Improving Activity Discovery with Automatic Neighborhood Estimation

David Minnen, Thad Starner, Irfan Essa, and Charles Isbell

Georgia Institute of Technology

GVU / College of Computing

{dminn,thad,irfan,isbell}@cc.gatech.edu

Abstract

A fundamental problem for artificial intelligence is identifying perceptual primitives from raw sensory signals that are useful for higher-level reasoning. We equate these primitives with initially unknown recurring patterns called motifs. Autonomously learning the motifs is difficult because their number, location, length, and shape are all unknown. Furthermore, nonlinear temporal warping may be required to ensure the similarity of motif occurrences. In this paper, we extend a leading motif discovery algorithm by allowing it to operate on multidimensional sensor data, incorporating automatic parameter estimation, and providing for motif-specific similarity adaptation. We evaluate our algorithm on several data sets and show how our approach leads to faster real world discovery and more accurate motifs compared to other leading methods.

1 Introduction

This paper addresses the problem of unsupervised discovery of perceptual primitives from time-varying sensor data. Such data includes sensor readings collected by mobile robots, on-body sensors, surveillance video, econometric data, and industrial diagnostic readings.

Perceptual primitives can be thought of as recurring patterns or *motifs* of the data; that is, they are sets of subsequences with high similarity. The goal of this work is to discover the motifs and the corresponding occurrences in real-valued, multivariate data. The problem is complicated because little is known about the motifs *a priori*. The number of motifs, number of occurrences, their length, shape, scale, and the location of each occurrence are all unknown. Furthermore, motifs may only appear similar after nonlinear warping, and the data may be corrupted by noise. Finally, any discovery algorithm must be efficient, since even small data sets may have anywhere from a few thousand to several hundred thousand sensor readings. Thus, algorithms with quadratic space or time complexity are unacceptable for real applications.

We are influenced by the philosophical position articulated by Cohen *et al.* [2002]. In that work, the authors argue that "most of the intellectual work in AI is done not by programs but by their creators, and virtually *all* the work involved in specifying the meanings of representations is done by people, not programs." They adopt Fred Dretske's theory of meaningful representation, positing that an internal state must both indicate a condition and have the function of indicating that condition. Applying this theory to robotics, Cohen *et al.* conclude that learning meaningful representations requires the robot to analyze its sensor readings to learn abstractions that inform action. Similarly, we seek to design an activity discovery system that learns endogenously meaningful representations via unsupervised analysis of sensor data.

Our contributions are: (1) to extend the applicability of current discovery algorithms, (2) to reduce the domain knowledge required of users, and (3) to improve performance of both discovery accuracy and real-world execution time. Applicability is extended by allowing discovery to work directly with multidimensional data (see Section 4). Our approach achieves two additional improvements by automatically estimating the size of the neighborhood for each motif. First, the user is no longer required to specify a parameter that is vital to locating correct motif occurrences but which is highly dependent on the domain and the distance metric used to measure subsequence similarity. Second, automatic estimation allows a different neighborhood size for each motif, something that is extremely difficult to specify manually but which leads to higher quality motifs in empirical tests (see Sections 4.1 and 6). We also present several evaluation methods for motif discovery and discuss many of the semantic and application-specific issues involved in comparing different approaches (see Section 5).

2 Related Work

Activity recognition is a supervised learning problem related to activity discovery. Approaches often use complex models, including hidden Markov models (HMMs), switching linear dynamic systems, and stochastic grammars [Starner *et al.*, 1998; Pavlovic *et al.*, 2000; Minnen *et al.*, 2003]. Bao and Intille [2004] provide a broad survey, comparing several methods for recognizing activities captured by on-body sensors.

