

## MDL-based time series clustering

Thanawin Rakthanmanon · Eamonn J. Keogh ·  
Stefano Lonardi · Scott Evans

Received: 11 December 2011 / Revised: 23 March 2012 / Accepted: 7 April 2012  
© Springer-Verlag London Limited 2012

**Abstract** Time series data are pervasive across all human endeavors, and clustering is arguably the most fundamental data mining application. Given this, it is somewhat surprising that the problem of time series clustering from a single stream remains largely unsolved. Most work on time series clustering considers the clustering of *individual* time series that have been carefully extracted from their original context, for example, gene expression profiles, individual heartbeats, or individual gait cycles. The few attempts at clustering time series *streams* have been shown to be objectively incorrect in some cases, and in other cases shown to work only on the most contrived synthetic datasets by carefully adjusting a large set of parameters. In this work, we make two fundamental contributions that allow for the first time, the meaningful clustering of subsequences from a time series stream. First, we show that the problem definition for time series clustering from streams currently used is inherently flawed, and a new definition is necessary. Second, we show that the minimum description length framework offers an efficient, effective, and essentially parameter-free method for time series clustering. We show that our method produces objectively correct results on a wide variety of datasets from medicine, speech recognition, zoology, gesture recognition, and industrial process analyses.

**Keywords** Time series · Clustering · MDL

---

T. Rakthanmanon (✉) · E. J. Keogh · S. Lonardi  
Department of Computer Science and Engineering, University of California, Riverside,  
Riverside, CA 92521, USA  
e-mail: rakthant@cs.ucr.edu

E. J. Keogh  
e-mail: eamonn@cs.ucr.edu

S. Lonardi  
e-mail: stelo@cs.ucr.edu

S. Evans  
GE Global Research, Niskayuna, NY, USA  
e-mail: evans@ge.com

## 1 Introduction

Time series data are pervasive across almost all human endeavors, including medicine, finance, science, and entertainment. As such it is hardly surprising that it has attracted significant attention in the research community [1,4,25,33]. Given the ubiquity of clustering both as a data mining application in its own right and as a subroutine in other higher-level data mining applications (i.e., summarization, outlier discovery, rule-finding, preprocessing for some classification algorithms etc.), it is surprising that the problem of time series clustering from a *single* time series stream remains largely unsolved, in spite of significant efforts by the community [1,4,33]. Most work on time series clustering considers the clustering of individual time series that have been carefully extracted from their original context, say, gene expressions, or extracted signals such as individual heartbeats. The few attempts at clustering the contents of a *single* time series stream have been shown to be objectively incorrect in some cases [18], and in other cases shown to work only on the most contrived datasets by carefully adjusting a large set of parameters. In this work, we make two fundamental contributions. First, we show that the problem definition for time series clustering from streams currently used is inherently flawed. Any meaningful algorithm must avoid trying to cluster *all* the data. In other words, the subsequences of a time series should only be clustered if they are clusterable. This seems to open up a “chicken and egg” paradox. However, our second contribution is to show that the minimum description length (MDL) framework offers an efficient, effective, and essentially parameter-free solution to this problem. MDL has had a significant impact in bioinformatics and data mining of discrete objects such as natural language [15] but has yet failed to have a significant impact on real-valued data mining [14,22,28].

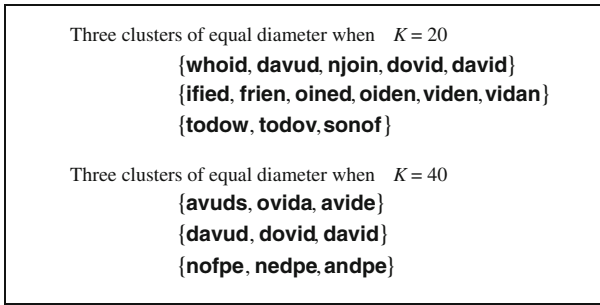
We begin by giving the intuition behind the fundamental observation that motivates and informs our work that clustering of time series from a single stream of data requires ignoring some of the data.

### 1.1 Why clustering time series streams requires ignoring some data

The observation motivating our efforts to cluster time series is that any attempt that insists on trying to explain *all* the data is doomed to failure. To see this, consider one of the most obviously “clusterable” time series data sources: motion-captured sign language, such as American Sign Language (ASL). There has been much recent work on nearest-neighbor classification of such data, with accuracies greater than 90% frequently reported [1]. This suggests that a long data stream of ASL might be amenable to clustering, where each cluster maps to a distinct “word” or “phrase.”

However, all such data contain movement epenthesis (ME) [8,33]. During the production of a sign language sentence, it is often the case that a movement segment needs to be inserted between two consecutive signs to move the hands from the end of one sign to the beginning of the next. These ME segments can be as long as—or even longer than—the true signs and are typically not performed with the precision or repeatability of the actual words, since they have no meaning. Recent sophisticated sign language recognition systems for continuous streams have begun to recognize that “*automated sign recognition systems need a way to ignore or identify and remove the movement epenthesis frames prior to translation of the true signs*” [33].

What we observed about ASL as a concrete and intuitive example matches our experience with dozens of other datasets and indicates that this is a pervasive phenomenon. We believe that almost all datasets have sections of data that do not represent a discrete underlying behavior, but simply a transition between behaviors or random drifts where no behavior is



**Fig. 1** Representative partitional clusters from dataset  $D$  for two settings of  $K$

taking place. In most datasets that we have examined, such sections constitute the *majority* of the data. If we are forced to try to model these in our clusters, they will swamp the true significant clusters. We can best demonstrate this effect and hint at our proposed solution by an experiment on a discrete analog of ASL time series, in this case English *text*.

We emphasize that this is *just a expository example*, and if we were really assigned to cluster such text data, we could do better than the attempt shown below.

Consider the following string  $D$ , which from left to right mentions three versions of the name David (English, Persian, Yiddish) and three versions of the name Peter (English, Croatian, Danish). Note that all names have five letters each, making this problem *apparently* simple.

David enjoined Peter who identified Davud son of Petar friend to Dovid and Peder, to do what...

Here, the words between the names are *exactly* the epenthesis previously referred to. To make it more like our time series problem, we can strip out the punctuation and spacing, leaving us:

davidenjoinedpeterwhoidentifieddavudsonofpetarfriendtodovidandpedertodowhat

The discrete analog of the time series clustering algorithm in [10] would begin by extracting all the subsequences of a given fixed length. Let us assume for simplicity the length five is used, and thus, the data are transformed into:

david  
 avide  
 viden  
 idenj  
 ...  
 owhat

In Fig. 1, we show representative clusters for two values of  $K$ , if we perform partitional clustering as in [10] on this extracted data.

Note that while the cluster of the name variants of David is discovered, we find that under any setting of  $K$ , there are equally significant *meaningless* clusters, for example {nofpe, nedpe, andpe}. This is in spite of the fact that this can be considered a particularly easy task. Exactly 40% of the signal consists of data we hope to recover, and we deliberately avoided name variants of different lengths (i.e., pieter, pere). In more realistic settings, we expect much less of the data to contain meaningful signals. Note also that the problem is not mitigated by using other clustering variants. The problem is inherent in the false assumption that a clustering of a single stream that must explain *all* such data could ever produce meaningful results [18].

## 1.2 How MDL can help

In contrast to the previous section, it is instructive to see what our proposed algorithm will do in this case. While the details of our algorithm are not introduced until Sect. 4, we can still outline the basic intuition here.

The original string  $D$  has a bit-level representation whose length we denote as  $DL(D)$ . Our algorithm can be imagined as attempting to losslessly compress the data by finding repeated structure in it. As there is little *exactly* repeated structure, we must find approximately repeated structure and encode the differences. For example, if we find the approximately repeated versions of the name “david,” we can think of one version as being a model or *hypotheses* for the data and encode only the difference between the other occurrences:

$$H_1 = \{1:\mathbf{david}\}$$

`1__enjoinedpeterwhoidentified1__u_sonofpetarfriendto1_o__andpertodowhat`

In terms of MDL, we can see **david** as a partial hypothesis  $H_1$  or description of the data. This model has some size, which is simply the length in bits of the word  $DL(H_1) = DL(\mathbf{david})$ . In addition, the size of the remaining data was both *reduced* by factoring out the common structure and (slightly) *increased* by the overhead of the pointers to the dictionary, etc.<sup>1</sup> When encoded with the hypothesis, the length (in bits) of the description of the data is given as  $DL(D|H_1)$ . The total cost of both the hypothesis and the data encoded using the hypothesis is just  $DL(H_1) + DL(D|H_1)$ .

Because this sum is less than the length of the original data  $DL(D)$ , we feel that we are making progress. Perhaps, however, there is more structure we can exploit. A brief inspection of the data suggests another model,  $H_2$ , that exploits *both* repeated names:

$$H_2 = \{1:\mathbf{david} \ 2:\mathbf{peter}\}$$

`1__enjoined2__whoidentified1__u_sonof2__a_friendto1_o__and2_d__todowhat`

Because  $DL(H_2) + DL(D|H_2) < DL(H_1) + DL(D|H_1)$ , we prefer this new hypothesis as a model of the data.

Are we now done? We can try other hypotheses. For example, we could consider the hypothesis  $H_3 = \{1:\mathbf{david} \ 2:\mathbf{peter} \ 3:\mathbf{ono}\}$ , attempting to exploit the two occurrences of a pattern “**o\*o**” (i.e.,...**sonof..** and **..to do..**). However, because this pattern is short, and only has two occurrences, we cannot break even with the cost of the overhead:

$$DL(H_2) + DL(D|H_2) < DL(H_3) + DL(D|H_3)$$

Because we cannot find any other hypotheses that produce a smaller model, we invoke the MDL principle to claim that  $H_2 = \{1:\mathbf{david} \ 2:\mathbf{peter}\}$  is the best model of the data  $D$ . Here, *best* means something beyond simply achieving the greatest compression. We can claim that MDL approach has achieved the most parsimonious explanation of the data, recovering the true underlying structure [9,13,15,16,21]. In at least this case, where the sentence was contrived as an excuse to use two names trice, MDL *did* recover the true underlying structure.

