2011 IEEE International Conference on Pervasive Computing and Communications (PerCom), Seattle (March 21-25, 2011)

On-body Device Localization for Health and Medical Monitoring Applications

Alireza Vahdatpour, Navid Amini, Majid Sarrafzadeh
Computer Science Department
University of California, Los Angeles
{alireza, amini, majid}@cs.ucla.edu

***Abstract*—We present a technique to discover the position of sensors on the human body. Automatic on-body device localization ensures correctness and accuracy of measurements in health and medical monitoring systems. In addition, it provides opportunities to improve the performance and usability of ubiquitous devices. Our technique uses accelerometers to capture motion data to estimate the location of the device on the user's body, using mixed supervised and unsupervised time series analysis methods. We have evaluated our technique with extensive experiments on 25 subjects. On average, our technique achieves 89% accuracy in estimating the location of devices on the body.**

***Keywords*-On-body device localization, Unsupervised activity discovery, Motion analysis**

I. Introduction

Recent advances in embedded systems technology have resulted in utilization of wearable and non-intrusive systems for remote health and activity monitoring. Pedometers and smart phones with motion sensors are among devices that are used to continuously monitor life characteristics such as the daily activity and exercise. With the production of ultra low power sensor interfaces such as Sensium [3], it is expected that smart band aids become widely used for health and medical monitoring applications. Smart band aids can be used for monitoring several metrics such as the daily energy expenditure, body temperature, skin moisture, heart rate, and other human vital signs. The proliferation of these wearable devices results in higher diversity of their usage. In addition, being non-intrusive in daily activities requires these devices to be adaptable to individual users habits. There are at least four reasons why it is necessary or beneficial to automatically discover the location and the placement of these devices on the user's body.

1) Correctness: Wearable non-intrusive embedded systems are used to remotely and ubiquitously monitor subjects conditions such as basal body temperature, skin moisture, and ultraviolet exposure. Correct functionality of these monitoring devices highly depends on the correct placement of them on the subject's body. For example, the following table summarizes the normal temperature range variation for normal adults according to the location of the thermometer on the body [8]:

Location	Ear	Axilla	Rectal	Oral
Range ($^\circ C$)	35.5-38	35.5-37	34.5-38	33-38

Table I
Normal temperature readings for different regions on the body

It is clear that the temperature varies in different body areas. If in each use a user attaches his temperature monitoring device to different regions on the body, clinicians and caregivers will be misguided by the data variation. Therefore, it is essential for the clinicians to know the placement of the sensing systems on the body, so that the data collected from the user becomes valuable and its analysis results in correct hypotheses.

2) Accuracy: The operation of motion and energy expenditure monitoring systems heavily relies on the placement and orientation of motion sensors, which usually include accelerometers and gyros. The functionality of pedometers is an example of the relation between the accuracy and the device placement. The step counting accuracy changes if the pedometer is attached to anywhere other than the waist, because the user's movements will be projected differently on the accelerometer. In an observation, we attached 6 pedometers to different regions on a subject body and asked him to take 220 steps. The following table shows how the location of the device greatly impacts the number of steps counted by each pedometer:

It is clear that in a location-aware pedometer, a simple dynamic threshold control technique can dramatically enhance the accuracy of the step counting. Energy expenditure estimation via motion sensing [13] is another application where the outcome accuracy highly depends on the accurate classification of the type of the movement performed by the user (e.g., walking, jogging, running, jumping, etc.) and the placement of the accelerometer on the body [6]. If the user wears an accelerometer on his/her foot instead of the waist, the accuracy of the activity detection degrades dramatically. This is because the activity detection models implemented in the device are tuned for waist motions. As a result of wrong activity detection, the system will overestimate or underestimate the user's caloric expenditure.

3) Communication Optimization: It is known that the human body affects the performance of radio transceivers. Shah et al. [16] have studied the performance of IEEE 802.15.4 radios when they are placed on different regions

Pedometer location	waist	shin	thigh	forearm	upper arm	chest
Steps counted	230	238	148	181	177	0
Error ($\frac{real-est.}{real}$)	4%	8%	32%	17%	19%	100%

Table II
COUNTED STEPS VS. LOCATION OF THE PEDOMETER

of the human body. According to the study, the wireless communication loss rate highly depends on the location of the sensor. Automatic on-body localization enables devices to control the communication power according to the anticipated channel quality.

