

Tight lower bounds for Dynamic Time Warping

Geoffrey I. Webb

Monash University, Clayton, Victoria, 3800, Australia

François Petitjean

Monash University, Clayton, Victoria, 3800, Australia

Abstract

Dynamic Time Warping (DTW) is a popular similarity measure for aligning and comparing time series. Due to DTW’s high computation time, lower bounds are often employed to screen poor matches. Many alternative lower bounds have been proposed, providing a range of different trade-offs between tightness and computational efficiency. LB_KEOGH provides a useful trade-off in many applications. Two recent lower bounds, LB_IMPROVED and LB_ENHANCED, are substantially tighter than LB_KEOGH. All three have the same worst case computational complexity—linear with respect to series length and constant with respect to window size. We present four new DTW lower bounds in the same complexity class. LB_PETITJEAN is substantially tighter than LB_IMPROVED, with only modest additional computational overhead. LB_WEBB is more efficient than LB_IMPROVED, while often providing a tighter bound. LB_WEBB is always tighter than LB_KEOGH. The parameter free LB_WEBB is usually tighter than LB_ENHANCED. A parameterized variant, LB_WEBB_ENHANCED, is always tighter than LB_ENHANCED. A further variant, LB_WEBB*, is useful for some constrained distance functions. In extensive experiments, LB_WEBB proves to be very effective for nearest neighbor search.

Keywords: dynamic time warping, lower bound, time series

1. Introduction

Dynamic Time Warping (DTW) is a time series similarity measure. From its origins in speech recognition [1], it has spread to a broad spectrum of further uses, recent examples of which include gesture recognition [2], signature verification [3], shape matching [4], road surface monitoring [5], neuroscience [6] and medical diagnosis [7]. DTW measures similarity by summing pairwise-distances

Email addresses: geoff.webb@monash.edu (Geoffrey I. Webb),
francois.petitjean@monash.edu (François Petitjean)

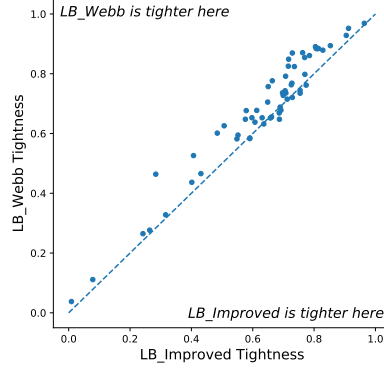
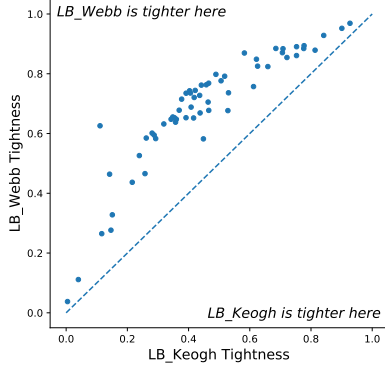


Figure 1: Relative tightness of LB-WEBB and LB-KEOGH Figure 2: Relative tightness of LB-WEBB and LB-IMPROVED

between points in two series. However, it allows flexibility in which points are aligned, to adjust for the way related events may unfold at different paces.

In such applications, lower bounding is often employed to discard potential candidate matches without need to compute the full DTW measure [10]. Numerous such lower bounds have been derived [11, 12, 13, 14, 15, 16]. These provide differing trade-offs between computational cost and tightness, ranging from loose constant time LB_KIM [12] to the relatively tight LB_NEW [14], which requires $O(\ell \cdot w)$ memory and $O(\ell \cdot \log w)$ time to compute a lower bound for a pair of series, where ℓ is the length of the series and w is the window size. These differing trade-offs will each be useful in different applications. The tighter the bound, the less frequent the need to compute the full DTW distance. However, the more compute resource needed to compute the bound, the lower the saving if DTW is not computed.