<sup>1</sup> In this toy example, we are deliberately glossing over the concrete details of how the pointers are represented and how the amount compression achieved is measured, etc. [21]. We will formalize these details in Sect. 3.

Note that while our informally stated algorithm does manage to recover the two embedded clusters, it *does not* attempt to explain all of the data. This is a critical observation, in order to cluster a single stream of data, be it discrete or real-valued, we *must* be able to represent and rank solutions that ignore some of the data.

## 2 Related work

The tasks of clustering *multiple* time series streams, or many individual time series (i.e., gene expressions) have received significant attention, but the solutions do not inform the problem we consider here, the task of clustering a *single* time series stream. The most commonly referenced technique for clustering a *single* time series stream is presented in [10] as a sub-routine for rule discovery in time series. In essence, the method slides a fixed length window across the stream, extracting all subsequences that are then clustered with  $K$ -Means. The reader may have already spotted a flaw here; the algorithm tries to explain *all* the data. In [18] (and follow-up works by more than twenty other authors [5, 6, 11]), it was shown that this method can only produce cluster centers that are sine waves, and the output of the algorithm is essentially *independent of the input*. Note that even if the algorithm did not have these fatal flaws, it assumes the cluster all have equal length and that we know the correct value of  $K$ . As we shall show, our method requires neither assumption.

Since the problem with [10] was pointed out in 2005 [18], at least a dozen solutions have been proposed. In Sect. 5.3, we show that the most referenced of these works [6] does not produce objectively correct results, even after extensive parameter tuning by the original authors on a relatively simple problem.

While there have been some efforts to use MDL with time series [26, 30], they all operate on a quantized representation of the data. This has the disadvantage of requiring three parameters (cardinality, dimensionality, and window size), eliminating the greatest advantage of MDL, its intrinsically parameter-free nature.

While MDL has had surprisingly little impact in data mining, it is a tool of choice for many bioinformatics problems. For example, working with RNA data, Evans et. al. have proposed a method using data compression and the MDL principle that is capable of identifying motif sequences, some of which were discovered to be miRNA target sites implicated in breast cancer [13]. Moreover, the authors showed the generality of their ideas by applying them, unmodified, to the problem of network traffic anomalies [14]. There is also a significant work on using MDL to mine graphs [16, 25], dating back to classic work by Cook et al. [9].

Finally, we note that the task was informed by, and may have implications for many other time series problems, including time series segmentation<sup>2</sup> [4]. To see why, let us revisit the technique of text analogy. It is not obvious how one should segment the three concatenated words “hisabasiais,” Perhaps the best we could do is to exploit the known frequencies of bigrams and trigrams, etc. In fact, most time series segmentation algorithms essentially do the real-valued equivalent of this [4]. However, if we see another such triplet of three concatenated words from later in the same stream, for example “withoutabasiais,” we can immediately see that “abasia” must be a word.<sup>3</sup>

<sup>2</sup> The phrase “*time series segmentation*” is unfortunately overloaded. It can mean approximating the data with the smallest number of piecewise polynomial segments for a given error threshold, or as here; extracting small, discrete, *semantically* meaningful segments of data [4].

<sup>3</sup> *Abasia* is the inability to walk due to impaired muscle coordination.

### 3 Background and notation

#### 3.1 Definitions and notation

We begin by defining the data type of interest, *time series*:

**Definition 1** A *time series*  $T$  is an ordered list of numbers.  $T = t_1, t_2, \dots, t_m$ . Each value  $t_i$  can be any finite number (e.g., for two-byte values, they could be integers in range  $[-32, 768, 32, 767]$ ) and  $m$  is the length of time series  $T$ .


Before continuing, we must make and justify a choice. The MDL technique that is at the heart of our algorithm requires *discrete* data, but most time series datasets use four or eight bytes per value and are thus *real-valued* [23]. Our solution is simply to cast the *real-valued* numbers into a reduced cardinality version. Does such a reduction lose meaningful information? To test this, we did one nearest-neighbor classification on eighteen public time series datasets, for cardinalities from the original four bytes down to a single bit. Figure 2 shows the results. As we can see, we can drastically reduce cardinality without reducing accuracy. The original four-byte cardinality is typically a by-product of file format convention or hardware specification, and not a claim as to the intrinsic cardinality of the data.

We note that there may be other things we *could* have done. For example, the MML framework [32] that is closely related to MDL would allow us to work in original continuous space. However, we choose MDL because it is more familiar and it allows for a more intuitive explanation of our algorithms. Likewise, we have at least a dozen choices of how to discretize the time series (adaptive binning, uniform binning, SAX, etc.); however, after testing all published algorithms and finding it made little or no difference, we settled on the simple idea shown below in Definition 3.

Based on the observations in Fig. 2, we will simply use 64-value (6-bit) cardinality in the rest of this work.

While the source data are one long time series, we ultimately wish to cluster it into sets of shorter *subsequences*:

**Definition 2** A *subsequence*  $T_{i,k}$  of a time series  $T$  is a short time series of length  $k$  that starts from position  $i$ . Formally,  $T_{i,k} = t_i, t_{i+1}, \dots, t_{i+k-1}$ ,  $1 \leq i \leq m - k + 1$ .



**Fig. 2** Classification accuracy on 18 time series datasets as a function of the data cardinality. Even if we reduce the cardinality of the data from the original 4,294,967,296 to a mere 64 (vertical bar), the accuracy does not decrease

As we previously noted, we are working in a space of reduced cardinality. Because comparing time series with different offsets and amplitudes is meaningless [18], we must (slightly) adapt the normalization process for our *discrete* representation:

**Definition 3** A *discrete normalization* function  $\text{DNorm}$  is a function to normalize a real-valued subsequence  $T$  into  $b$ -bit discrete value of range  $[1, 2^b]$ . It is defined as followings:

$$\text{DNorm}(T) = \text{round} \left( \left( \frac{T - \min}{\max - \min} \right) \times (2^b - 1) \right) + 1$$

where  $\min$  and  $\max$  are the minimum and maximum value in  $T$ , respectively.

Based on the results in Fig. 2,  $b$  is fixed at 6 for all experiments. We need to define a distance measure; we use the ubiquitous *Euclidean distance* measure:

**Definition 4** The distance between two subsequences  $T_{i,k}$  and  $T_{j,k}$  is the *Euclidean distance* (ED) between  $T_{i,k}$  and  $T_{j,k}$ . Both subsequences must be in the same length. Hence, it is:

$$\text{Dist}(T_{i,k}, T_{j,k}) = \sqrt{\sum_{l=0}^{k-1} (t_{i+l} - t_{j+l})^2}$$

As we shall see later, the Euclidean distance is not general enough to support clustering from time series streams; nevertheless, it is still a useful subroutine to speed up our more general measures. As generally noted [3, 19, 31, 34], the Euclidean distance is a fast and robust distance measure.

For both the full time series  $T$  and any subsequences derived from it, we are interested in knowing how many bits are necessary to represent it. Normally, the number of bits depends solely on the data format, which is typically a reflection of some arbitrary choices of hardware and software. In contrast, we are interested in knowing the *minimum* number of bits to exactly represent the data. In the general case, this number is not calculable, as it is the Kolmogorov complexity of the time series [22]. However, there are numerous ways to approximate this, using Huffman coding, Shannon–Fano coding, etc. Because entropy is a lower bound on the average code length from any such encoding, we can use the *entropy* of the time series as its description length:

**Definition 5** The *entropy* of a time series  $T$  is defined as following. For special case when  $P = 0$ ,  $P \log_2 P$  is defined as 0.


$$H(T) = - \sum_t P(T = t) \log_2 P(T = t)$$

We can now define the description length of a time series.


**Definition 6** A *description length*  $DL$  of a time series  $T$  of length  $m$  is the total number of bits required to represent it, that is  $DL(T) = m * H(T)$ .

The  $DL$  of a time series using entropy clearly depends on the data itself, not just arbitrary representational choices. Figure 3 shows four time series, which all require 250 bytes to characterize in the original representation, but which have differing *entropies* and thus different description lengths.

The reader may have anticipated the following observation. While the (slightly noisy) straight line  $B$  has high entropy, we would subjectively consider it a simple shape. It is



**Fig. 3** Four time series of length 250 and with a cardinality of 256. Naively all require 250 bytes to represent, but they have different *description lengths*



**Fig. 4** Time series B can be represented exactly as the sum of the straight line  $H$  and the difference vector  $B'$

simple given our belief (hypothesis) that it is a slightly corrupt version of a straight line. If  $H$  is this hypothesis, then we can consider instead the entropy of a time series  $B'$ , which as shown in Fig. 4 is simply  $B$  encoded using  $H$  and written as  $B' = (B|H)$ . As a practical matter, to use  $H$  to encode  $B$ , we simply subtract  $H$  from  $B$  to get a difference vector  $B'$  and encode this simpler vector  $B'$ .

While the vector  $B'$  is also of length **250**, it has only 10 unique values, all of which are small in magnitude; thus, its entropy rate is only **2.51** bits. In contrast,  $B$  has 172 unique values and an entropy rate of 7.29. Note that if we are given only  $B'$ , we cannot reconstruct  $B$ ; we also need to know the slope and mean of the line. Thus, when reporting the overall number of bits for  $B'$ , we must also consider the number of bits it takes to encode the hypothesis (the line  $H$ ). We can encode the line simply by recording the heights' **two** locations, the first and last points,<sup>4</sup> each of which requires a single **byte**. Thus, the number of bits required to represent  $B$  using our hypothesis is:

$$DL(B) = DL(H) + DL(B|H) = (2 * 8) + (250 * 2.51) = 643.5 \text{ bits}$$

which is significantly less than the 1,822 bits required for the naive encoding of  $B$  without any hypothesis.