In addition, study in [14] evaluated the impact of the human body and where sensors are located on the body on the transmission range of 802.15.4 transmitters. It shows that particularly in outdoor experiments *the body factor* halves the transmission range of a device.

4) Usability and Adaptability: Pervasive use of monitoring systems requires enabling the users to change sensors location and placement according to their life style, clothing, and activities. Not only does the on-body device localization enhance the usability of these devices, but also it provides the opportunity for software developers to design smarter devices, which can automatically adapt to users and their environment. Automatic device localization allows the smart phones to rapidly and autonomously adapt to users needs. Different placements of the device on the body can be interpreted as the users desire for an application or a major change in the environment. For example, as soon as the user mounts the smart phone on her upper arm, the smart phone should activate its gym/music mode.

In this paper, we propose a technique for automatic on-body device localization using acceleration data. Accelerometer is one of the most widely used types of motion sensors, which is used for variety of applications such as device orientation detection, game controlling, shock protection, and activity discovery. Given that accelerometers are relatively cheap and consume low energy, nowadays they are embedded in most of the high-end and portable systems. Our developed technique allows both on-line and off-line discovery of the device location on the body. Our technique consists of two main phases. Unsupervised activity discovery and Supervised location estimation. In the rest of this paper, we first review the related work. Later, the description of the technique is presented. Next, we present the experimental results, and finally the last section provides the concluding remarks.

II. RELATED WORK

The problem of localization on the human body has been addressed by a few number of researchers before. The most well known study is done by Kunze in [11] and [10]. While these studies address the problem of localizing devices on the body, the results of their classification is limited to four exact locations (wrist, breast pocket, trousers pocket, and right eye). As denoted by the authors, the locations that have been chosen represent typical location of appliances and accessories. In contrast, our approach is designed to discover device location on any region on the body. As a result not is it only applicable in in accessories and wearable systems (with higher diversity in discovering the location) but also it can be leveraged in smart band aids and electronic implants.

A number of studies in mobile and smart phone area have also addressed simplistic device localization problem, where for example it is desired to detect whether a cellphone is on the table or in the pocket [9].

Accelerometer data has also been used in other studies to improve the utility of portable devices. For example, the study by Lester et al. [12] uses motion data to determine if two portable devices are carried by the same person.

Supervised techniques for activity detection using on-body sensors have been studied widely in the last decade [5]. All of these studies leverage body-fixed sensors and predefined models for activities. In the recent years, a few studies have addressed unsupervised classification of human activities [15], [17], [18]. We use the unsupervised most frequent activity detection technique proposed in [18] as part of our technique for on-body device localization.

III. OVERALL APPROACH

The essential idea behind the on-body localization is to analyze the acceleration data when the user is performing a specific activity. On the other hand, requiring the subjects to perform a specific activity will degrade the practicality of the solution and makes it user dependent. The study by Welker et al. [19] shows that on average people take 8265 steps per day (performing only normal daily activities). Such a large number of the steps clearly shows that walking is the most frequent and consistent activity people perform throughout the day (excluding non active states such as sleeping and sitting). Hence, the walking activity could be used as a characterizing action for our device localization technique. However, since the location, placement, and orientation of the accelerometer on the body is not known a priori (and may change in each use), the projection of the body movement patterns on the sensor is not predictable. Hence, it is impossible to define a generic model for detecting the walking activity which is valid for all placements and locations of the accelerometer. Because of the same reason, all of the studies in activity detection assume that the location of

