Published in final edited form as:

J Neurosci Methods. 2014 July 15; 231: 22–30. doi:10.1016/j.jneumeth.2013.09.015.

Hand, belt, pocket or bag: Practical activity tracking with mobile phones

Stephen A. Antos^{1,2}, Mark V. Albert^{2,3,4}, and Konrad P. Kording^{2,3,5,6}

¹Department of Biomedical Engineering, Northwestern University, Evanston, IL, USA

²Sensory Motor Performance Program, Rehabilitation Institute of Chicago, Chicago, IL, USA

³Department of Physical Medicine and Rehabilitation, Northwestern University, Chicago, IL, USA

⁴Center for Bionic Medicine, Rehabilitation Institute of Chicago, Northwestern University, Chicago, IL, USA

⁵Department of Physiology, Northwestern University, Chicago, IL, USA

⁶Department of Physiology, Northwestern University, Chicago, IL, USA

Abstract

For rehabilitation and diagnoses, an understanding of patient activities and movements is important. Modern smartphones have built in accelerometers which promise to enable quantifying minute-by-minute what patients do (e.g. walk or sit). Such a capability could inform recommendations of physical activities and improve medical diagnostics. However, a major problem is that during everyday life, we carry our phone in different ways, e.g. on our belt, in our pocket, in our hand, or in a bag. The recorded accelerations are not only affected by our activities but also by the phone's location. Here we develop a method to solve this kind of problem, based on the intuition that activities change rarely, and phone locations change even less often. A Hidden Markov Model (HMM) tracks changes across both activities and locations, enabled by a static Support Vector Machine (SVM) classifier that probabilistically identifies activity-location pairs. We find that this approach improves tracking accuracy on healthy subjects as compared to a static classifier alone. The obtained method can be readily applied to patient populations. Our research enables the use of phones as activity tracking devices, without the need of previous approaches to instruct subjects to always carry the phone in the same location.

Keywords

Activity recognition; Parkinson's disease; smartphone; accelerometer; context awareness; classification

© 2013 Elsevier B.V. All rights reserved.

Corresponding Author: Name: Stephen Antos, Phone: 011-1-585-406-6223, Fax: 312-238-2208, santos@u.northwestern.edu, Postal Address: Sensory Motor Performance Program, 345 E Superior St, Chicago, IL 60611.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

1. Introduction

Activity tracking uses wearable sensors, combined with computer algorithms to quantify what we are doing during our everyday lives. Such tracking has been applied to ambulatory monitoring (Albert et al., 2012b; Mathie et al., 2004) and automatic fall detection (Albert et al., 2012a; Kangas et al., 2008; Lee and Carlisle, 2011) to inform clinical providers about patient activity at home and in the community. Precise, objective measures of patient activity can be used to create, modify, and evaluate individualized treatment plans to improve a patient's outcomes and quality of life.

There are not many ways of evaluating patient mobility in everyday life. For clinical evaluations the patient needs to travel to the health care provider, and testing there is expensive in terms of money and clinician time. This cost prohibits frequent evaluations. Consequently, most prospective studies use patient journaling. However, such self-reporting is associated with various problems. First, the approach is necessarily subjective, leading to changes based on the mental state of the subject. Journaling can also have low compliance rates, in some cases as low as 11% (Stone et al., 2003). It would be helpful to develop a measure of evaluating patient mobility that is both frequent and convenient for the subject.

Few studies have applied activity tracking systems to people with neurodegenerative diseases. In a study involving patients with Duchenne muscular dystrophy, a single monitoring device worn on the chest reliably quantified walking parameters and time spent performing different activities (Jeannot et al., 2011). Patients with Parkinson's disease wore three body-fixed inertial sensors which were used to accurately identify activities including walking, sitting, standing, lying, sit to stand, and stand to sit transitions (Salarian et al., 2007). It has also been shown that a mobile phone placed in a pocket can be used to recognize activities in people with Parkinson's disease (Albert et al., 2012b). Each of these systems can quantify physical activity, but they require participants to wear sensors in a fixed location, and generally wearing more sensors leads to more accurate tracking (Bao and Intille, 2004). However, wearing sensors can be a burden to the wearer and not practical for everyday use.

