Optimal Subsequence Bijection

Longin Jan Latecki, Qiang Wang, Suzan Koknar-Tezel, and Vasileios Megalooikonomou Temple University Department of Computer and Information Sciences Philadelphia, PA 19094, USA {latecki, wang32, tezel, vasilis}@temple.edu

Abstract

We consider the problem of elastic matching of sequences of real numbers. Since both a query and a target sequence may be noisy, i.e., contain some outlier elements, it is desirable to exclude the outlier elements from matching in order to obtain a robust matching performance. Moreover, in many applications like shape alignment or stereo correspondence it is also desirable to have a one-to-one and onto correspondence (bijection) between the remaining elements. We propose an algorithm that determines the optimal subsequence bijection (OSB) of a query and target sequence. The OSB is efficiently computed since we map the problem's solution to a cheapest path in a DAG (directed acyclic graph). We obtained excellent results on standard benchmark time series datasets. We compared OSB to Dynamic Time Warping (DTW) with and without warping window. We do not claim that OSB is always superior to DTW. However, our results demonstrate that skipping outlier elements as done by OSB can significantly improve matching results for many real datasets. Moreover, OSB is particularly suitable for partial matching. We applied it to the object recognition problem when only parts of contours are given. We obtained sequences representing shapes by representing object contours as sequences of curvatures.

1. Introduction

Sequences of real numbers are commonly used in all research fields. Due to historical reasons, they are also called time series in the data mining community. Time series are a ubiquitous and increasingly prevalent type of data, therefore, there has been much research effort devoted to time series data mining in recent years. Many data mining algorithms have similarity measurements of sequences at their core. Examples include motif discovery [3], anomaly detection [7, 18], rule discovery [6], classification [16] and clustering [1]. In this paper we deal with computation of time series distances based on elastic matching.

To solve the problem of alignment of sequences, Dynamic Time Warping (DTW) [2, 17] has been used. The DTW distance between two sequences is the sum of distances of their corresponding elements. Dynamic programming is used to find corresponding elements so that this distance is minimal. The DTW distance has been shown to be superior to the Euclidean distance in many cases [1, 4, 20, 21]. However, DTW is particularly sensitive to outliers, since it is not able to skip any elements of the sequences. In DTW, each element of the query sequence must correspond to some element of the target sequence and vice versa. Thus, the optimal correspondence computed by DTW is a relation on the set of indices of both sequences, i.e., a one-to-many and many-to-one mapping. The fact that outlier elements must participate in the correspondence optimized by DTW often leads to an incorrect correspondence of other sequence elements. This fact is illustrated in Fig. 1b, where the query sequence on top has one outlier element (spike), and the target sequence (at the bottom) has two outlier elements. The correspondence of the elements is illustrated with straight lines. Observe that the outliers corrupt the correspondence computed by DTW. In particular, this is reflected in the fact that single elements of one sequence correspond to a large number of elements of the other sequence. For comparison, the result of the proposed Optimal Subsequence Bijection (OSB) is shown in Fig. 1a.

The Longest Common Subsequence (LCSS) measure has been used in time series [5, 19] to deal with the alignment and outlier problems. Given a query and a target series, LCSS determines their longest common subsequence, i.e., LCSS finds subsequences of the query and target that best correspond to each other. The distance is based on the ratio between the length of longest common subsequence and the length of the whole sequence. The subsequence does not need to consist of consecutive points, the order of points is not rearranged, and some points can remain unmatched. However, when LCSS is applied to sequences



Figure 1: The top and bottom sequences represent parts of contours of two different but very similar bone shapes. The correspondence obtained by the proposed OSB is shown in (a). Observe how outliers corrupt the result of DTW in (b). The correspondence obtained by LCSS with two different thresholds in (c) and (d) are also mis-aligned.

of numeric values, one needs to set a threshold that determines when values of corresponding points are treated as equal [19]. The performance of LCSS depends heavily on the correct setting of this threshold, which may be a particularly difficult problem for many applications.

The proposed OSB computes the distance value between two sequences based directly on the distances of corresponding elements, just as DTW does, and it allows the ignoring of outlier points on both the query and target sequences to find the best match, just as LCSS does. The main difference between DTW and OSB is that, unlike DTW, OSB can skip outlier elements of the query and target sequences when computing the correspondence. This makes the performance of OSB robust in the presence of outliers. Moreover, OSB defines a bijection on the matched subsequences, which means that we have a one-to-one correspondence of the matched elements.

