

Toward Unsupervised Activity Discovery Using Multi-Dimensional Motif Detection in Time Series

Alireza Vahdatpour, Navid Amini, and Majid Sarrafzadeh

Department of Computer Science
University of California, Los Angeles
{alireza, amini, majid}@cs.ucla.edu

Abstract

This paper addresses the problem of activity and event discovery in multi dimensional time series data by proposing a novel method for locating multi dimensional motifs in time series. While recent work has been done in finding single dimensional and multi dimensional motifs in time series, we address motifs in general case, where the elements of multi dimensional motifs have temporal, length, and frequency variations. The proposed method is validated by synthetic data, and empirical evaluation has been done on several wearable systems that are used by real subjects.

1 Introduction

Activity discovery and classification using wearable sensor data has always been in the interest of several research communities. Supervised and unsupervised methods have been proposed to discover signal patterns that represent certain activities [Bao and Intille, 2004]. Although supervised methods have more accurate outcome in detecting known patterns, they require knowledge from a domain expert to be fed into the system. In addition, due to the sensitivity of such methods to the quality and quantity of the training data, they are only able to discover activities that are defined and known by the domain expert a priori.

Nowadays, the proliferation of convenient handheld devices, wearable sensors, and broadband wireless services for monitoring and guidance has lead to generation of tremendous amount of data. There is a great interest in designing signal mining algorithms that can discover frequent and non-frequent unknown phenomena in such data. This information is useful in higher level reasoning and analysis of subject behavior during a short or long period of time. Time series motifs are approximately repeated subsequences in a longer time series data [Chiu *et al.*, 2003]. Motifs can be thought as constructing elements or perceptual primitives [Minnen *et al.*, 2007a] of activities. Several studies have shown the potential of using motifs for detecting and classifying activities and events [Minnen *et al.*, 2007b],[Tanaka *et al.*, 2005] in the past few years. A significant advantage of using unsupervised activity discovery methods is that it eliminates the training and setup phase required by other methods previously used

such as decision tree and hidden Markov model. This is especially valuable in nowadays increasing proliferation of different wearable sensors to monitor human activities.

The notion of motifs in time series data has been first defined by [Lin *et al.*, 2002] and addresses recurring patterns in single time series. Several attempts have been made to generalize this notion to multi-dimensional time series data [Tanaka *et al.*, 2005] [Minnen *et al.*, 2007a]. Since most of the real world applications deal with several data sources and therefore, require methods to analyze multi dimensional data. However, none of these methods considers the reality of motifs in systems which include a number of independent sensors, especially wearable systems used for activity monitoring, where motifs representing an activity have different length and timing characteristics in each signal (a signal is data gathered by a sensor). To clarify more, as it is shown in Figure 1.b activities are usually projected in different sensors (signal data) with small timing differences. However, as depicted in figure 1.a, current state-of-the-art multi dimensional motif discovery [Minnen *et al.*, 2007a] [Tanaka *et al.*, 2005] and activity discovery [Minnen *et al.*, 2007b] approaches only consider synchronous motifs distributed on each dimension of the signal. As Figure 1.b suggests, when motifs in each dimension are not synchronized, discovering them and constructing activity primitives are not a naive task anymore. In this paper, we present and evaluate an approach for activity discovery in multi time series data using multi signal motifs. The major contribution of this study is to propose a new motif definition and discovery algorithm in multi time series data.

Section 2 will cover some previous leading work in motif detection and activity discovery. Section 3 introduces our multi dimensional motif discovery and optimization method, which is an extension of the well-known approach proposed by [Chiu *et al.*, 2003]. In section 4, we present the evaluation results of our approach for activity discovery and sub dimensional motif detection. Finally, conclusion and future work are presented in the last section.

