challenges to implementation. We include some preliminary work we have completed on a sensor system for classifying treadmill usage.

Introduction

In this paper we compare a selection of recent studies within the field of human activity recognition. We also outline challenges for future work in this area.

Two main categories for implementation for human activity recognition exist. The first, and most common seeks to perform data classification on the data collection device itself [1,4,5,6]. The second category, which we have implemented in our own work [10], separates data collection and data classification across more than one [2,3,7,8,9,10].

Single device implementations are increasingly viable due to the widespread adoption of smartphones, which commonly contain both the sensors required, and a processor powerful enough for data classification. Alternatively, separating the sensor unit from the computing device can provide additional computing power by utilizing a much larger and more powerful computer or server for classification, and ideally reducing power on the sensor device. However, in practice, the power consumption from continual data transfer for real time classification may negate the benefits of increased computing power [6,10].

The outline of this paper is as follows: Four major concerns of human activity recognition are reviewed. We focus on current results and conclusions from existing studies including work previously performed by the author. The concerns are:

1. Power consumption due to the strict limitations provided by mobile sensor platforms such as smartphones.
2. Device location due to the wide variety present in phone location in real world usage.

3. Sensor fusion, utilizing additional sensors to potentially provide a more accurate classification.
4. The difficulty of utilizing labeled data versus unlabeled data due to additional time and preparation requirements.

Finally, we present a table comparing and contrasting the methods and results of the selection of studies examined in this article.

Energy Consumption

One major concern of human activity recognition is the effect it has on the battery life on a mobile device. Traditional methods of human activity classification involve post-processing of sensor data, which can introduce delay in data analysis. On device classification can provide rapid results which can provide advantages in, for example, preventative healthcare [3], medical monitoring [1,10], and athletic performance [10]. Post-processing limits implementations to either on-device analysis [1,4,5,6] or to syncing data to an additional system for analysis [2,3,7,8,9,10]. It has been noted in several studies, separating data gathering and analysis across different devices introduces additional power drain due to wireless data transfer [6,10].

Khan et al illustrates that a tradeoff between power consumption and accuracy. Increasing the time window or sampling rate of the gathering device directly correlates to power consumption of the device [4]. Khan provides the most directly measurable information concerning these effects of the examined studies. Khan states that across 4 different feature sets formed from a single data source, there is a marked increase in accuracy as sampling rate is increased, but this is also paired with considerable increases in power consumption. As such, additional thought must be placed into feature selection methods based on power consumption and required accuracy.

From Khan's work [4], he finds the usage of Auto-regression Analysis coefficients to provide the highest accuracy for the lowest power consumption at all sampling rates, providing nearly 90% accuracy at a 20 Hz sampling rate with approximately 100 Joules of energy consumption. While these results do aid in illustrating the challenges of power consumption, extending these specific results to other studies proves difficult due to the various differences in power consumption for both hardware and various classifiers [6].

In addition to sampling rate, another major influence on power consumption is the specific sensors utilized, as the usage of additional sensors increase the power drain on the device [1,6]. To minimize power drain from sensors, a subset of the examined papers investigated activity classification via a single sensor, the Accelerometer [1,3,4,6]. This resulting data set can be comparatively smaller than that of a multi-sensor approach, saving power in classification. Additional steps, such as disabling sensors when not in use [2], can provide minor benefits, but fails to solve this issue for systems intended for continuous usage.

The focus on limiting the number of high-power requirement sensors provides additional limitations, as the most common sensor used for calculating user speed, the Global Positioning System (GPS), is an exceptionally power hungry device [5]. The addition of GPS data allows for both the detection of user speed and location, which can be valuable data in determining user activity when a broader set of similar motion activities is examined, such as driving a car, riding a bus, or simply sitting at home [5].

Furthermore, although power consumption is a critical issue in mobile activity analysis, a majority of the current studies in this field currently lack the tools required to provide concrete comparable data on the power consumption of alternative methods. As an example, one paper specifies that the usage of time domain features provides significant power savings over that of the commonly used frequency domain features, but fails to provide empirical data on the power usage, limiting the utility of this conclusion [5].

Device Location

Earlier works that featured dedicated sensor devices or that focused on specific classification techniques primarily utilized fixed sensor locations [1,8,10]. As smart phones became more prevalent, studies began to focus more on the sensors included in these devices, as they are willingly worn by the subject in real world scenarios and throughout the day. Unfortunately, these smart phones are also multipurpose and as such are prone to switch location based on activity and user preference. A majority of the studies focused on the most common phone locations, including pants and jacket pockets [2,3,4], while a few allowed for more potential usage scenarios [5,6].

With more possible locations introduced, including the possibility of the phone being placed on a nearby table during some actions, additional information must be utilized to allow for accurate human activity recognition, including the usage of additional sensors. For example, microphones, pressure, light sensors, GPS, and more were used to increase accuracy. Of particular note, Martin et al. [6] examined a two stage process where the classifier first attempted to recognize the location of the phone, then afterwards classify the activity with pre-trained data for that location. This provided an average 10% increase in classification accuracy when all sensor data was utilized across three different classifiers. This shows that as additional usage scenarios are introduced, more data must be gathered to provide more accurate classifications.