The main difference between LCSS and OSB is that LCSS optimizes the length of the longest common subsequence and requires a distance threshold, while OSB optimizes directly the sum of distances of corresponding elements. Moreover, the OSB penalty for skipping consecutive elements is proportional to the number of elements skipped, thus skipping one outlier costs less than skipping a consecutive subsequence of several elements. There is no penalty for skipping elements in LCSS, which often leads to accidental matches as illustrated in Fig. 1c. As can be seen also in Figs. 1c and 1d, the threshold chosen greatly affects the resulting correspondence. The main difference between OSB and assignment algorithms, like the Hungarian algorithm [10, 15], is that the Hungarian algorithm, as well as other assignment algorithms, do not preserve the order of sequences.

OSB is a natural extension of Minimum Variance Matching (MVM) [14]. MVM aims at finding a subsequence of the target sequence that best matches the query sequence. While MVM can avoid noise on the target sequence, every element in the query sequence must be matched. This is a serious limitation of MVM. In particular, it implies that the length of the query sequence must be shorter than the length of the target sequence. This problem has been corrected in OSB. Moreover, the theoretical framework for OSB is significantly different from the framework introduced for MVM.

The paper is organized as follows. In Section 2 we introduce our new method OSB together with its underlying theoretical framework. In Section 3 we propose a method to automatically compute the main parameter for OSB which is the cost of jumping over an element of the query and/or target sequence. In section 4 we discuss the time complexity of OSB. Finally, in Section 5 we present the experimental evaluation of the performance of OSB on two types of datasets: benchmark time series datasets and a shape contour dataset. In particular, we show that OSB outperforms DTW, LCSS, and the Euclidean distance.

2. Optimal Subsequence Bijection

The new algorithm, called *Optimal Subsequence Bijection* (*OSB*), works for the elastic matching of two sequences of different lengths m and n:

$$a = (a_1, \ldots, a_m)$$
 and $b = (b_1, \ldots, b_n)$

The goal of OSB is to find subsequences a' of a and b' of b such that a' best matches b'. Skipping (not matching) some elements of a and b is necessary because both sequences may contain some outlier elements. However, skipping too many elements of either sequence increases the chance of accidental matches. To prevent this from happening, we introduce a penalty for skipping. We call this penalty the *jump cost* and denote it with C. We describe one method to compute the jump cost in Section 3.

We assume that the distance function d used to compute the dissimilarity value between elements of sequences a and b, i.e., $d(a_i, b_j)$, is given for $(i, j) \in \{1 \dots m\} \times \{1 \dots n\}$. We do not have any restrictions on the distance function d, and therefore, any distance function is possible. Usually, for sequences of real numbers we simply have the distance $d(a_i, b_j) = (a_i - b_j)^2$, which is also the case for our experimental results reported in Section 5.

The optimal correspondence can be found with a shortest path algorithm on a DAG (directed acyclic graph). The nodes of the DAG are all index pairs $(i, j) \in \{1 \dots m\} \times \{1 \dots n\}$ and the edge cost w is defined as

$$w((i,j)(k,l)) = \begin{cases} \sqrt{(k-i-1)^2 + (l-j-1)^2} \cdot C + d(a_k,b_l) \\ & \text{if } i < k \land j < l \\ \infty & \text{otherwise} \end{cases}$$
(1)

Thus, the cost of an edge from (i, j) to (k, l) is the Euclidean distance of vertices (i, j) and (k, l) in the matrix $\{1 \dots m\} \times \{1 \dots n\}$ times the jump cost plus the dissimilarity measure of elements a_k and b_l .

OSB can be viewed as extension of DTW. To see this, observe that the edge weight dtw for DTW is defined as

$$dtw((i, j)(k, l)) = \begin{cases} (a_k, b_l) & \text{if } i < k \land j < l \land ((k - i) + (l - j) = 1 \lor 2) \\ \infty & \text{otherwise} \end{cases}$$

This means that if *i* maps to *j*, then either k = i maps to l = j + 1 or k = i + 1 maps to l = j + 1 or k = i + 1 maps to l = j in the DTW correspondence. Thus, in comparison to DTW, OSB allows penalized jumps. We view as our main contribution the definition of the edge cost *w* in DAGs with a jump penalty that includes the jump cost defined in Section 3.