Note the straight line would not help in reducing the number of bits required to represent time series  $C$ , but using a sine wave as the hypothesis *would* significantly help. This observation inspired one of the principle uses of MDL, model section [28]. Statisticians use this principle to decide whether some noisy observations suggest an underlying physical model is produced by, say, a piecewise linear model as opposed to a sinusoidal model. However, our work leverages off a simple but unexploited observation. The hypotheses are not limited to well-defined functions such as sine waves, wavelet basis functions, polynomial models, etc. The hypothesis model can be *any arbitrary time series*. We will see how this observation can be exploited in detail later, but in brief: if  $k$  subsequences of a stream truly form a cluster, then it should be possible to store them in less space by encoding them as a set of difference

<sup>4</sup> If we know the time series is z-normalized, we only need one byte to record the line.



vectors to the mean of all of them. Thus, we have a potential test to guide our search for clusters.

Having seen this intuition, we can now formalize the notion of *hypothesis* as it pertains to our problem:

**Definition 7** A *hypothesis*  $H$  is a subsequence used to encode one or more other subsequences of the same length.

As a practical matter, the encoding we use is the one visualized Fig. 4, we simply subtract hypothesis  $H$  from the target subsequence(s) and encoded the difference vector(s). We could encode the difference vector(s) with, say, Huffman encoding, but as we noted in Definition 5, we really only care about the *size* of the encoding, so we simply measure the entropy of the difference vector(s) to get a lower bound of the size of encoding.

A necessary (but not sufficient) condition to place two subsequences  $H$  and  $B$  into the same cluster is:

$$DL(B) > DL(B|H)$$

This inequality requires that the subsequence  $B$  takes fewer bits to represent when  $H$  is used as a basis to encode it, encoding the intuition that the two subsequences are related or similar.

We can hint at the utility of thinking about our data in terms of hypothesis encoding by revisiting our text example. When a clustering text stream, would it be better to merge  $\mathbf{A}$  or  $\mathbf{B}$ ?

$$\mathbf{A} = \{\text{david, dovid}\}, \mathbf{B} = \{\text{petersmith, petersmidt}\}$$

The first case allows a tight cluster of two short words, is that better than a looser but longer cluster  $\mathbf{B}$ ? The problem is exacerbated when we consider the possibility of clusters with more than two members: how would we rank the relative utility of the tentative cluster  $\mathbf{C} = \{\text{bob, rob, hob}\}$ ?

Normally, clustering decisions are made by considering *Euclidean distance* (or its text counterpart, *Hamming distance*); however, Euclidean distance only allows meaningful comparisons when all the subsequences are the same length. The solution for text, to use the length-normalized Hamming distance, *cannot* be generalized here. The reason is subtle and underappreciated, suppose we have two subsequences of length  $k$  that are distance  $d$  apart. If we truncate the end points and measure the distance again, we might find it has increased! This is because we should only compare z-normalized time series when using Euclidean distance,<sup>5</sup> and after (re)z-normalizing the slightly shorter subsequences, we may find they have grown further apart. Thus, the z-normalized Euclidean distance function is not linear in length and is not even monotonic.

We have already hinted at the fact that the  $DL$  function can use extra information, by using “given,” that is,  $DL(B'|H)$  is the  $DL$  of  $B'$  given  $H$ . We can now formalize this notion:

**Definition 8** A *conditional description length* of a subsequence  $A$  when a hypothesis  $H$  is given is

$$DL(A|H) = DL(A - H)$$

Recall from Figs. 3 and 4 that the  $DL$  of a subsequence depends on the structure of the data. For example, a constant line has a very low  $DL$ , whereas a random vector has a very high

<sup>5</sup> The solution of not normalizing the time series would mitigate this problem, but measuring the Euclidean distance between two time series with different offsets or amplitudes produces meaningless results [18].

*DL*. If  $A$  and  $H$  are very similar, their difference ( $A - H$ ) will be close to a constant line and thus have a tiny *DL*. In essence then, the *DL* function gives us a parameter-free test to see whether two subsequences should be clustered together.

We generalize the notion of *DL* to multiple sequences next. We can apply the same spirit by using a hypothesis to calculate the minimum number of bits required to keep a cluster. We call this *description length of a cluster*:

**Definition 9** A *Description Length of a Cluster (DLC)*  $\mathbf{C}$  is the number of bits needed to represent all subsequences in  $\mathbf{C}$ . In this special case,  $H$  is the center of the cluster. Hence, the description length of cluster  $\mathbf{C}$  is defined as:

$$DLC(\mathbf{C}) = DL(H) - \max_{A \in \mathbf{C}} DL(A|H) + \sum_{A \in \mathbf{C}} DL(A|H)$$

The above *DLC* gives us a primitive to measure the reduction in bits achieved by encoding data with a hypothesis. The two right terms in the equation record the number of bits needed to represent the cluster  $\mathbf{C}$ . Concretely, if the cluster  $\mathbf{C}$  contains  $n$  subsequences, we need to keep at most  $n$  subsequences for reconstructing  $\mathbf{C}$ . We choose to keep the hypothesis  $H$  and the  $n - 1$  smallest differences. Therefore, the description length of the cluster  $\mathbf{C}$  is the combination of the description length of the hypothesis  $H$  and the minimum of any  $n - 1$  conditional description lengths with respect to  $H$ .

Our clustering algorithm is essentially a search algorithm. Three operators avail of the *DLC* definition to test how many bits a particular choice can save. Thus, these three operators fall under the umbrella definition of *bitsave*:

**Definition 10** A *bitsave* is the total number of bits saved after applying an operator that creates a new cluster, adds a subsequence to an existing cluster, or merges two existing clusters together. It is the difference in the number of bits *before* and *after* applying a given action:

$$bitsave = DL(Before) - DL(After)$$

In detail, the *bitsave* for each operator is defined as following:

- 1) *Creating* a new cluster  $\mathbf{C}'$  from subsequences  $A$  and  $B$

$$bitsave = DL(A) + DL(B) - DLC(\mathbf{C}')$$

- 2) *Adding* a subsequence  $A$  to an existing cluster  $\mathbf{C}$

$$bitsave = DL(A) + DLC(\mathbf{C}) - DLC(\mathbf{C}')$$

where  $\mathbf{C}'$  is the cluster  $\mathbf{C}$  after including subsequence  $A$ .

- 3) *Merging* cluster  $\mathbf{C}_1$  and  $\mathbf{C}_2$  to a new cluster  $\mathbf{C}'$ .

$$bitsave = DLC(\mathbf{C}_1) + DLC(\mathbf{C}_2) - DLC(\mathbf{C}')$$

Note that, as we discussed earlier, we do *not* use Euclidean distance to make decisions about which subsequences to place into which clusters. We use only use Euclidean distance in two subroutines: motif discovery and finding the closest subsequence from a given cluster center. Next, we will define closest subsequence or the *nearest neighbor*:

**Definition 11** A *nearest neighbor* of the given subsequence  $A$  is the subsequence  $B$  such that

$$\text{Dist}(A, B) \leq \text{Dist}(A, X) \quad \text{for any subsequences } X$$

Another definition, which we borrow from the literature [24], is *time series motifs*.

**Definition 12** A *time series motif* is pair of subsequence  $A$  and  $B$  such that

$$\text{Dist}(A, B) \leq \text{Dist}(X, Y) \quad \text{for any subsequences } X \neq Y, A \neq B$$


Similar to the nearest neighbor, the time series motif contains two most similar subsequences in the given time series. We use the *simple-but-robust* Euclidean distance as *dist* function (cf. Definition 4) for finding both time series motif and the nearest neighbor of the given subsequence. Note that if subsequence  $X$  and  $Y$  are overlapped, it may lead to the discovery of trivial matches. For more details about the time series *motif*, we refer the reader to [24]. In next section, we will explain our algorithm in detail.

## 4 Clustering algorithm

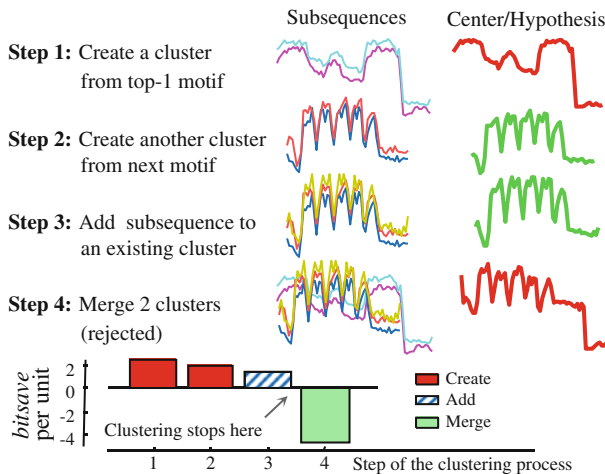
Having introduced the necessary notation, we are finally in a position to introduce our algorithm. We begin by giving a simple text and visual intuition in the next section and follow by giving detailed and annotated pseudo code in Sect. 4.2.

### 4.1 The intuition behind stream clustering

Recall that our input is a single time series like the one shown in Fig. 5 (bottom) and our required output is a set of clusters—possibly of different lengths and sizes. Recall that the union of all the subsequences in this set of clusters may only cover a fraction of the input time series. Indeed, for pathological cases we are given a pure noise time series, we want our



**Fig. 5** Two interwoven bird calls featuring the *Elf Owl*, and *Pied-billed Grebe* are shown in the original audio space (top), and as a time series extracted by using MFCC technique (middle) and then clustered by our algorithm (bottom)



**Fig. 6** A trace of our algorithm on the bird call data shown in Fig. 5(bottom)

algorithm to return a *null* set of clusters. In Fig. 5, we show our running example. It contains the interwoven calls of two very different species of birds.

Our proposed clustering algorithm is a bottom-up greedy search over the space of clusters. For the moment, we will ignore the computational effort that it requires and simply explain *what* is done, leaving the *how* it is (efficiently) done for the next section.

Our algorithm is an iterative merging algorithm similar in spirit to an agglomerative clustering algorithm [17]. However, the differences are telling and worth enumerating:

- Our algorithm typically stops merging before explaining all the data, thus producing a *partitioning* a subset of the data, not producing a *hierarchy* of all the data.
- Agglomerative clustering algorithms are typically implemented such that they require quadratic space; our algorithm has only linear space requirements.<sup>6</sup>
- Most critically, agglomerative clustering algorithms assume the  $K$  items of a *fixed* dimensionally (subsequence length) to be clustered are inputs to the algorithm. However, we do not know how many items will ultimately be clustered, or even how long the items will be.