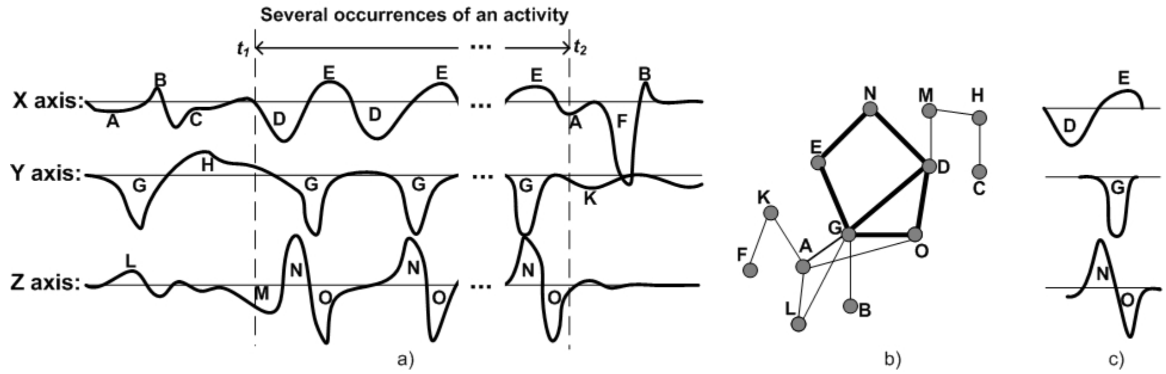

Figure 1. An illustrating example of the clustering algorithm for the recurrent activity discovery; a) three dimensions of accelerometer data, with several occurrences of an activity between t_1 and t_2. b) The interval coincidence graph, representing the primitives in the time series data. The thickness of edges show higher coincidence between the primitives. c) Primitives with high coincidence are clustered and the recurrent activity is discovered. t_1 and t_2 are sent to output as the start and end time of the most frequent activity.

the accelerometer is known and is restricted to one or a few places. In this study, we use an unsupervised activity discovery method to discover time intervals that include walking patterns in long time series data. An important advantage of unsupervised techniques is that these methods derive the activity models during the run-time, and hence they do not require predefining activity models and patterns in the setup phase [17].

Our unsupervised technique discovers walking activity occurrences based on their frequency and consistency during long time intervals. Once the time series subsequences representing walking activity are discovered, they are analyzed with a supervised classification technique. We use a support vector machine (SVM) to estimate the location of the device based on the frequency and time domain properties of the time series subsequences that represent the walking activity. In the following, we discuss each phase of the technique separately.

A. *Most frequent activity discovery*

Discovering most frequent activities has been in the interest of the research community in the recent years. Motif detection in time series data is a technique to discover the frequent subsequences in a longer time series [7]. A recent study by Vahdatpour et al. [17] has extended the motif discovery mechanism to multi-dimensional time series data for discovering the most frequent activities in wearable systems. However, as the study by Chiu et al. [7] suggests, motif discovery is computationally intensive and is not suitable for embedded computing and applications with a huge amount of data. Hence, in this paper, we use an approach similar to the method presented in [18] for discovering the most frequent activity (walking) in 3 -dimensional accelerometer data. Our approach for most frequent activity discovery consists of two steps. In the first step, each dimension of the accelerometer data is analyzed separately to discover frequent patterns. In contrast to [17], we use a binning approach to classify similar patterns in single dimensional accelerometer data. In the second step, classified patterns are grouped together with a clustering approach to form the activities. Next, we describe each step of the frequent activity discovery method:

1) Activity primitive discovery in single-dimensional time series: Numerous approaches have been used to discover similar patterns in a time series data. Although using motif discovery results in accurate classification of patterns, it is computationally intensive and sensitive to noise in the data. Here, we classify patterns in single dimensional accelerometer data based on the total movement that is resulted from the activity in the corresponding axis. We define an activity primitive to be a subsequence between two consecutive stable regions (or default sensor value) in the accelerometer time series data. Such subsequences are mostly observed as bell shape patterns (with local minimum or maximum points) or linear slopes between two stable values in human actions, which depending on the speed and nature of the activity have different length and frequency. The acceleration data can be converted to displacement information, using the following equation on a subsequence of length m $(a_1 a_2 .. a_m)$.

$$displacement = D \cdot \sum_{k=1}^{m} \sum_{i=1}^{k} |a_i - a_1|$$

where D is the calibration constant factor, which depends on the placement of the sensor and other parameters such as the frequency of data sampling. Although the location of the sensor in not known in our application, since displacement information is used to compare patterns in the same accelerometer, it is constant for all patterns and can be eliminated from the calculations. Figure 1.a depicts three accelerometers and several activity primitives classified in them (labeled by alphabets).

Algorithm 1 presents the method for extracting primitives of activity subsequences projected on single-dimensional time series. The algorithm extracts the patterns from the time series, classifies them based on the physical attribute of the data and according to the discretization cardinality, and assigns a symbol to them. It also assigns starting time to each symbol. There are variety of methods in literature regarding effective discretization, especially in time series. Here, in order to discretize the calculated $displacement$ and assign symbols to it, we consider using the probability distribution of the calculated $displacement$s and assigning symbols to primitives such that primitives are classified fairly considering the variation of $displacement$ in the application runtime.