Mobile phones have become a practical and popular activity tracking platform because people are accustomed to carrying phones on a daily basis. This makes data easier to collect as compared to custom tracking systems. In addition, mobile phones eliminate the need to design custom hardware, because they have built in sensors, memory, and computing power. Furthermore, software creation and distribution is easier because open source tools allow anyone to create applications and deploy them on mobile phones. Thus, mobile phones conveniently contain all of the hardware and software capabilities to create a stand-alone activity tracking system, with the practical benefit that people wear them every day.

Most activity tracking studies using mobile phone accelerometers instruct their participants to wear the phone in a predetermined location and orientation, the most common location being the pocket (Bieber et al., 2009; Kwapisz et al., 2010; Lau and David, 2010). In everyday life, however, people wear their mobile phones in many different ways. In a cross cultural study on phone carrying, 60% of men carried the phone in their pocket while 61%

of women carried the phone in their bag (Cui et al., 2007). The next most common locations were the belt clip, upper body, and hand. Instructing people to wear their phone in a specific location disregards their motivations for why they prefer carrying the phone in a particular way, such as ease of carrying, comfort, or fashion. For people with any type of movement disorder, restrictions for how they carry the phone can limit compliance, and consequently both the quality and amount of data collected by the tracking system. To be both accurate and practical, we must ensure our classifiers work when the phone is worn in different ways.

Here we track activities and phone locations using Support Vector Machines (SVM), a Hidden Markov Model (HMM), and the idea that activities change rarely, and the location of the phone changes even less often. This idea is made explicit through the HMM with an appropriately constructed transition matrix. Previous work in speech and activity recognition has shown improved prediction accuracy when combining SVMs and HMMs (Ganapathiraju et al., 2000; Lester et al., 2005; Yang, 2009). Incorporating time and transitions into our model improves activity and location tracking accuracy by smoothing out temporal irregularities. With this model, we conduct several analyses to understand the impact of phone location on tracking accuracy, and which phone locations are best for activity tracking. We then use this model in a pilot study to continuously track activities for two people with Parkinson's disease. By improving tracking accuracy when the phone is worn in multiple locations, people using activity tracking can wear their phone however they like.

2. Materials and methods

2.1 Protocol

Experimenters guided twelve able-bodied subjects (29.25 ± 6.27 years, mean $\pm$ SD) through a predetermined sequence of activities (see Figure 1). When prompted, subjects would walk, sit, or stand, and periodically change how they were wearing the phone. For each sequence of activities, the subject wore a phone in a different location: either in their pocket, in a belt clip, in their hand, or in a bag. Sit to stand and stand to sit transitions were included as activities to capture transitory states. The sequence was created to include as many transitions and locations as possible, while also leaving enough time for the person to perform each activity naturally. We used a T-Mobile G1 phone with a built in tri-axial accelerometer capable of measuring $\pm 2.8g$ and operating on Android 1.6 OS. Accelerations were recorded using a custom Android application with a sampling rate between 15 and 25 Hz depending on the movement speed (Fernandes et al., 2011). The experimenter used a stopwatch to ensure each activity was performed for a minimum amount of time, and accelerations were recorded continuously for the entire experiment. This study was approved by the Northwestern Institutional Review Board, and all subjects provided written, informed consent.

To prevent constraining how the phone was worn, subjects could place the phone in any orientation during the pocket, hand, or bag sequences. The belt clip, because of its design, constrained the phone to an upright orientation. We recorded a total of eight sequences, two per location. For the pocket and belt clip, one sequence was recorded from the right side of the body (e.g. right pocket), and the other sequence was recorded from the left side of the

body. To carry the phone in their hand, subjects first pretended to make a phone call, and then pretended to text message or glance at the screen. Subjects brought their own bag, which included backpacks, purses, and in one case, a plastic grocery bag. The bag was carried in a two different ways; for example, a backpack can be worn on the back or carried in the hand. Carrying the phone for eight sequences resulted in approximately 35 minutes of data per subject, or 7 hours of data over all subjects.