2 Background and Related Work

2.1 Motif Detection

Motif discovery is an active research topic in bioinformatics. The large data size in such applications has inspired research community to design several algorithms to increase the speed




Figure 1: Two different types of multi dimensional motifs in a set of time series a) 2-dimensional motif with synchronous occurrence of motif elements b) 2-dimensional motif with temporal variation of motif occurrences

of motif detection in biosequences (e.g. DNA). One of the most efficient algorithms is proposed by [Buhler and Tompa, 2001], where random projection concept is introduced as a mean of reducing computational complexity (and therefore, increasing the speed of the process). Later, [Chiu *et al.*, 2003] introduced the concept of motif for time series data and used projection technique to propose a probabilistic approach that is able to find motifs in single time series data in sub quadratic time. In [Tanaka *et al.*, 2005], authors use Principle Component Analysis (PCA) to transform multidimensional time series to one dimensional time series to detect motifs that happen on all dimensions of a set of time series. More recently, [Minnen *et al.*, 2007a] extended the motif discovery method for single time series to detect motifs that happen in some dimensions of a multi dimensional signal. Although this generalization increases the applicability of motif discovery methods, it still assumes that motifs along all the dimensions are happening synchronously (e.g., see figure 1.a), which restricts the applicability of the method to data coming from systems with highly accurate actions (e.g., space shuttles with accurate and almost noise free sensors).

2.2 Activity Discovery and Classification

Activity recognition and classification has been of special interest among research communities such as computer vision and body sensor networks for years. One of the major uses of activity recognition is to discover life habits and trends, which is especially attractive to medical and healthcare communities. A remarkable number of studies have been published addressing activity recognition targeting applications that each has different considerations and assumptions [Bao and Intille, 2004]. Most of these studies, propose using of supervised and semi-supervised methods for activity discovery and classification, meaning that a phase of setup and training is necessary. Consequently, the result of the classifier heavily depends on the quality and quantity of the training data set.

As an outstanding example, the study [Tapia *et al.*, 2004] uses hierarchical hidden semi-Markov model to track daily activities of assisted living community residents. The method

requires a training phase and works on data coming from simple *state-change* sensors which are massively installed in the environment. The presented method is supervised in the sense that a professional user must define activities and determine the sensors corresponding to each activity. In another study [Parkka *et al.*, 2006] evaluated the performance of custom decision trees, automatically generated decision trees, and artificial neural networks in recognizing activities using realistic data from wearable sensors. It is shown that custom decision tree has the most stable performance for all activities while the automatically generated decision tree has the best overall activity classification. It is also stated that artificial neural network does not perform well since gathered data is noisy. It is obvious that in such methods, the more effort the field expert user puts on customizing the discovery system, will result in better performance of discovering certain activities and events, however, this human customization is a costly task that is desired to be eliminated in less critical situation.

Beside the supervised methods for activity classification, recently, Minnen *et al.* [Minnen *et al.*, 2007b] proposed the idea of using motif discovery algorithms for automatic classification of activities. In their study, authors propose a method to estimate the length of motifs in a signal, which eliminates the requirement of a domain expert user to supervise the motif discovery algorithm performance. They also show how automatic motif length estimation can increase the performance of motif discovery algorithms in discovering exercise activities captured by inertial sensors. A great advantage of applying motif discovery algorithms for activity classification is its automatic operation which eliminates the requirement for supervision of domain experts (which is generally costly, time consuming, and difficult).

In this study we continue the idea of using motif discovery as an unsupervised approach for activity discovery by extending the definition of motifs in multi time series. The extended definition of motifs is a more realistic representation of activities and events in time series. Our proposed approach will leverage the benefits of *random projection* method and extends it by designing an activity clustering algorithm that constructs activity perceptual primitives from time series motifs.

3 The Unsupervised Method to Discover Activities in Multi Dimensional Time Series

3.1 Single Dimensional Motif Discovery

For better presentation of the main contribution of this study, which is enhancing the applicability of motif discovery in multi dimensional time series, we will briefly overview the state-of-the-art time series motif discovery mechanism that is the basis of this work. However, since these methods have been reviewed by many previous work, we encourage the readers to read [Chiu *et al.*, 2003] [Minnen *et al.*, 2007a] for the detailed description and evaluations.