In the bioinformatics community, researchers have developed many approaches for discovering recurring patterns that correspond to binding sites in nucleotide sequences. MEME uses expectation maximization (EM) to estimate the parame-

ters of a probabilistic model for each motif [Bailey and Elkan, 1994]. GEMODA unifies several previous discovery methods, but the algorithm has an expected run time quadratic in the size of the data [Jensen *et al.*, 2006]. Finally, the greedy EM method developed by Blekas *et al.* [2003] uses incremental mixture learning to discover motifs represented by probabilistic models similar to MEME.

The data mining and artificial intelligence communities have developed methods that work with real-valued sequences. PERUSE finds recurring patterns in multidimensional data by estimating the parameters of an HMM-like model [Oates, 2002]. Tanaka *et al.* [2005] apply principal component analysis (PCA) to multidimensional data, using a minimum description length criterion to select motifs found in the principal component. Lin *et al.* [2002] use hashing and lower-bounding to efficiently search for motifs. Chiu *et al.* [2003] extend Lin *et al.*'s approach by using random projections to quickly find candidate motifs in noisier data.

3 Motif Discovery via Random Projection

In this section, we focus on Chiu *et al.*'s probabilistic discovery method [2003], discussing extensions to improve its applicability, effectiveness, and real-world run time. The algorithm operates in five steps (see Table 1 for a summary of the user-specified parameters):

1. Extract all subsequences of a given length (n) from the time series using a sliding window.
2. Convert each subsequence to a string using SAX with length w and a unique symbols.
3. Build a collision matrix by comparing the strings via several iterations of random projection.
4. Select motif seeds as the two subsequences corresponding to the largest entry in the collision matrix.
5. Extract motifs by detecting other windows in the neighborhood of each seed. The neighborhood is defined as all sequences within a fixed distance (R) of either seed.

Subsequence Extraction: Subsequences are extracted using a sliding window of user-specified length. Each window location corresponds to one subsequence of the time series.

Subsequence Discretization: Each subsequence is discretized using the <u>s</u>ymbolic <u>a</u>ggregate appro<u>x</u>imation (SAX) method developed by Lin *et al.* [2003]. SAX is a local quantization method that first computes a piecewise aggregate approximation (PAA) [Keogh *et al.*, 2000] of the normalized window data and then replaces each PAA segment with a symbol. Normalization ensures that the subsequence has zero mean and unit variance. PAA represents a sequence by dividing it into equal length segments and then storing the mean within each segment. The final SAX quantization is based on precomputed breakpoints that divide the data range into equiprobable regions assuming an underlying standard normal distribution. Each PAA segment is represented by the symbol of the corresponding bin (see Figure 1).

As presented, SAX appears applicable only to univariate data; however, one can use the algorithm with multivariate data by applying it to each dimension independently
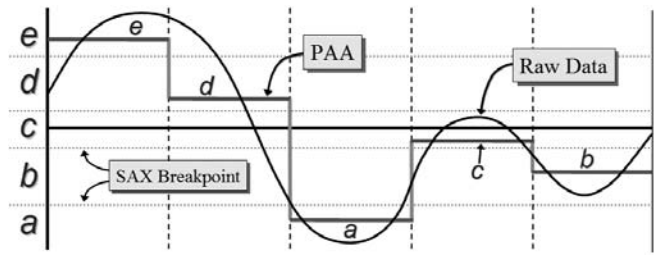

Figure 1: A real-valued, 1D signal is converted into a SAX string (`edacb`) using five PAA segments (bounded by the vertical dotted lines) and five SAX symbols.

and then concatenating the resulting strings. This leads to a constant number of symbols and linear increase in string length [Keogh, 2006].

Building the Collision Matrix: Once each subsequence has been converted to a string, the similarity between each pair is efficiently estimated using random projection. The key insight is that while direct comparison is infeasible since there are $O(T^2)$ pairs given T strings, similar sequences can be identified in $O(I \cdot T)$ time using I iterations of random projection (see Figure 2). Each iteration involves selecting a random subset of the string positions and building a hash table with the corresponding characters. After all strings are hashed, collisions (equivalent projections from different strings) are taken as evidence of similarity and the corresponding positions in a collision matrix are incremented.