2.2 Data preprocessing

One experimenter performed annotations for the activity and phone location after data collection. Transitions between phone locations were labeled as a *transition* activity in the *hand* location because subjects used their hands to transfer how the phone was worn. Data before the subject picked up the phone at the beginning of the experiment and after the subject set the phone down at the end of the experiment were removed prior to analysis. One subject dropped the phone as they were changing the phone's location; this event was noted and removed from the data as well. These were the only data removed from the time series. After annotations, accelerometer data was linearly interpolated to match 20 Hz, and then divided into two second clips. Clips were included in the SVM training set only if at least 80% of the clip was marked as a single activity to prevent mixed training data samples. We used all clips to test our classifiers (ignoring the 80% threshold) so the resulting predictions were continuous in time. The ground truth label for each clip was the activity taking up the largest percentage of time in those two seconds. For each two second clip, we calculated 106 features (Table 1), and each feature was linearly scaled to a range between 0 and 1.

2.3 Classification algorithms

We used SVM classifiers from the LIBSVM package (Chang and Lin, 2011). The SVM kernel functions were radial basis functions requiring a soft slack variable (C) and Gaussian kernel size (γ). Using a 10^x grid search, where x is an integer from -5 to 5 , we found the hyper-parameters ($C = 10$, $\gamma = 1$) that minimized the 4-fold cross validation error when the location of the phone was known. Each SVM classifier produced probabilistic predictions, which were then used to calculate the mean of Gaussian emission distributions for each hidden state when the ground truth is known. All emission distributions were assigned a constant standard deviation $\sigma = 0.05$ using a grid search to minimize cross validation error. We created HMMs using the pmtk3 package (Murphy and Dunham, 2011). The prior distribution was set to be uniform across all analyses, and we manually specified a transition matrix for each analysis. The transition matrix reflects the idea that activities change rarely and the phone location changes even less often. That is, the $\text{Prob}(\text{Activity}_i(t+1) | \text{Activity}_i(t))$ and the $\text{Prob}(\text{Location}_k(t+1) | \text{Location}_k(t))$ is closer to one, and the $\text{Prob}(\text{Activity}_i(t+1) | \text{Activity}_j(t))$ and the $\text{Prob}(\text{Location}_k(t+1) | \text{Location}_l(t))$ is closer to zero (where $i \neq j$, $k \neq l$). To test our models, we first made probabilistic predictions with the SVM and then applied the forward-backward algorithm on our HMM to infer the most likely activity and/or location.

2.4 Activity tracking

We conducted four types of analyses using SVMs and HMMs to examine how accurately we can track activities when the phone is worn in different locations and orientations. Each analysis differs from the others in one or more of the following: the training set, the test set, or the transition matrix for the HMM.

In the first case, we assume the exact location of the phone is known (location-known). To evaluate this model, SVMs were trained with clips from one of the four locations (pocket, belt, hand, or bag) and then tested with clips from that same location. The transition matrix for the HMM was constructed based on the idea that activities change rarely and the phone location does not change (Figure 2A). Transitory activities like sit-to-stand or stand-to-sit were biased toward their assumed states before and after the transition.

In the next analysis we train our models with data from one phone location (e.g. pocket), but we test on data from all four phone locations (location-assumed). The training set and transition matrix were the same as the previous analysis (Figure 2A). However, in this analysis, the test set was expanded to include all phone locations.

The third analysis trains the SVM and HMM with data from all locations, and then predicts activities that were performed in all locations (global model). This analysis has a transition matrix that includes an extra transition activity that reflects the subject changing the phone's location, resulting in a total of six activities. For this type of analysis, we also report the accuracy of transitioning from one state to another (transition accuracy). We define a correct transition as beginning in the correct state, ending in the correct state, and occurring within 6 seconds of the labeled transition.

2.5 Location tracking

In addition to activity tracking, we also evaluated our ability to track the phone's location using the same three types of analyses outlined above. First, we predicted the location of the phone when the training set consisted of one activity and the testing set consisted of the same activity (activity-known). Next, we predicted the phone location when the training set consisted of one activity, but the testing set consisted of all five activities (activity-assumed). Finally, we predicted the phone location with a training and test set consisting of all activities (global model).