This paper presents four new bounds. To the best of our knowledge, the first of these new bounds, LB_PETITJEAN, is the tightest known DTW lower bound that has linear complexity with respect to series length and constant complexity with respect to window size. The second, LB_WEBB, belongs to the same complexity class, but is nonetheless substantially faster. Less tight than LB_PETITJEAN, LB_WEBB is always tighter than LB_KEOGH, often substantially so (see Figure 1), and usually tighter than the previous tightest lower bound in the complexity class, LB_IMPROVED (see Figure 2). The third, LB_WEBB_ENHANCED, is a variant of LB_WEBB that may be useful in the context of large window sizes. The fourth, LB_WEBB*, is a variant of LB_WEBB suited to some specific pairwise distance functions.

The paper is organized as follows. Section 2 describes DTW. Section 3 describes key existing bounds. Section 4 introduces the first of the new bounds, LB_PETITJEAN and Section 5 the second, LB_WEBB and its variants, LB_WEBB* and LB_WEBB_ENHANCED. We provide proofs that they are DTW lower bounds and algorithms for calculating them. Section 6 presents experimental evaluation of the utility of these bounds. We first compare their tightness to

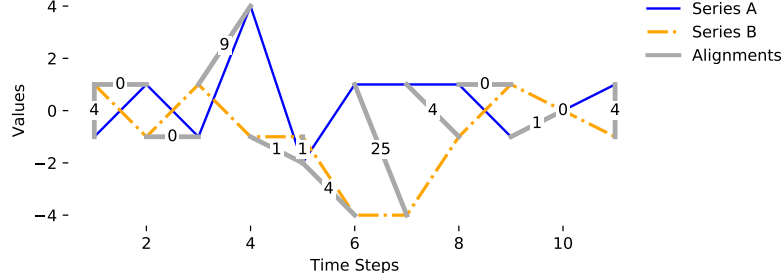


Figure 3: DTW warping path for time series $A = \langle -1, 1, -1, 4, -2, 1, 1, 1, -1, 0, 1 \rangle$ and $B = \langle 1, -1, 1, -1, -1, -4, -4, -1, 1, 0, -1 \rangle$ with window $w = 1$ and $\delta(A_i, B_j) = (A_i - B_j)^2$. Each alignment is labeled with the distance between the elements that are aligned. The DTW distance is the sum of these distances (52).

that of key existing bounds. We next compare their practical value for nearest neighbor search. We end with discussion and conclusions.

2. Problem description

DTW is a similarity measure for aligning and comparing time series [1]. DTW finds the global *alignment* of time series $A = \langle A_1, \dots, A_\ell \rangle$ and $B = \langle B_1, \dots, B_\ell \rangle$, as illustrated in Figure 3. For ease of exposition, we assume A and B are of the same length. However, it is trivial to extend this work to the case of different length series. A *warping path* of A and B is a sequence $\mathcal{A} = \langle A_1, \dots, A_P \rangle$ of *alignments*. Each alignment is a pair $\mathcal{A}_k = (i, j)$ indicating that A_i is aligned with B_j . \mathcal{A} must obey the following constraints:

- **Boundary Conditions:** $\mathcal{A}_1 = (1, 1)$ and $\mathcal{A}_P = (\ell, \ell)$.
- **Continuity and Monotonicity:** for any $\mathcal{A}_k = (i, j)$, $1 < k \leq P$, $\mathcal{A}_{k-1} \in \{(i-1, j), (i, j-1), (i-1, j-1)\}$.

The cost $\text{DTW}(A, B)$ for series A and B is the minimal cost of any warping path and is given in Equation 1, where $\delta(A_i, B_j)$ represents the cost of aligning the two elements. Two common such functions are $\delta(A_i, B_j) = |A_i - B_j|$ and $\delta(A_i, B_j) = (A_i - B_j)^2$. $\text{DTW}(A, B) = D(A_\ell, B_\ell)$.

$$D(A_i, B_j) = \begin{cases} \delta(A_i, B_j) & \text{if } i = 1 \wedge j = 1 \\ \delta(A_i, B_j) + D(A_i, B_{j-1}) & \text{if } i = 1 \wedge 1 < j \leq \ell \\ \delta(A_i, B_j) + D(A_{i-1}, B_j) & \text{if } 1 < i \leq \ell \wedge j = 1 \\ \delta(A_i, B_j) + \min \begin{bmatrix} D(A_{i-1}, B_{j-1}), \\ D(A_i, B_{j-1}), \\ D(A_{i-1}, B_j) \end{bmatrix} & \text{if } 1 < i \leq \ell \wedge 1 < j \leq \ell. \end{cases} \quad (1)$$

[illegible]

Figure 4: A cost matrix for calculating DTW with window $w = 1$.