Using random projection to estimate similarity has several important properties. First, it is effective in the presence of shot noise. Even if some of the symbols are drastically different, when random projection does not select these positions, the string will still appear similar. Second, although the collision matrix could require $O(T^2)$ time and space in the worst case, realistic data sets produce a linear number of non-zero entries leading to linear time and space requirements with a sparse matrix implementation.. Furthermore, the complexity can be controlled online by dynamically changing the number of string positions included in each iteration of random projection [Keogh, 2006]. In a pathological data set, as when the distance between all subsequences is effectively zero, the matrix will still grow prohibitively large despite dynamic projection adaptation. When such a condition is detected, the algorithm can switch to storing row sums rather than the full matrix. Our experiments show that this simplification does reduce accuracy, but the penalty may be tolerable to achieve strict linear time and space complexity.

Selecting Seed Motif Occurrences: Once the collision matrix is built, seed motif occurrences are extracted by locating maximal entries. Each entry in the matrix corresponds to a pair of subsequences that can be used to find the other occurrences of the motif. The distance between the seeds must be less than R, they must not temporally overlap previously identified occurrences, and they must not overlap each other (called a *trivial match*). Search for additional motif seeds terminates when the largest entry in the collision matrix is smaller than the expected score, which can be estimated by
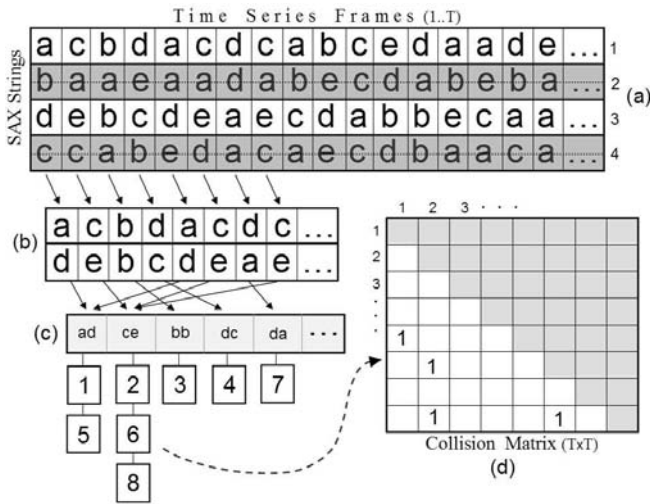
Figure 2: (a) For each iteration of random projection, a subset of string positions are selected (here, positions one and three). (b,c) The remaining symbols are hashed, and (d) equivalent projections are tallied in a collision matrix.

Equation 9 in [Chiu *et al.*, 2003].

Locating Additional Motif Occurrences: The final step of the discovery algorithm is to locate all other occurrences of the motif identified by the seed locations. This is achieved by scanning the original subsequences and selecting those within the neighborhood of the seeds (*i.e.,* $min(d(w_i, seed_1), d(w_i, seed_2)) \leq R$, where $w_i : i \in \{1..T\}$ represents each of the original subsequences).

Discussion: In our experiments (see Section 6), Chiu's approach works very well despite the erroneous assumption that all motifs and all motif occurrences have the same length. It is, however, quite sensitive to the various user-specified parameters summarized in Table 1. In particular, the neighborhood radius, R, is vital for accurate motif discovery, yet it is very difficult to specify since it depends on the domain, the distance metric, the motif length, and the features selected for processing.

4 Motif Discovery using Automatic Neighborhood Estimation

Before discussing our approach for automatically estimating a per-motif neighborhood radius, we address the issue of applying the discovery algorithm to multidimensional data.

