

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/320692949>

Speeding up dynamic time warping distance for sparse time series data

Article in Knowledge and Information Systems · January 2018

DOI: 10.1007/s10115-017-1119-0

CITATIONS

11

READS

668

6 authors, including:



Abdullah Mueen

University of New Mexico

90 PUBLICATIONS 3,224 CITATIONS

[SEE PROFILE](#)



Nikan Chavoshi

University of New Mexico

13 PUBLICATIONS 379 CITATIONS

[SEE PROFILE](#)



Noor Abu-El-Rub

University of Kansas Medical Center

9 PUBLICATIONS 50 CITATIONS

[SEE PROFILE](#)



Hossein Hamooni

Visa Inc.

13 PUBLICATIONS 336 CITATIONS

[SEE PROFILE](#)

Some of the authors of this publication are also working on these related projects:



machine learning [View project](#)



On URL Changes and Handovers in Social Media [View project](#)

Speeding up dynamic time warping distance for sparse time series data

Abdullah Mueen¹ · Nikan Chavoshi¹ · Noor Abu-El-Rub¹ · Hossein Hamooni¹ · Amanda Minnich¹ · Jonathan MacCarthy²

Received: 26 January 2017 / Revised: 6 June 2017 / Accepted: 10 October 2017
© Springer-Verlag London Ltd. 2017

Abstract Dynamic time warping (DTW) distance has been effectively used in mining time series data in a multitude of domains. However, in its original formulation DTW is extremely inefficient in comparing long sparse time series, containing mostly zeros and some unevenly spaced nonzero observations. Original DTW distance does not take advantage of this sparsity, leading to redundant calculations and a prohibitively large computational cost for long time series. We derive a new time warping similarity measure (AWarp) for sparse time series that works on the run-length encoded representation of sparse time series. The complexity of AWarp is quadratic on the *number of observations* as opposed to the *range of time* of the time series. Therefore, AWarp can be several orders of magnitude faster than DTW on sparse time series. AWarp is *exact* for binary-valued time series and a *close approximation* of the original DTW distance for any-valued series. We discuss useful variants of AWarp: bounded (both upper and lower), constrained, and multidimensional. We show applications of AWarp to three data mining tasks including clustering, classification, and outlier detection, which are otherwise not feasible using classic DTW, while producing equivalent results. Potential areas of application include bot detection, human activity classification, search trend analysis, seismic analysis, and unusual review pattern mining.

Keywords Sparse time series · Dynamic time warping · Run-length encoding

1 Introduction

Sparse temporal data are becoming more common with the advancement of active sensing techniques to monitor discrete events. Consider a motion sensor, which continuously senses the environment for object movement, and actions (i.e., turn on the light) are only taken

✉ Abdullah Mueen
mueen@cs.unm.edu

¹ Department of Computer Science, University of New Mexico, Albuquerque, NM, USA

² Los Alamos National Laboratory, Los Alamos, NM, USA

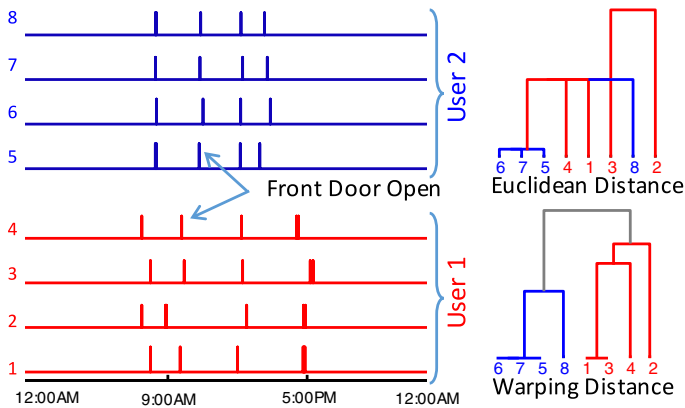


Fig. 1 (Left) Day-long signals generated from the front doors of two single-resident apartments of two users. (Right) Euclidean distance cannot capture the difference between the two users, while DTW distance can

when there is a change in the environment surrounding the sensor. Thus, intelligent sensors hide complexities related to signal processing, report a set of related events with precise time information, and produce sparse time series data. Sparse time series are different from traditional time series in having arbitrary gaps between real observations. The data mining community has been developing mining techniques for traditional time series data to discover patterns [1, 2] and their associations with real events. A successful adaptation of these techniques to sparse time series will create opportunity to mine meaningful patterns from sparse time series data.

In this paper, we particularly consider time warping (i.e., stretching and squeezing of time) sensitive pattern mining. Time warping naturally appears in many domains, especially in the activities of humans and animals. For example, humans can produce the same motion [3] or speech [4] at a different pace and acceleration and, have the speech or signal be still recognizable. Time warping is also present in sparse time series. For example, Fig. 1 shows the 24-h time series of the front door statuses of two single-resident apartments. A spike in the data shows a door-opening event. Most of the time, front doors are closed. Each day shows a unique schedule of the resident in that apartment. Note the time warping across days. A simple hierarchical clustering of the data shows that the daily patterns of a person can be clustered well if we use dynamic time warping (DTW) distance instead of the widely used Euclidean distance.

DTW is a distance measure that has been used in dozens of research works on mining equally sampled time series data [5]. However, new sensor technologies (both soft and hard) can capture a sequence of discrete events that forms a sparse time series (as in Fig. 1). In its original form, DTW distance does not take advantage of temporal sparsity. For example, Twitter records discrete activities of more than 300 million users at a resolution of ms. Comparing the activities of two users for a day at this resolution requires $86,400,000^2$ arithmetic operations, which is equivalent to more than a day in an off-the-shelf machine. In contrast, the number of activities performed by average users are on the order of tens or hundreds. Clearly, the amount of computation required to calculate DTW distance using existing algorithms is unnecessarily excessive.

We develop a time warping distance measure, AWarp, for sparse time series data that works on run-length encoded time series. Run-length encoded time series are much shorter than their

versions before encoding; for example, in Fig. 1 the run-length encoded time series for the seventh time series will have only eight numbers, as opposed to 86,400,000 observations for a day. Thus, AWarp will require around 8^2 arithmetic operations to calculate DTW distance between two such run-length encoded time series. By just adopting this simple strategy, we achieved several orders of magnitude speedup in calculating warping distance on sparse data. We show that AWarp is *exact* for binary-valued time series and closely approximates the DTW distance for any-valued time series. AWarp is extensible to constrained warping and multidimensional warping. We demonstrate applications of AWarp on bot discovery, human activity classification, search trend analysis, seismic analysis, and unusual review pattern discovery. We experimentally demonstrate

We give necessary background (Sect. 2) on sparse time series and their various representations, and on dynamic time warping. Next we describe the core AWarp algorithm and its variants in Sect. 4. We show performance analysis of the algorithm in Sect. 5 and demonstrate potential applications in Sect. 6. We conclude in Sect. 7.

2 Encoding sparse time series

2.1 Definition

A *time series* is defined as a vector $T = \langle v_1, v_2, \dots, v_n \rangle$ of observations made at equal intervals. Most distance measures and mining algorithms are invariant to the absolute start time and sampling interval of the time series [6, 7].

For two series $x = x_1, x_2, \dots, x_n$ and $y = y_1, y_2, \dots, y_m$ of length n and m , where $n > m$ without losing generality, the classic dynamic time warping distance is defined as below.

$$DTW(x, y) = D(n, m) \tag{1}$$

$$D(i, j) = (x_i - y_j)^2 + \min \begin{cases} D(i - 1, j) \\ D(i, j - 1) \\ D(i - 1, j - 1) \end{cases} \tag{2}$$

$$D(0, 0) = 0, \forall_{ij} D(i, 0) = D(0, j) = \infty \tag{3}$$

We intentionally skip taking the square root of $D(n, m)$, as it does not change the relative ordering of pairs and makes it efficient for speedup techniques. A dynamic programming algorithm to populate the DTW matrix and calculate the DTW distance is well known. An example DTW matrix for two time series is given in Fig. 2a.

Constrained DTW distance is a variant that limits the allowed time gap between two aligned observations. In effect, the DTW matrix is populated partially around the diagonal (readers can find details about DTW in many online resources such as Wikipedia and also in [5]).

2.2 Sparse time series and representations

A time series is simply a sequence of observations made in temporal order. The phenomena that we observe can be continuous or discrete in time. For example, the temperature of a sea surface at specific point on earth is a continuous phenomenon. In contrast, the activities of a user on social media are discrete because the user can be inactive at times. When observing a discrete phenomena, a sparse time series is produced, which is the focus of this work.