Algorithm 1 Single-dimensional activity primitive extraction

```
Input: Time series T = t_1..t_n
Output: List of activity primitives along with their occurrences location p_occurrence(1..j), start(1..j)
i ← 1
while true do
    while t_i = default_value | t_i is stable region do
        increment i
    end while
    start(j) ← i
    while t_i ≠ default_value & t_i is not stable region do
        d ← d + Σ_{i=1}^{start(j)} t_i − t_{start(j)}
    end while
    length(j) ← i − start(j)
    displacement(j) ← D · d
    d ← 0
    if i = n then
        p_occurrence(1..j) ←
                discretize(displacement(1..j), length(1..j))
        return p_occurrence(1..j), start(1..j)
    end if
    j ← j + 1
end while
```

Note that in contrast to motif discovery method is used in [17] for detecting recurrent patterns in time series, this algorithm leverages domain specific attributes of data (here, being the motion sensing data, which is used frequently in wearable systems). Regardless of the method used to discover recurrent activity patterns in the single-dimensional time series, the output of this phase is passed to the multi-dimensional activity and abnormality detection algorithm, which is described in the next section.

2) Activity in multi-dimensional time series: By using 3-dimensional accelerometers, activity motions are projected on one, two, or three dimensions of the accelerometer. Since the original orientation and placement of the accelerometer is unknown, all dimensions should be considered together in order to discover recurrent activity. If only a subset of dimensions is used for counting the occurrence of activities, activities with similar projection will be classified together and over counted. As a result, we first use a graph clustering approach to construct activity structures in multi-dimensional data. This step basically groups several activity primitives from different dimensions together according to their temporal characteristics.

In order to construct the multi-dimensional activity structures, we first convert the list of discovered activity primitives into a weighted directed graph, so that a graph clustering mechanism can be applied to construct model of activities in multi-dimensional data. In the proposed graph, each vertex represents an activity primitive, and the weight on the vertices represents the number of occurrences of the primitive in the corresponding time series. The weight on each directed edge in the graph is calculated by the following equation:

$$e(i,j) = \frac{coincidence(i,j)}{total\_occurrences(j)}$$

The $coincidence(i, j)$ denotes the number of times there is temporal overlap between occurrences of primitives i and j. As a result, $e(i, j)$ will be at most 1, when all the occurrences of primitive j have overlap with occurrences of primitive i, and it is at least 0, when there is no overlap between occurrences of primitives i and j.

Next, a graph clustering approach is used to construct multi-dimensional activities structure from the primitive coincidence graph. Algorithm 2 presents our method for constructing the activity structures from the coincidence graph. This method is similar to the clustering mechanisms proposed in [4]. Clustering the primitives starts with sorting the list of vertices based on the number of occurrences of the primitives. Then, the most frequent primitive is selected as a candidate core for activity structure, and the graph is searched for primitives with high occurrence correlation with the candidate core primitive. If there are several primitives with equal frequency, the one with the larger number of highly correlated neighbors is selected as the core primitive. The threshold for the correlation is set to $1 - \beta$, where β is the abnormality frequency constant. Upon construction of an activity by clustering correlated primitives, selected primitives are removed from the graph, and the algorithm continues for discovering new activity structures by selecting a new core activity primitive.

Figure 1.c illustrates the results of applying the algorithm to the coincidence graph in Figure 1.b (the weights in the graph are omitted to make the figure readable). As it is denoted in the figure, the activity structure extracted contains primitives E, D, G, O, and N (G is the activity core primitive, due to its frequent number of occurrences).

Algorithm 2 Activity structure construction from the primitives coincidence graph

Input: $G(V, E)$ the primitives coincidence graph, $total\_occurrence(v)$ for all vertices and $e(v_1, v_2)$ for all directed edges
Output: $S = \{S_i\}$ The set of most recurrent activity structures in the time series, (each S_i is an activity)
Sort the vertices list in G in descending order of $total\_occurrence$
$\{$if some primitives have equal frequency, then the primitive with the highest number of highly correlated neighbors has the highest priority$\}$
for each vertex v_k in the sorted list of vertices **do**
    Add v_k to activity S_i
    for all neighbors of v_k, if $e_{(j,k)} > 1 - \beta$ **do**
        Add v_j to activity S_i
        remove v_j from the graph
    end for
    $i \leftarrow i + 1$
    Update the sorted list of vertices
end for
return $S = \{S_1, ..., S_{i-1}\}$

Note that the effective construction of activity structures highly depends on the frequency of primitives, and in a real world application the frequency and coincidences are significantly higher (β will be set to less than 0.1 in real world applications). The depicted graph is formed from a set of short time series and is selected to be small to better illustrate the logic behind the algorithm.