As discussed previously, we can apply a univariate discretization algorithm to each dimension of a multidimensional subsequence and concatenate the result. If this combined string is used directly for random projection, however, the space requirements will likely be quadratic in the length of the input data. This increase is understandable because projecting columns of the concatenated string is equivalent to defining multidimensional similarity as a logical OR between dimensions. Thus, two strings are deemed equivalent for a particular iteration if they match in *any* dimension. Instead we implement similarity using a logical AND by requiring

Table 1: User-Specified Parameters for Chiu's Algorithm

Symbol	Description
n	window length
w	number of PAA segments (string length)
a	SAX alphabet size
R	radius of motif neighborhood

that *all* dimensions agree for a particular string position. This drastically reduces the effective size of the collision matrix and requires motifs to span all of the dimensions. To our knowledge, no existing algorithm addresses the more general problem of discovering motifs that exist in only a subset of the dimensions.

4.1 Automatic Estimation of Motif-Specific Neighborhoods

In order to address the theoretical and practical issues surrounding a user-specified neighborhood radius, we have developed a method for automatically estimating an appropriate value from data. Because the radius is not used during the random projection phase, that part of the algorithm proceeds as usual. Once a pair of motif seeds has been selected, the distance from each subsequence to the closer of the two seeds is computed using the original, real-valued data. Next, the smallest portion of the distances are selected for further analysis[1]. In the experiments presented below, the smallest 10% are used based on the assumption that 10% is a very loose upper bound on the number of motif occurrences. Our empirical results show that this precise size of the portion only has a minor impact on the final estimate (see Figure 6).

Finally, to determine the appropriate neighborhood radius, we sort the distances and seek a "knee in the curve." The knee corresponds to the change from distances within the motif to those outside. This change is an inflection point of the curve. Inflection points can only occur at a zero of the second derivative, which corresponds to an optima of the first derivative. Thus, estimating the radius is equivalent to finding the maxima of the derivative (see Figure 3), and because we can directly compute the discrete first derivative, no assumptions about the parametric form of the distance curve are necessary.

The inherent undersampling of the true distribution of motif occurrences causes the observed distances to be quite noisy. Therefore, the mode is estimated as the weighted mean of the derivative. This procedure is summarized below:

1. Calculate the minimum distance from each subsequence, $w_i : i \in \{1..T\}$, to a motif seed:
   $d_i = min(d(w_i, seed_1), d(w_i, seed_2))$
2. Compute, v_k, the k^{th} order statistic for $k = \frac{T}{10}$
3. Select the k smallest distances: $d_i : d_i \leq v_k$
4. Sort the distances in ascending order
5. Calculate the discrete first derivative: $d'_i = d_{i+1} - d_i$
6. Treat the derivative as weighted votes for the best radius and compute the expected value:
   $E(i) = (\sum_i d'_i \cdot (i + \frac{1}{2}))/(\sum_i d'_i)$

---

[1]See [Cormen *et al.*, 1997] for a linear time algorithm for finding arbitrary order statistics.
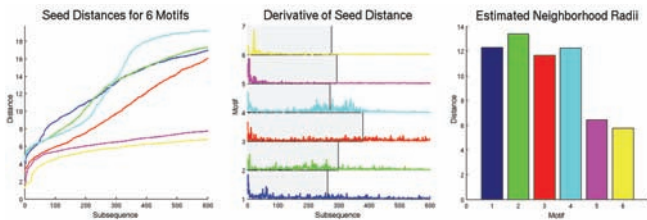
Figure 3: The radius of each motif neighborhood is estimated from the distance between each potential member and the motif seeds (*left*). It is equated with the inflection point, which is estimated by the weighted mean of the derivative of the distances (*center*). The result is automatic motif-specific neighborhood sizes (*right*).

5 Evaluation Methods

Before evaluating the performance of our approach, we must discuss what it means for one set of discovered motifs to be "better" than another set. As with activity recognition, the efficacy of a discovery system depends on how the results will be used. We have identified three broad approaches for evaluating the quality of discovered motifs: comparison to ground truth, primary utility, and interpretability.