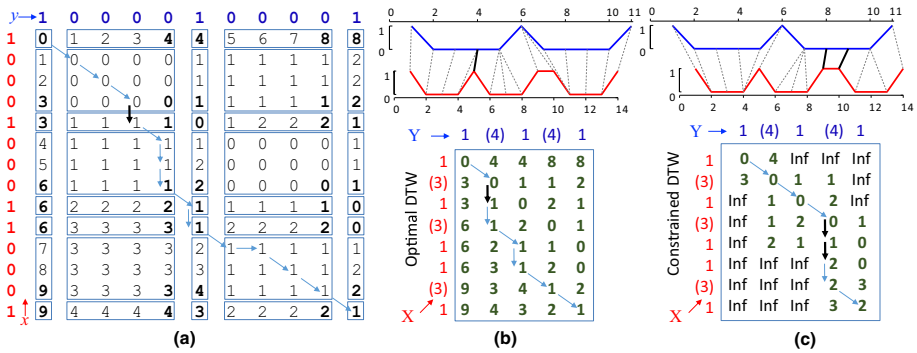


Fig. 2 **a** Two sparse time series x (red) and y (blue) and their DTW matrix. **b** The AWarp matrix for their encoded versions, X and Y . **c** The AWarp matrix for a constraint window of size 5. The bold faced arrows and their corresponding alignments show cost accumulations on the warping paths (color figure online)

A sparse time series has many more zero-valued observations than nonzero observations. We define the *sparsity factor*, s , of a time series as the ratio between the length of the time series and the number of nonzero observations. The higher the sparsity factor, the more sparse a time series is. Representing a sparse time series in the traditional vector format wastes significant amount of space. For example, the REFIT [8] datasets are stored in this format. A more optimal way to store sparse time series is as a sequence of time-value pairs.

2.2.1 Time-value sequence

Each observation is stored as a (t, v) pair and a sparse time series is an ordered set $T_v = \{(t_i, v_i) | t_i < t_{i+1}, i = 1, \dots, n - 1\}$. For example, the CASAS datasets [9] are represented in this format. This is the most common representation of sparse time series.

Example The time series $T = \langle 7, 0, 0, 9, 6, 0, 0, 0, 1 \rangle$ can be represented equivalently as $T_v = \{(1, 7), (4, 9), (5, 6), (9, 1)\}$ if the start time is 1.

In this paper, we use a well known compression technique, *run-length encoding* [10], to represent sparse time series. We differ from the classic run-length encoding as we only encode the runs of zeros and leave the runs of nonzero observations as they are.

2.2.2 Length encoded series

Let us assume we have a time series T . A length encoded time series is T_e where we replace a run of k zeros in T with a (k) . Here we use the parenthesis to represent the duration of zeros.

Example For the same sparse time series, $T = \langle 7, 0, 0, 9, 6, 0, 0, 0, 1 \rangle$, the length encoded series is $T_e = \langle 7, (2), 9, 6, (3), 1 \rangle$.

We can also define length encoded series in a rather complex way from the time-value sequence T_v as $T_e = \langle v_1, (t_2 - t_1 + 1), v_2, (t_3 - t_2 + 1), \dots, v_{n-1}, (t_n - t_{n-1} + 1), v_n \rangle$. In other words, we insert the duration between each pair of observations in between the observations to create a length encoded series. In the rest of the paper, we use “encoded series” to denote length encoded series.

Note that a time series of four observations, such as T_v , needs eight integers for storage in the time-value sequence representation. In traditional representation, T could require any number of integers larger or equal to eight to store the series because the lengths of the runs of zeros can arbitrarily vary in size. In an encoded series, T_e needs at most eight integers. Thus, for a fixed sparsity factor, the encoded series require the lowest amount of space.

Our length encoding is different from classic run-length encoding. Consider an example string WWWB BBBBCWWW. The run-length encoding of the string is W4B4C1W3, where every run of a symbol is compressed to only two items: symbol and length. Our length encoding compresses based on a preselected symbol (i.e., zero). In the example string, there can be two possible encodings: 4BBBB3 if W is selected and WWW4CWWW if B is selected. Classic run-length encoding has higher compression factor, yet we use our length encoding based on zeros for two reasons: our encoding is more suitable for sparsity related to *absence* of observations as opposed to *repetitions* of observations, and the principle of speeding up DTW distance by length encoding one symbol can easily be extended to encoding multiple symbols.

Run-length encoding compresses a run of zeros by the length of the run. There is no better compression than just one number. In that sense, run-length encoded series are *fully encoded series*. We can also define partially encoded series, which will be useful to calculate multidimensional DTW distance.

2.2.3 Partially encoded series

Given an encoded series T_e , a partially encoded series T_{pe} is an equivalent series where one or more of the runs of zeros are split into parts.

Example $T_{pe} = \langle 7, (2), 9, 6, (2), (1), 1 \rangle$ is a partially encoded series of T_e from the previous example. If we keep *splitting* the runs of zeros in a partially encoded series, we reach the same length as the traditional series, with zero being represented by (1) and no more possible splits.

If a time series starts with a run of zeros, we treat the first zero as an observation and encode the rest of the run. This ensures that an encoded series always starts with an observation, and not with a run of zeroes. Similarly, we ensure that the series ends with an observation. Since T_e and T_{pe} are equivalent, their DTW distances to any other series remain identical. The conversion between the three representations of sparse time series can be performed in time linear to the length of the time series.

2.3 Motivating example

We now present an example to motivate AWarp. In Fig. 2a, we show two toy time series x and y of lengths 14 and 11, respectively. The DTW distance between the two time series is 1. The DTW matrix is a 14×11 matrix as shown in Fig. 2a. If we encode the time series x and y , the two time series shrink to X (length 8) and Y (length 5), respectively. The AWarp matrix calculated on these encoded time series is only of size 8×5 (shown in (b)). The AWarp distance is the same as the DTW distance, 1. The computation in each boxed sub-matrix of the DTW matrix is replaced by a one cell in the AWarp matrix. The value in the bottom-right corner of a sub-matrix is identical to the corresponding cell in the AWarp matrix. Note that a sub-matrix is not always a constant matrix with identical values. Some of the sub-matrices are monotonically increasing sequences. To complete the example, we also show the constrained AWarp matrix for a constraint window of size 5 in Fig. 2c. The constrained warping distance is always larger than the optimal DTW distance. In this example, the constrained AWarp

distance is 2, which is exactly the same as the constrained DTW distance under the same constraint window.

3 Related work

Dynamic time warping is a long-studied algorithm in many research communities, including signal processing [11], speech recognition [4, 12], data mining [13], and image processing [14]. One of the earliest research on using dynamic time warping to discover patterns in time series data is by Berndt and Clifford [15]. We adopt warping distance for sparse time series. Although many human activity datasets are publicly available, warping-invariant mining has not been applied to sparse time series from generated discrete human activities (to the best of our knowledge). Our work is the first to exploit sparsity for time efficiency in warping-invariant mining.

Some works exploit other forms of sparsity in DTW calculations [16, 17]. In [16], the authors reduce space complexity by unfolding only required cells in the DTW matrix by exploiting local correlation; however, there is no reduction in time complexity. In contrast, our method reduces both time and space complexity with negligible difference in accuracy. In [17], the authors have not used the sparsity of the time series or the sparsity of the DTW matrix, rather sparsity is used when combining features that are independently calculated without using DTW. We claim our work as the first to calculate warping similarity on an encoded representation of sparse time series data.

A significant body of research exists on efficient DTW calculation [18–20]. In all of these work, calculation of one global DTW distance has a worst-case time complexity of $O(n^2)$, where n is the length of the time series. AWarp has a worst-case complexity of $O(m^2)$, where m is the number of nonzero observations. This makes a significant difference in performance for sparse time series.

DTW-based similarity search in streaming or database settings has been made efficient by indexing [5], hybrid bounding [21], admissible pruning [22], and filter-and-refine [23] approaches. These approaches are equally applicable for sparse time series and can use AWarp, instead of DTW, for un-pruned distance comparisons. We leave it as a future work to adopt these techniques to perform similarity search under AWarp. In [24], the authors have shown that locally relevant constraints learned from salient features of the comparing time series are better than a fixed constraint for the entire time series. We will evaluate this approach on constrained AWarp in future.

4 AWarp distance measure

We start by describing the AWarp algorithm for simple binary-valued series, on which AWarp *exactly matches* the classic DTW distance. We will then relax this simplification by considering the general case of any-valued time series, on which AWarp *closely approximates* the DTW distance. Finally, we show the constrained and multidimensional versions of AWarp.