To clarify more, it should be noted that in this paper, **single dimensional** and **multi dimensional** are used to describe time series and motifs which consist of one and more than one data dimensions, respectively (Some previous work have used terms such as *univariate* and *multivariate* time series.)

Motif discovery is basically the process of finding recurring subsequence in a time series. Since the shape of the patterns is unknown, $O(n^2)$ comparisons are required to compare every two subsequences and find subsequences that iterate in the time series (we assume n is the length of the time series and also note that the length of the subsequences is a constant $m \ll n$). Because the size of the time series is usually large in most applications, studies have been done to reduce the quadratic run time of the algorithm. The most well know approach is proposed by [Chiu 2003], where discretization and random projection are used to reduce the comparison space size and the number of candidate motif subsequences, respectively. In the discretization step, SAX [Lin *et al.*, 2003] method is used. SAX is a symbolic representation of time series that reduce the size of a time series by discretizing a signal using piecewise aggregation approximation [Keogh *et al.*, 2000]. It is shown that although applying this method to time series data decreases the accuracy of signal representation, information loss is bounded and by choosing appropriate parameters (such as the number of discrete levels) adequate information about the behavior of the signal can be preserved.

After SAX conversion, all subsequences in the time series are compared to each other using random projection [Buhler and Tompa, 2001]. Random projection is the process of randomly comparing some values from two subsequences. Since, when two subsequences match, their random projection also must match, this method is useful to avoid performing of full comparison between subsequences that highly differ. Result of several iterations of random projection is saved in a collision matrix that represents the similarity of subsequences. Generally, element i, j of collision matrix represents the number of times random projections of subsequence i and subsequence j have been matched to each other. Figure 2 depicts a block diagram representation of the motif discovery process.

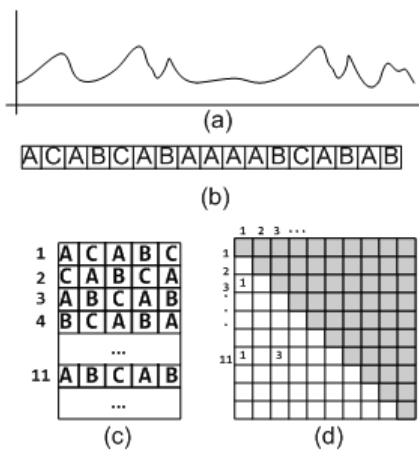


Figure 2: General overview of single dimensional motif discovery algorithm (a) A time series (b) Time series SAX conversion (c) SAX words construction (d) Collision table via random projection

3.2 Multi Dimensional Motif Discovery via Graph Clustering

In order to extend the motif discovery algorithm to multi dimensional time series, considering the practical problems associated with current methods, we have developed a two-stage algorithm that extracts non-synchronous multi dimensional motifs from a set of time series. Currently the state of the art studies addressing detecting multi dimensional motifs in time series [Tanaka *et al.*, 2005] [Minnen *et al.*, 2007a] [Minnen *et al.*, 2007b] use Logical AND and Logical OR to concatenate collision matrices resulted from subsequence matching in each of time series dimensions. Although these approaches can discover synchronous motifs, in reality, they show poor performance because motifs in all dimensions do not occur completely synchronous and their relative timing differs time to time (e.g. the timing between moving right and left feet in walking activity slightly varies). In addition, while two motifs in two time series, can represent the same activity, the length of these subsequences can be completely different, making it impossible to relate them using simple methods such as Logical AND and Logical OR. Since in these methods, the length of motifs are considered to be the same in all dimensions of the time series. To overcome the two major problems stated above (motif synchronization and motif length variation in different dimensions of time series) we have designed a two phase algorithm.

In the first phase of our algorithm, single dimensional motifs in all dimensions of the time series are discovered. To do so, the method in [Chiu *et al.*, 2003] is implemented and optimized by the complementary mechanism introduced in [Yankov *et al.*, 2007]. While the optimization mechanism makes the discovery method resilient to uniform scaling, since the discovery algorithm requires m (length of the motif) as an input, it is required to perform the discovery task on the time series considering all possible motif lengths. It should be noted that we aimed to use these methods, since they are the most efficient methods in the literature up to now, however, our algorithm for the second phase only depends on the result of single dimensional motif discovery and these methods can be replaced by any valid algorithm.