2.6 Joint activity-location tracking

We trained the SVM and HMM to predict both the location of the phone and the activity. The transition matrix is a block matrix that uses the transition matrices from previous analyses (Figure 2B). The SVM predicts joint activity-location pairs, and the HMM simultaneously transitions between both the activity of the subject and location of the phone (joint classifier).

2.7 Subject-wise cross validation

To determine how well phone locations generalize to a new person, we performed subject-wise cross validation and predicted the person's activities. For this type of validation, the

classifier was trained on 11 of the subjects, and tested on the one remaining subject. This was repeated twelve times, so that each subject was tested once.

2.8 Pilot application to patients with Parkinson's disease

We applied our method to track activities for two subjects with Parkinson's disease. Both subjects were able to walk without an assistive device, and could perform sit to stand and stand to sit transitions independently. The subjects both carried the phone in either their hand or pocket and performed the following activities: walking, sitting, standing, or they set the phone down (not wearing). The sequence was determined in a similar way to controls – incorporating as many transitions as possible while also leaving enough time in each state for the person to do the activity naturally. An experimenter followed the subjects and marked the sequence of activities for annotation after data collection. Data was collected continuously for approximately 10 minutes for each subject, and all classification was performed offline. The joint SVM+HMM was constructed using data from only the subjects with Parkinson's disease. In this pilot application we observed a subset of the locations and activities, so the transition matrix was constructed using a subset of the previous transition matrix and appended to include the new “not wearing” condition.

3. Results

Twelve subjects carried a mobile phone in four different locations and performed five different activities while their accelerations were recorded (Figure 3). The accelerations were divided into clips and features were calculated (see table 1). These features were input into SVMs to probabilistically predict the activity or phone location. The probabilistic outputs from the SVM were then input into a HMM. In addition to models which separately predict the activity or location, we also created a joint activity-location SVM+HMM classifier to simultaneously make predictions for both the activity and location of the phone.

3.1 Activity tracking

In most mobile phone activity tracking studies, the location of the phone is known. We first perform an analysis with a similar validation method. SVM+HMM activity classifiers were trained and tested on the same phone location while different activities were performed (location-known). We predicted the activity with 90.8% accuracy (Figure 4A). The pocket location achieved the highest activity tracking accuracy (95.2%), followed by the belt (93.6%), the bag (91.3%) and the hand (83.2%). Knowing how someone wears the phone and applying the classifier that uses this information results in the most accurate activity tracking (Figure 5A).

However, it is often the case that we do not know how the phone is being worn in everyday life. If an activity classifier is trained using only one phone location but tested on all four locations (location-assumed), we predict activities with only 56.8% accuracy (Figure 4A). Training with data from the bag location resulted in the highest activity tracking accuracy (69.8%), followed by the pocket (63.4%), the hand (48.7%), and the belt (45.2%). Assuming someone wears the phone in only one location results in poor tracking accuracy if the person wears the phone in other locations (Figure 5B).

Even if we do not know how the phone will be worn, we can collect data from all potential phone locations to train the activity classifiers with all locations simultaneously (global model). By using this method, we can predict activities with 88.1% accuracy (Figure 4A). The best tracking accuracy was while the phone was in the pocket (94.2%), followed by the belt (93.1%), the bag (89.1%), and the hand (79.0%). The transition accuracy for activities was 81.5%. The pocket achieved the highest transition accuracy (91.6%), followed by the belt (90.9%), the bag (84.4), and the hand (70.1%). When we collect data from all potential locations we can track activities almost as well as when we know the location of the phone.

3.2 Location tracking

We trained and tested location classifiers with data from the same activity (activity-known) while the phone was worn in different locations. We found an overall location tracking accuracy of 94.4% (Figure 4B); walking achieved the highest accuracy (99.5%), followed by sitting (99.4%), standing (94.3%), stand to sit transitions (61.5%), and sit to stand transitions (48.5%). If we know the activity being performed, we can accurately predict the phone location in non-transitory states (Figure 5D).