4.1 Binary-valued series

Algorithm 2 is the AWarp distance function for run-length encoded time series. The inputs to the algorithm are two run-length encoded time series. The algorithm fills in a matrix D of

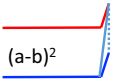
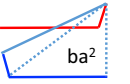

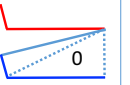
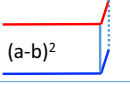
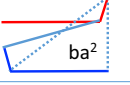
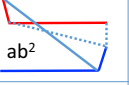
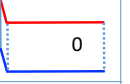
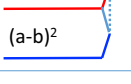
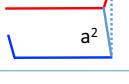
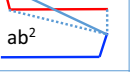
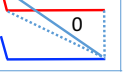
	a: OBS b: OBS	a: OBS b: ROZ	a: ROZ b: OBS	a: ROZ b: ROZ
Top	 $(a-b)^2$	 ba^2	 b^2	 0
Diagonal	 $(a-b)^2$	 ba^2	 ab^2	 0
Left	 $(a-b)^2$	 a^2	 ab^2	 0

Fig. 3 Twelve cases covered by the Algorithm 1. *OBS* observation, *ROZ* run of zeros

size $l_x \times l_y$ in the same way as the DTW algorithm. Here l_x and l_y are the lengths of the two encoded series x and y , respectively. The algorithm has two loops in lines 4 and 5 that go over all the cells of the AWarp matrix. The algorithm calculates three costs for a cell based on three other cells: (*diagonal*, *left*, and *top*) relative to the cell being populated. Finally, in line 9, the algorithm takes the minimum of the costs as per the definition of DTW (see Eq. 1).

While calculating the cost of a pair of values x_i and y_j , Algorithm 1 treats various mutually exclusive cases differently based on the values of x_i and y_j (i.e., a real observation or a run of zeros), and the direction of the cell (i.e., $D_{i-1,j-1}$, $D_{i-1,j}$ or $D_{i,j-1}$) to which the cost will be added to. The following facts describe the cases in *UBCosts*, one by one.

Observation 1 AWarp (Algorithm 2) is identical to DTW for any traditional time series, although it is designed for encoded series.

It is a trivial observation. If x and y are traditional vectors, there is no run of zeros in x and y by definition. Therefore, the *UBCosts* algorithm must always execute the first case in line 1, which is the squared error between the values, as in the definition of DTW.

Observation 2 AWarp distance of encoded binary-valued series is identical to the DTW distance of their traditional representations.

Algorithm 1 describes the cases we need to treat separately for binary-valued encoded series. The case in line 1 is the trivial case when both of the inputs a and b are real observations. The value v is simply the squared error. In line 2, we have one observation ($a = 1$) and one run of zeros (b). There can be two inner cases: the run of zeros has already been aligned (*left*) or it is being aligned for the first time (*right* or *diagonal*). If the run of zeros is being aligned for the first time, we have no choice other than aligning all of the zeros with some real observation(s). In the case of a binary-valued series, the real observation(s) are always identical and their values are one, no matter where they are located. Thus the term ba^2 aligns the zeros. If the run of zeros has already been aligned to previous value(s) of the real observation a , we just align a with the last zero of the run, hence the term $a^2 = 1$. The case in line 4 is the mirror of the case in line 2. The default case in line 6 is triggered when both a and b are runs of zeros, which can only result into a distance of zero. In Fig. 3, we show twelve cases, which are all of the possible cases in binary-valued time series, and we illustrate how *UBCosts* calculates the optimal alignment. The solid lines (aligning the red and blue time series) represent the so-far-alignment, and the dotted lines show the new alignment for which *UBCosts* is calculating the cost.

As shown in Fig. 2, if we take the DTW matrix of the traditional binary-valued time series and remove the rows and columns corresponding to zeros that are followed by other zeros, we obtain the matrix calculated by the AWarp algorithm.

Algorithm 1 $UBCosts(a, b, c)$

Require: $a \leftarrow$ an observation, $b \leftarrow$ another observation, $c \leftarrow$ a case identifier
Ensure: Output the distance value v between a and b
 1: case: a and b are observations: $v \leftarrow (a - b)^2$
 2: case: a is an observation and b is a run of zeros:
 3: **if** $c = left$ $v \leftarrow a^2$ **else** $v \leftarrow ba^2$
 4: case: a is a run of zeros and b is an observation:
 5: **if** $c = top$ $v \leftarrow b^2$ **else** $v \leftarrow ab^2$
 6: case *default*: $v \leftarrow 0$
 7: **return** v

Algorithm 2 $AWarp(x, y)$

Require: $x, y \leftarrow$ two encoded time series for comparison
Ensure: Output warping distance between x and y
 1: $l_x \leftarrow length(x), l_y \leftarrow length(y)$
 2: $D(0 : l_x, 0 : l_y) \leftarrow \infty$
 3: $D_{0,0} \leftarrow 0$
 4: **for** $i \leftarrow 1$ to l_x **do**
 5: **for** $j \leftarrow 1$ to l_y **do**
 6: $a_d \leftarrow D_{i-1, j-1} + UBCosts(x_i, y_j, diagonal)$
 7: $a_t \leftarrow D_{i, j-1} + UBCosts(x_i, y_j, top)$
 8: $a_l \leftarrow D_{i-1, j} + UBCosts(x_i, y_j, left)$
 9: $D_{i,j} \leftarrow \min(a_d, a_t, a_l)$
 10: **return** D_{l_x, l_y}

4.2 Any-valued series

As we have described the exactness of AWarp in case of binary-valued time series, the natural question is if the exactness holds for any-valued time series. The answer is no.

Observation 3 AWarp on any-valued encoded series approximates the DTW distance between their traditional representations.

We first discuss why AWarp is not exact for any-valued time series. Although the encoded representation is not lossy, the optimal alignment, which is similar to classic DTW, is not possible for any-valued encoded series. This is because run-length encoding treats all zeros as identical, while an optimal warping alignment may treat zeros in the same run differently.

Example In Fig. 4, two time series $x = \langle 1, 2, 3, 0, 1 \rangle$ and $y = \langle 1, 0, 0, 4, 1 \rangle$ are shown in red and blue, respectively. Note that these time series contain various positive observations as opposed to just one. The optimal DTW aligns the first zero of y with the first one of x and the second zero of y is aligned with the two of x . Such a scenario of aligning part of a run of zeros to one observation and the remaining part of the run to another observation is not possible in the encoded representation, where we treat all the zeros as one entity. If we encode

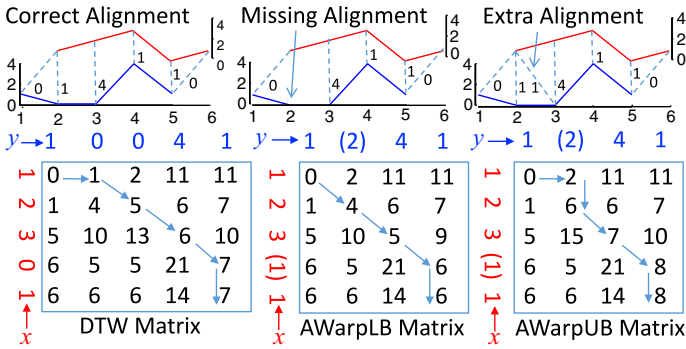


Fig. 4 An example demonstrating that optimal alignment in the encoded representation is not possible

x and y and calculate the AWarp distance, the *UBCosts* function aligns the run of two zeros of y to the first one of x . Therefore, AWarp accumulates a higher distance than the optimal DTW and forms an upper-bounding function of the DTW distance measure. Similarly, if in the *UBCosts* algorithm, we skipped aligning the run of two zeros of y with the first one of x , AWarp would have accumulated a smaller distance than the optimal DTW and formed a lower-bounding function of the DTW distance.

Algorithm 3 *LBCosts*(a, b, c)

Require: $a \leftarrow$ an observation, $b \leftarrow$ another observation, $c \leftarrow$ a case identifier
Ensure: Output the distance value v between a and b
 1: case: a and b are observations: $v \leftarrow (a - b)^2$
 2: case: a is an observation and b is a run of zeros:
 3: **if** $c = top$ $v \leftarrow ba^2$ **else** $v \leftarrow a^2$
 4: case: a is a run of zeros and b is an observation:
 5: **if** $c = left$ $v \leftarrow ab^2$ **else** $v \leftarrow b^2$
 6: case *default*: $v \leftarrow 0$
 7: **return** v

We define the lower-bounding cases in Algorithm 3, where the term ba^2 is applied to only the *top* case and the term ab^2 is applied to only the *left* case. The difference between the *UBCosts* and *LBCosts* is that the *diagonal* cost in the former is always equal or larger (ab^2 or ba^2) than the latter (b^2 or a^2). From now on, we will use AWarp_UB and AWarp interchangeably to refer to Algorithm 2 and AWarp_LB to refer to the same algorithm where *UBCosts* are replaced with *LBCosts*.

At this point, the most important question is: *how good are these bounding functions?* To test them, we generate a comprehensive set of synthetic datasets in the following way. Each dataset has a sparsity factor from the following: 2, 4, 8, 12, 16, 24, 32. Each dataset is associated with a distribution (uniform, normal, binomial and exponential) to generate random numbers from. To generate a dataset, we create 1000 pairs of zero vectors of length 128. We insert random values between one and five in the zero vectors at random locations drawn from the associated distribution. The number of values that are inserted depends on the associated sparsity factor.

For each pair of time series in a dataset, we calculate the upper bound (i.e., AWarp), the lower bound as described above, and the DTW distance in the traditional representation.