5.1 Comparison to Ground Truth

When operating on well-understood data, domain experts may be able to manually identify motifs. Thus, they can create a set of ground truth labels that identify which motifs exist and where they occur in the data. A discovery system can then be evaluated by comparing the discovered motifs to the "correct" motifs identified by the experts in the same way that supervised recognition systems are evaluated. This approach is often used in the bioinformatic literature where algorithms are tested against nucleotide data for which binding sites area already known.

Although ground truth labels allow for a quantitative, objective measure of performance, evaluation is still complicated by issues such as different kinds of errors, "bounded" versus "fluent" activities, and different levels of error analysis (see [Minnen *et al.*, 2006b] for a detailed discussion). Furthermore, the number of discovered motifs may not be the same as what was identified by a domain expert. Different error metrics may prefer methods that discover fewer but more accurate motifs (high precision) or those that identify all of the expected motifs but also find unlabeled motifs (high recall). Note that these additional motifs are neither correct nor incorrect in general. They may be errors, or they may correspond to real patterns in the data that were simply unexpected.

5.2 Primary Utility

The best measure of performance is the utility of the discovered motifs in accomplishing a primary task. The assumption here is that motifs are not discovered for their own sake but rather in the service of a primary goal. For example, when analyzing speech data, a discovery algorithm may identify phonemes as motifs if a short temporal scale parameter is specified. Speech recognition systems, however, are not typically evaluated based on their phoneme recognition rates. In-

stead, they are only scored according to their ability to accurately detect higher-level structures such as words, phrases, or sentences. The accuracy of phoneme-level motifs is not important *per se*: they only matter insofar as the motifs represent useful features for improving word recognition rates.

5.3 Interpretability

Discovered motifs can also be qualitatively evaluated by presenting them to a domain expert. The expert can score the results by inspecting either the original data or some other time-synchronized recording such as video or audio. Evaluation may consist of manually checking that all occurrences of a motif are semantically similar and that similar events are not split across multiple motifs. The key difference between this approach and a comparison to ground truth is that interpretability is evaluated by inspecting the motifs that were actually discovered, whereas ground truth labels encode an expert's *a priori* expectations.

5.4 Further Issues for Evaluation

The quality of the discovered motifs is not the only factor that contributes to the value of a discovery algorithm. Other factors matter, such as run time characteristics (including asymptotic performance and expected real world run time), online versus batch processing, and memory requirements. In addition, the overall usability of the system is a major component of its value and is a fundamental source of motivation for our approach. Although we will show that motif-specific neighborhood size estimation improves results, even if the difference was negligible, automatic estimation would still be a contribution. The benefit arises because automatic estimation eases the demands on the user to understand the domain before using the discovery system and reduces the need to manually explore the space. A similar idea is discussed in more detail in [Keogh *et al.*, 2004]; however, the compression-based distance metric proposed in that work proved insufficient for detecting motifs in the test data.

6 Empirical Results

In this section, we evaluate our approach to motif discovery and compare it to existing methods. We compare the discovered motifs to known motifs (for the manually labeled exercise data set), and we also visually present the motifs discovered in public data sets that have been used for empirical evaluation by other researchers (the shuttle and fetal ECG data).

6.1 Exercise data

The exercise data set was captured from an exercise regime composed of six different dumbbell exercises. An Xsens MT9 inertial motion sensor was attached to the subject's wrist by fitting it into a pouch sewn to the back of a thin glove. The MT9 was sampled at 100Hz and three-axis accelerometer and gyroscope readings were recorded. In total, approximately 27.5 minutes of data was captured over 32 sequences. For the experiment, the data was resampled at 12.5Hz leading to 20,711 frames. The data contains six different exercises and 864 total repetitions (roughly 144 occurrences of each exercise). Each frame is composed of the raw accelerometer and

gyroscope readings leading to a six-dimensional feature vector.