Similar to agglomerative clustering, we have a search problem that uses operators, in our case, *create*, *add*, and *merge* (Definition 10). When the algorithm begins, only *create* is available to us.

We begin by finding the best initial pair of subsequences to combine so that we may *create* a cluster of two items. To find this best pair, we treat one as a hypothesis and see how well it encodes the other (Definition 8). The pair that reduces the bit cost the most is the pair of choice. This is shown in Fig. 6 as **Step 1**.

There is a potential problem here. Even if we fix the length of subsequences to consider to a constant  $s$ , the number of candidate pairs to consider is quadratic in the length of the time series, approximately  $O((m - s)^2/2)$ . Furthermore, there are no known shortcuts that let us search this space in sub-quadratic time.

<sup>6</sup> Linear space agglomerative clustering algorithms *do* exist, but require highly multiply redundant calculations to be performed, and are thus rarely used due to their lethargy.

The solution to this problem is to note that Euclidean distance and conditional description length are highly correlated where either of them is small. We can leverage off this fact because there exist very fast algorithms to find the closest pair of subsequences (which are known as *time series motifs* [24]) under Euclidean distance. So rather than a brute force search using the conditional description length, we do a fast motif search and then test the motif pair's conditional description length.

In the next stage of the algorithm, there are two operators available to us. We can either *add* a third item to our existing cluster of size two, or we can *create* a new cluster, possibly of a different length. In Fig. 6 in **Step 2**, we can see that in this case, our scoring functions suggest creating a new cluster is the better option in.

In the subsequent phase of the algorithm, it happens that *all* operators are available to us: we could try to *create* a new cluster, we could *merge* our two existing clusters, or we could *add* a subsequence to one of our two clusters. As we can see in Fig. 6, **Step 3**, the last option is chosen.

In the next iteration, the cheapest operator was to merge our two existing clusters as shown in **Step 4**. However, doing this does not decrease the size of the representation—it *increases* it. As such, our algorithm terminates after returning the two clusters it had created up to **Step 3**. The only other way that our algorithm can terminate is if it simply runs out of data to cluster.

## 4.2 Our algorithm in detail

As we noted in the last section, our algorithm is a bottom-up search algorithm. The input is a single time series, and the output is a set of clusters of subsequences. Our algorithm can cluster subsequences of different lengths, and it does not require the number of clusters to be specified.

There are three operators in our search algorithm: *create*, *add*, and *merge*. In each step, we do all (legal) operations and choose the operator that maximizes the number of bits saved as measure by *bitsave* (Definition 10). The current clusters are updated with respect to that choice.

The algorithm can terminate in just two ways; either the best possible choice cannot save any bits, or all data are used up.

Most attempts to cluster time series [6, 10, 11] suffer from a surfeit of parameters. Our algorithm allows essentially none. However, if we allow subsequences that are too short, we can get pathological results in some cases. For example, there are only two possible  $z$ -normalized subsequences of length two. Moreover, a user may wish to bias the algorithm toward certain clustering. For example, for electrical power demand load, we may be interested in weekly *or* daily patterns. Thus, as shown in Table 1, we allow the user the option of suggesting an approximate length  $s$ .

The algorithm begins by initializing the cluster set to *empty*, then it enters a loop until no more bits can be saved (line 2), or it runs out of data. Within each iteration the loop, we perform three operators *create* (line 4–9), *add* (line 10–15), and *merge* (line 16–22), and we keep the results of the most parsimonious operator.

For the *create* process, we call a subroutine to find time series motifs under Euclidean distance using the fastest currently known technique [24]. Because we do not know how long the subsequences in the cluster should be, the algorithm runs `MotifDiscovery` multiple times on different lengths of motif (line 5). If the new cluster is created, then the number of bits saved is calculated (line 7). The temporary version of updated clusters are kept (line 8) and used if the algorithm eventually chooses to *create* this cluster (line 24). Recall that the details of function `ComputeBitsave` are provided in Definitions 9 and 10.

**Table 1** Main time series stream clustering algorithm

<b>Input:</b>	<i>ts</i> : time series, <i>s</i> : approximate length
<b>Output:</b>	<i>cluster</i> : final cluster of subsequences
1	<i>cluster</i> = {}
2	<b>while</b> <i>bitsave</i> >0
3	<i>bitsave</i> =-∞, <i>i</i> =0 // create new cluster
4	<b>for</b> <i>len</i> = <i>s</i> to 2 <i>s</i>
5	( <i>A</i> , <i>B</i> ) = MotifDiscovery( <i>ts</i> , <i>len</i> )
6	<i>C'</i> = CreateCluster( <i>A</i> , <i>B</i> )
7	<i>bs.append</i> (ComputeBitsave( <i>C'</i> , <i>A</i> , <i>B</i> ))
8	<i>cluster'</i> .append( <i>cluster</i> ∪ { <i>C'</i> })
9	<b>end for</b>
	// add subsequence to an existing cluster
10	<b>for</b> <i>C</i> ∈ <i>cluster</i>
11	<i>A</i> = NearestNeighbor( <i>ts</i> , <i>C</i> )
12	<i>C'</i> = AddToCluster( <i>C</i> , <i>A</i> )
13	<i>bs.append</i> (ComputeBitsave( <i>C'</i> , <i>C</i> , <i>A</i> ))
14	<i>cluster'</i> .append( <i>cluster</i> ∪ { <i>C'</i> } - { <i>C</i> })
15	<b>end for</b>
	// merge 2 clusters
16	<b>for</b> <i>C</i> <sub>1</sub> ∈ <i>cluster</i>
17	<b>for</b> <i>C</i> <sub>2</sub> ∈ <i>cluster</i> and <i>C</i> <sub>1</sub> ~ <i>C</i> <sub>2</sub>
18	<i>C'</i> = MergeClusters( <i>C</i> <sub>1</sub> , <i>C</i> <sub>2</sub> )
19	<i>bs.append</i> (ComputeBitsave( <i>C'</i> , <i>C</i> <sub>1</sub> , <i>C</i> <sub>2</sub> ))
20	<i>cluster'</i> .append( <i>cluster</i> ∪ { <i>C'</i> } - { <i>C</i> <sub>1</sub> } - { <i>C</i> <sub>2</sub> })
21	<b>end for</b>
22	<b>end for</b>
	// update the result
23	[ <i>bitsave index</i> ] = max( <i>bs</i> );
24	<i>cluster</i> = <i>cluster'</i> ( <i>index</i> );
25	<b>end while</b>

It is possible to *add* a subsequence into an existing cluster (line 10–15). We first find the most similar subsequence in the input time series with respect to the center of a given cluster (line 11); we can achieve this task by using any nearest-neighbor search algorithm [12], including brute force search. After the search, the cluster is updated to include that nearest subsequence (line 12), the number of bits saved is calculated, and the temporary clusters are recorded (line 13–14).

For our last operator, any pair of clusters is allowed to *merge* (line 18); we then compute the number of bits saved for each pair and record the temporary cluster.

After the algorithm measures the number of bits saved from all possible choices, the final cluster is updated with respect to the choice that maximizes the number of bits saved (line 23–24).

We have glossed over an important detail: the two items being combined by the merge/add operators may be of different lengths. To allow this critical flexibility, we use a simple data structure to record a cluster. For any given cluster *C*, *C.size* records the number of subsequences in the cluster, *C.cen* is the center of the cluster, *C.seq* is a set of subsequences in the cluster, and *C.shift* is a set of shift positions (i.e., offsets) of each subsequence in *C* when it aligned to the *C.cen*. Note that to compute the conditional description length (Definition 8), a subsequence and its hypothesis must be of the same length.

**Table 2** Create operator

Function $C = \text{CreateCluster}(A, B)$	
1	$C.size = 2;$
2	$C.cen = (A+B)/2$
3	$C.seq = [A; B]$
4	$C.shift = [0; 0]$

**Table 3** Add operator

Function $C = \text{AddToCluster}(C, A)$	
1	$C.cen = (C.cen * C.size + A * 1) / (C.size + 1)$
2	$C.size = C.size + 1$
3	$C.seq = [C.seq; A]$
4	$C.shift = [C.shift; 0]$

**Table 4** Merge operator

Function $C' = \text{MergeClusters}(C_1, C_2)$	
1	$C'.seq = [C_1.seq; C_2.seq]$
2	$C'.size = C_1.size + C_2.size$
3	$n_1 = C_1.size, \quad m_1 = \text{length}(C_1.cen)$
4	$n_2 = C_2.size, \quad m_2 = \text{length}(C_2.cen)$
5	$i = 0$
6	<b>for</b> $off = 0$ to $m_2$
7	$cen1 = [C_2.cen(1, off), C_1.cen]$
8	$cen2 = [C_2.cen, C_1.cen(1, m_1 + off - m_2)]$
9	$C'.cen = (cen1 * n_1 + cen2 * n_2) / (n_1 + n_2)$
10	$C'.shift = [C_1.shift + off; C_2.shift]$
11	$bs.append(\text{ComputeBitsave}(C', C_1, C_2))$
12	$Ctmp.append(C')$
13	<b>end for</b>
14	<b>for</b> $off = 1$ to $m_1$
15	$cen1 = [C_1.cen, C_2.cen(1, m_2 + off - m_1)]$
16	$cen2 = [C_1.cen(1, off), C_2.cen]$
17	$C'.cen = (cen1 * n_1 + cen2 * n_2) / (n_1 + n_2)$
18	$C'.shift = [C_1.shift; C_2.shift + off]$
19	$bs.append(\text{ComputeBitsave}(C', C_1, C_2))$
20	$Ctmp.append(C')$
21	<b>end for</b>
22	$[bitsave\ index] = \max(bs);$
23	$C' = Ctmp(index);$

Table 2 shows how a new cluster can be created from two subsequences of the same size. Because those two subsequences are from motif discovery under Euclidean distance, their align position is set to 0 (line 4). The center of new cluster is the average of those two subsequences. In Table 3, when we want to add a subsequence  $A$  to an existing cluster  $C$ , the new center is created by the weighted average of the current center and the subsequence (line 1). Because  $A$  is the nearest neighbor of  $C.cen$ , no offset alignment is needed for  $A$ .