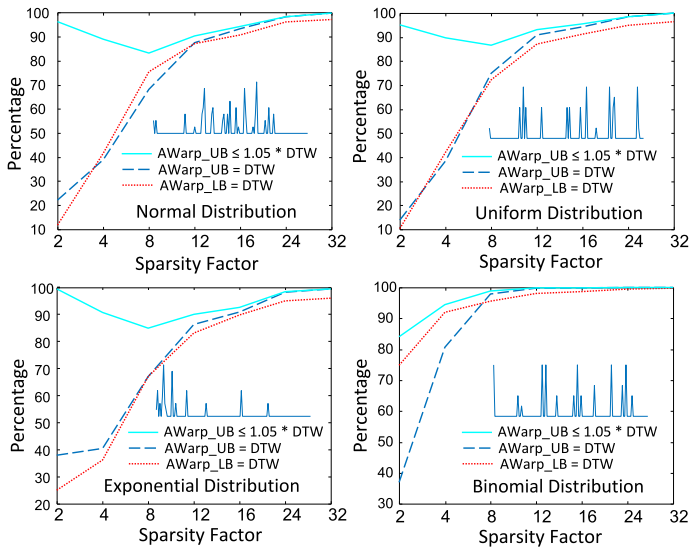


Fig. 5 AWarp_LB and AWarp_UB on encoded series with respect to DTW on vector representation. On average, 90% of the times the upper bound is within 5% of the true distance. Sample time series are shown inside

We calculate the percentage of exact and approximate matches (up to 5% error) between the bounds and DTW distances. The results are shown in Fig. 5. AWarp_UB, approximately 90% of the times, is within 5% of the true distance value. The accuracy converges to 100% as data becomes sparser. These results empirically support that AWarp distance for sparse time series in the encoded form is almost identical to the DTW distance in the traditional form.

The cup-shapes of the approximate matches in Fig. 5 can be explained. For low sparsity factor, the number and length of the runs of zeros are smaller than that when sparsity factor is high. Thus, for low sparsity factor, high accuracy is achieved by exploiting the Observation 1.

Although AWarp is not exactly identical to DTW, there is a simple way to test if AWarp distance is exact. we can calculate AWarp_LB and check if it is equal to AWarp. If they are the same, the distance must be exactly equal to the DTW distance. Thus, we can validate the exactness without calculating the expensive DTW distance by just two AWarp calculations on encoded series, and use AWarp as a preprocessing step ahead of the exact DTW calculation on sparse data.

4.3 Invariance to partial encoding

As mentioned before, a partially encoded series is a longer version of an encoded series where a run of zeros can follow another run of zeros. Let us informally define *order* of partially encoded series as the number of zeros that have been encoded.

Observation 4 AWarp is invariant to the order of partial encoding.

Let us first give an example. If $x = \langle 7, (2), 9, 6, (3), 1 \rangle$ is an encoded series and $x' = \langle 7, (2), 9, 6, (2), (1), 1 \rangle$ is a partially encoded series of x , then the above fact ensures $\text{AWarp}(x, y) = \text{AWarp}(x', y)$. This observation can be easily explained by the *UBCosts* algorithm, which solely depends on the two values, a and b , and is not impacted by prior or later values in the series. Since x and x' are equivalent series, the distance values must be




Fig. 6 (Left) The exactness of constrained AWarp_LB and AWarp_UB for various windows. (Right) The error and exactness of partially encoded representation as we split runs of zeros into halves iteratively

identical. Optimality in substructures is a classic property of dynamic programming. This fact is simply an alternative description of the optimal substructure of the AWarp algorithm that we will exploit in the multidimensional version.

AWarp(x' , y') is always closer to the DTW distance on traditional representations than AWarp(x , y), where x' and y' are partial encodings of x and y , respectively. The reason is that the more runs of zeros are split, the closer the partial encoding is to the traditional representation. To test this statement, we define an operation, *split*, on an encoded series that splits every run of two or more zeros into half. If we iteratively split an encoded series, the series is eventually converted to the traditional version. The impact of such iterative splits on exactness is shown in Fig. 6 (right). As we split more, the error decreases and the exactness increases.

4.4 Multidimensional warping

We have so far discussed the one dimensional algorithms for calculating AWarp. We consider the multidimensional extension of AWarp using approaches similar to those developed for traditional DTW in [25]. There are three general ways to extend DTW to multidimensional time series:

- Independent:** Calculate the individual optimal distances and sum them after normalization by the path length.
- Aggregate:** Sum up the individual dimensions into one superposed time series and encode them to calculate the AWarp distance using Algorithm 2.
- Dependent:** Calculate the global optimal distance assuming that all of the observations at a timestamp must be aligned together to the observations of another timestamp.

Extending AWarp to multidimensional-encoded time series is trivial for the independent scenario. In the aggregate scenario, we sum up the individual dimensions. A simple way to sum two encoded sparse time series is to convert them to traditional time series, add the series, and encode them back to obtain the aggregated time series. It is even more simple to aggregate two time-value sequences. We concatenate the two sequences, sort the concatenated sequence based on time, and add observations which appear at the same time. The time cost is linear in both the cases.

In the dependent scenario, it is nontrivial to calculate the global optimal distance. The recursive step of the dependent version of multidimensional warping distance is given below.

$$D(i, j) = \sum_{k=1}^d (X_{ik} - Y_{jk})^2 + \min \begin{cases} D(i-1, j) \\ D(i, j-1) \\ D(i-1, j-1) \end{cases} \quad (4)$$

The above definition of the multidimensional DTW does not work on encoded series directly. For example, if a two-dimensional series is $(x_1, x_2) = ((1, 0, 0, -1, 0, 0, 0, 1), (1, 0, 0, 0, 0, 1, 0, 1))$, then the encoded representation is $(x_1, x_2) = ((1, (2), -1, (3), 1), (1, (4), 1, (1), 1))$. Clearly, the locations of real observations are not aligned in x_1 and x_2 . In order to convert them to a workable representation, we partially encode x and y in a way that runs of zeros always end at an observation in one of the dimensions. For example, $(x'_1, x'_2) = ((1, (2), -1, (1), (1), (1), 1), (1, (2), (1), (1), 1, (1), 1))$ is an equivalent representation of x and y where the values are time aligned. On sequences of different lengths, aligning them requires managing the ends carefully. We provide the Algorithm 4 that describes the alignment process for two run-length encoded sequences corresponding to two dimensions. The algorithm aligns every positive observation with another observation or a zero in the other dimension. When there are more than two dimensions, the process will be to align pairs of dimensions until no change is needed.

The AWarp algorithm will need to calculate the sum of *UBCosts* over all of the dimensions in lines 8–10 to accommodate the recursion specified above. We skip the details due to lack of space and will explain in detail in an extended version of this paper. In [25], the authors have shown that a combination of the dependent and independent algorithms can beat both of them individually. We will consider such extensions for multidimensional AWarp in future.

Algorithm 4 *AlignDimensions*(x, y)

Require: $x, y \leftarrow$ run length encoded dimensions of a multidimensional time series

Ensure: $dfx, dfy \leftarrow$ aligned run length encoded time series

```

1: while  $x$  is not empty or  $y$  is not empty do
2:   case:  $x$  empty
3:   while  $y$  is not empty do
4:     Append ( $head(y)$ ) to  $dfx$  if  $isRun(head(y))$ 
5:     Append (1) to  $dfx$  if  $isValue(head(y))$ 
6:   case:  $y$  empty
7:   while  $x$  is not empty do
8:     Append ( $head(x)$ ) to  $dfy$  if  $isRun(head(x))$ 
9:     Append (1) to  $dfy$  if  $isValue(head(x))$ 
10:  case:  $isValue(head(x))$  and  $isValue(head(y))$ 
11:  Append  $head(x)$  to  $dfx$  and Append  $head(y)$  to  $dfy$ 
12:  Move to next  $x$  and  $y$ 
13:  case:  $isRun(head(x))$  and  $isValue(head(y))$ 
14:  Append (1) to  $dfx$  and Append  $head(y)$  to  $dfy$ 
15:  Move to next  $y$  and set  $head(x) \leftarrow -(|head(x)| - 1)$ 
16:  case:  $isValue(head(x))$  and  $isRun(head(y))$ 
17:  Append  $head(x)$  to  $dfx$  and Append (1) to  $dfy$ 
18:  Move to next  $x$  and set  $head(y) \leftarrow -(|head(y)| - 1)$ 
19:  case:  $isRun(head(x))$  and  $isRun(head(y))$ 
20:   $m \leftarrow \min(|head(x)|, head(y))$ 
21:  Append ( $m$ ) to  $dfx$  and  $dfy$ 
22:  if  $m = |head(x)|$  then
23:    Move to next  $x$  and set  $head(y) \leftarrow -(|head(y)| - m)$ 
24:  else
25:    Move to next  $y$  and set  $head(x) \leftarrow -(|head(x)| - m)$ 
26: return  $dfy, dfy$ 

```

4.5 Constrained warping

It is widely accepted that constraining the warping between two time series in a user-given window not only helps data mining algorithms to run more quickly, but also enforces physical laws in the matching process [5, 12, 21]. Figure 2 (right) shows an example of a constrained (Sakoe–Chiba band) AWarp matrix. The constrained AWarp algorithm for encoded time series is shown in Algorithm 5. This algorithm is identical to Algorithm 2 except the lines 6–9. In line 6, the absolute difference between the timestamps of x_i and y_j is calculated. We assume that the timestamp of every observation in the encoded series is available to us. It takes linear time to calculate these absolute timestamps if we know t_0 , and the overhead is minimal compared to the overall computational cost.