The edge cost can be easily extended to impose a warping window constraint, i.e., we can limit the number of elements that can be jumped over in one step by setting a maximal value for the index differences k - i - 1 and l - j - 1.

We illustrate the proposed method on a simple example. Fig. 2a demonstrates the correspondence found with OSB as a shortest path for two sequences a = (1, 2, 8, 6, 8) and b = (1, 2, 9, 15, 3, 5, 9) with the distance between two elements being the squared difference. The corresponding sequence indices computed by OSB (not the values) are (1, 1), (2, 2), (3, 3), (4, 6), and (5, 7). Observe that the outliers $b_4 = 15$ and $b_5 = 3$ are skipped. The automatically computed jump cost is C = 1.15 (see Section 3). For comparison, DTW (Fig. 2b) yields the sequence indices (1, 1), (2, 2), (3, 3), (3, 4), (4, 5), (4, 6), and (5, 7). Thus, $a_3 = 8$ is matched to $b_3 = 9$ (correct) and the outlier $b_4 = 15$, while $a_4 = 6$ is forced to match the outlier $b_5 = 3$ and the correct match $b_6 = 5$.

The output of OSB gives us a *correspondence* defined as a mapping on sequence indices $f : \{1 \dots m\} \rightarrow \{1 \dots n\}$ that is a monotonic injection, i.e., f(i) < f(i+1) for $i \in \{1 \dots m\}$. The sets of indices (i_k) and $(f(i_k))$ $i_k \in \{1 \dots m\}$ define the subsequences a' of a and b' of b, such that f restricted to (i_k) is a bijection. This explains the phrase "subsequence bijection" in *Optimal Subsequence Bijection (OSB)*.



Figure 2: Time series alignment example

3. Penalty for skipping elements

In this section we describe how to determine the jump cost C. As we stated in Section 2, while in most applications $d(a_i, b_j)$ is given for $(i, j) \in \{1 \dots m\} \times \{1 \dots n\}$, the jump cost should be carefully selected. We propose to compute the jump cost in two phases. Let B be a set of all target sequences to which query sequence a should be compared. In the first phase, the query sequence a is compared to a target sequence $b \in B$, and we define

$$C(a,b) = \underset{i}{\operatorname{mean}}(\underset{j}{\operatorname{min}}(d(a_i,b_j))) + \underset{i}{\operatorname{std}}(\underset{j}{\operatorname{min}}(d(a_i,b_j)))$$

Thus, for every element a_i we find the closest element b_j , and then we take the mean plus one standard deviation (std) of the distances to the closest elements. So, for example, if sequences a and b are similar with the exception of one outlier element, call it a_k , then for every a_i with $i \neq k$ we find an element b_j with a small distance $d(a_i, b_j)$. Consequently, C(a, b) will be small, so that the distance to the closest element in b for a_k will be greater than C(a, b), and the element a_k will be excluded from matching by this relatively small penalty, i.e., we will jump over it.

In order to ensure a fair comparison of sequence a to all target sequences $b \in B$, we need to fix the jump cost for all the comparisons. This is obtained in the second phase by simply taking the mean of all the jump costs:

$$C(a) = \operatorname{mean}\{C(a,b) : b \in B\}$$

4. Time complexity

It is well known that the time complexity of DTW is O(mn) for two sequences with m and n elements respectively. Since OSB considers all possible jumps, its time complexity is $O(m^2n^2)$.

It is also well known that in practice DTW is run with a warping window restriction, which means that element a_i from the query sequence can only match to an element b_j from the target sequence if |i - j| is smaller than a predefined constant, say r. The time complexity of DTW with this warping window restriction is O(m). We can also impose the warping window restriction on OSB. In addition, we can limit the number of elements that can be jumped over in one step by setting maximal values for the index differences k - i - 1 and l - j - 1 in formula (1). With these restrictions, the time complexity of OSB is also O(m).

As our experimental results in Section 5 demonstrate, these restrictions do not influence the matching results of OSB. In all experimental results, we set all three index differences to be bounded by five, i.e., |i - j| < 5, k - i - 1 < 5, and l - j - 1 < 5. In contrast, the performance of DTW is evaluated with respect to the best possible warping window restriction.