Evaluation was performed by comparing the discovered motifs to the hand labeled occurrences of each exercise. Both discovery methods located extra motifs in addition to the six exercises. We do not consider this to be an error, however, as the extra motifs are legitimate clusters of similar subsequences. To quantify performance, we performed analysis over events (individual occurrences) and frames (each of the 20,711 sensor readings). Accuracy, precision, and recall were calculated in both cases considering only the six motifs that most closely correspond to the exercises. For all experiments, the dynamic time warping (DTW) distance was used with a Sakoe-Chiba band of 10%, which follows standard practice in the field [Ratanamahatana and Keogh, 2005].

Figure 4 shows the result of running Chiu's algorithm over a range of values for the neighborhood radius. The chart shows performance for each of the six evaluation metrics along the horizontal axis. For each metric, performance is measured for 18 different radius values ranging from 8.0 to 14.0. Note that event-level accuracy, which we consider to be the most important overall summary statistic, peaks at 61.1% with a radius of 9.5, but it is very unstable as accuracy drops to 39.1% with a radius of 9.75. This instability is due to the use of a single neighborhood size across motifs when the motifs actually have different scales (see Figure 3).

The performance of our approach is shown in Figure 6. For each of the six evaluation metrics, performance is presented across a range of estimation portion sizes. Each portion size corresponds to the percentage of the smallest distances used to estimate the neighborhood size as explained in Section 4.1. Although the portion size does impact overall performance, the results are very stable between the relatively wide range from 5.0% to 15.0%. For all of the experiments performed in this section, the generic value of 10.0% was selected since it is in the center of this stable range. Note that a value of 7.5% leads to superior results on this data set but will likely fail to generalize as well since it is closer to the boundary of the stable region.

Figure 5 shows a direct comparison of our proposed approach to Chiu's algorithm on the exercise data. For Chiu's algorithm, a radius of 9.5 was used since it was found to be the optimal choice after manually comparing a range of

Figure 5: Comparison of the proposed approach to Chiu's algorithm for the exercise data showing superior performance of the proposed approach across all six evaluation metrics.

values. The graph shows that the proposed approach outperforms the previous method according to all of the metrics. For event accuracy, the proposed approach achieves 78.2% compared to 61.1%, which is a 44% reduction in error. For comparison, another discovery method, which relies on global discretization and a minimum description length criterion to find motif seeds, achieves an event accuracy of 86.7% [Minnen *et al.*, 2006a]. This approach, however, requires considerably more time to locate the motifs, and it includes a step that is quadratic in the number of motif occurrences.

6.2 Shuttle and Fetal ECG Data

In order to demonstrate the generality of our approach, we tested on two public data sets commonly used in the motif discovery literature [Keogh and Folias, 2002]. In the case of the shuttle telemetry data (see Figure 7), the algorithm discovers a single motif with two occurrences. For the fetal ECG data, the algorithm finds thirteen occurrences of the heartbeat (see Figure 8). In both cases, the motif neighborhood size was automatically estimated, and the results are qualitatively equivalent to those presented in the literature (see, for example, [Tanaka *et al.*, 2005]).

7   Future Work

We are currently investigating several ways of extending our approach. We are particularly interested in generalizing the motif length parameter by incorporating automatic estimation, allowing variable length occurrences within a motif, and
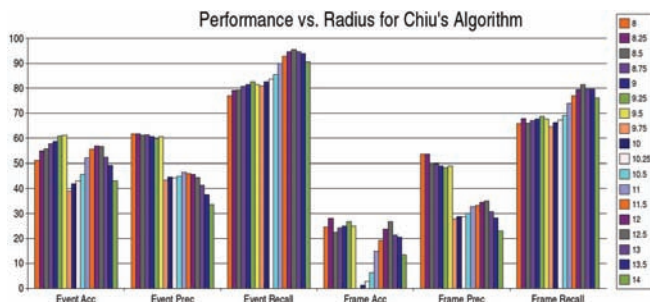