The path with minimal cost can be found using dynamic programming by building a cost matrix D . Each cell (i, j) of the matrix records the minimum cost of aligning $\langle A_1, \dots, A_i \rangle$ and $\langle B_1, \dots, B_j \rangle$.

Windowing adds a further constraint, that A_i may only be aligned with B_j if $i - w \leq j \leq i + w$, where $w \in \mathbb{N}$ is the *window*. $\text{DTW}_w(A, B) = D_w(A_\ell, B_\ell)$ where,

$$D_w(A_i, B_j) = \begin{cases} \delta(A_i, B_j) & \text{if } i = 1 \wedge j = 1 \\ \delta(A_i, B_j) + D(A_i, B_{j-1}) & \text{if } i = 1 \wedge 1 < j \leq w + 1 \\ \delta(A_i, B_j) + D(A_{i-1}, B_j) & \text{if } 1 < i \leq w + 1 \wedge j = 1 \\ \delta(A_i, B_j) + \min \begin{bmatrix} D(A_{i-1}, B_{j-1}), \\ D(A_i, B_{j-1}) \end{bmatrix} & \text{if } i = j + w \wedge 1 < j \leq \ell \\ \delta(A_i, B_j) + \min \begin{bmatrix} D(A_{i-1}, B_{j-1}), \\ D(A_{i-1}, B_j) \end{bmatrix} & \text{if } 1 < i \leq \ell \wedge j = i + w \\ \delta(A_i, B_j) + \min \begin{bmatrix} D(A_{i-1}, B_{j-1}), \\ D(A_i, B_{j-1}), \\ D(A_{i-1}, B_j) \end{bmatrix} & \begin{array}{l} \text{if } 1 < i < j + w \\ \quad \wedge 1 < j < i + w. \end{array} \end{cases}$$

Figure 4 shows the cost matrix corresponding to the warping path with window $w = 1$ illustrated in Figure 3.

The time complexity of calculating DTW with window w is $O(\ell \cdot w)$. While this is linear on both ℓ and w , when both are relatively large the total computation can be prohibitive for the many repetitions that are entailed by operations such as nearest neighbor search. To this end, it is often desirable to employ lower bounds, such as the popular LB_KEOGH with complexity $O(\ell)$. These allow some potential nearest neighbors to be discarded without recourse to the expensive process of computing DTW.

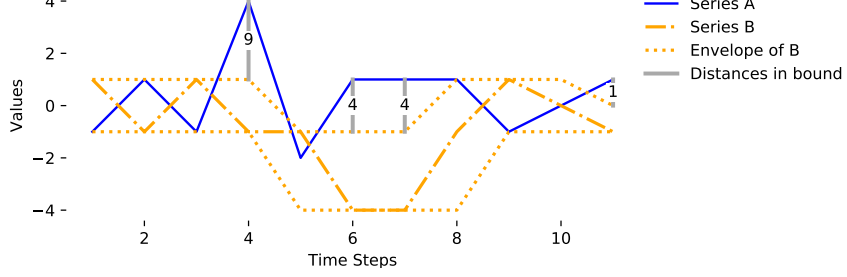


Figure 5: Illustration of LB_KEOGH with $w = 1$ and $\delta(A_i, B_j) = (A_i - B_j)^2$. The gray lines represent the distances LB_KEOGH captures.

3. Key existing bounds

LB_KEOGH [11] employs a pair of derived series called the *envelopes*. Given window w , the upper, \mathbb{U}^S , and lower, \mathbb{L}^S , envelopes of time series S are series representing the maximum and minimum values of S within the window for each point in S .

$$\begin{aligned}\mathbb{U}_i^S &= \max_{\max(1, i-w) \leq j \leq \min(\ell, i+w)} (S_j) \\ \mathbb{L}_i^S &= \min_{\max(1, i-w) \leq j \leq \min(\ell, i+w)} (S_j) \\ \text{LB_KEOGH}_w(A, B) &= \sum_{i=1}^{\ell} \begin{cases} \delta(A_i, \mathbb{U}_i^B) & \text{if } A_i > \mathbb{U}_i^B \\ \delta(A_i, \mathbb{L}_i^B) & \text{if } A_i < \mathbb{L}_i^B \\ 0 & \text{otherwise} \end{cases}\end{aligned}$$

This bound is illustrated in Figure 5.