It was also noted that in the experiments recorded, males were more likely to place their phone in a front trouser pocket, while females commonly preferred a bag or purse [6]. Focusing classification on the assumption that the device will be in a specific location depending on user gender may provide some statistical advantages, however it would fail to properly classify the alternative usage locations, such as the usage of arm band phone holders while exercising.

The difference in stride generated by users of different heights and habits can present additional complications. To mitigate differences between users, multiple papers [5.6] investigated user-calibrated systems, which would take user habits in phone location into consideration. In practice, Khan et al. found that a subject dependent classification system increased accuracy by 9.3% [5].

Some of the above issues can be avoided by specifying fixed sensor locations. In our own work on human activity recognition for treadmill usage [10], sensors were applied to fixed locations on the treadmill itself. Data then was transferred via Bluetooth. User activity classification was achieved independent of the specific user.

Unfortunately, sensors fixed in a remote location require additional power for wireless communication.

As companies build products and devices for the “Internet of Things,” sensors may become part of everyday items such as user clothing [1]. These sensors would possess the advantages originally noted to smartphones as always being available, but in addition having a fixed known location.

Sensor Fusion

Accelerometer sensors were used in all of the studies examined. In some studies, [1,3,4,6], Accelerometer data was the only sensor utilized, yielding accuracies up to 98% [1,3]. As noted in Device Location section, as usage scenarios are added, more information is needed to separate activities with similar movements, such as recognizing the difference between driving and sitting in a bus [5]. To provide more versatile classification, a majority of the studies examined utilized additional sensors [2,5,6,8,10], either across multiple devices [2], or within the same device [5,6,8,10].

Some examples of this include the usage of both a smartphone and smartwatch [2], allowing for multiple sensor points to classify user activity. For example, Guiry et al. found a wrist mounted accelerometer to be particularly good at separating walking from running, and in recognizing stair climbing. It is of interest to note that in this particular study, they found that the majority of subjects reached out and used support railings when climbing up or down stairs, creating a separation in the wrist accelerometer data of subjects who used and did not use the railing. The presence of a second sensor can identify the presence of this subclassification and enhance general accuracy.

Another study of note in multisensory usage included the addition of a chest mounted accelerometer [3]. This fixed location sensor provided essential data to separate activities such as sitting and lying in their study, but the author does note that classifying both activities as “Sedentary” also overcomes this shortcoming without the need of an additional sensor. In this way, we can see that limiting classification to upper level activities instead of more specific subsets can reduce the amount of data required.

Labeling Training Data

Obtaining usable labeled data is a time intensive task, often requiring the inclusion of an additional user to mark and label data as it is obtained [2,3,6,10], or additional custom software must be made to facilitate user self-labeling [4,5,7]. This in turn increases time required for cleaning and preparing the ground truth data set for proper classification [10].

One relatively unexplored method of human activity recognition is utilizing unsupervised methods, or using unlabeled data. Of the studies examined, only one [9] investigated the usage of unsupervised methods. They examined five activities: walking, running, sitting, standing, and lying down; across a series of clustering algorithms. They found when the number of expected clusters were known, Gaussian Mixed Models provided exact classification in their tests, though without an expected number of clusters accuracy of 90% was still obtainable via hierarchical clustering.

CONCLUSIONS

From the examined works we safely conclude that although much work has been performed in the field of human activity recognition, many questions and issues still exist worthy of future work. Battery life has been nearly universally noted as an important point of consideration, but current studies lack the tools required to provide comparable results on energy usage. Furthermore, the direct link between potential accuracy via increased sensor and power consumption still lacks a definitive answer between studies.

The introduction of smartphones as a sensor platform has provided an easily obtainable system for human activity recognition, but difficulties with sensor location and user-independent classification provide additional fields of investigation. Additionally, the introduction of sensors beyond that of an accelerometer provides solutions to classifying otherwise similar activities and increase accuracy, but do introduce considerable power considerations depending on the sensors employed.

Finally, unsupervised training provides a relatively untapped field of study that would minimize many of the difficulties currently associated with training and preparing existing classifiers.

REFERENCES

[1] Álvarez de la Concepción, M., Soria Morillo, L., Gonzalez-Abril, L., & Ortega Ramírez, J. (2014). Discrete techniques applied to low-energy mobile human activity recognition. A new approach. Expert Systems With Applications, 41(14), 6138-6146. http://dx.doi.org/10.1016/j.eswa.2014.04.018

[2] Guiry, J., van de Ven, P., & Nelson, J. (2014). Multi-Sensor Fusion for Enhanced Contextual Awareness of Everyday Activities with Ubiquitous Devices. Sensors, 14(3), 5687-5701. http://dx.doi.org/10.3390/s140305687

[3] Guiry, J., van de Ven, P., Nelson, J., Warmerdam, L., & Riper, H. (2014). Activity recognition with smartphone support. Medical Engineering & Physics, 36(6), 670-675. http://dx.doi.org/10.1016/j.medengphy.2014.02.009