Table 4 shows how two clusters of different lengths can be merged into the same cluster. The new cluster contains all subsequences from both clusters (line 1–2). Because two clusters may be different lengths, we need to align them before finding the new center. As the Fig. 6 example shows, **Step 4** merges two clusters from **Step 1** and **Step 3**. The red center in **Step 1** is longer than the green center in **Step 3**. We align the red center at all possible offsets (line 6) and then create a new center by averaging two current centers. Parts of the

new centers are created by weighted averaging from all (one *or* two) centers that cover that part (line 7–9). To make a decision among all possible offsets, MDL plays an important role again; at each offset, *bitsave* is calculated (line 11), and we choose the offset that can save the maximum number of bits (line 22–23). Similar to the code in line 6–13, which evaluates all offsets when  $C_1.cen$  moves *into*  $C_2.cen$ , the code in line 14–21 evaluates the inverse when  $C_1.cen$  moves *out of*  $C_2.cen$ .

To summarize, our algorithm contains three operators (*create*, *add*, and *merge*), which all use MDL to decide the best choice at each step of the clustering. As there are no known indexing/motif discovery algorithms for MDL, we avail ourselves of two fast external modules that use Euclidean distance for motif discovery [24] and nearest-neighbor search [20]. Using Euclidean distance as a fast proxy for MDL is possible because they are highly correlated when both are small.

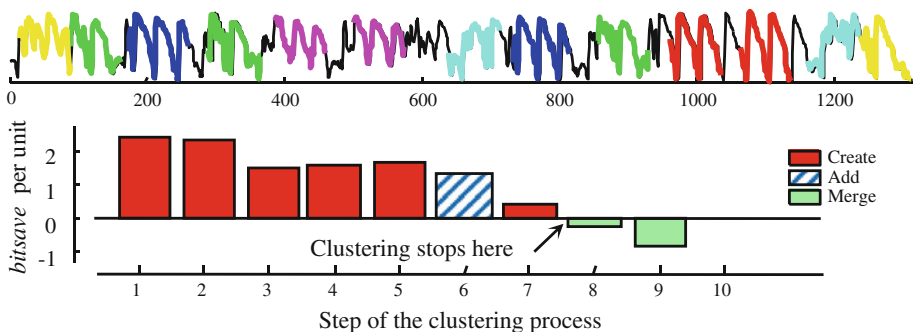
## 5 Experimental results

We begin by stating our experimental philosophy. To ensure our experiments are reproducible, all codes/data are available at [29]. In addition, the site contains many more experiments omitted due to space limitations. Furthermore, the web site contains video animations of the clustering process for each dataset.

### 5.1 Comparison to ground truth

We begin by considering a time series for which we have access to the ground truth (albeit indirectly). Consider the time series shown in Fig. 7(top). A visual inspection gives a hint of some structure, but even on this tiny example, it is not clear exactly what the clustering should be—or even what is the natural length for potential clusters. This dataset was obtained by taking an audio snippet of a recording of Edgar Allen Poe’s poem “The Bells” and transforming it in to the Mel-Frequency Cepstral Coefficients (MFCC) retaining only the first coefficient.

The clustering we obtained looks subjectively intuitive; however, because of the original source material, we are in a unique position to do a more objective test. Table 5 shows the original source text brushed with the colors reflecting the clustering obtained.



**Fig. 7** Top 29.8 s of an audio snippet, represented by the first coefficient in MFCC space, and then annotated with colors to reflect the clusters. Bottom a trace of the steps use to produce the clustering



**Table 5** The text corresponding to the time series shown in Fig. 7, annotated by color/font

Original Order	Grouped by Clusters
In a <i>sort of Runic rhyme</i> ,	<b>bells,bells, bells</b>
To <b>the throbbing of the bells</b> --	<b>Bells, bells, bells</b>
Of the bells, bells, bells,	Of the bells, bells, bells
To <b>the sobbing of the bells</b> ;	Of the bells, bells, bells
Keeping <i>time, time, time</i> ,	<b>the throbbing of the bells</b>
As he <i>knells, knells, knells</i> ,	<b>the sobbing of the bells</b>
In a happy Runic rhyme,	<b>the tolling of the bells</b>
To <i>the rolling of the bells</i> --	<b>the tolling of the bells</b>
Of the bells, bells, bells--	To the rolling of the bells
To <b>the tolling of the bells</b> ,	To the moaning and the
Of the <b>bells, bells, bells</b> , bells,	<b>time, time, time</b>
<b>Bells, bells, bells</b> --	<b>knells, knells, knells</b>
To the moaning and the <i>groan-</i>	<i>sort of Runic rhyme</i>
<i>ing of the bells.</i>	<i>groaning of the bells.</i>

The results are not perfect with reference to the text version. Recall that we are only considering *one* of the MFCC coefficients, instead of the ten plus typically used in speech processing. This allows some collisions, such as “*t i m e*” and “*k n e l l s*.” However, the structure recovered by our algorithm is significant. Note that our clusters are of different sizes (three items, and two items) and of different lengths (from 55 to 70 points). Also note that we could have had a single cluster of eighteen occurrences of the word “bells.” However, that would have obfuscated the information that this word tends to be repeated in this work, as in “bells, bells, bells.” These longer clusters are arguably more parsimonious.

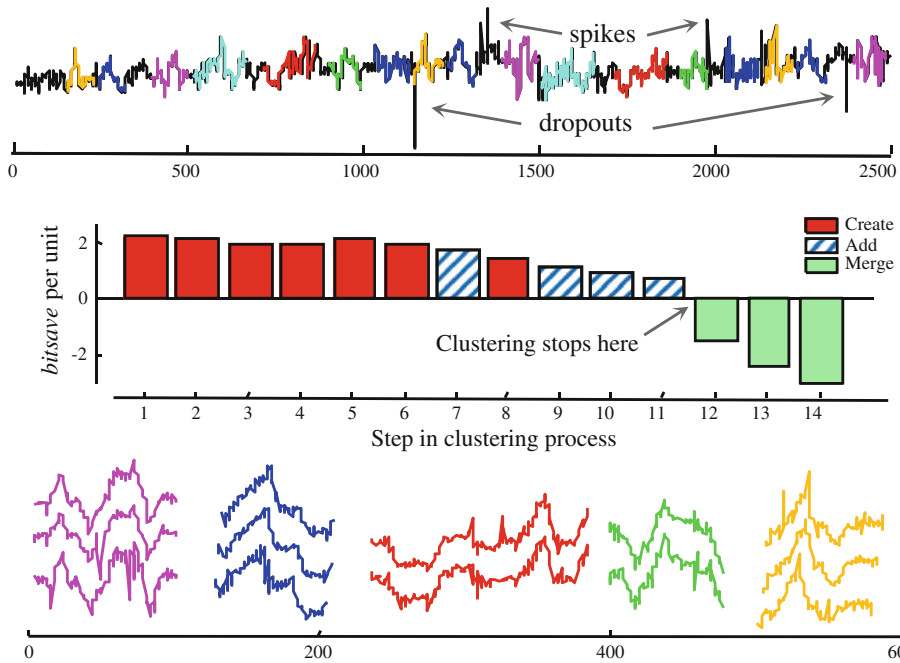
## 5.2 Clustering a noisy dataset

In Fig. 8, we show the results of clustering a noisy industrial dataset. The data come from an industrial wire winding process [2]. The original data consist of seven dimensions; here, we show only the results of clustering the noisiest channel, labeled U1 (the results on the other channels are at [29]). Note that the data have significant non-uniform noise, including spikes and dropouts. While we do not have access to the ground truth here, the clusters, which have different sizes and length, clearly have the property of being similar *within* a cluster and dissimilar *between* clusters. Note that approximately 26% of the data remain unclustered.

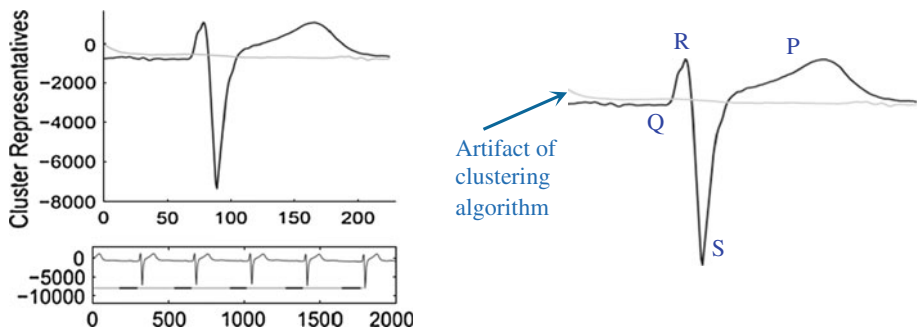
## 5.3 Comparison to other methods

As we noted above, there are few candidate strawmen to compare our work to. Here, we compare our work to the most referenced work in the literature. In a sequence of papers, Chen proposes a series of fixes for the stream clustering problem [5–7]. He demonstrates his ideas mostly on synthetic data; however, as shown in Fig. 9 (right), he also tests on short section of the Koski heartbeat dataset.