The condition on line 7 ensures that if $tx_i > ty_j + w$ then $tx_{i-1} > ty_j + w$ must be true to set a cell to infinity. If $tx_i > ty_j + w$ and $tx_{i-1} < ty_j + w$, then x_i is a run of zeros, which contains the timestamp $ty_j + w$ (boundary of the Sakoe–Chiba band). As mentioned before, AWarp cannot align a run of zeros in parts, therefore, when a run of zeros contains the boundary of Sakoe–Chiba band, we extend the band until the next observation after the run of zeros. This forces us to calculate some extra cells that would have been infinity if we used the traditional representation. However, constrained AWarp ensures that no cell within the band is skipped, as Line 7 also checks the mirror case for $ty_j > tx_i + w$.

In Fig. 6 (left), we show the correctness of the *AWarp_{LB}* and *AWarp_{UB}* algorithms as we increase constraint window size. We generate a time series of length 200 with 50% sparsity and normally distributed observations. We calculate 10,000 random distances using Algorithm 5 and check what percentage of the distances match the exact constrained DTW distance. We find that the accuracy increases as the window grows. *AWarp_{UB}* converges quickly to 100%, while *AWarp_{LB}* show some variance. Note that the exactness is always above 96.5% for *AWarp_{LB}* and above 99% for *AWarp_{UB}*.

Algorithm 5 *Constrained_AWarp*(x, y, w)

Require: $x \leftarrow$ a sequence of timestamps, $y \leftarrow$ another sequence of timestamps

Ensure: Output warping distance between the two sequences x and y

```

1:  $l_x \leftarrow \text{length}(x), l_y \leftarrow \text{length}(y)$ 
2:  $D(0 : l_x, 0 : l_y) \leftarrow \infty$ 
3:  $D_{0,0} \leftarrow 0$ 
4: for  $i \leftarrow 1$  to  $l_x$  do
5:   for  $j \leftarrow 1$  to  $l_y$  do
6:      $gap \leftarrow |tx_i - ty_j|$ 
7:     if  $gap > w$  and
       ( $ty_{j-1} - tx_i > w$  or  $tx_{i-1} - ty_j > w$ ) then
8:        $D_{i,j} \leftarrow \infty$ 
9:     else
10:       $a_d \leftarrow D_{i-1,j-1} + \text{UBCosts}(x_i, y_j, \text{diagonal})$ 
11:       $a_l \leftarrow D_{i,j-1} + \text{UBCosts}(x_i, y_j, \text{left})$ 
12:       $a_t \leftarrow D_{i-1,j} + \text{UBCosts}(x_i, y_j, \text{top})$ 
13:       $D_{i,j} \leftarrow \min(a_d, a_l, a_t)$ 
14: return  $D_{l_x, l_y}$ 

```

4.6 Conversion of representation

The best sparse representation for time series data depends on sparsity. Time-value sequence is space saver if more than half of the sequence contains zeros. Length encoding can save even

more when the sequence is *very* sparse. It is clear that conversion between representations is useful to harvest benefits of various representations. We provide two algorithms to convert the two common representations (traditional series and sequence of time-value pairs) of sparse time series into run-length encoded series. The conversion algorithms work in linear time and linear space. Both of the algorithms are implemented and shared in our project page.¹

Algorithm 6 *EncodeTimeSeries*(x)

Require: $x \leftarrow$ a uniformly-sampled continuous time series

Ensure: Output the run length encoded time series y

```

1:  $y \leftarrow$  empty
2: append the positive prefix of  $x$  to  $y$ 
3: while end of  $x$  is not reached do
4:   if current value of  $x$  is zero then
5:     keep track of the run length
6:   else
7:     append the previous run-length (if any) to  $y$ 
8:     append the current  $x$  to  $y$ 
9:   append any trailing run-length to  $y$ 
10: return  $y$ 

```

Algorithm 7 *EncodeTimeValueSequence*(x)

Require: $x \leftarrow$ a discrete time series of time-value pairs

Ensure: Output the run length encoded time series y

```

1:  $y \leftarrow$  a one
2: append the negative of the first timestamp to  $y$ 
3: append the value of the first timestamp to  $y$ 
4: while end of  $x$  is not reached do
5:    $t \leftarrow$  the time since the previous timestamp
6:   append ( $t$ ) to  $y$ 
7:   append the current value to  $y$ 
8: return  $y$ 

```

4.7 Normalization of sparse time series

AWarp, as a distance measure, is agnostic to any preprocessing such as normalization. Many recent articles demonstrate the need for normalization when mining patterns from time series data to achieve scale and shift invariance [21, 26–28]. We evaluate the impact of normalization in sparse time series mining.

A sparse time series contains zeros to represent the absolute absence of a phenomenon, which serves as a reference for the observations that present. Thus sparsity is inherently absolute as opposed to being relative to a dynamic baseline. We, therefore, argue that mining sparse time series should not consider shift invariance. However, the values of the observations can be in different units (miles, meters, etc.) requiring scale invariance. We, therefore, propose to normalize (i.e., scale) sparse time series with respect to the absolute maximum value of the observations.

$$\hat{x}_i = \frac{x_i}{\max(x_i)} \quad (5)$$

¹ <http://www.cs.unm.edu/~mueen/Projects/AWarp/>.

Table 1 Dataset summary

Dataset	Instances	Length	Resolution	Duration
TA	4170	36,799	1 s	1 day
AR	3755	1334	1 day	Years
HA	1628	288	5 min	1 day
PW	3089	288	5 min	1 day

Note that, rescaling does not change the zeros in the sparse time series. This ensures that normalization can be done in both representations: length encoded series and sequence of (time-value) pairs in linear scan.

5 Experiments

We validate AWarp empirically on several real datasets, evaluate speed-accuracy trade-off against a baseline and demonstrate tractability on all-pair comparisons.

5.1 Reproducibility statement

We share *everything* related to this paper in our supporting webpage [29]. We share code for AWarp in two languages (C++ and MATLAB), presentation slides, datasets, experimental results, additional experiments, and additional data.

5.2 Datasets

We use four real datasets from diverse domains to demonstrate the scalability of AWarp. The datasets are: Twitter user activity time series (TA), app review time series (AR), human activity time series (HA) and power usage time series (PW). In Table 1, we briefly describe the datasets. The resolutions of the datasets are very carefully chosen to be relevant for the respective domains. In human behavioral activity and electric power usage, a resolution of 5 min is reasonable. In online reviewing activity, a resolution of a day is enough. In Twitter activity time series, a resolution of a second is required because many actions in Twitter only need mouse clicks (e.g., follow, retweet). Detailed descriptions of the datasets are given in the subsequent application sections.

5.3 Speedups

We generate 100,000 pairs of sparse time series for various sparsity factors and lengths where the activities are uniformly distributed. We calculate the average speedup achieved by AWarp over DTW for these pairs and show the results in Fig. 7.

As data becomes more sparse, speedup increases. As data gets larger, the speedup increases even more. This is an incredible feature of AWarp that can *enable applications of warping distance to datasets where DTW cannot run on the uncompressed sparse time series*.

5.4 Tractability

A valid question at this point is: are the sizes and sparsity factors of real datasets large enough to require a method like AWarp? We first validate the major motivation of AWarp. We test




Fig. 7 Speed and accuracy with respect to the sparsity and size of the datasets

Table 2 Speedup achieved on real datasets

Dataset	s	DTW	AWarp	SpeedUp
TwitterActivity	746	180h	0.3 h	557×
AppReviews	3	46h	21 h	2×
HumanActivity	42	907s	34 s	27×
PowerUsage	28	1170s	40 s	29×


the speed of AWarp by comparing the running time of AWarp in the encoded representation with that of DTW in the traditional representation. The gain in speed naturally depends on the resolution of the time series. The higher the resolution, the more sparse the data becomes and the more speedup we gain. We use reasonable resolutions for all of our datasets as shown in Table 1.

We perform *all-pair distance calculations* on each of the datasets using DTW and AWarp. All-pair distance calculations is a basic operation for many data mining task including: hierarchical clustering [30], outlier detection [31], and nearest neighbor classification [32]. We record the speedup and the respective sparsity factors for four real datasets in Table 2. The sparsity factors in our real datasets are large enough to extract at least 2×, and up to 557×, speedup. In each of these domains, the data owners (e.g., Twitter, Google Play) have several orders of magnitude more data than what we use for this experiment. AWarp will be very useful at that scale for performing many basic data mining tasks under warping similarity. We describe four such data mining tasks in the next section.