For detailed description and evaluation of unsupervised multi-dimensional activity detection algorithm, we refer the reader to [17]. However, it should be noted here that the execution overhead of this method is linearly correlated to the execution overhead of single dimensional motif discovery algorithm. The computational intensity of algorithms for motif discovery can be reduced by introducing assumptions about the input signals such as the minimum and maximum length of subsequences and the frequency of activities.

B. Location discovery via SVM

Once the walking interval is discovered via the unsupervised method, the walking subsequence is analyzed to discover the location of the accelerometer on the body. We extract several frequency and time domain features from the walking subsequences. Frequency domain features are analyzed to evaluate the impact of each step on the accelerometer. In each step, the energy of the impact between the foot and the floor is distributed in the body. The closer the sensor is to the foot (the origin of the strike in each step), the stronger is the energy sensed by the accelerometer. Time domain features are evaluated to analyze the motion range of the sensor in all three directions (we are using tri-axial accelerometer). It is clear that some limbs (e.g., forearms) move more freely than other limbs (e.g., chest).

Activity	Feature	Characteristic
Walking	Max energy in power spectrum (D)	Impact of strides
Walking	Sum of energy in power spectrum (F)	on acceleration
Walking	Maximum of Max amplitudes for all axes ($A > B, A > C$)	Motion range
Walking	Ratio of amplitude in different axes ($\frac{A}{C}$, $\frac{A}{B}$)	Degree of freedom
Non-walk	Number of orientation variations	in movement

Table III
THE INTUITION BEHIND THE SIX MAJOR EXTRACTED FEATURES. LETTERS IN PARENTHESIS REPRESENT THE FEATURES IN FIGURE 2

In addition to analyzing the walking patterns, the time series subsequences representing general activity (non-walking) were also analyzed. In this regard, we designed an algorithm to count the number of times the orientation of the device (body part) changes over time. The algorithm first approximates the orientation of the sensor in each 10 second interval, via calculating the DC value of the Fourier transform. Then, the total number of orientation variations is calculated for a long time series by comparing the consecutive approximate orientations. Table III and Figure 2 summarize six features extracted from walking and non-walking subsequences and the intuition behind using them for location classification. Figure 2.a depicts a time series representing multiple walking steps in a 3 second interval (captured on the forearm). The power spectrum depicted in Figure 2.b is the sum of the frequency power spectrum of the 3-dimensions. As illustrated, feature A is the maximum amplitude among all the dimensions. $\frac{A}{C}$ and $\frac{A}{B}$ are the ratio of the maximum amplitudes in different axes. D and F represent the maximum energy (which is at 1.85 Hz) and the overall energy captured by the accelerometer respectively. Note that although Table III presents 6 features, we used 17 features in our SVM, which are the variations of the presented features in the table.

It is noteworthy that for each placement of the sensor on the body, the relation between coordinate systems of the sensor and the subject's body varies. Hence, no assumption can be made during the feature extraction regarding the initial orientation of the sensor. As shown in Table III, features either represent an accumulation value calculated among all the directions or depict the maximum of a property among the three dimensions.

Figure 3 illustrates the variation of features A (maximum amplitude), D (maximum energy), and F (sum of the energy) for a subject, when the accelerometer is mounted on different regions on his body. While as a result of strides variation, the features span over the normalized y axis, however, it is clear that different locations exhibit different ranges for each