A tighter bound is provided by LB_IMPROVED [13]. This bound augments LB_KEOGH by capturing not only distances from series A to the envelope of B , but also some distances from B that are not captured by LB_KEOGH. It uses the envelopes of a further derived series called the *projection* of A . The projection $\Omega_w(A, B)$ of A onto B is a sequence such that for all i , $1 \leq i \leq \ell$,

$$\Omega_w(A, B)_i = \begin{cases} \mathbb{U}_i^B & \text{if } A_i > \mathbb{U}_i^B \\ \mathbb{L}_i^B & \text{if } A_i < \mathbb{L}_i^B \\ A_i & \text{otherwise} \end{cases}$$

$$\text{LB_IMPROVED} = \text{LB_KEOGH}_w(A, B) + \sum_{i=1}^{\ell} \begin{cases} \delta(B_i, \mathbb{U}_i^{\Omega_w(A, B)}) & \text{if } B_i > \mathbb{U}_i^{\Omega_w(A, B)} \\ \delta(B_i, \mathbb{L}_i^{\Omega_w(A, B)}) & \text{if } B_i < \mathbb{L}_i^{\Omega_w(A, B)} \\ 0 & \text{otherwise} \end{cases}.$$

This bound is illustrated in Figure 6.

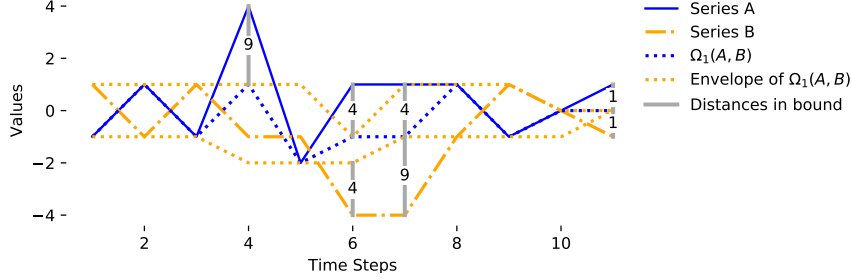


Figure 6: Illustration of LB_IMPROVED with $w = 1$ and $\delta(A_i, B_j) = (A_i - B_j)^2$. The gray lines represent the distances LB_IMPROVED captures.

A recently derived bound, LB_ENHANCED [16], melds two strategies for establishing a lower bound. It uses the strategy of summing distances to an envelope employed by both LB_KEOGH and LB_IMPROVED. It adds to this a constant time operation applied to the start and end of the series, where alignments are more constrained. It employs the concept of a *band*. This is a continuous path through the cost matrix from a cell at the top or right of the accessible region to a cell at the bottom or left. The sum of the minimum values in any collection of non-overlapping bands forms a DTW lower bound. Two key forms of band are the *left bands*

$$\mathcal{L}_i^w = \{(\max(1, i-w), i), (\max(1, i-w) + 1, i), \dots, (i, i), (i, i-1), \dots, (i, \max(1, i-w))\}$$

and the *right bands*

$$\mathcal{R}_i^w = \{(\min(\ell, i+w), i), (\min(\ell, i+w) + 1, i), \dots, (i, i), (i, i-1), \dots, (i, \min(\ell, i+w))\}.$$

The use of each of these types of bands to calculate a lower bound in isolation is illustrated in Figures 7 and 8.