[4] Khan, A., Siddiqi, M., & Lee, S. (2013). Exploratory Data Analysis of Acceleration Signals to Select Light-weight and Accurate Features for Real-time Activity Recognition on Smartphones. Sensors, 13(10), 13099-13122. http://dx.doi.org/10.3390/s131013099

[5] Khan, A., Tufail, A., Khattak, A., & Laine, T. (2014). Activity Recognition on Smartphones via Sensor-Fusion and KDA-Based SVMs. International Journal Of Distributed Sensor Networks, 2014, 1-14. http://dx.doi.org/10.1155/2014/503291

[6] Martín, H., Bernardos, A., Iglesias, J., & Casar, J. (2012). Activity logging using lightweight classification techniques in mobile devices. Pers Ubiquit Comput, 17(4), 675-695. http://dx.doi.org/10.1007/s00779-012-0515-4

[7] Reyes-Ortiz, J., Oneto, L., Samà, A., Parra, X., & Anguita, D. (2016). Transition-Aware Human Activity Recognition Using Smartphones. Neurocomputing, 171, 754-767. http://dx.doi.org/10.1016/j.neucom.2015.07.085

[8] San-Segundo, R., Montero, J., Barra-Chicote, R., Fernández, F., & Pardo, J. (2016). Feature extraction from smartphone inertial signals for human activity segmentation. Signal Processing, 120, 359-372. http://dx.doi.org/10.1016/j.sigpro.2015.09.029

[9] Kwon, Y., Kang, K., & Bae, C. (2014). Unsupervised learning for human activity recognition using smartphone sensors. Expert Systems With Applications, 41(14), 6067-6074. http://dx.doi.org/10.1016/j.eswa.2014.04.037

[10] Blank, N., Buckner, M., Owen, C. Scott, A. (2016). Real time activity recognition of treadmill usage via machine learning. Rose-Hulman Institute of Technology, Terre-Haute, Indiana, USA

Paper	Activities	Sensors	Examined Locations	Top Classifier(s)	Classification on/off device	Accuracy	Comments
Álvarez: Discrete techniques	Walk, Jump, Immobile, Run, Climb, Descend, Cycle, Drive	Accelerometer	Hip	Custom Classifier	On Device	98%	
Guiry: Multi-Sensor Fusion	Sitting, Standing, Walking, Running, Cycling, Stair Decent, Stair Ascent, Elevator Descent, Elevator Ascent	Phone: Accelerometer, Magnetometer, Gyroscope, GPS, light sensor, pressure sensor Smartwatch: Accelerometer	Phone: Pants pocket (any), Smartwatch: Wrist (any)	Support Vector Machine (For balanced datasets)	Off Device	72.63%	
				CART based decision trees (For unbalanced datasets)		94.73%	
				Multi-Layer Perceptrons (Overall principle component analysis on smartphone)		92.89%	
				C4.5 decision trees (Overall principle component analysis on smartwatch)		56.89%	
Guiry: Activity recognition	Varied based on trial; Sitting, Standing, Walking (various), Cycling (various), Jogging, Running	Accelerometer	Phone: Pants pocket (any) PLUX Sensor: Chest	C4.5 decision trees	Off Device	98%	

Paper	Activities	Sensors	Examined Locations	Top Classifier(s)	Classification on/off device	Accuracy	Comments
Khan: Exploratory Data Analysis	Standing, Walking, Running, Upstairs, Downstairs, Hopping	Accelerometer	Pants pocket (any), Jacket inner pocket	Artificial Neural Network with Autoregressive Modeling Features	On Device	87.1%	Examined effects of various sampling rates and feature selection methods, including energy consumption and accuracy
Khan: Activity Recognition on Smartphones	Walking (on/off treadmill), Running (on/off treadmill), Upstairs, Downstairs, Elevator up/down, Biking, Hopping, Idle, Watching TV, Vacuuming, Driving, Riding a Bus, Others	Accelerometer, Pressure sensor, Microphone	Pants pocket (any), Jacket pocket (any), in hand (while idle or watching TV), on table (while idle or watching TV)	Kernel Discriminant Analysis based Support Vector Machine	On Device	94%	
Martín: Activity logging	Sit, Stand, Walk, Slow Walk, Rush Walk, Run	Accelerometer, Magnetometer, Orientation Sensor, Light Sensor, Proximity Sensor, Gyroscope, Rotation Sensor	Pants Pocket (any), Shirt pocket, Hand (Texting or “Talking” positions), Waist Case, Backpack, Jacket Pocket, In Short/Long Strap Bag, Armband	Decision Tree	On Device	92.92%	Separated classification into recognizing location of the phone and recognizing activity

Paper	Activities	Sensors	Examined Locations	Top Classifier(s)	Classification on/off device	Accuracy	Comments
San-Segundo: Feature extraction	Walking, Walking Upstairs, Walking Downstairs, Sitting, Standing, Lying	Accelerometer, Gyroscope	Waist	Custom Classifier	Off Device	99.3%	Applies speech processing methods to Human Activity recognition
Buckner: Real time activity recognition	Walking, Running, Not Walking (No activity)	Accelerometer, Gyroscope, Magnetometer	On Treadmill	Random Forests	Off Device	98.1%	