While the results are perhaps reasonable, it is not clear why we should have two clusters here since there is clearly just one heartbeat. In addition, there is a subtle artifact noticed by cardiologist, Dr. Helga Van Herle, whom we asked to examine this. The slight slope on the light-gray cluster show in Fig. 9 (left) is not in the data; it comes from the fact that the input data are not an integer multiple of beats, instead being roughly 5.2 beats. Since the



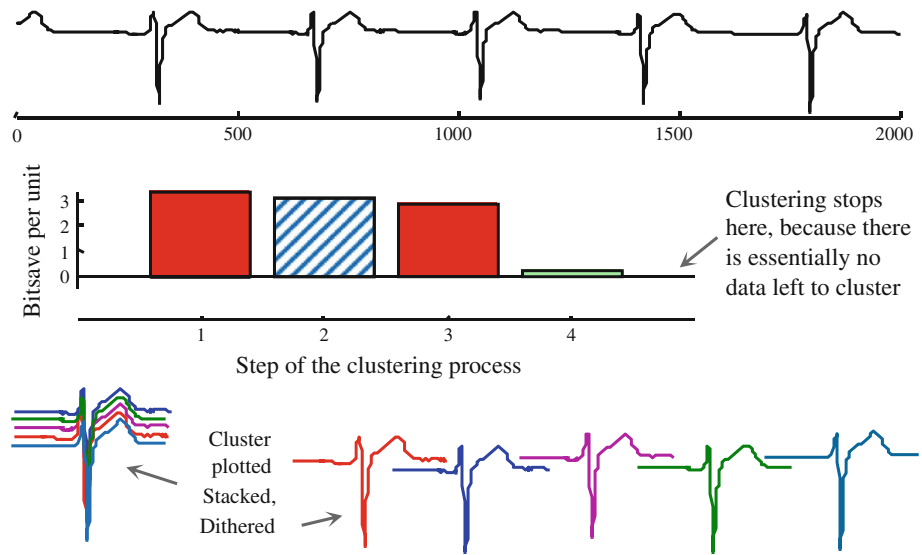
**Fig. 8** Top dimension U1 of the Winding dataset. Middle a trace of the clustering steps produced by our algorithm. Bottom representative clusters obtained



**Fig. 9** Left a screen dump of Fig. 11 from [6]. The original caption read “TF Clustering: Koski-ECG result.” Right an annotation of the clusters by a USC cardiologist

algorithm is trying to explain *all* the data, it must explain the *extra* P-wave by averaging it into a place where it does not belong. Furthermore, as acknowledged in the original paper, the algorithm requires the setting of several parameters and “magic numbers” (i.e., “we chose  $p$  as the number of points in the time series divided by 15.”). Finally, we note in passing that the algorithm requires multiple calls to a *quadratic* space and time (in the length of the time series) algorithm, which would make it impractical for many real data mining problems. Our algorithm requires linear space.

In Fig. 10, we show the clustering we achieved on *exactly* the same dataset. We believe the results here are intuitively correct, discovering a complete *single* heartbeat as the cluster.



**Fig. 10** *Top* the same 2,000 data points from Koski-ECG as used in Fig. 9. *Middle* a trace of the clustering steps produced by our algorithm. *Bottom* the single cluster discovered has five members

Note that our algorithm explains 87.5 % of the data; it does not try to explain the extra P-wave “bump” caused by the fact that we do not have an integer number of heartbeats.


### 5.4 Scalability

From our algorithm in Sect. 4.2, assume that `MotifDiscovery` takes time  $O(T)$ . In each *create* step, `MotifDiscovery` is called multiple times to find motifs of different length; we run it at most  $O(s)$  times. Because each subsequence is of length at least  $s$ , there are at most  $O(m/s)$  new clusters to be created. This is why the running time for creating new clusters is  $O(T*s*m/s) = O(mT)$ .


Assume that `NearestNeighbor` can be finished in time  $O(ms)$ . The maximum number of clusters we can have is  $O(m/s)$ , and the original time series can be updated only when a new motif is discovered, so the number of clustering steps (cf. line 2 in Table 1) is at most  $O(m/s)$ . Thus, for add steps, we have  $O(ms * m/s * m/s) = O(m^3/s)$ .

For *merge* steps, if a cluster is created by merging  $k$  clusters so far, the number of subsequences in that cluster is at most  $O(k)$ . The length of its center is at most  $O(ks)$ ; therefore, the number of possible offsets is  $O(ks)$ , and *bitsave* calculation is finished in time  $O(k^2s^2)$ . The maximum number of clusters we can have is at most  $O(m/s)$ , so we can have cluster of size  $k$  at most  $O(m/sk)$  clusters, and there are at most  $O(m/s)$  steps in our algorithm. This means that the running of *merge* steps is at most  $O(m/s*(m/sk)^2*k^2s^2) = O(m^3/s)$ . Hence, the total running time of our algorithm is at most  $O(mT + m^3/s)$  where  $T$  is a running time for a motif discovery. The empirical behavior is shown in Fig. 11.

To put these results in perspective, the ornithology lab we are working with has spent *months* collecting data in the field (cf. Fig. 5), so they are willing to wait the hour we require to cluster several minutes of audio. Nevertheless, we believe that a 100×speedup will soon be possible simply by caching some near redundant motifs calculations.



**Fig. 11** Running time of our algorithm on Koshi data when  $s = 350$



**Fig. 12** The relationship between Euclidean distance (ED) of pairs of subsequences in a random walk time series and MDL of their difference. Euclidean distance is calculated in original continuous space but MDL is calculated in discrete space (64 cardinality)

### 5.5 Discussion of the MDL choice

Now that the reader has gleaned some intuition for our algorithm and its utility for clustering data; we will briefly revisit a discussion of why MDL on a discretized time series is our choice of measure to steer the clustering search.

We cannot use Euclidean distance (or the related *correlation* or *Dynamic Time Warping*, etc. [12,20]) directly because it does not allow us to compare the relative merits of clusters of different lengths or different sizes. In contrast, MDL does allow such meaningful comparisons. Moreover, in the limited case when MDL and Euclidean distance can be compared (when time series lengths are the same), we find that the two measures are *highly* correlated so long as they are small (if both are destined to be large, it does not really matter how correlated they are). The relationship between Euclidean distance and MDL is shown in Fig. 12.

We work in the discrete space rather than the original continuous space because MDL requires it, and because working with the discretized time series makes *no perceptible difference* in classification (as shown in Fig. 2) or in similarity search, indexing, motif discovery, or outlier discovery.

Because of their relationship especially when the distance is small, to make an intractable problem solvable (in term of acceptable running time), we can apply ED-based techniques to speed up the algorithm in some modules. For example, instead of finding the pair of subsequences whose difference has the smallest MDL, we can use the fast motif discovery algorithm from [24] to find the most similar pair using Euclidean distance. We also can

apply some techniques, such as early abandoning and lower bounding, in searching a nearest subsequence of the existing clusters.

Although Euclidean distance can dramatically speed up the running time of algorithm, sometimes the most similar subsequences using Euclidean distance and MDL are *not* the same. This makes our final score, bitsave, non-monotonically decreasing as the readers may have noticed in Fig. 7 (bottom) and in Fig. 8 (middle).

## 6 Multi-dimensional clustering

In addition to a *single* dimensional time series clustering, our algorithm can be extended to cluster multi-dimensional time series data. In many applications that contain more than one feature, the quality of the clustering can be better if we can do clustering across different features or dimensions. For example, in motion capture, some activities have key features only in the upper parts of the body, but lower parts are keys in some activities.

### 6.1 Notation

For the sake of clarity, this section will define some necessary notations related to multi-dimensional time series clustering; however, we note that many of these definitions are obvious analogs of the single-dimensional case discussed in Sect. 3. We begin by defining *multi-dimensional time series*:

**Definition 13** A *d-dimensional time series*  $T$  is a  $d$ -dimensional ordered list of numbers.  $T = \langle T_1, T_2, \dots, T_d \rangle$  where a time series in dimension  $i$ ,  $T_i = t_{i,1}, t_{i,2}, \dots, t_{i,m}$ . Each value  $t_j$  can be any finite number and  $m$  is the length of the time series, which is equal for all dimensions.

When the original time series contains multiple dimensions, its subsequence is called a *d-dimensional subsequence*:

**Definition 14** A *d-dimensional subsequence*  $T_{i,j,k}$  of a  $d$ -dimensional time series  $T$  is a short time series of length  $k$  of dimension  $i$  that starts from position  $j$ . Formally,  $T_{i,j,k} = t_{i,j}, t_{i,j+1}, \dots, t_{i,j+k}$  where  $1 \leq i \leq d$  and  $1 \leq j \leq m - k$ .

As we mentioned before in previous sections, Euclidean distance and MDL are highly correlated when the distances are small. As in the single version algorithm, many techniques using the Euclidean distance are used for speeding up our algorithm. Hence, we define the *Euclidean distance* for two  $d$ -dimensional subsequences as following:

**Definition 15** The *Euclidean distance* between two  $d$ -dimensional subsequences  $A$  and  $B$  of the same length,  $k$ , is:

$$\text{Dist}(A, B) = \sqrt{\sum_{i=1}^d \sum_{j=1}^k (A_{i,j} - B_{i,j})^2}$$

The description length for a multi-dimensional time series is calculated based on the entropy of the given time series. We define *entropy* for  $d$ -dimensional time series as:

**Definition 16** The *entropy* of a  $d$ -dimensional time series  $T$  is defined as following equation. For special case when  $P = 0$ ,  $P \log_2 P$  is defined as 0.

$$H(T) = - \sum_{i=1}^d \sum_t P(T_i = t) \log_2 P(T_i = t)$$

For simplicity, we define the entropy of  $d$ -dimensional time series as the simple summation of the entropy from each dimension. This means we treat each dimension independent from others. While we do not preclude other methods for calculating  $d$ -dimensional entropy, this method works *very* well empirically.

However, there may be some room here to improve the quality of the clustering if we can exploit use the relationship between different dimensions. We leave this as the future improvement.

Similar as in Definition 6 in Sect. 3.1, the description length of a multi-dimensional time series  $T$  is defined as  $m * H(T)$  where  $m$  is the length of the time series  $T$ . We can reuse Definition 8, 9 and 10 in our multi-dimensional clustering algorithm.

We invite the readers to revise some definitions in Sect. 3.1 before we move forward to the algorithm in next section.

## 6.2 Multi-dimensional clustering algorithm

In general, the idea of clustering multi-dimensional time series is similar to the idea of clustering a single-dimensional time series in Sect. 4.2.