When the location classifiers are trained using one activity and then tested over all activities (activity-assumed), the location tracking accuracy is 72.0% (Figure 4B). Walking achieved the highest location tracking accuracy (77.2%), followed by sit to stand transitions (74.7%), stand to sit transitions (72.8%), standing (70.0%), and sitting (65.4%). Assuming the person is performing only one activity, when they perform several activities reduces the location tracking accuracy as compared to knowing the activity (Figure 5E).

To account for many activities, we can train our location classifiers using data from all of the potential activities (global model). We track the phone location with 95.0% accuracy using this method (Figure 4B). Again, walking had the highest tracking accuracy (99.6%), followed by sitting (97.1%), then standing (94.5%), sit to stand transitions (89.9%), and then stand to sit transitions (87.3%). The transition accuracy for phone locations was 82.1%.

3.3 Joint Activity/Location Classifier

The previous approaches use a location classifier to predict locations, and an activity classifier to predict activities. Here, we combine these approaches to create a joint activity-location classifier that simultaneously predicts both the phone location and activity. The accuracy of correctly predicting both the activity and phone location was 87.1%. This method correctly classified 88.0% of activities (Figure 5C) and 96.4% of phone locations (Figure 5F). The transition accuracy was 77.4% for activities, 84.1% for locations, and 76.0% for both activities and locations combined. This combined approach yields tracking accuracies for each activity and location that are comparable to the individual approaches (Figure 4A, B).

3.4 Subject-wise cross validation

We conducted the same analysis as the previous sections using subject-wise cross validation (Figure 6). There is a decrease in SVM+HMM accuracy across all analyses as compared to the four-fold cross validation. When the classifier is trained and tested with data from the

same location, the pocket performs the best (91.2%), followed by the belt (89.0%), the hand (66.0%), and the bag (44.0%). When the classifier is trained with data from one location and tested with data from all locations, the pocket performs the best (61.8%), followed by the bag (53.9%), the hand (44.5%), and the belt (41.1%). When the classifier is trained and tested with data from all locations, the belt performs the best (87.5%), followed by the pocket (79.7%), the hand (66.8%), and the bag (48.6%). For the joint classifier, the belt performs is best (80.9%), followed by the pocket (73.2%), the hand (69.0%), and the bag (39.0%).

3.5 Pilot application to patients with Parkinson's disease

Two subjects with Parkinson's disease carried a phone while walking, sitting, or standing and we used this data to evaluate our joint SVM+HMM method. The activity tracking accuracy was 94.4% and the location tracking accuracy was 98.1% (Figure 7), and the transition accuracy was 97.7%. These preliminary results show that the joint SVM+HMM method can be easily extended to a new population with various activity and location sets.

4. Discussion

We tracked subjects' activities and how they were wearing the phone using SVMs, HMMs, and the idea that activities change rarely and the phone location changes even less often. We found that knowing the location of the phone results in the best activity tracking accuracy. Assuming the phone is worn in one location when its location is not known results in the worst performance. We built a classifier that simultaneously predicts both the activity and location of the phone. Finally, we successfully applied this method to track activities for two people with Parkinson's disease.

Numerous activity tracking studies have been able to accurately track activities using a variety of classification methods (Preece et al., 2009; Ravi et al., 2005). The use of smartphones has made activity tracking systems more convenient (Brezmes et al., 2009; Ravi et al., 2005; Yang, 2009). We found that augmenting SVM classifiers with an HMM improves accuracy as compared to using only an SVM, which agrees with previous work on hybrid models in activity recognition (Lester et al., 2006; Lester et al., 2005; Suutala et al., 2007). The main innovation of our approach is that we use the HMM to simultaneously predict activities and locations, and we incorporate the concept that activities change rarely, and the phone location changes even less often. Additionally, an HMM can smooth out misclassifications due to temporary irregularities from spastic movements or brief bouts of tremor that may otherwise deceive an SVM classifier from correctly inferring activities. Therefore, a hybrid SVM+HMM provides a more robust approach for people with movement disorders.

Orientation independent features have been used to improve classification accuracy when the phone is worn in different ways. This method estimates the direction of gravity using the accelerometer signal and, in turn, the orientation of the phone. This enables estimation of the vertical and horizontal components of the user's motion, which can be used to calculate orientation independent features (Mizell, 2003). This method has been applied successfully to activity tracking while the phone is in a pocket (Sun et al., 2010). However, this method

assumes the phone is still always in the same location. Accelerations differ if the phone is placed in a belt clip, hand, or bag so this approach would not be expected to apply to new locations.