LB_ENHANCED uses just the k leftmost left bands and rightmost right bands, as these are the smallest bands and hence will usually contribute most to the lower bound. The distance between these bands is bridged using LB_KEOGH, as illustrated in Figure 9.

$$\begin{aligned} \text{LB_ENHANCED}_w^k(A, B) &= \sum_{i=1}^k [\min(\mathcal{L}_i^w) + \min(\mathcal{R}_{\ell-i+1}^w)] \\ &\quad + \sum_{i=k+1}^{\ell-k} \begin{cases} \delta(A_i, \mathbb{U}_i^B) & \text{if } A_i > \mathbb{U}_i^B \\ \delta(A_i, \mathbb{L}_i^B) & \text{if } A_i < \mathbb{L}_i^B \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

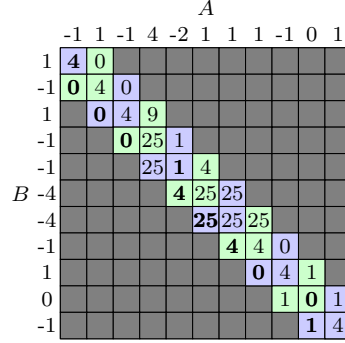


Figure 7: The cost matrix for calculating a lower bound using *left* bands with $w=1$ and $\delta(A_i, B_j) = (A_i - B_j)^2$. Alternating colors distinguish successive bands. The sum over all bands of the minimum distances for that band provides a lower bound (39).

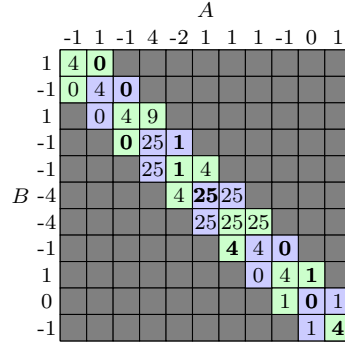


Figure 8: The cost matrix for calculating a lower bound using *right* bands with $w=1$ and $\delta(A_i, B_j) = (A_i - B_j)^2$. Alternating colors distinguish successive bands. The sum over all bands of the minimum distances for that band provides a lower bound (36).

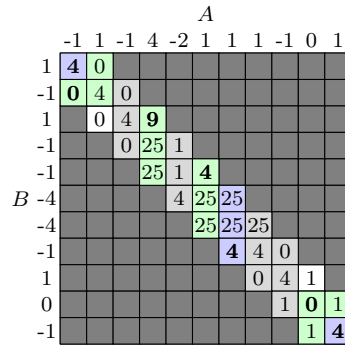


Figure 9: The cost matrix for calculating $\text{LB_ENHANCED}_1^2(A, B)$ with $k=2$, $w=1$ and $\delta(A_i, B_j) = (A_i - B_j)^2$. Alternating colors distinguish successive bands. The vertical bands represent the regions bridged by LB_KEOGH . For the vertical bands in gray $\mathbb{L}_i^B \leq A_i \leq \mathbb{U}_i^B$ and hence the values cannot contribute to the bound. The minimum value in each other band is set in bold. The sum of all these minimum distances provides a lower bound (25).

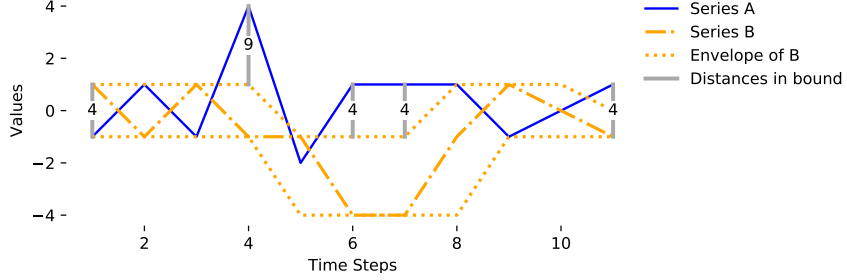


Figure 10: Illustration of LB_ENHANCED with $w = 1$, $k = 1$ and $\delta(A_i, B_j) = (A_i - B_j)^2$. The gray lines represent the distances LB_ENHANCED captures.

4. The LB_Petitjean lower bound

For some elements of A and B , LB_IMPROVED can identify a boundary on the region that LB_KEOGH can reach (the envelope of the projection) and add distances from elements in B to this boundary. Consider an ideal case where the warping path connects A_i and B_j and LB_IMPROVED can incorporate $\delta(A_i, \mathbb{U}_i^B) + \delta(\mathbb{U}_i^B, B_j)$, such as alignment (A_6, B_7) in Figure 3 and the LB_IMPROVED illustration in Figure 6. With $\delta(A_i, B_j) = (A_i - B_j)^2$, the warping path for these points costs $\delta(A_i, B_j) = 25 > \delta(A_i, \mathbb{U}_i^B) + \delta(\mathbb{U}_i^B, B_j) = 4 + 9$, which is the allowance by LB_IMPROVED. We present LB_PETITJEAN, a tighter variant of LB_IMPROVED that uses a stronger strategy for adding to LB_KEOGH an allowance for points in B that cannot be reached by the distances included in LB_KEOGH. It also exploits the constraints on alignments at the start and end of the series, employing a strategy similar to LB_ENHANCED. LB_PETITJEAN uses the following observations.