The algorithm composes of three operations—*create*, *add*, and *merge*. The general idea is the same as in the previous algorithm. All possible operators will be considered, and their *bitsave* will be computed. Then, the choice whose *bitsave* is maximum will be selected and the algorithm performs the operation corresponding to that choice. The algorithm for clustering multi-dimensional time series is shown in Table 6.

We allow user to set the approximate length of subsequences, contained inside the final clusters. The set of the final clusters is initialed to an empty set (line 1). The algorithm will be terminated on only two conditions, first, when no possible choices can reduce the number of bits for representing the clusters and, second, when there is no data left and all clusters have been merged.

In the *create* operation (line 4–10), we find a most similar pair of subsequences. The Euclidean distance is used here instead of MDL for speeding up the process and makes this process much faster. A trivial extension of the motif discovery algorithm from [24] using multi-dimensional Euclidean distance (cf. Definition 15) is used here to find the most similar subsequences in multi-dimensional time series (line 6). Then, a new cluster is created, and its *bitsave* is calculated (line 7–8). For convenience, the temporary cluster is collected (line 9).

In *add* process (line 11–17), the nearest neighbor of each existing cluster is discovered (line 13). The rest of the algorithm, including *merge* process (line 18–25), is the same as in the algorithm in Sect. 4.2. After all choices are calculated, the choice, which maximizes *bitsave*, will be selected to perform (line 27–28).

Because each subsequence contains multi-dimensional data, the cluster representative or the cluster's center also contains multi-dimensional data. Table 7 explained how to create the cluster. The new cluster is always created from two subsequences so it will contain only two multi-dimensional subsequences.

Similarly, how to add a multi-dimension subsequence to an existing cluster is explained in Table 8. The idea of all three operators is same as in the single-dimension clustering algorithm.

**Table 6** Multidimensional stream clustering algorithm

<b>Input:</b> $ts$ : multidimensional time series, $s$ : approximate length	
<b>Output:</b> $cluster$ : final cluster of subsequences	
1	$cluster = \{ \}$
2	<b>while</b> $bitsave > 0$
3	$bitsave = -\infty, i = 0$
4	// create new cluster
5	<b>for</b> $len = s$ to $2s$
6	$(A, B) = \text{MultiDimMotifDiscovery}(ts, len)$
7	$C' = \text{CreateMultiDimCluster}(A, B)$
8	$bs.append(\text{ComputeBitsave}(C', A, B))$
9	$cluster'.append(cluster \cup \{C'\})$
10	<b>end for</b>
11	// add subsequence to an existing cluster
12	<b>for</b> $C \in cluster$
13	$A = \text{MultiDimNearestNeighbor}(ts, C)$
14	$C' = \text{AddToMultiDimCluster}(C, A)$
15	$bs.append(\text{ComputeBitsave}(C', C, A))$
16	$cluster'.append(cluster \cup \{C'\} - \{C\})$
17	<b>end for</b>
18	// merge 2 clusters
19	<b>for</b> $C_1 \in cluster$
20	<b>for</b> $C_2 \in cluster$ and $C_1 \sim C_2$
21	$C' = \text{MergeMultiDimClusters}(C_1, C_2)$
22	$bs.append(\text{ComputeBitsave}(C', C_1, C_2))$
23	$cluster'.append(cluster \cup \{C'\} - \{C_1\} - \{C_2\})$
24	<b>end for</b>
25	<b>end for</b>
26	// update the result
27	$[bitsave \text{ index}] = \max(bs)$
28	$cluster = cluster'(index)$
29	<b>end while</b>

**Table 7** Create operator

<b>Function</b> $C = \text{CreateMultiDimCluster}(A, B)$	
1	$C.size = 2$
2	<b>for</b> all dimension $d$
3	$C.cen_d = (A_d + B_d) / 2$
4	$C.seq_d = \{A_d, B_d\}$
5	$C.shift_d = \{0, 0\}$
6	<b>end for</b>

**Table 8** Add operator

<b>Function</b> $C = \text{AddToMultiDimCluster}(C, A)$	
1	$C.size = C.size + 1$
2	<b>for</b> all dimension $d$
3	$C.cen_d = (C.cen_d * (C.size - 1) + A_d * 1) / C.size$
4	$C.seq_d.append(A_d)$
5	$C.shift_d.append(0)$
6	<b>end for</b>

**Table 9** Merge operator

Function $C' = \text{MergeMultiDimClusters}(C_1, C_2)$	
1	$n_1 = C_1.size, \quad m_1 = \text{length}(C_1.cen)$
2	$n_2 = C_2.size, \quad m_2 = \text{length}(C_2.cen)$
3	$C'.size = n_1 + n_2$
4	$bs = \{\}$
5	$Ctmp = \{\}$
6	<b>for</b> all dimension $d$
7	// add all sequences from $C_1$ and $C_2$ to $C'$
8	$C'.seq_d = [C_1.seq_d; C_2.seq_d]$
9	<b>end for</b>
10	// align $C_1.cen$ by moving to all positions
11	<b>for</b> $off = 0$ to $m_2$
12	<b>for</b> all dimension $d$
13	// create centers of the same length
14	$cen1 = [C_2.cen_d(1, off), C_1.cen_d]$
15	$cen2 = [C_2.cen_d, C_1.cen_d(1, m_1 + off - m_2)]$
16	$C'.cen_d = (cen1 * n_1 + cen2 * n_2) / (n_1 + n_2)$
17	$C'.shift_d = [C_1.shift_d + off; C_2.shift_d]$
18	<b>end for</b>
19	$bs.append(\text{ComputeBitsave}(C', C_1, C_2))$
20	$Ctmp.append(C')$
21	<b>end for</b>
22	// move $C_1.cen_d$ to opposite direction
23	<b>for</b> $off = 1$ to $m_1$
24	<b>for</b> all dimension $d$
25	// create centers of the same length
26	$cen1 = [C_1.cen_d, C_2.cen_d(1, m_2 + off - m_1)]$
27	$cen2 = [C_1.cen_d(1, off), C_2.cen_d]$
28	$C'.cen_d = (cen1 * n_1 + cen2 * n_2) / (n_1 + n_2)$
29	$C'.shift = [C_1.shift_d; C_2.shift_d + off]$
30	<b>end for</b>
31	$bs.append(\text{ComputeBitsave}(C', C_1, C_2))$
32	$Ctmp.append(C')$
33	<b>end for</b>
34	$[bitsave\ index] = \max(bs);$
35	$C' = Ctmp(index);$


Table 9 explains in detail how to *merge* two clusters. As in the single-dimension version, when two clusters are being merged, we have to align their center to find the best position, whose *bitsave* is maximized, after merging. We can align the center of the first cluster's center when it moves *into* the center of another cluster (line 10–21), and also when it moves *out of* another center (line 22–33). Each position in alignment is kept as a local variable inside the corresponding cluster (line 20). However, only the cluster that maximum *bitsave* is return (line 34–35).

In next section, we will demonstrate that our algorithm can handle the time series that contain more than one dimension.

### 6.3 Experimental result

We demonstrate our multi-dimensional clustering algorithm on a people activity dataset, *Physical Activity Monitoring for Aging People* (PAMAP) from [27]. In this dataset, eight subjects (people) perform both indoor and outdoor activities such as normal walk, Nordic walk, cycle, run, ascend stairs, and descend stairs. Each subject has sensors placed on three locations on their bodies, which are hand, chest, and ankle. We pick one time series





**Fig. 13** Three time series generated from z-accelerometer of sensors at hand, chest and shoe from PAMAP [27]. The subject performs three activities: descending stairs, ascending stairs, and descending stairs again

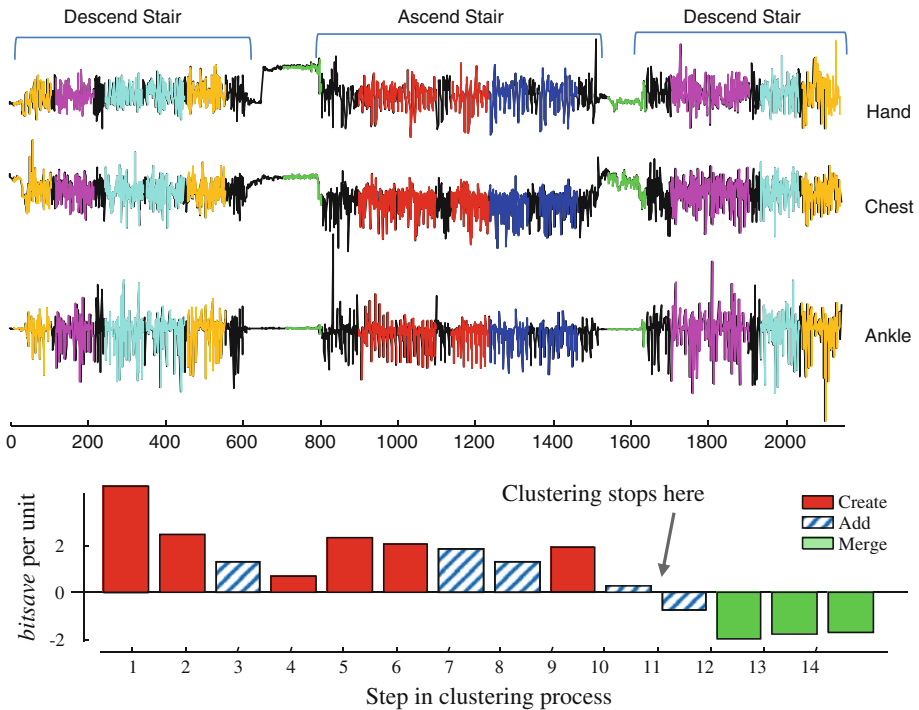
from each position; we choose to use only three time series generated by z-accelerometers instead of using all data from all 45 sensors because experience from indexing ASL [1] and motion capture data suggest that three or four time series is enough to represent most activities/behaviors.

A snippet of a three dimensional time series showing an individual walking up and down stairs is shown in Fig. 13. Each time series is generated from a z-accelerometer. The data are from subject number 1 of indoor activities in PAMAP dataset [27]. In this example, the subject performed the sequence, walking up stair, walking down stair, and walking up stair.