Figure 4: Performance of Chiu's algorithm on the exercise data set over a range of radii. Note the instability (for example, between the radius values of 9.5 and 9.75 for event and frame-based accuracy).
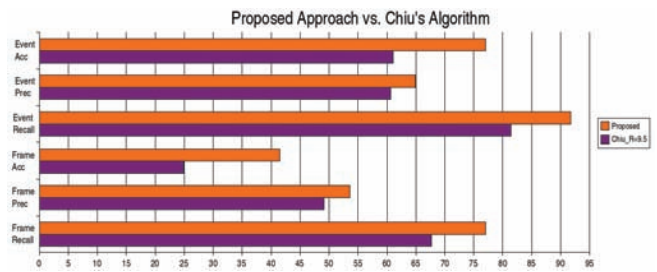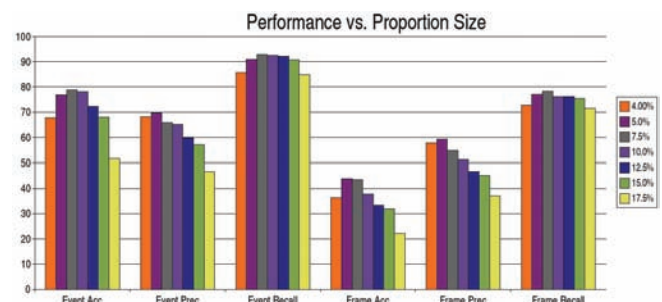

Figure 6: Performance of the proposed approach on the exercise data set showing stability over a relatively wide range of portion sizes used to estimate the neighborhood (5%-15%).
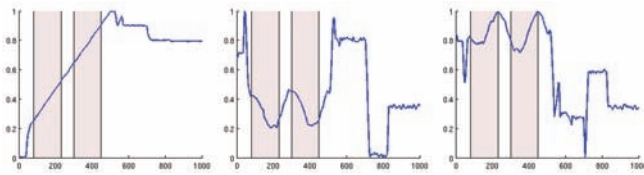
Figure 7: The motif discovered in the 6D shuttle data (the first three dimensions shown). The vertical bars represent the location of the two occurrences of the discovered motif.

detecting motifs at different temporal resolutions. In addition, we are extending our algorithm for discovering motifs that only exist in a subset of the dimensions.

Our other major research focus is on learning higher-level structure from the discovered motifs. Our long-term goal is to learn "from signals to scripts"; that is, we wish the discovery system to be able to take raw sensory data from an activity, discover semantically meaningful primitives, and then learn relationships between them that characterize the activity.

8 Conclusion

Discovering perceptual primitives from sensory data is a fundamental task for intelligent agents. Towards this goal, we have presented an approach to real-valued, multivariate motif discovery that incorporates automatic estimation of the size of each motif neighborhood. Our approach improves on existing methods by removing a difficult-to-specify parameter and improving overall accuracy. We demonstrated the efficacy of our approach in multiple domains including activity discovery from on-body accelerometer and gyroscope data where the event-level error rate was reduced by 44% compared to a leading method.

References

[Bailey and Elkan, 1994] T.L. Bailey and C. Elkan. Fitting a mixture model by expectation maximization to discover motifs in biopolymers. In *Second International Conference on Intelligent Systems for Molecular Biology*, pages 28–36. AAAI Press, 1994.

[Bao and Intille, 2004] Ling Bao and Stephen S. Intille. Activity recognition from user-annotated acceleration data. In A. Ferscha and F. Mattern, editors, *Pervasive*, volume 3001 of *Lecture Notes in Computer Science*, pages 1–17. Springer, 2004.

[Blekas et al., 2003] K. Blekas, D. Fotiadis, and A. Likas. Greedy mixture learning for multiple motif discovery in biological sequences. *Biinformatics*, 19(5), 2003.

[Chiu et al., 2003] B. Chiu, E. Keogh, and S. Lonardi. Probabilistic disovery of time series motifs. In *Conf. on Knowledge Discovery in Data*, pages 493–498, 2003.