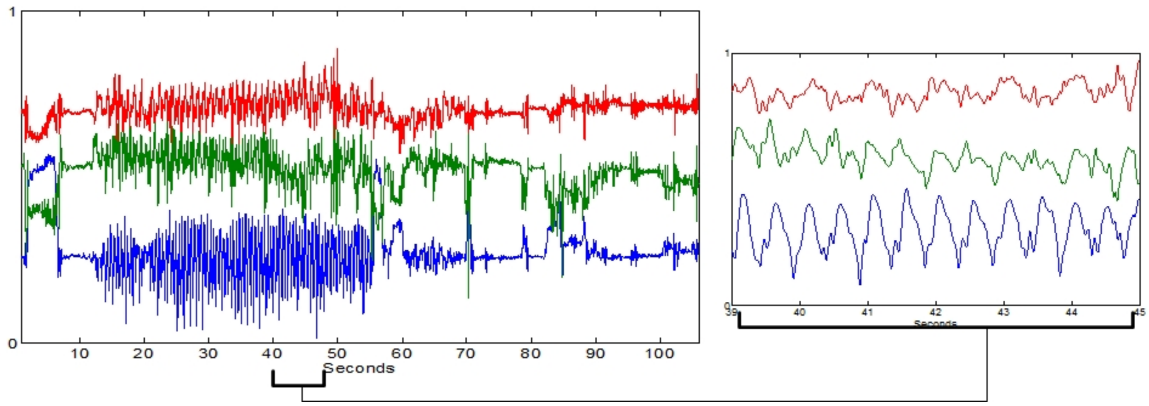

Figure 4. Acceleration data collected from a sensor positioned on the arm during a 105 seconds interval (left). The magnified interval (right) depicts several occurrences of a walking activity. Y axis is normalized and X, Y, and Z acceleration dimensions are shifted to enhance the presentation.

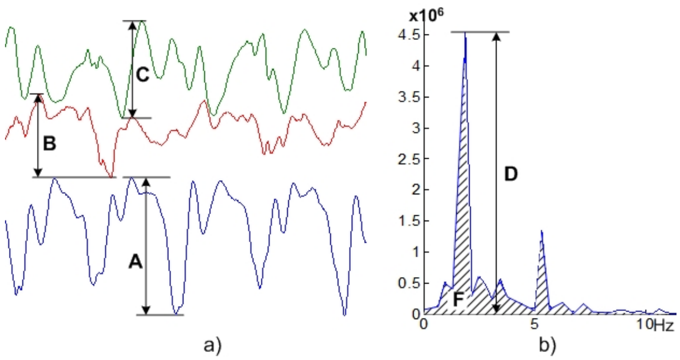

Figure 2. a) Accelerometer time series captured during walking in a sensor mounted on the forearm (3 second interval). b) Power spectrum.

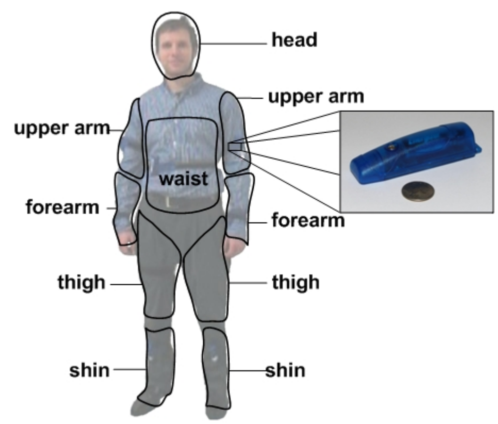

Figure 5. Sensors are placed on 10 different regions on the body. The localization algorithm has 6 different outputs: {forearm, upper arm, head, thigh, shin, waist}.

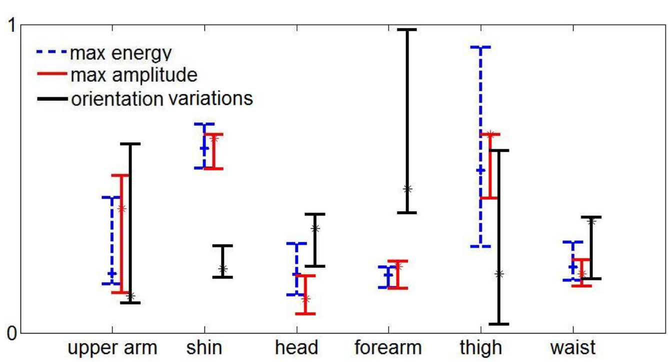

Figure 3. Variation of three features for a subject in different sensor locations. Data is captured in a 5 minute time frame. Y axis is normalized. Blue is the maximum energy of the power spectrum, red is the maximum amplitude among all direction, and green is the average number of orientation variations in a 5 minute time frame.

feature. In the next section, we evaluate the performance of SVM in classifying the location of the sensors into six different regions depicted in Figure 5. Each of these regions have semi-independent movements during human activities. We performed several experiments to evaluate the effect of dividing each of these six regions into smaller areas. However, due to the motion characteristics of the human body, no meaningful statistical difference was observed between the features of the detailed areas. Hence we conducted the rest of our experiments based on these six areas on the body. In the next section, we present the experimental evaluation of our classification technique for the six regions depicted in Figure 5.