The multidimensional clustering result is shown in Fig. 14. Each color represents a cluster; hence, the subsequences of the same color are clustered into the same cluster. The result demonstrates that our algorithm can capture the similar actions inside the same activities. For example, three clusters appeared in both side of the time series, which are in the same action, that is, walking down stair. Some data in the time series have been ignored as we can see in black color. From this example, we conjecture that the subject may occasionally have a special behavior because it looks like there are patterns inside the clusters. For example, it appear that she initially starts to climb stairs briskly (red) but begins to slow down as she was tired out after a few flights (blue). Likewise, the sequence of three clusters in the same order (pink, cyan, yellow) in *both* descending stairs sections is suggested.

We also ran our single-dimensional clustering algorithm on each dimension separately. The result shows that there are incorrect clusters across the different activities in the final clusters and, in some dimension, the algorithm is terminated very early and covers less than 25 % of the data. Hence, in this example, the clustering result from a single dimension is worse than the result from the multi-dimensional clustering algorithm. We believe that the multi-dimensional time series can create the better clusters than the clusters created from just one dimension, especially when the dataset contains many related features such as motion capture.

We can speed up the algorithm using Euclidean distance and make the problem solvable in an acceptable time, but we have to trade of the quality of the clustering somehow. Our parameter-lite algorithm can be extended to cluster multi-dimensional time series and it works quite well; however, some readers may notice that the *bitsave* score [in Fig. 14(bottom)] is not much stable as the result of a single-dimensional clustering algorithm in Sect. 5. This is



**Fig. 14** *Top* the multi-dimensional time series clustering result. Two clusters are detected in ascending stair, and three clusters are detected in descending stair. *Bottom* a trace of the multi-dimensional clustering steps produced by our algorithm

because the Euclidean distance is used instead of MDL in many core modules. By the curse of dimensionality, when the number of dimensionality is increase, the Euclidean distance performs worse and the difference between MDL and Euclidean distance is increase. We believe that if the multi-dimensional time series clustering using MDL has been well study in the future, some researchers will provide the distance that is fast enough and very close to MDL to get the better result.

## 7 Conclusions

In this work, we have shown that any attempt to cluster a single time series stream that insists on explaining *all* the data is almost certainly doomed to failure. We introduced a clustering representation that has the expressive power to ignore some of the data and can have clusters with different length subsequences. We further showed an efficient and parameter-lite MDL-based algorithm to perform the clustering. We have shown on our algorithm is effective on a wide variety of datasets, for both single and multi-dimensional problems.

Currently, our algorithm only works on batch time series. In ongoing work, we are attempting to generalize it to the online case.

**Acknowledgments** We would like to acknowledge the financial support for our research provided by the Royal Thai Government and NSF grants 0803410 and 0808770.

## References

1. Athitsos V, Wang H, Stefan A (2010) A database-based framework for gesture recognition. *Pers Ubiquitous Comput* 14(6):511–526
2. Bastogne T, Noura H, Richard A, Hittinger JM (1997) Application of subspace methods to the identification of a winding process. In: *Proceeding of the 4th European control conference, Brussels, Belgium*
3. Batista GEAPA, Wang X, Keogh EJ (2011) A complexity-invariant distance measure for time series. In: *SDM*, pp 699–710
4. Bouchard D, Badler NI (2007) Semantic segmentation of motion capture using laban movement analysis. In: *IWA*, pp 37–44
5. Chen JR (2005) Making subsequence time series clustering meaningful. In: *ICDM*, pp 114–121
6. Chen JR (2007) Useful clustering outcomes from meaningful time series clustering. In: *The Australasian data mining conference*
7. Chen JR (2007) Making clustering in delay-vector space meaningful. *Knowl Inf Syst* 11(3):369–385
8. Chuang ZJ, Wu CH, Chen WS (2006) Movement epenthesis generation using NURBS-based spatial interpolation. *IEEE Trans Circuit Syst Video Technol* 16(11):1313–1323
9. Cook DJ, Holder LB (1994) Substructure discovery using minimum description length and background knowledge. *J Artif Intell Res* 1:231–255
10. Das G, Lin K, Mannila H, Renganathan G, Smyth P (1998) Rule discovery from time series. In: *Proceeding of the 3rd KDD*, pp 16–22
11. Denton AM, Basemann CA, Dorr DH (2009) Pattern-based time-series subsequence clustering using radial distribution functions. *Knowl Inf Syst J* 18(1):1–27
12. Ding H, Trajcevski G, Scheuermann P, Wang X, Keogh EJ (2008) Querying and mining of time series data: experimental comparison of representations and distance measures. *PVLDB* 1(2):1542–1552
13. Evans SC et al (2007) MicroRNA target detection and analysis for genes related to breast cancer using MDLcompress. *EURASIP J Bioinform Syst Biol* 2007:1–16
14. Evans SC, Eiland E, Markham TS, Impson J, Laczó A (2007) MDLcompress for intrusion detection: signature inference and masquerade attack. *MILCOM*, Orlando, Florida
15. Grünwald PD, Myung IJ, Pitt MA (2005) *Advances in minimum description length: theory and applications*. MIT Press, Cambridge
16. Jonyer I, Holder LB, Cook DJ (2004) MDL-based context-free graph grammar induction and applications. *J Artif Intell Tools* 13(1):65–79
17. Kamvar SD, Klein D, Manning CD (2002) Interpreting and extending classical agglomerative clustering algorithms using a model-based approach. In: *ICML*, pp 283–290
18. Keogh EJ, Lin J (2005) Clustering of time-series subsequences is meaningless: implications for previous and future research. *Knowl Inf Syst* 8(2):154–177
19. Keogh EJ, Lin J, Lee SH, Herle HV (2007) Finding the most unusual time series subsequence: algorithms and applications. *Knowl Inf Syst* 11(1):1–27
20. Keogh EJ, Kasetty S (2003) On the need for time series data mining benchmarks: a survey and empirical demonstration. *Data Mining Knowl Discov* 7(4):349–371
21. Li H, Abe N (1996) Clustering words with the MDL principle. In: *Proceeding of the 16th international conference on computational linguistics*, pp 5–9
22. Li M, Vitanyi P (1997) *An introduction to Kolmogorov complexity and its applications*, 2nd ed. Springer, Berlin
23. Molkov YI, Mukhin DN, Loskutov EM, Feigin AM (2009) Using the minimum description length principle for global reconstruction of dynamic systems from noisy time series. *Phys Rev E* 80:046207
24. Muen A, Keogh EJ, Shamlo NB (2009) Finding time series motifs in disk-resident data. In: *ICDM* pp 367–376
25. Papadimitriou S, Sun J, Faloutsos C, Yu PS (2008) Hierarchical, parameter-free community discovery. In: *PKDD* pp 170–187
26. Pednault E (1998) Some experiments in applying inductive inference principles to surface reconstruction. In: *IJCAI*, pp 1603–1609
27. Reiss A, Weber M, Stricker D (2011) Exploring and extending the boundaries of physical activity recognition. In: *IEEE SMC workshop on robust machine learning techniques for human activity recognition*
28. Stine RA (2004) Model selection using information theory and the MDL principle. *Sociol Methods Res* 33(2):230–260
29. Supporting webpage. <http://www.cs.ucr.edu/~rakthant/TSEpenthesis>
30. Tanaka Y, Iwamoto K, Uehara K (2005) Discovery of time-series motif from multi-dimensional data based on MDL principle. *Mach Learn* 58(2):269–300

31. Ueno K, Xi X, Keogh EJ, Lee DJ (2006) Anytime classification using the nearest neighbor algorithm with applications to stream mining. In: ICDM, pp 623–632
32. Wallace CS, Boulton DM (1968) An information measure for classification. *Comput J* 11(2):185–194
33. Yang R, Sarkar S, Loeding BL (2010) Handling movement epenthesis and hand segmentation ambiguities in continuous sign language recognition using nested dynamic programming. *IEEE PAMI* 32(3):462–477
34. Yankov D, Keogh EJ, Rebbapragada U (2008) Disk aware discord discovery: finding unusual time series in terabyte sized datasets. *Knowl Inf Syst* 17(2):241–262

## Author Biographies



**Thanawin Rakthanmanon** is currently a Ph.D. candidate of Computer Science at the University of California, Riverside. His research interests are data mining and machine learning, especially in efficient algorithm development, motif discovery, time series clustering/classification, and document analysis. He receives the Royal Thai Government Scholarship for studying his doctoral program.



**Eamonn J. Keogh** is a full professor of Computer Science at the University of California, Riverside. His research interests are data mining, machine learning, and information retrieval. Several of his papers have won best paper awards, including papers at SIGKDD, ICDM, and SIGMOD. Dr. Keogh is the recipient of a 5-year NSF Career Award for “Efficient Discovery of Previously Unknown Patterns and Relationships in Massive Time Series Databases.”



**Stefano Lonardi** is Professor and Vice Chair of the Department of Computer Science and Engineering at University of California, Riverside, CA. He is also a faculty member of the Graduate Program in Genetics, Genomics and Bioinformatics, the Center for Plant Cell Biology, the Institute for Integrative Genome Biology, and the Graduate Program in Cell, Molecular and Developmental Biology. Stefano’s recent research interest includes design of algorithms, computational molecular biology, data compression and data mining. He has published over 40 papers in major theoretical computer science and computational biology journals and has over 50 publications in referred international conferences. In the year 2005, he received the CAREER award from National Science Foundation. He has received funding from NSF, NIH, DARPA, and USDA.



**Scott Evans** holds a Ph.D. in Electrical Engineering from Rensselaer Polytechnic Institute, an MS in Electrical Engineering from the University of Connecticut and a BS in Electrical Engineering from Virginia Tech. He is a Senior Research Engineer in the Machine Learning lab at GE Global Research in Niskayuna, NY. His research interests include machine learning, data mining, and algorithmic information theory for cyber-security, prognostics and health monitoring, and remote monitoring and diagnostics.