Observation 1. For any alignment (A_i, B_j) , if $B_j > \mathbb{U}_j^\Omega$ then either $A_i < \mathbb{L}_i^B \leq B_j \leq \mathbb{U}_i^B$ or $\mathbb{L}_i^B \leq A_i \leq B_j \leq \mathbb{U}_i^B$.

Proof 1. For A_i to be aligned with B_j it is necessary that $i - w \leq j \leq i + w$. Hence $\mathbb{L}_i^B \leq B_j \leq \mathbb{U}_i^B$ (note, B indexed by j , the index with which A_i is aligned, but bounds indexed by i). It cannot be that $A_i > B_j$ because that would require that $\Omega_i < A_i$, which can only happen if $A_i > \mathbb{U}_i^B$, in which case $\Omega_i = \mathbb{U}_i^B$ which entails that $\mathbb{U}_j^\Omega \geq \mathbb{U}_i^B$ which contradicts that $B_j > \mathbb{U}_j^\Omega$.

Observation 2. By the same reasoning as Observation 1, for any alignment (A_i, B_j) , if $B_j < \mathbb{L}_j^\Omega$ then either $A_i > \mathbb{U}_i^B \geq B_j$ or $\mathbb{U}_i^B \geq A_i \geq B_j \geq \mathbb{L}_i^B$.

LB_PETITJEAN uses these observations to derive tighter bounds than LB_IMPROVED. Returning to the ideal case, where the warping path connects A_i and B_j and LB_IMPROVED can incorporate $\delta(A_i, \mathbb{U}_i^B) + \delta(\mathbb{U}_i^B, B_j)$, LB_PETITJEAN can instead incorporate the greater amount of $\delta(A_i, \mathbb{U}_i^B) + \delta(\mathbb{L}_j^A, B_j) - \delta(\mathbb{U}_i^B, \mathbb{L}_j^A)$. For alignment (A_6, B_7) in Figure 3 this is 21. $B_7 = -4$ must align with one of A_6, A_7 or A_8 , all of which have value 1. Thus, the value of

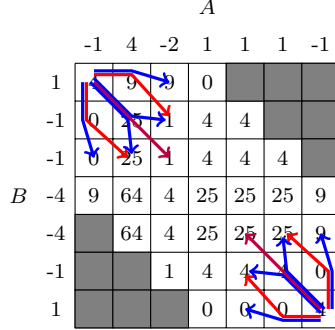


Figure 11: The limited number of potential minimal cost start and end paths of length three.

its alignment must equal $(1 - 4)^2 = 25$. However, the LB_KEOGH bound may already have incorporated an allowance for this alignment of up to the distance between the furthest point in A to which B_7 might be aligned, $\mathbb{U}_7^A = 1$, and the closest point in the envelope of B that is within the window of any point to which B_7 could be aligned, $\mathbb{L}_7^{\Omega_7(A,B)} = -1$. This allowance for the largest possible value from LB_KEOGH for an alignment with B_7 , $\delta(\mathbb{U}_7^A, \mathbb{L}_7^{\Omega_7(A,B)}) = 4$ is subtracted from the distance that is added to the bound, resulting in $25 - 4 = 21$.