5. Experimental results

We tested our approach on the benchmark time series datasets published on the UCR Time Series Classification/Clustering Page [9] and on the MPEG-7 shape dataset [11]. These datasets are online to serve the data mining/machine learning community as an effort to encourage reproducible research for time series classification and clustering, and for shape matching. Section 5.1 will discuss the results for the UCR time series datasets, and Section 5.2 will discuss the results for partial sequence similarity on the MPEG-7 dataset.

In addition, we entered the Workshop and Challenge on Time Series Classification at SIGKDD 2007 [8]. In this competition, we were given twenty time series datasets of unknown taxonomy and we had 24 hours to run our algorithm, classify the data, and submit the results. Due to technical problems, we could finish only twelve of the twenty datasets, but using OSB, we had the highest accuracy of all competitors on three of the datasets.

5.1. The UCR time series datasets results

We were able to download 19 datasets from [9] that came from various application domains. The description of these datasets can be found on the UCR webpage.

The table published on [9] reports 1-Nearest Neighbor (1-NN) classification results for three methods: Euclidean Distance, DTW (Dynamic Time Warping) without warping window, and DTW with best scoring warping window. The overall accuracy is reported as the percentage of misclassified queries (i.e., the smaller the value, the better is the result). We followed the same 1-NN classification procedure as reported on [9] to obtain our results and compare OSB to the other methods.

To summarize, OSB yields the best results on 10 datasets (best on 7 datsets and tied for best on 3), followed by DTW without warping window with best results on 8 datasets (best on 4 and tied for best on 4), followed by DTW with best scoring warping window with best results on 6 datasets (best on 2 and tied for best on 4). The Euclidean Distance obtained the best score on only two datasets (tied for best on both datasets). Please see [12] for a complete discussion and the full results.

We do not claim that OSB is always superior to DTW, we claim only that skipping outlier elements as done by OSB can significantly improve matching results for many real datasets, and this fact is clearly demonstrated by our results (see Table 1 in [12]).

5.2. The MPEG7 dataset results

The goal of this section is to test OSB's ability as a measure for partial sequence similarity. Our query sequences represent significant contour parts of shapes in the MPEG-7 Shape 1 Part B dataset [11]. The MPEG-7 dataset is a standard dataset for testing shape-matching algorithms [13]. The dataset consists of 70 object classes (e.g. "Birds") and within each class there are 20 shapes for a total of 1400 shapes. Each shape is represented with 100 equidistant sample points on the contour. We manually selected 10 query sequences as contour segments representing shapes from 10 different classes. They are represented as the black parts of the contours in the first column in Fig. 3. Their lengths range from 30 to 57 points. We then converted both the target shapes and the query segments into sequences by calculating the curvature of each point with respect to its 5 neighbors on each side. This particular transformation makes the sequence representation invariant to rotation and scale changes, which is particularly relevant for finding similar contour parts. In other words, the shape of a cell phone with its antenna pointing up can still match with the same cell phone shape scaled and rotated so that its antenna is now smaller and pointing down. This curvature conversion was done for each target shape (yielding target sequences of length 100), and on each query segment (giving query sequences with lengths 10 less than the original number of points for each segment).

We ran the OSB algorithm for each query against each target to find the optimal subsequence between each query segment and target shape. For each query/target pair we used a step of 5 points to move the starting point of the target. First we compared the query against target elements 1 - 100. Then we compared the query against target elements 6 - 105 (doing a circular shift of the elements), then against target elements 11 - 110, and so on (incrementing by 5).

The final score for a query/target pair is the dissimilarity value of the best matching starting value. Since we have only 10 query sequences, and we want a statistically significant comparison, we used 1, 5, 10 and 20 nearest neighbors' classification. For each query we looked at the one/five/ten/twenty targets that yielded the smallest distance measure among 1400 shapes and calculated the rate of cor-

bird:05.17	bird:05.17	bird:05.19	bird:05.16	dog:33.08	dog:33.04
bone:06.01	bone:06.01	bone:06.07	bone:06.11	bone:06.04	bone:06.13
5	C A	J A	A		5
cell:14.15	cell:14.15	cell:14.16	cell:14.19	cell:14.12	cell:14.10
crown:20.16	crown:20.16	dev1:24.05	crown:20.01	crown:20.11	dev1:24.19
glas:42.13	glas:42.13	glas:42.05	glas:42.01	glas:42.12	glas:42.02
fish:36.09	fish:36.09	fish:36.05	fish:36.11	fish:36.19	fish:36.02
rat:59.16	rat:59.16	rat:59.01	rat:59.11	rat:59.18	rat:59.06
S S S	555	S	S S	S S	53
fount:40.17	fount:40.17	fount:40.02	fount:40.06	fount:40.19	fount:40.15
≪()	≮()	\$	≪)	≪)	€
watch:70.16	watch:70.16	watch:70.17	watch:70.15	watch:70.02	watch:70.14
	}	{}	}	{}	{}
stef:65.01	stef:65.01	stef:65.13	stef:65.12	stef:65.17	stef:65.07
(Z	(m	So the second se	S	L	X