5.5 Comparison with a baseline

As described earlier, the purpose of AWarp is to calculate the warping similarity of sparse time series much more quickly than the classic dynamic time warping algorithm while retaining the accuracy of a warping distance measure. There are other methods (e.g., FastDTW) that achieve the same for arbitrary time series data, as opposed to sparse time series. We compare AWarp to FastDTW [19] on 1000 pairs of sparse time series for different values of the *radius* parameter. We measure total execution times and percentages of exact distances produced by FastDTW and show the results in Fig. 8. On the same chart, we point to the worst and median accuracy achieved by AWarp (implemented in MATLAB) and the corresponding execution time for various sparsity factors. Note that AWarp has no input parameters. Also note that FastDTW does not vary on sparsity. For completeness, we point to the timings of two classic DTW implementations. FastDTW (Python) is completely dominated by our

Fig. 8 Speed-accuracy trade-off for various methods and implementations



implementations. We show a hypothetical $10\times$ accelerated curve for FastDTW, which is also dominated by our implementations of AWarp and DTW.

Dozens of techniques are available to speedup similarity search [5], subsequence search [21], and indexing time series [33] data. These techniques are equally applicable to sparse time series and can benefit from AWarp's speedup just by replacing DTW with AWarp when calculating true distances to eliminate false positives. Comparing AWarp, DTW, and FastDTW in searching or indexing algorithms is out of scope of this work.

6 Data mining applications

AWarp is a distance measure that nearly optimally aligns two discrete time series much more quickly than DTW aligns them in their traditional representation. However, this work needs to be justified by showing the utility of this speedup in real data mining tasks. In this section, we show four cases of important data mining tasks that *require time warping* and could not have been performed using time warping distance functions without the speedup provided by AWarp.

6.1 Bot discovery in Twitter

We evaluate the performance of AWarp for clustering the Twitter activities of thousands of users. We assemble a dataset of every activity, including tweet, retweet and delete, from 4170 randomly chosen users for a day. We form activity time series for each of the users at a resolution of seconds (the data is available at ms resolution).

Activity time series can be very useful for finding surprisingly correlated user groups that are mostly bot operated [34]. To find such correlated user groups, we hierarchically cluster the users based on their AWarp distances. We use the single linkage technique and a threshold of 1 to create the clusters.

We find ten clusters that are very dense groups of ten or more users with highly synchronous activities. Several of these clusters can be further merged to form four semantically coherent clusters. One of the clusters was spreading pornographic content and is now mostly suspended by Twitter. Another cluster is spreading news, videos, and images about Selena Gomez (wedselena13, wedselena, wedselena12). The remaining two clusters were spreading identical content in two specific languages: Portuguese (patetamos, IndiretasMusica, LoucoDeVodka) and Malaysian (elzmn01, _ItSy4mimi, zazaizzaty96).




Fig. 9 (Left) Time series of a cluster of 35 bots. Each spike is one tweet. Note the warping in time axis. (Right) Dendrogram of the Twitter accounts using constrained (60s) AWarp. Most of the random users are outliers and several clusters of bots are formed




Fig. 10 Example of time series motif in bot activities. x axis is in ms, y axis shows number of tweets

We show some of the activity time series from the cluster of Portuguese language in Fig. 9 (left). The time series show arbitrary shifts in tweet timestamps because of queuing delay, transmission delay, tweet registration delay, geographically separated data centers, and many other reasons. Such unstructured delay between synchronous tweets breaks Euclidean distance- and lagged Euclidean distance-based methods and prevents this bot group from being detected and suspended. Since AWarp is two orders of magnitude faster on Twitter data, we could perform the clustering under warping distance and discover such a cluster.

6.2 Temporal patterns in Bot activities

Twitter bots are very active agents. It is interesting analyze temporal patterns in these bots to understand their dynamics. With that objective, we select a group of 1500 bots, and collect 100% of their activities in Twitter for five consecutive days. We then perform two temporal pattern mining algorithms (motif discovery and discord discovery) to identify repeating and outlying structure in the activities.

Time series motif is a repeating subsequence in a long time series [35]. Motif can be very simply defined as the most similar pair of subsequence. Motif discovery is an important data mining tool to identify preserved structure in the underlying dynamics of the data source. We use our time warping distance measure, AWarp, to extract the most similar repeated segments for each bot.

In Fig. 10, we show the activity series of the user DSGuarico for 5 days. Visually there is no periodicity in the activity other than some long pauses. However, the user has a motif that occurs many times (two occurrences are shown in the Fig. 10). The motif is simply a sequence of tweets made at about 500 ms interval (exact interval varies). Clearly it is impossible for a human being to post tweets at this rate even if the tweets are identical.




Fig. 11 Example of discord in bot activities, x axis is in ms, y axis shows number of tweets

Upon further investigation, we observe that all of these tweets are copied from the President of Venezuela, Nicols Maduro. *DSGuarico* was synchronous with at least fifty other bots engaged in similar kind of proliferation of political tweets.

Time series discord is the most anomalous subsequence in a long periodic time series [36]. Discord is defined as the subsequence whose nearest neighbor is the furthest among other nearest neighbors. A good segment of Twitter bots are periodic. For example The Count (@countforever), is a harmless bot that just counts periodically. Another example is Red Swingline (@RedSwingline1), which posts political content periodically. A discord in such bots is unusual and potentially indicates downtime in bot master. In Fig. 11, we show the bot *m_and_e_2* that is periodically posting at every 4 s. We discover a discord of 32 s long pause.

Both motifs and discords are computationally expensive tasks requiring quadratic number of distance computation in the worst-case. A 5-day-long time series at ms resolution contains 4.32×10^8 samples in the traditional representation. AWarp on Length encoded sequences makes it feasible to discover motifs and discords by considering only the timestamps of the tweets. Note that the motifs and discords described above requires high resolution (s or ms) data to be discovered as patterns. Aggregated tweet counts over minutes would not require AWarp, and fail to discover the patterns.

6.3 Pseudo-sparse time series analysis

AWarp is motivated to exploit sparsity. Many real-world time series are not sparse in their raw forms, while can easily be converted to sparse time series without losing much information. For example, seismic recordings are typically stationary having mostly noise and only infrequent signatures of seismic activities. We can very simply use a cut-off threshold to increase sparsity of the signal. Thus, AWarp can be applied on the converted sparse times series to mine patterns in an efficient manner.

We show a simple application of motif discovery in a pseudo-sparse time series. We collect digital seismic data recorded at a station near Yellowstone, WY (station SM06 of network ZH). The station is strategically picked with a hope to contain seismic signals of both natural and human generated activities. In Fig. 12, we show a 10-min long segment of time series. We convert the time series by reducing observations with absolute value less than 5×10^3 . This conversion preserves all high amplitude data, while allowing a sparsity factor of over seven. Dynamic time warping (DTW) alignment can produce valuable insights in seismic data, for example, linking wells to their seismic activities [37,38]. We perform motif discovery on the compressed seismic signal using AWarp and identify a motif that periodically appears in a short window of 10 s. The constant periodicity of the motif within the window is more likely to be human generated, although the signal shape does not confirm anything more specific. Nevertheless, the process of efficiently finding motifs in pseudo-sparse time series potentially can improve seismic data analysis methodologies.

Reducing low magnitude observations is a relatively straight-forward technique to add sparsity. Clearly, it works when the expected mean of the time series is zero, as in some seismic data. When a signal has nonzero mean, we can extend the technique to reduce




Fig. 12 Example of a motif discovered in seismograph after conversion to sparse time series




Fig. 13 Clustering Google trends with AWarp

observations with values in an arbitrary range about the mean. For example, in an extreme scenario, we can convert all the values less than the mean to zeros. Adding sparsity in such way can be useful in search engine trend analysis.

In Fig. 13, we show trends of some keywords as search query in Google [39]. Most trends contain periodicity (i.e., annual, monthly, etc.) or sudden bursts. Ignoring the vast amount of small observations does not change the periodic or bursty patterns much, while provides significant performance boost via algorithms such as AWarp. We collect trends for




Fig. 14 Multidimensional power usage data from two households. Each time series is 1 day long at 5 min resolution starting at midnight. There is neither a fixed schedule nor a fixed load to these appliances

Table 3 Accuracies of different distance functions

Euclidean	DTW	DTW_100	AWarp	AWarp_100
59.89%	62.71%	78.19%	76.78%	78.50%

Maximum accuracy is bold faced

two groups of keywords related to the holiday season and tax season. The keywords are: Christmas, Turkey, Gift, Black-Friday, W2, 1040, H&R and Tax. We convert the trends to sparse time series by replacing observations lower than the mean with zeros. We use constrained AWarp with a window size of a month (i.e., 30 days) to *perfectly* cluster the trends and show the dendrogram in Fig. 13. Note that the grouping within clusters are also meaningful: Christmas is more related to Gift than to Black-Friday or Turkey. Computationally, AWarp has captured the shape of the periodic patterns. Holiday keywords have single spikes whereas tax keywords have double spikes denoting the start and ending of the season (Fig. 13).