Let's assume the input to the algorithm is a set of k time series, where n is the length of the time series. By applying the single dimensional motif discovery of the first phase, a set $R = \{r_i\}$, $|R| = W$ of motifs are discovered, where r_i represents a single dimensional motif with several occurrences in one of the time series. Each occurrence of motifs in r_i is denoted by $r_{i,j}$, where $r_{i,1}$ is the first occurrence and $r_{i,size_i}$ is the last occurrence of motif r_i in its associated time series ($size_i$ = the total number of occurrences of motif in r_i).

The next phase of the algorithm is to construct multi dimensional time series motifs via combining discovered single dimensional motifs. In this study, we have used a graph clustering approach to group the single dimensional motifs, and will show its efficiency in the next section. To perform the motif clustering, first we construct a coincidence graph G as follows: In directed graph $G = (V, E)$, each motif r_i is represented by a vertex v_i . Edge $e_{i,j}$ (edge connecting vertex i to vertex j) has weight $w_{i,j}$ which is calculated as:

$$w_{i,j} = \text{coincident}(r_i, r_j) / \text{size}_i$$

Where $\text{coincident}(r_i, r_j)$ is the total number of times that there is a temporal overlap between occurrences of r_i and r_j . As $w_{i,j}$ gets closer to one, it denotes a higher coincidence factor between two single dimensional motifs. Figure 3 depicts an example of three time series and the coincidence adjacency matrix generated from it. Note that the table in the figure includes the weights of the edges of the coincidence graph.

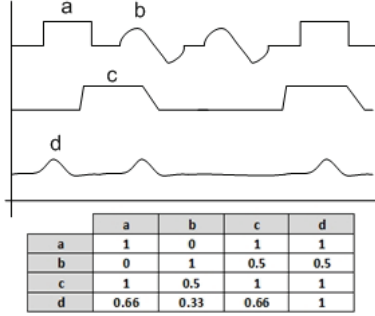


Figure 3: Three example time series and the weight of the edges in their adjacency matrix

Algorithm 1 Multi dimensional motif construction

- 1: Input: $G(V, E)$, α , size_i and $w_{i,j}$ for all i and j
 - 2: Output: S_i set of clustered motifs, each representing a multi dimensional time series motif
 - 3: Sort the vertices v_i in G based on the number of incidents (size_i)
 - 4: **for** each vertex v_k in the sorted list of vertices **do**
 - 5: **for** all neighbors of v_k , if $e_{k,j} > \alpha$ **do**
 - 6: Add v_j to S_k
 - 7: Update the weight of edges connected to v_j by removing the motif occurrences of r_j that has coincident with r_k
 - 8: **end for**
 - 9: Update the sorted list of vertices
 - 10: **end for**
 - 11: **return** all nonempty S_k
-

The procedure to cluster the graph and construct multi dimensional motifs is stated in Algorithm 1. Starting from the motifs with most occurrences in the time series, the algorithm groups motifs that have high coincidence with each other. after performing each grouping iteration, the coincidence graph is updated by eliminating the motif occurrences that are associated to an activity (multi dimensional motif). Note that α ($0 < \alpha < 1$) is the threshold to determine the minimum correlation two motifs should have to construct a multi dimensional activity. This factor can be determined statistically by a field analyzer or can be determined dynamically by statistical analysis of weights of the coincidence graph (We have used the dynamic approach in our experiments).

The key feature of the above method for discovering multi dimensional motifs, comparing to previous studies [Minnen

et al., 2007b] [Tanaka *et al.*, 2005] are *a*) Since the relation between two single dimensional motifs is based on any temporal overlap, slight variation of timing does not affect the results. *b*) Each single dimensional motif is discovered independently, which enables them to be of any size and frequency. *c*) Discovered multi dimensional motifs can span on any subset of time series.