Figure 3: OSB results on the MPEG-7 dataset for our 10 queries. Each contour is identified by label:class:id. The first column shows the query part (in black) in its original contour (red). Columns 2 through 6 are the top five matches out of the 1400 target shapes. The black points on the targets indicate the corresponding points as computed by OSB

rect classification, i.e., how many of the target contours in the top one/five/ten/twenty were from the same class as the query segment. The accuracy reported is the recall rate (sometimes also called the retrieval rate), i.e., of all the possibly correct matches, how many were actually found in the top one/five/ten/twenty.

For comparison, we used the same query and target sequences in three other algorithms: DTW, DTWCW (DTW with Correspondence Window, described below), and LCSS. As shown in Table 1, OSB performed signifi-

	OSB	DTW	DTWCW	LCSS
1NN	100%	0%	90%	90%
5NN	92%	2%	72%	42%
10NN	84%	2%	67%	34%
20NN	67%	3%	59%	26%

Table 1: Recall rates on the MPEG-7 dataset.

cantly better than any of the other algorithms. OSB was the only algorithm to achieve 100% recall rate in the 1NN classification; it had a 92% recall rate in the 5NN classification.

The extremely low performance of DTW is due to the fact that we did not use any window restriction, e.g., if the query sequence is of length 20, we still matched it to the target sequences of length hundred 1-100, 6 - 105, and so on. In order to eliminate this problem, we also considered DTW with Correspondence Window (DTWCW), where we matched the query sequence against a target sequence of the same length, i.e., the correspondence window has the same length as the query sequence. For example, if the query sequence is of length 20, we matched it to subsequences of the target with indices 1-20, 6-25, and so on. The results for DTWCW improved dramatically, but still did not match the performance of OSB. It may be possible to further improve the performance of DTW, by varying the size of the correspondence window, but this would add one more parameter and increase the time complexity (from the linear time of DTW with a warping window to quadratic time if the size of the correspondence window is added). The excellent performance of OSB illustrates that such a window size parameter is not needed, since OSB computes a bijective embedding of the query sequence. This fact clearly demonstrates that it is more suitable for partial sequence similarity.

For DTW, DTWCW, and OSB we used the same warping window of size five. Thus, all three algorithms have the same linear time complexity. LCSS requires a threshold that decides whether two sequence elements are matching or not. We used LCSS with various thresholds and selected the results that gave the best performance (0.3 times the value of the smallest element.) For OSB, we used a fixed jump cost of 10 for each query, which was determined experimentally. (OSB was not sensitive to the value of the jump cost, since we obtained the same results for a large range of jump cost values.)

Fig. 3 shows the retrieval results of OSB. It can be observed that the performance of OSB is actually better than is reflected by the classification based on class labels. For example, the query contour segment representing a bird head in the first row matches to dog heads, which actually have very similar shape.

The best match for each segment (shown in the second column) is the same shape from which we took the query segment. This holds for all the queries and is not a trivial result. Because of the step of 5 we used to position the starting point of the target sequence, the query and target sequences are not identical. In fact, DTW was not able to find any of the queries' original shapes, and DTWCW and LCSS both could find only 9 of the 10.

To summarize, the OSB algorithm performed significantly better than DTW, DTWCW, and LCSS. This difference clearly shows that the ability of skipping outlier elements is essential in partial sequence similarity.

6. Conclusions and future work

The proposed sequence matching method (OSB) can be viewed as an extension of DTW. In comparison to DTW, OSB allows penalized skipping of outlier elements, which we view as one of our main contributions. As demonstrated in the test results in Section 5.1, the ability of skipping outlier elements leads to improved retrieval performance for most datasets. However, for some datasets without significant outliers it may lead to reduced retrieval performance. As demonstrated in the test results in Section 5.2, when dealing with partial shape similarity in the presence of noise, the ability of skipping outlier elements is essential. Here the performance of OSB in the 5NN classification was 90 percentage points higher than the performance of DTW, 20 percentage points higher than that of DTWCW, and 50 percentage points higher than that of LCSS.