[Cohen et al., 2002] Paul R. Cohen, J. Tim Oates, and Carole R. Beal. Robots that learn meanings. In *submitted to First Joint Conference on Autonomous Agents and Multiagent Systems*, 2002.

[Cormen et al., 1997] T. Cormen, C. Leiserson, and R. Rivest. *Introduction to Algorithms*. The MIT Press, 1997.

[Jensen et al., 2006] K.L. Jensen, M. P. Styczynski, Isidore Rigoutsos, and G.N. Stephanopoulos. A generic motif discovery algorithm for sequential data. *Bioinformatics*, 22(1):21–28, 2006.
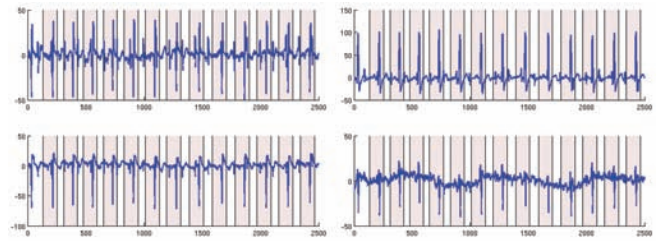
Figure 8: The motif discovered in the 8D fetal ECG data (the first four dimensions are shown) using automatic radius estimation.

[Keogh and Folias, 2002] E. Keogh and T. Folias. UCR time series data mining archive, 2002.

[Keogh et al., 2000] E. Keogh, K. Chakrabarti, M.J. Pazzani, and S. Mehrotra. Dimensionality reduction for fast similarity search in large time series databases. *Journal of Knowledge and Information Systems*, 3(3):263–286, 2000.

[Keogh et al., 2004] E. Keogh, S. Lonardi, and C. Ratanamahatana. Towards parameter-free data mining. In *ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining*, 2004.

[Keogh, 2006] E. Keogh. Personal communication, 2006.

[Lin et al., 2002] J. Lin, E. Keogh, S. Lonardi, and P. Patel. Finding motifs in time series. In *Proceedings of the Second Workshop on Temporal Data Mining*, Edmonton, Alberta, Canada, July 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 *8th ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery*, June 2003.

[Minnen et al., 2003] D. Minnen, I. Essa, and T. Starner. Expectation grammars: Leveraging high-level expectations for activity recognition. In *CVPR 2003*, June 2003.

[Minnen et al., 2006a] D. Minnen, T. Starner, I. Essa, and C. Isbell. Discovering characteristic actions from on-body sensor data. In *International Symposium on Wearable Computers*, pages 11–18, October 2006.

[Minnen et al., 2006b] D. Minnen, T. Westeyn, T. Starner, J.A. Ward, and P. Lukowicz. Performance metrics and evaluation issues for continuous activity recognition. In *Performance Metrics for Intelligent Systems*, 2006.

[Oates, 2002] T. Oates. PERUSE: An unsupervised algorithm for finding recurring patterns in time series. In *Int. Conf. on Data Mining*, pages 330–337, 2002.

[Pavlovic et al., 2000] V. Pavlovic, J.M. Rehg, and J. MacCormick. Learning switching linear models of human motion. In *NIPS*, pages 981–987, 2000.

[Ratanamahatana and Keogh, 2005] C.A. Ratanamahatana and E. Keogh. Three myths about dynamic time warping. In *SIAM Int. Conf. on Data Mining*, pages 506–510, 2005.

[Starner et al., 1998] T. Starner, J. Weaver, and A.P. Pentland. Real-time american sign language recognition using desk and wearable computer based video. *PAMI*, 20(12):1371–1375, Dec. 1998.

[Tanaka et al., 2005] Yoshiki Tanaka, Kazuhisa Iwamoto, and Kuniaki Uehara. Discovery of time-series motif from multi-dimensional data based on mdl principle. *Machine Learning*, 58(2-3):269–300, 2005.