IV. Experimental Results

We used the acceleration logging device from Gulf Coast Data Concepts to conduct our experiments [1]. The tri-axial accelerometer in the device has $\pm 6g$ gain and therefore, its sensed acceleration does not saturate during normal

		Inferred location					
		upper arm	forearm	waist	shin	thigh	head
Real location	upper arm	88%	6%	3%	3%	0%	0%
	forearm	11%	85%	0%	0%	4%	0%
	waist	5%	0%	95%	0%	0%	0%
	shin	0%	0%	0%	100%	0%	0%
	thigh	0%	0%	0%	12%	88%	0%
	head	18%	10%	9%	0%	0%	63%

Table IV
DETAILED CLASSIFICATION RESULTS FOR EACH BODY LOCATION

activities.We set the sampling rate to 160Hz. The device is suitable for long intervals of data collection because it has 1 GB of memory. In addition, the size of the device is relatively small ($1in \times 3.5in \times .75in$) and does not interfere with the normal activities of the subjects.

25 subjects (20 males, age between 18 and 35) participated in our study; each subject wore 10 sensors on 10 different parts of his/her body, according to the sample setup shown in Figure 5. No instruction was given to the users regarding the exact placement and the orientation of the sensors and users were allowed to place the devices anywhere in the requested body areas. Subjects were allowed to choose to attach the sensors to the skin or to put them on their clothes. Each subject had sensors on for 30 minutes, during which he/she was asked to perform a sample set of his/her own daily activities in the university environment. We required the users to include a number of walking intervals in their activities. The collected time series data was divided into 3 minute subsequences. The total of 2500 time series samples were collected in this phase. We used an open source SVM Library (LIBSVM [2]) for implementation of the classifier. A radial basis function (RBF) was used for the SVM kernel function. Parameter selection and tuning of of the kernel function was done via cross-validation and grid-search techniques.

1) Walking interval discovery: For each sample, the walking intervals were discovered via the unsupervised technique for the most frequent activity detection. 96% of the time series samples were detected to include walking periods. This rate is high due to the fact that subjects were asked to include walking periods in their activity. The walking discovery algorithm had a consistent accuracy in detecting walking occurrences in the time series. On average 93% of all walking strides were discovered and classified correctly (versus 84% reported by [11]). In discovering the walking occurrences, the highest accuracy was in the foot area with 95% accuracy and the lowest was in hand area (89%) due to the high variation in the motion patterns. Figure 4 depicts an example of data collected from a sensor placed on a subject's arm. On average, 2% false positive detection was observed which denotes the number of non-walking activity occurrences that were labeled as walking (versus 2% reported by [11]). Since we were interested in finding relatively longer walking time intervals, the effect of false positives was eliminated as these occurrences were sparsely distributed in the time series data.

	upper arm	forearm	waist	shin	thigh	head
Precision	83%	93%	100%	88%	92%	84%
Recall	91%	100%	63%	85%	87%	94%

Table V
CLASSIFICATION PRECISION AND RECALL FOR DIFFERENT REGIONS OF THE BODY, WHEN THE TRAINING AND TESTING WERE DONE ACROSS SEVERAL SUBJECTS

2) Sensor location detection: In the first experiment, we trained and tested the localization SVM for each participant separately. Of the 100 time series samples collected from each subject, 20 samples were used for training and the remaining 80 for the testing. As expected, the classification result was near perfect, since the motion characteristics for each subject is consistent during his/her activity. The minimum, mean, and maximum overall classification precision for 25 subjects were 88%, 94%, and 100%, respectively. However, having training data from each user is not a valid assumption in real life scenarios, where the algorithm is supposed to work on off-the-shelf devices, for unknown users. Hence in the second experiment, we trained the system using 500 randomly selected samples out of 2500 time series. The average classification accuracy in this scenario was 89%. Table V presents the classification precision and recall rate for each region separately. Note that direct comparison to the results of sensor localization in [11] is not possible since the set of locations in [11] is more limited and different from ours. However, the average classification accuracy reported by authors in [11] is in the same range reported here ($\sim 90\%$) and confirms the applicability of localization methods in improving the utilization of wearable (our study) or portable ([11]) devices.

Table IV presents the details of false classifications for each area on the body. Head has the least classification accuracy, because it generally mimics the motion characteristics of the upper extremities. Classification of the shin area is performed perfectly, since the impact of each step on the shin area is significantly higher than other locations.

Finally we performed analysis to study the feasibility of classification of left limbs from right limbs (e.g., left arm vs. right arm). No meaningful classification was observed using the SVM. This is due to the fact that the sensors were mounted on different sides of the limbs by different users, and hence the accelerometer coordinate system changed in each experiment.