References

- Aach and Church. Aligning gene expression time series with time warping algorithms. *Bioinformatics*, 17:495–508, 2001.
- [2] Berndt and Clifford. Using dynamic time warping to find patterns in time series. In *Proceedings AAAI-94 W.* on Knowledge Discovery and Databases, pages 229–248, 1994.
- [3] Chiu, Keogh, and Lonardi. Probabilistic discovery of time series motifs. In Proceedings ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining, Washington, 2003.
- [4] Chu, Keogh, Hart, and Pazzani. Iterative deepening dynamic time warping for time series. In *Proceedings SIAM International Conference on Data Mining*, 2002.
- [5] Das, Gunopulos, and Mannila. Finding similar time series. In *Principles of Data Mining and Knowledge Discovery*, pages 88–100, 1997.
- [6] Hoeppner. Discovery of temporal patterns. learning rules about the qualitative behavior of time series. In *Proceedings* of the 5th European Conference on Principles and Practice of Knowledge Discovery in Databases, Freiburg, pages 192– 203, 2001.

- [7] Keogh, Lonardi, and Ratanamahatana. Towards parameterfree data mining. In *Proceedings ACM SIGKDD Int. Conf.* on Knowledge Discovery and Data Mining, Seattle, 2004.
- [8] Keogh, Shelton, and Moerchen. Workshop and challenge on time series classification at SIGKDD 2007. Website, 2007. http://www.cs.ucr.edu/~eamonn/ SIGKDD2007TimeSeries.html.
- [9] Keogh, Xi, Wei, and Ratanamahatana. Ucr time series classification/clustering page. Website. http://www.cs. ucr.edu/~eamonn/time_series_data/.
- [10] Kuhn. The hungarian method for the assignment problem. Naval Research Logistic Quarterly, 2:83–97, 1955.
- [11] Latecki. Shape data for the MPEG-7 Core Experiment CE-Shape-1. Website. http://www.cis.temple.edu/ ~latecki/TestData/mpeg7shapeB.tar.gz.
- [12] Latecki, Koknar-Tezel, Wang, and Megalooikonomou. Sequence matching capable of excluding outliers. In Proceedings of Workshop on Time Series Classification at ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining (KDD), San Jose, CA, August 2007.
- [13] Latecki, Lakaemper, and Eckhardt. Shape descriptors for non-rigid shapes with a single closed contour. In *IEEE Conf.* on Computer Vision and Pattern Recognition (CVPR), pages 424–429, 2000.
- [14] Latecki, Megalooikonomou, Wang, Lakaemper, Ratanamahatana, and Keogh. Partial elastic matching of time series. In *Proceedings IEEE Int. Conf. on Data Mining (ICDM)*, *Houston*, pages 701–704, 2005.
- [15] Munkres. Algorithms for the assignment and transportation problems. *Journal of the Society of Industrial and Applied Mathematics*, 5(1):32–38, 1957.
- [16] Rafiei. On similarity-based queries for time series data. In Proceedings of the Int. Conf. on Data Engineering, Sydney, pages 410–417, 1999.
- [17] Sakoe and Chiba. Dynamic programming algorithm optimization for spoken word recognition. *IEEE Trans. on Acoustics, Speech, and Signal Processing*, 26(1):43–49, 1978.
- [18] Salvador, Chan, and Brodie. Learning states and rules for time series anomaly detection. In *Proceedings of the 17th Intl. Florida Artificial Intelligence Research Society Conference, Florida*, pages 306–311, 2004.
- [19] Vlachos, Hadjieleftheriou, Gunopulos, and Keogh. Indexing multi-dimensional time-series with support for multiple distance measures. In *Proceedings of ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining, Washington*, pages 216–225, 2003.
- [20] Vlachos, Kollios, and Gunopulos. Discovering similar multidimensional trajectories. In *Proceedings of 18th ICDE, San Jose, CA*, pages 673–684, 2002.
- [21] Yi, Jagadish, and Faloutsos. Efficient retrieval of similar time sequences under time warping. In *Proceedings Int. Conf. on Data Engineering (ICDE98)*, pages 201–208, 1998.