3.3 Improvement of the Single Dimensional Motif Discovery Algorithm

As mentioned before, the presented method for discovering multi dimensional motifs is based on clustering the single dimensional motifs discovered using the probabilistic approach of [Chiu *et al.*, 2003]. Due to the existence of the noise in the time series, it is well known that motif discovery algorithms fail to discover some of the occurrences of motif seeds. Although by relaxing the threshold distance in which two substrings match together, one can increase the probability of discovering all motif occurrences, this will lead to possible detection of wrong subsequences as motifs. An advantage of the two phase algorithm to discover multi dimensional motifs is the ability to increase the accuracy of the results of the first phase (single dimensional motif discovery) by the results of the second phase (motif clustering). Generally, when two motifs have high temporal coincidence, an exclusive occurrence of one of them suggests high probability for discovering an additional occurrence of the second motif (with increasing the acceptable noise threshold). Algorithm 2 uses this idea to discover potential motif occurrences that are not discovered in the first phase.

Algorithm 2 Single dimensional motif discovery optimization

- 1: Input: $R = r_i$ initial set of discovered single dimensional motifs
 - 2: Output: $R = r_i$ modified set of discovered single dimensional motifs
 - 3: **for** each r_k and r_l in R , if $1 > w_{k,l} > \beta$ **do**
 - 4: **for** all $r_{k,x}$ (occurrences of r_k) **do**
 - 5: **if** there is no coincidence between occurrences of r_l and $r_{k,x}$ **then**
 - 6: Re-evaluate r_l for additional occurrences of motifs
 - 7: Update r_l
 - 8: **end if**
 - 9: **end for**
 - 10: **end for**
 - 11: **Return** R
-

β ($0 < \beta < 1$) is the precision factor, which determines the revisiting criteria of the algorithm. As β gets closer to one, the algorithm only considers highly correlated motifs and evaluates their missing occurrences while a small β requires more reevaluations of the motifs occurrences. It is obvious that small β values result in larger execution time, however, it is recommended that in dealing with noisy data, which increases the error probability in motif detection, a smaller β is used, to increase the accuracy of the motif discovery.

4 Empirical Results

In this section, we will provide the evaluation results of our multi dimensional motif discovery algorithm on several data sets. Here, to show the applicability and timing characteristics of the method, we first evaluate the method using synthetic data, however, in the most of the experiments our emphasis was on the data collected from inertial sensors such as accelerometers and gyroscopes worn on the body or embedded in several devices to monitor activities of subjects.

4.1 Synthetic Data

To generate the synthetic data, motifs were planted in a set of randomly generated time series. Figure 4 shows a small snapshot of the discovered multi dimensional motifs using our method and the sub dimensional motif discovery method presented in [Minnen *et al.*, 2007a]. As depicted in figure 4.b, the previous work only detects motifs that are well synchronized. In this test case, since motifs have different temporal relations in each occurrence, it fails to discover some of the dimensions and occurrences correctly. It is important to note that it is possible to increase the distance error threshold in the previous work, to able the algorithm to discover the time shifted motifs, however, increasing the threshold will result in increasing rate of falsely detected motifs.

Figure 5 shows how the algorithms execution times scale over the number of time series dimensions, when the length of all time series and motifs are 10000 and 60 respectively. As graph shows, both algorithms are scaled linearly, however, the linear factor is bigger in our algorithm since the single dimensional motif discovery should be performed against each time series, while in the previous work, increasing the dimension only affects the calculations required for constructing the collision table. Note that the in our approach, although the execution time of the clustering algorithm (phase 2) is quadratic, the time required to calculate motif clusters is negligible versus the time required to discover single dimensional motifs (phase 1), since the size of the input to the clustering algorithm is the number of detected single dimensional motifs which is substantially smaller than the size of the time series.