Knowing the location of a phone is a useful capability. Depending on the location of the phone on the body or the activity of the user, phone displays and sensors can respond differently (e.g. turn off if not wearing). The location of a sensing device has been tracked during walking using only an accelerometer (Kunze et al., 2005), and a microphone was able to determine whether or not a phone is in the pocket (Miluzzo et al., 2010). We found that walking produced the best location tracking results compared to any other activity. By predicting the location of the phone with near perfect accuracy, we can provide an estimate of the error associated with our activity tracking. We show that location tracking, or context awareness, is easily incorporated into our activity tracking algorithm. Although this study evaluated activity tracking in mobile phones across multiple locations, we included only five activities. This was to focus on the most common daily activities for people with movement disorders. Even though only five activities were included, we still see a significant reduction in tracking accuracy if all potential phone locations are not taken into account. Further studies are needed to examine the effect of increasing the number of activities across multiple potential phone locations. Each subject followed the same predetermined sequence for all locations and was assigned the same transition matrix, which may not be representative of that individual's day to day activities. Ideally, the probability of these transitions would be known from natural behavior, but our supplied transition matrix was used in absence of such information. If the probability of such transitions is better understood, this could be extended to a semi-Markov approach, where the probability of changing activities is dependent on how long a person has been engaged in a given activity. For example, it may be difficult for a patient to walk. The longer they are in a walking state, the more likely it is for them to stop walking.

Incorporating additional features may improve tracking accuracy. Anthropomorphic parameters such as height, weight, and leg length may affect a person's gait or posture, thereby influencing the accelerations recorded and tracking accuracy. Assistive devices such as walkers and canes may also have an impact on tracking accuracy if the patients gait changes, or they lean or change posture with the assistive device. Furthermore, patterns in patient activities could be influenced by the time of day, day of the week, the season or weather. However, in this study we have limited the features used to those derived directly from the accelerometer signal, and see the use of additional features as a possible contribution in future work.

We contribute an analysis demonstrating the effect of multiple phone locations on activity tracking accuracies. If it is possible to control the phone location, we recommend placing the phone in the pocket or on a belt because these locations achieved the highest activity tracking accuracies for sitting, standing, and walking, when tested on new subjects. A joint activity/location SVM+HMM can simultaneously track both the activity and location of the phone with high accuracy for both healthy subjects and Parkinson's patients.

5. Conclusions

Mobile phones are not worn in the same location throughout a person's daily activities, and it is important that activity tracking platforms take this into account. If a person wears their phone in multiple locations, we recommend collecting data for every activity in each location. Additionally, we have shown that activity tracking and location tracking can be combined into a joint classifier. Most importantly, this work better enables people to carry their phones more naturally.

Acknowledgments

We would like to thank Petra Conaway, Kevin Greene, Aaron Yang, and Andrew Cichowski. This work was supported by the National Parkinson Foundation and the U.S. National Institutes of Health under grants P01NS044393, R01NS063399, and T32EB009406. This work was also supported by the Washington Square Health Foundation.