V. Conclusion and Future Work

We presented a technique for discovering the location of the devices on the human body, using acceleration data captured during daily activities. Our technique first leverages unsupervised activity discovery to detect time intervals when the user is walking. Then it uses a support vector machine (SVM) to analyze the patterns and estimate the position of the device on the body. We have presented extensive experimental results, which have been performed on a diverse set of participants. We are currently developing a set of applications for smart phones to enhance the adaptability and location sensitivity of these devices. In addition, we are investigating the use of acoustic sensors (to sense heartbeat sound) to improve the granularity of the location discovery.

References

[1] Gulf coast data concepts, http://gcdataconcepts.com/x6-2.html/ retrieved on 1/25/2011.

[2] Libsvm, http://www.csie.ntu.edu.tw/˜cjlin/libsvm/ retrieved on 1/25/2011.

[3] Ultra low power sensor interface, http://www.toumaz.com/ retrieved on 1/25/2011.

[4] J. Aslam, K. Pelekhov, and D. Rus. A practical clustering algorithm for static and dynamic information organization. In *SODA '99: Proceedings of the tenth annual ACM-SIAM symposium on Discrete algorithms*, pages 51–60, Philadelphia, PA, USA, 1999. Society for Industrial and Applied Mathematics.

[5] L. Bao and S. S. Intille. Activity recognition from user-annotated acceleration data. In *PERVASIVE '04*, pages 1–17. Springer, 2004.

[6] C. V. C. Bouten and et al. Effects of placement and orientation of body-fixed accelerometers on the assessment of energy expenditure during walking. *Medical and Biological Engineering and Computing*, 35:50–56, 1997.

[7] B. Chiu, E. Keogh, and S. Lonardi. Probabilistic discovery of time series motifs. In *KDD '03*, pages 493–498, 2003.

[8] D. C. Crawford and et al. Which thermometer? factors influencing best choice for intermittent clinical temperature assessment. *Journal of Medical Engineering & Technology*, 30:199–211, 2006.

[9] H. W. Gellersen, A. Schmidt, and M. Beigl. Multi-sensor context-awareness in mobile devices and smart artifacts. *Mob. Netw. Appl.*, 7(5):341–351, 2002.

[10] K. Kunze and P. Lukowicz. Using acceleration signatures from everyday activities for on-body device location. In *ISWC '07: Proceedings of the 2007 11th IEEE International Symposium on Wearable Computers*, pages 1–2, Washington, DC, USA, 2007. IEEE Computer Society.

[11] K. Kunze, P. Lukowicz, H. Junker, and G. Trster. Where am i: Recognizing on-body positions of wearable sensors. In *In: LOCA04: International Workshop on Location and Context-Awareness*, pages 264–275. Springer-Verlag, 2005.

[12] J. Lester, B. Hannaford, and G. Borriello. Are you with me? using accelerometers to determine if two devices are carried by the same person. In *In Proceedings of Second International Conference on Pervasive Computing (Pervasive 2004*, pages 33–50, 2004.

[13] J. Lester, C. Hartung, L. Pina, R. Libby, G. Borriello, and G. Duncan. Validated caloric expenditure estimation using a single body-worn sensor. In *UBICOMP '09*, pages 225–234, 2009.

[14] E. Miluzzo, X. Zheng, K. Fodor, and A. T. Campbell. Radio characterization of 802.15.4 and its impact on the design of mobile sensor networks. In *In Proc. 5th European Conf. on Wireless Sensor Networks (EWSN 08), pp 171188*, 2008.

[15] D. Minnen and et al. Improving activity discovery with automatic neighborhood estimation. In *IJCAI '07*, pages 2814–2819, 2007.

[16] R. C. Shah and et al. On the performance of bluetooth and ieee 802.15.4 radios in a body area network. In *BODYNETS '08*, pages 1–9, 2008.

[17] A. Vahdatpour, N. Amini, and M. Sarrafzadeh. Toward unsupervised activity discovery using multi-dimensional motif detection in time series. In *IJCAI '09*, pages 1261–1266, 2009.

[18] A. Vahdatpour and M. Sarrafzadeh. Unsupervised discovery of abnormal activity occurrences in multi-dimensional time series, with applications in wearable systems. In *SDM '10: Siam Conference on Data Mining*, 2010.

[19] G. Welk and J. Differding. The utility of the digi-walker step counter to assess daily physical activity patterns. *Medicine & Science in Sports & Exercise*, 32(9):S481S488, 2000.