6.4 Behavioral classification

We evaluate the classification performance of AWarp in a real-world setting. We use two human activity datasets (HH102 and HH104) from the WSU CASAS repository [9]. Each dataset is from a single-resident apartment recording the activities (e.g., door open, light on, etc.) of the resident. The datasets are partially annotated by labeling the beginning and end of some day-to-day activities, such as toilet, dress, sleep, cook, leave_home, etc. Instead of using the annotations to classify the activities, we ask an alternate question: *can we identify a person based on the status (e.g., opened or closed) of the front door of his apartment?* We pick the daily time series of the front door of the two apartments for over 2 years and create a balanced two-class classification problem of 1628 instances of daily time series of length 288 (i.e., one observation every 5 min). A sample of the dataset is shown in Fig. 1.

We use a 1-NN classifier under Euclidean distance, DTW distance (global and constrained), and our proposed AWarp distance (global and constrained). We evaluate the leave-one-out accuracy for each of these classifiers (see Table 3).

It is interesting to note that there is a big gap between the accuracy of global DTW distance (62.71%) and the accuracy of the global AWarp distance (78.19%). Although global DTW finds *the* optimal alignment between the two series, AWarp penalizes a run of zeros being aligned with some real observations more than DTW does. The difference goes away when we use constrained versions of both of the measures with 100-min widows. Because long runs of zeros are broken into at most 100 min runs, the difference between the global versions is reduced.

Table 4 Accuracy of different distance functions

	Eucl. (%)	DTW (%)	DTW _{1h} (%)	AWarp (%)	AWarp _{1h} (%)
DW	79.56	82.16	76.95	83.57	77.15
CW	81.96	87.58	82.77	85.37	81.16
Both	82.16	88.98	85.77	87.58	71.34

CW cloth washer, *DW* dish washer

Maximum accuracy in each row is bold faced

Irrespective of the difference noted above, a 1-NN classification using AWarp is $26\times$ faster than the DTW-based classifier. This is a substantial difference for large datasets. We estimate that if we use all of the fourteen CASAS datasets of single-resident apartments, it would take 50 min to perform these experiments using AWarp, versus 23 h using a DTW-based classifier.

6.5 Power usage classification

We also evaluate the performance of AWarp on a dataset of the power usage of appliances from two different houses. This dataset has been collected from [8]. Instead of considering all the appliances, we first consider only the power usage of the dishwasher appliance. Typically a dishwasher consumes more than 2000 watts at regular operation. We discretize the power usage time series to on–off time series at a resolution of 5 min. In total we have 500 days of on–off time series for the dishwashers. The two classes have 214 and 286 instances of days. These data are very sparse because dishwashers are not often in use. We consider classifying households by using their dishwashing pattern (See Figure 14).

We use a 1-NN classifier under Euclidean distance, DTW distance (global and constrained), and our proposed AWarp distance (global and constrained). We evaluate the leave-one-out accuracy for each of these classifiers and report the results in Table 4.

We also evaluate the classification accuracy of the same two houses based on the power usage of washing machines. We finally evaluate the accuracy considering both of the appliances together using the multidimensional extension of AWarp. In all three cases, global DTW or AWarp has the highest accuracy compared to constrained DTW, constrained AWarp and Euclidean distances. Classic DTW on full resolution data is expected to produce the best accuracy. However, AWarp beats DTW on DishWasher data because of the longer runs of zeros compared to the ClothWasher data. To perform a leave-one-out cross-validation, DTW took 4.5 h while AWarp took 9 min with a tiny reduction in accuracy of 1.4%.

6.6 Unusual review pattern discovery

We collect a dataset of app reviews from the Google Play Marketplace. This dataset contains the review time series for 3755 mobile apps. To form review time series, we collect the number of reviews an app receives in a day since the beginning of data availability. The time series are therefore of varying lengths, with an average length of 1334 days.

We perform discord discovery [36] on these data to identify the most anomalous review time series. The discord is the object in a dataset whose nearest neighbor is the farthest among all other nearest neighbors. We use AWarp as a distance measure to identify the discord. We find a pair of apps that are “far” from every other app while they are reasonably similar to each other. These apps are `com.facebook.katana` and `com.supercell.clashofclans`, which are two of the most popular apps in the Google




Fig. 15 Review time series found as outliers illustrate the capacity hit and subsequent 2 days cycle in the data collection system

Play Marketplace [40]. These apps have received more than 20 million reviews each and they receive several thousands of reviews every day, which is much greater than the average number of reviews an app receives in the store.

However, the success of AWarp is not catching the popular apps, which can easily be found in Wikipedia, but in efficiently identifying anomalous patterns. The patterns that cause AWarp to detect these two apps as outliers are shown in Fig. 15. These patterns show that the apps receive thousands of reviews in 1 day and do not receive any on another day, which is an impossible scenario. The data collection system has a dynamic limit on the number of reviews it can collect and the system works in a 2-day cycle. If an app is highly popular, the number of reviews it receives in a day exceeds the dynamic limit. For the two outlier apps, the limit is exceeded every day and the collection system gets reviews written in 1 day every 2 days, which is why the pattern appears. Thus, the outliers represent the overloaded scenarios of the data collection system.

7 Conclusion

The goal of our work is to develop a time warping distance measure for sparse time series to exploit sparsity for efficiency. We develop AWarp, which is orders of magnitude faster than DTW and calculates a close approximation of DTW, if not a more accurate measure in some cases, such as in human activity datasets. We show applications of AWarp to four domains where DTW is unusable and AWarp can produce interesting results. We discover new bot behavior in Twitter, and we classify human activity much more quickly than with DTW-based classifiers. In future work, our goal is to use AWarp in large-scale data analysis tasks in distributed platforms. We will consider utilizing AWarp to discover meaningful patterns such as motifs [35], discords [41], and shapelets [42] in sparse time series data.

Acknowledgements This work was supported by the NSF CCF Grant No. 1527127 and the NSF Graduate Research Fellowship under Grant No. DGE-0237002.

References

1. Mueen A, Keogh E (2010) Online discovery and maintenance of time series motifs. In: Proceedings of the 16th ACM SIGKDD international conference on knowledge discovery and data mining—KDD'10, number C in KDD'10. ACM Press, p 1089
2. Ye L, Keogh E (2009) Time series shapelets: a new primitive for data mining. In: Proceedings of the 15th ACM SIGKDD international conference on knowledge discovery and data mining, KDD, pp 947–956