References

1. Albert MV, Kording K, Herrmann M, Jayaraman A. Fall Classification by Machine Learning Using Mobile Phones. *PloS one*. 2012a; 7
2. Albert MV, Toledo S, Shapiro M, Kording K. Using mobile phones for activity recognition in Parkinson's patients. *Frontiers in neurology*. 2012b; 3:158. [PubMed: 23162528]
3. Bao L, Intille SS. Activity recognition from user-annotated acceleration data. *Lect Notes Comput Sc*. 2004; 3001:1–17.
4. Bieber G, Voskamp J, Urban B. Activity Recognition for Everyday Life on Mobile Phones. *Universal Access in Human-Computer Interaction, Pt II, Proceedings*. 2009; 5615:289–296.
5. Brezmes T, Gorricho JL, Cotrina J. Activity Recognition from accelerometer data on mobile phones. *IWANN '09: Proceedings of the 10th International Work Conference on Artificial Neural Networks*. 2009:796–799.
6. Chang C-C, Lin C-J. Libsvm. *ACM Transactions on Intelligent Systems and Technology*. 2011; 2:1–27.
7. Cui YQ, Chipchase J, Ichikawa F. A cross culture study on phone carrying and physical personalization. *Usability and Internationalization, Pt 1, Proceedings*. 2007; 4559:483–492.
8. Fernandes HL, Albert MV, Kording KP. Measuring generalization of visuomotor perturbations in wrist movements using mobile phones. *PloS one*. 2011; 6:e20290. [PubMed: 21629659]
9. Ganapathiraju A, Hamaker J, Picone J. Hybrid SVM/HMM architectures for speech recognition. *INTERSPEECH*. Citeseer. 2000:504–507.
10. Jeannot PY, Aminian K, Bloetzer C, Najafi B, Paraschiv-Ionescu A. Continuous monitoring and quantification of multiple parameters of daily physical activity in ambulatory Duchenne muscular dystrophy patients. *Eur J Paediatr Neuro*. 2011; 15:40–47.
11. Kangas M, Konttila A, Lindgren P, Winblad I, Jamsa T. Comparison of low-complexity fall detection, algorithms for body attached accelerometers. *Gait Posture*. 2008; 28:285–291. [PubMed: 18294851]
12. Kunze K, Lukowicz P, Junker H, Troster G. Where am I: Recognizing on-body positions of wearable sensors. *Location- and Context-Awareness, Proceedings*. 2005; 3479:264–275.
13. Kwapisz JR, Weiss GM, Moore SA. Activity recognition using cell phone accelerometers. *4th ACM SIGKDD International Workshop on Knowledge Discovery from Sensor Data*. 2010:1.
14. Lau, SL.; David, K. *Proceedings of Future Network & MobileSummit*. Florence, Italy: 2010. Movement recognition using the accelerometer in smartphones. 2010.
15. Lee RYW, Carlisle AJ. Detection of falls using accelerometers and mobile phone technology. *Age Ageing*. 2011; 40:690–696. [PubMed: 21596711]
16. Lester J, Choudhury T, Borriello G. A practical approach to recognizing physical activities. *Lect Notes Comput Sc*. 2006; 3968:1–16.

17. Lester J, Choudhury T, Kern N, Borriello G, Hannaford B. A Hybrid discriminative/Generative Approach for Modeling Human Activities. 19th International Joint Conference on Artificial Intelligence (IJcai-05). 2005:766–772.
18. Mathie MJ, Coster AC, Lovell NH, Celler BG. Accelerometry: providing an integrated, practical method for long-term, ambulatory monitoring of human movement. *Physiological measurement*. 2004; 25:R1–R20. [PubMed: 15132305]
19. Miluzzo E, Papandrea M, Lane ND, Lu H, Campbell AT. Pocket, bag, hand, etc. - automatically detecting phone context through discovery. *Proc. of the First International Workshop on Sensing for App Phones (Phonesense '10)*. 2010
20. Mizell D. Using gravity to estimate accelerometer orientation. *Seventh IEEE International Symposium on Wearable Computers, Proceedings*. 2003:252–253.
21. Murphy K, Dunham M. pmtk3 - probabilistic modeling toolkit for matlab/octave version 3. 2011
22. Preece SJ, Goulermas JY, Kenney LP, Howard D, Meijer K, Crompton R. Activity identification using body mounted sensors--a review of classification techniques. *Physiological Measurements*. 2009; 4:1–33.
23. Ravi N, Dandekar N, Mysore P, Littman ML. Activity recognition from accelerometer data. *Association for the Advancement of Artificial Intelligence*. 2005:1541–1546.
24. Salarian A, Russmann H, Vingerhoets FJG, Burkhard PR, Aminian K. Ambulatory monitoring of physical activities in patients with Parkinson's disease. *Ieee T Bio-Med Eng*. 2007; 54:2296–2299.
25. Stone AA, Shiffman S, Schwartz JE, Broderick JE, Hufford MR. Patient compliance with paper and electronic diaries. *Controlled clinical trials*. 2003; 24:182–199. [PubMed: 12689739]
26. Sun L, Zhang DQ, Li B, Guo B, Li SJ. Activity Recognition on an Accelerometer Embedded Mobile Phone with Varying Positions and Orientations. *Ubiquitous Intelligence and Computing*. 2010; 6406:548–562.
27. Suutala J, Pirttikangas S, Rönning J. Discriminative temporal smoothing for activity recognition from wearable sensors. *Ubiquitous Computing Systems, Proceedings*. 2007; 4836:182–195.
28. Yang J. Toward physical activity diary: motion recognition using simple acceleration features with mobile phones. *1st international workshop on interactive multimedia for consumer electronics: Beijing, China*. 2009