Figure 5: The scalability of the two approaches with increasing the number of dimensions. A1 is our approach and L1 is [Minnen *et al.*, 2007a]

4.2 Real Data

In order to further evaluate the performance of the method, we tested the method for data gathered from sensors embed-

ded in two different wearable systems, SmartCane [Wu *et al.*, 2008] and SmartShoe [Dabiri *et al.*, 2008]. Both systems include accelerometers, gyros and pressure sensors. Two different scenarios were tested on the SmartCane system, one being normal use of the cane in walking and the other one was the cane falling on the ground (The scenarios were separated since the range and quality of the data differs significantly in the two scenarios). Eight time series were collected from the SmartCane in the normal walking and five were recorded in the falling scenario. Also, six signals were collected from the SmartShoe. Study in [Minnen2 2007] suggests that although comparing the results of such activity monitoring mechanisms to a hand labeled ground truth allows quantitative measurement of performance, still issues such as detecting additional motifs are not incorrect results. The additional motifs can be a result of actions that the field expert does not expect or they can be due to the noise in the data or error in the method, while less number of detected motifs can be due to the noise in the time series or the inaccuracy of the discovery method.

Table 1 summarizes the average accuracy of the activity discovery by applying multi dimensional motif discovery algorithm to the time series data collected from the above systems in 1000 seconds, comparing to hand-labeled classification of the data. Note that hand-labeled data of the SmartCane walking data comprises three different activities (each spans on some time series dimensions) while there were two and three activities for SmartShoe and SmartCane fall applications, respectively. Since the notion of the activity defined in this paper is an extension of the actions defined in [Minnen *et al.*, 2007a][Tanaka *et al.*, 2005], we could not conduct a direct accuracy and performance comparison between the methods (previous work do not consider time shifted multi dimensional motifs, which is the main scenario happened in our systems). It also includes the results of evaluating the single dimensional motif discovery improvement mechanism proposed in section 3.3. Comparing the results in Table 1 to the results provided in [Bao and Intille, 2004] (which has accuracy varying between 42%-96% for different supervised activity classification algorithms performed on different systems) the results of our unsupervised method seems to be promising. It is worth to mention that the performance of any method heavily relies on the test data and as it is presented in Table 1, our results accuracy varies between 62% and 85%. In addition, since our method leverages and expands the method of [Chiu *et al.*, 2003], parameters such as the motif lengths and scaling factor changes the outcome of the algorithm and we have manually optimized them in our evaluation scenario to effectively evaluate the contribution of this study.

5 Conclusion and Future Work

Using time series motifs to discover activities and events has been shown to be beneficial in the past [Minnen *et al.*, 2007b] [Tanaka *et al.*, 2005]. In this study, we described an approach to discover multi dimensional motifs in time series. The major contribution of the method is its ability to discover multi dimensional motifs that have time and value irregularities, which is a common case in activity monitoring applications.

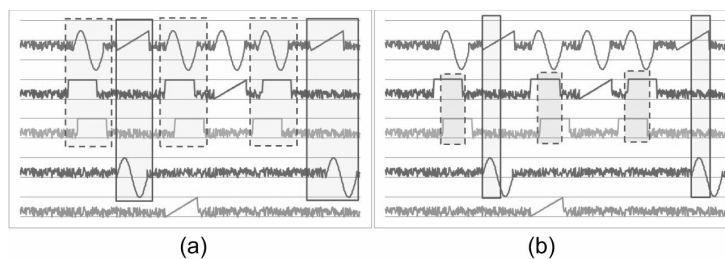


Figure 4: A simple set of synthetic time series and the result of multi dimensional motif discovery using (a) our approach and (b) the approach of [Minnen et al., 2007a].

Table 1: The results of performing the unsupervised activity discovery approach on three applications

	SmartCane Walk	SmartShoe	SmartCane Fall
Multi dim. activity discovery accuracy	71%	85%	62%
Single dim. motif discovery improvement	9%	17%	12%

The method is based on the well known single dimensional motif discovery approach of [Chiu et al., 2003]. We proposed to construct activity primitives based on clustering single dimensional motifs. The clustering was done on the coincidence graph which is based on the temporal coincidence of motifs in different time series dimensions. We also presented a feedback approach to improve the accuracy of single dimensional motif discovery process. In addition to evaluating the approach using synthetic data, we used real data collected from previously deployed wearable systems. Although the activity discovery method using motif concept is an unsupervised approach, which eliminates the need for a field expert, it still shows acceptable outcome. For the future work, we plan to evaluate using different clustering approaches for grouping the single dimensional motifs.