3. Shokoohi-Yekta M, Chen Y, Campana B, Hu B, Zakaria J, Keogh E (2015) Discovery of meaningful rules in time series. In: Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining—KDD'15. ACM Press, New York, pp 1085–1094
4. Hamooni H, Mueen A (2014) Dual-domain hierarchical classification of phonetic time series. In: ICDM 2014. ICDM
5. Keogh E (2002) Exact indexing of dynamic time warping. In: Proceedings of the 28th international conference on very large data bases, VLDB'02, pp 406–417
6. Faloutsos C, Ranganathan M, Manolopoulos Y (1994) Fast subsequence matching in time-series databases. ACM SIGMOD Rec 23(2):419–429
7. Wang X, Mueen A, Ding H, Trajcevski G, Scheuermann P, Keogh E (2013) Experimental comparison of representation methods and distance measures for time series data. Data Min Knowl Disc 26(2):275–309
8. Murray D, Stankovic L, Refit: electrical load measurements. <http://www.refitsmarthomes.org/>
9. Cook DJ, Crandall AS, Thomas BL, Krishnan NC (2013) CASAS: a smart home in a box. Computer 46(7):62–69
10. Run-Length Encoding. https://en.wikipedia.org/wiki/Run-length_encoding
11. Boulgouris N, Plataniotis K, Hatzinakos D (2004) Gait recognition using dynamic time warping. In: IEEE 6th workshop on multimedia signal processing. IEEE, pp 263–266
12. Sakoe H, Chiba S (1978) Dynamic programming algorithm optimization for spoken word recognition. IEEE Trans Acoust Speech Signal Process 26(1):43–49
13. Keogh EJ, Pazzani MJ (2000) Scaling up dynamic time warping for datamining applications. In: Proceedings of the 6th ACM SIGKDD international conference on knowledge discovery and data mining—KDD'00. ACM Press, New York, pp 285–289
14. Rath TM, Manmatha R (2003) Word image matching using dynamic time warping. In: 2003. Proceedings. 2003 IEEE computer society conference on computer vision and pattern recognition, vol 2. IEEE, p II—521
15. Berndt DJ, Clifford J (1994) Using dynamic time warping to find patterns in time series. In: KDD Workshop, pp 359–370
16. Al-Naymat G, Chawla S, Taheri J (2009) SparseDTW: a novel approach to speed up dynamic time warping. In: Proceedings of the Eighth Australasian data mining conference, vol 101. Australian computer society, Inc., Darlinghurst, Australia, pp 117–127
17. Tan LN, Alwan A, Kossan G, Cody ML, Taylor CE (2015) Dynamic time warping and sparse representation classification for birdsong phrase classification using limited training data. J Acoust Soc Am 137(3):1069–80
18. Chu S, Keogh E, Hart D, Pazzani M (2002) Iterative deepening dynamic time warping for time series, Chapter 12, pp 195–212
19. Salvador S, Chan P (2007) Toward accurate dynamic time warping in linear time and space. Intell Data Anal 11(5):561–580
20. Sart D, Mueen A, Najjar W, Niennattrakul V, Keogh E (2010) Accelerating dynamic time warping subsequence search with GPUs and FPGAs. ICDM 2010. In: Proceedings—IEEE international conference on data mining, ICDM, pp 1001–1006
21. Rakthanmanon T, Campana B, Mueen A, Batista G, Westover B, Zhu Q, Zakaria J, Keogh E (2012) Searching and mining trillions of time series subsequences under dynamic time warping. In: Proceedings of the 18th ACM SIGKDD international conference on knowledge discovery and data mining—KDD '12. ACM Press, New York, p 262
22. Begum N, Ulanova L, Wang J, Keogh E (2015) Accelerating dynamic time warping clustering with a novel admissible pruning strategy. In: Proceedings of the 21st ACM SIGKDD international conference on knowledge discovery and data mining- KDD'15. ACM Press, New York, pp 49–58
23. Assent I, Wichterich M, Krieger R, Kremer H, Seidl T (2009) Anticipatory DTW for efficient similarity search in time series databases. J Proc VLDB Endow 2(1):826–837
24. Candan KS, Rossini R, Sapino ML, Wang X (2012) sDTW: computing DTW distances using locally relevant constraints based on salient feature alignments. PVLDB 5(11):1519–1530
25. Shokoohi-Yekta M, Wang J, Keogh E, On the non-trivial generalization of dynamic time warping to the multi-dimensional case, Chapter 33, pp 289–297
26. Lines J, Davis L, Hills J, Bagnall A (2012) A shapelet transform for time series classification. In: Proceedings of the 18th ACM SIGKDD international conference on knowledge discovery and data mining, KDD, pp 289–297
27. Mueen A (2013) Enumeration of time series motifs of all lengths. In: Proceedings—IEEE international conference on data mining, ICDM. ICDM, pp 547–556

28. Zhu Y, Zimmerman Z, Senobari NS, Yeh CCM, Funning G, Mueen A, Brisk P, Keogh E (2016) Matrix profile II: exploiting a novel algorithm and GPUs to break the one hundred million Barrier for time series motifs and joins. In: 2016 IEEE 16th international conference on data mining (ICDM). IEEE, pp 739–748
29. AWarp: Warping Similarity for Sparse Time Series. <http://www.cs.unm.edu/~mueen/Projects/AWarp/>
30. Zhu Q, Batista G, Rakthanmanon T, Keogh E (2012) A novel approximation to dynamic time warping allows anytime clustering of massive time series datasets. In: Proceedings of the 2012 SIAM international conference on data mining, pp 999–1010
31. Yeh CCM, Zhu Y, Ulanova L, Begum N, Ding Y, Dau HA, Silva DF, Mueen A, Keogh E (2016) Matrix profile I: all pairs similarity joins for time series: a unifying view that includes motifs, discords and shapelets. In: 2016 IEEE 16th international conference on data mining (ICDM). IEEE, pp 1317–1322
32. Silva DF, Batista GEAPA (2016) Speeding up all-pairwise dynamic time warping matrix calculation. In: Proceedings of the 2016 SIAM international conference on data mining. Society for Industrial and Applied Mathematics, Philadelphia, pp 837–845
33. Shieh J, Keogh E (2009) ISAX: disk-aware mining and indexing of massive time series datasets. *Data Min Knowl Disc* 19(1):24–57
34. Chavoshi N, Hamooni H, Mueen A (2016) DeBot: Twitter Bot detection via warped correlation. In: 2016 IEEE 16th international conference on data mining (ICDM). IEEE, 12, pp 817–822
35. Mueen A, Keogh E, Zhu Q, Cash S, Westover B (2009) Exact discovery of time series motifs. In: Proceedings of the 2009 SIAM international conference on data mining, pp 473–484
36. Yankov D, Keogh E, Medina J, Chiu B, Zordan V (2007) Detecting time series motifs under uniform scaling. In: Proceedings of the 13th ACM SIGKDD international conference on knowledge discovery and data mining KDD 07, KDD'07, p 844
37. Anderson KR, Gaby JE (1983) Dynamic waveform matching. *Inf Sci* 31(3):221–242 12
38. Herrera RH, Fomel S, van der Baan M (2014) Automatic approaches for seismic to well tying. *Interpretation* 2(2):SD9–SD17
39. Google Trends. <https://www.google.com/trends/>
40. List of Most Downloaded Android Applications. https://en.wikipedia.org/wiki/List_of_most_downloaded_Android_applications
41. Yankov D, Keogh EJ, Rebbapragada U (2007) Disk aware discord discovery: finding unusual time series in terabyte sized datasets. In: ICDM, pp 381–390
42. Mueen A, Keogh E, Young N (2011) Logical-shapelets: an expressive primitive for time series classification. In: The 17th ACM SIGKDD international conference, pp 1154–1162



Abdullah Mueen is an Assistant Professor in Computer Science at University of New Mexico since 2013. Earlier he was a Scientist in the Cloud and Information Science Laboratory at Microsoft Corporation. His major interest is in temporal data mining with a focus on two unique types of signals: social networks and electrical sensors. He has been actively publishing in the data mining conferences including KDD, ICDM and SDM and journals including DMKD and KAIS. He has received runner-up award in the Doctoral Dissertation Contest in KDD 2012. He has won the best paper award in the same conference. Earlier, he earned Ph.D. degree at the University of California at Riverside and B.Sc. degree at Bangladesh University of Engineering and Technology.



Nikan Chavoshi is a Ph.D. candidate at the computer science department of the University of New Mexico. Earlier, she got a Master degree in Software Engineering, and a Bachelor degree in Computer Engineering from Amirkabir University of Technology in Tehran, Iran. Her research interests are in time series mining and temporal behavior analysis in social media. She has designed and implemented a bot detection system called DeBot. She has published in data mining conferences such as WWW, ICDM, and ASONAM. Nikan also did two internships at Visa Research Laboratory as a research intern.



Noor Abu-El-Rub is a Ph.D. student in the Computer Science department at the University of New Mexico. She received her B.S. and M.S. degrees in computer science from Jordan University of Science and Technology. Her Research Interests include Social Networks and Data Mining. More specifically, her focus is on finding fraudulent and unusual behaviors for abusing review systems, by applying statistical analysis and graph mining techniques. She has published articles in ICDM, ASONAM, and ICWE.



Hossein Hamooni is a research scientist of the data analytics team at Visa Research in Palo Alto, CA. He received his Ph.D. (with distinction) in Computer Science from the University of New Mexico. He also has a Master degree in Computer Networks, and a Bachelor degree in Computer Engineering both from Amirkabir University of Technology (Tehran Polytechnic) in Iran. His research interests are in distributed knowledge discovery, big data, and time series mining. During his Ph.D., he worked on scalable distributed systems that can process very big datasets in real time. He has published several papers in data mining conferences like ICDM, WWW, and CIKM. He also did two internships at NEC Labs (Princeton) and Visa Research.



Amanda Minnich received a BA in Integrative Biology from UC Berkeley and an M.S. and Ph.D. with Distinction in Computer Science from the University of New Mexico. Her thesis topic was “Spam, Fraud, and Bots: Improving the Integrity of Online Social Media Data.” While at UNM she was named an NSF Graduate Research Fellow, a PiBBs Fellow, a Grace Hopper Scholar, and the Outstanding Graduate Student for the C.S. Department in 2017. She has published her work at top data mining conferences including WWW, ASONAM, ICDM, and ICWE, and has a patent pending on her dissertation work. Amanda also has a passion for advocating for women in tech; she co-founded and served as President of UNM’s first chartered Women in Computing group, and she frequently volunteers at women in tech events. Amanda is currently a Research Scientist at Lawrence Livermore National Laboratory, where she applies Machine Learning techniques to biological data for drug discovery purposes.



Jonathan MacCarthy is a technical staff member at the Los Alamos National Laboratory. He earned Ph.D. degree in geophysics from New Mexico Institute of Mining and Technology in 2010. His research interest is in computational and mathematical techniques for seismic data analysis. He is a member of American Geophysical Union (AGU) and Seismological Society of America (SSA).