Highlights

We use the output of a support vector machine as input to a hidden Markov model to improve accuracy.

We make simultaneous predictions of a person's activity and how the phone is worn.

Method successfully applied to people with Parkinson's disease.

Wearing the phone in a pocket or belt results in the most accurate activity tracking.

Walking is the best activity to predict how the phone is worn.

Figure 1. Experimental protocol and analysis

A) Subjects carried a mobile phone in their pocket, belt, hand, or bag and B) performed a predetermined sequence of activities. After performing one full sequence of activities, the subject moved the phone to the next phone location. C) Data processing pipeline.

Accelerometer data was divided into clips that were two seconds long and features were calculated for each clip. The SVM used the features to make probabilistic estimates of the activity and/or location. These estimates were then smoothed over time using an HMM to predict the most likely activity and/or location.

Figure 2. HMM transition matrices

The SVM processes accelerometer features (Table 1) and outputs a probabilistic estimate of the state of the person. The transition matrix for the HMM gives the Prob(current hidden state | previous hidden state). A) The transition matrix for the activity-known and location-known models. B) The transition matrix for the joint activity-location classifier. The joint matrix is in a block diagonal form analogous to previous models with only slight variations in the parameters.

Figure 3. Raw data clips

Example data clips for each activity location pair for one subject. Note the accelerations recorded by the phone are different for each activity across the four phone locations. The blue, red and green lines correspond to the X, Y, and Z axes of the phone's accelerometer, respectively.

Figure 4. Four-fold cross validation

A) Activity tracking accuracy. The location-known analysis builds classifiers using clips from one phone location, and tests each classifier with clips only from its respective location. The location-assumed analysis builds classifiers using clips from one phone location, but tests each classifier with clips from all four locations. The global model builds a classifier with clips from all the phone locations, and tests the classifier with clips from all locations. The joint classifier is built and tested using activity-location pairs; the accuracies shown here are marginalized over each phone location. B) Location tracking accuracy. Each analysis is analogous to the activity tracking, except the classifiers are now trained and tested to predict the location of the phone.

Figure 5. Classification matrices

for A) activity-known B) activity-unknown C) joint classifier marginalized over location D) location-known E) location-unknown and F) joint classifier marginalized over activities. Each entry in the matrices is shaded proportionally to its number, divided by the sum of all entries in that same row. A one hundred percent accurate matrix is black along the diagonal and white elsewhere.

Figure 6. Subject-wise cross validation for activity tracking

Classifiers are trained on 11 subjects and then tested on the one remaining subject. The accuracies shown here are marginalized over phone location.

Figure 7. Method applied to a person with Parkinson's disease

Actual (blue) and predicted (orange) activity-location pairs for one subject with Parkinson's disease using a joint SVM+HMM classifier.

Table 1**Features Used for Activity Recognition**

Description	Total Number of Values
Mean, absolute value of the mean	6
Moments: standard deviation, skew, kurtosis	9
For the instantaneous change: mean, std dev, skew, kurtosis	12
Root mean square	3
Smoothed root mean square (5 pt kernel, 10 pt kernel)	6
Extremes: min, max, abs min, abs max	12
Histogram: includes counts for -3 to 3 z-score bins	21
Fourier components: 8 samples for each axis	24
Overall mean acceleration	1
Cross product means, xy, xz, yz	3
Abs mean of the cross products	3
Normalized mean cross products	3
Abs normalized mean cross products	3