6 Acknowledgments

The authors would like to thank Professor Eamonn Keogh for providing us with the implementation of the single dimensional motif discovery algorithm.

References

- [Bao and Intille, 2004] L. Bao and S. S. Intille. Activity recognition from user-annotated acceleration data. In *PERVASIVE*, pages 1–17. Springer, 2004.
- [Buhler and Tompa, 2001] Jeremy Buhler and Martin Tompa. Finding motifs using random projections. In *Journal of Computational Biology*, pages 69–76, 2001.
- [Chiu et al., 2003] B. Chiu, E. Keogh, and S. Lonardi. Probabilistic discovery of time series motifs. In *KDD '03*, pages 493–498, NY, USA, 2003. ACM.
- [Dabiri et al., 2008] F. Dabiri, A. Vahdatpour, H. Noshadi, H. Hagopian, and M. Sarrafzadeh. Ubiquitous personal assistive system for neuropathy. In *HealthNet '08*, 2008.
- [Keogh et al., 2000] E. Keogh, K. Chakrabarti, M. Pazzani, and S. Mehrotra. Dimensionality reduction for fast similarity search in large time series databases. *Journal of Knowledge and Information Systems*, 3:263–286, 2000.
- [Lin et al., 2002] J. Lin, E. Keogh, S. Lonardi, and P. Patel. Finding motifs in time series. In *Second Workshop on Temporal Data Mining, (KDD'02)*, pages 53–68, 2002.
- [Lin et al., 2003] J. Lin, E. Keogh, S. Lonardi, and B. Chiu. A symbolic representation of time series, with implications for streaming algorithms. In *DMKD '03: 8th ACM SIGMOD workshop on Research issues in data mining and knowledge discovery*, pages 2–11, NY, USA, 2003. ACM.
- [Minnen et al., 2007a] D. Minnen, C. Isbell, M. Essa, and T. Starner. Detecting subdimensional motifs: An efficient algorithm for generalized multivariate pattern discovery. *ICDM 2007*, pages 601–606, Oct. 2007.
- [Minnen et al., 2007b] D. Minnen, Th. Starner, I. A. Essa, and C. L. Isbell Jr. Improving activity discovery with automatic neighborhood estimation. In *IJCAI*, pages 2814–2819, 2007.
- [Parkka et al., 2006] J. Parkka, M. Ermes, Panu Korpipää, Jani Mäntyjärvi, Johannes Peltola, and Ilkka Korhonen. Activity classification using realistic data from wearable sensors. *IEEE Transactions on Information Technology in Biomedicine*, 10(1):119–128, 2006.
- [Tanaka et al., 2005] Y. Tanaka, K. Iwamoto, and K. Uehara. Discovery of time-series motif from multi-dimensional data based on mdl principle. *Mach. Learn.*, 58(2-3):269–300, 2005.
- [Tapia et al., 2004] E. Munguia Tapia, S. S. Intille, and Kent Larson. Activity recognition in the home using simple and ubiquitous sensors. In *Pervasive*, pages 158–175, 2004.
- [Wu et al., 2008] W. Wu, L. Au, B. Jordan, T. Stathopoulos, M. Batalin, W. Kaiser, A. Vahdatpour, M. Sarrafzadeh, M. Fang, and J. Chodosh. The smartcane system: an assistive device for geriatrics. In *BodyNets '08*, pages 1–4, Belgium, 2008. ICST.
- [Yankov et al., 2007] D. Yankov, E. Keogh, J. Medina, B. Chiu, and V. B. Zordan. Detecting time series motifs under uniform scaling. In *KDD*, pages 844–853, 2007.