LB_PETITJEAN further exploits the tight constraints on the first and last few alignments in any warping path. This additional mechanism, called the *left and right paths*, is inspired by LB_ENHANCED. It incorporates the minimum of each of the possible first and last three sequences of alignments in the path, an even tighter mechanism than that of LB_ENHANCED. The length of these initial and final paths are limited to three because there are just seven such options involving just six distances, as illustrated in Figure 11. If these paths are increased to length 4, the number of alternatives leaps to 21. The most efficient manner to compute start and end paths of any greater length than three would almost certainly be to use the same dynamic programming process used to find the full path, an operation of the same complexity as directly finding the path and hence of little practical utility for calculating a lower bound.

$$\begin{aligned}
\text{MinLRPaths}(A, B) = & \delta(A_1, B_1) + \delta(A_\ell, B_\ell) \\
& + \min[\delta(A_1, B_2) + \delta(A_1, B_3), \delta(A_1, B_2) + \delta(A_2, B_3), \\
& \quad \delta(A_2, B_2) + \delta(A_2, B_3), \delta(A_2, B_2) + \delta(A_3, B_3), \\
& \quad \delta(A_2, B_2) + \delta(A_3, B_2), \delta(A_2, B_1) + \delta(A_3, B_2), \\
& \quad \delta(A_2, B_1) + \delta(A_3, B_1)] \\
& + \min[\delta(A_\ell, B_{\ell-1}) + \delta(A_\ell, B_{\ell-2}), \delta(A_\ell, B_{\ell-1}) + \delta(A_{\ell-1}, B_{\ell-2}), \\
& \quad \delta(A_{\ell-1}, B_{\ell-1}) + \delta(A_{\ell-1}, B_{\ell-2}), \delta(A_{\ell-1}, B_{\ell-1}) + \delta(A_{\ell-2}, B_{\ell-2}), \\
& \quad \delta(A_{\ell-1}, B_{\ell-1}) + \delta(A_{\ell-2}, B_{\ell-1}), \delta(A_{\ell-1}, B_\ell) + \delta(A_{\ell-2}, B_{\ell-1}), \\
& \quad \delta(A_{\ell-1}, B_\ell) + \delta(A_{\ell-2}, B_\ell)]
\end{aligned}$$

$$\text{LB_PETITJEAN assumes that } \forall_{x,y: A_i \leq x \leq y \leq B_j \vee A_i \geq x \geq y \geq B_j} \delta(A_i, B_j) \geq$$

$\delta(A_i, y) + \delta(B_j, x) - \delta(x, y)$. This is true of both $\delta = |A_i - B_j|$ and $\delta(A_i, B_j) = (A_i - B_j)^2$.

Theorem 1. *If $\forall x, y: A_i \leq x \leq y \leq B_j \vee A_i \geq x \geq y \geq B_j$ $\delta(A_i, B_j) \geq \delta(A_i, y) + \delta(B_j, x) - \delta(x, y)$,*

$$\text{LB_PETITJEAN}_w(A, B) = \text{MinLRPaths}(A, B)$$

$$+ \sum_{i=4}^{\ell-3} \begin{cases} \delta(A_i, \mathbb{U}_i^B) & \text{if } A_i > \mathbb{U}_i^B \\ \delta(A_i, \mathbb{L}_i^B) & \text{if } A_i < \mathbb{L}_i^B \\ 0 & \text{otherwise} \end{cases} + \sum_{j=4}^{\ell-3} \begin{cases} \delta(B_j, \mathbb{U}_j^A) - \delta(\mathbb{U}_j^\Omega, \mathbb{U}_j^A) & \text{if } B_j > \mathbb{U}_j^\Omega > \mathbb{U}_j^A \\ \delta(B_j, \mathbb{L}_j^A) - \delta(\mathbb{L}_j^\Omega, \mathbb{L}_j^A) & \text{if } B_j < \mathbb{L}_j^\Omega < \mathbb{L}_j^A \\ \delta(B_j, \mathbb{U}_j^\Omega) & \text{if } B_j > \mathbb{U}_j^\Omega \leq \mathbb{U}_j^A \\ \delta(B_j, \mathbb{L}_j^\Omega) & \text{if } B_j < \mathbb{L}_j^\Omega \geq \mathbb{L}_j^A \\ 0 & \text{otherwise} \end{cases}$$

is a lower bound on $\text{DTW}_w(A, B)$, where \mathbb{U}_i^Ω denotes $\mathbb{U}_i^{\Omega_w(A, B)}$ and \mathbb{L}_i^Ω denotes $\mathbb{L}_i^{\Omega_w(A, B)}$.

Proof 2.

$$\text{DTW}_w(A, B)$$

$$= \sum_{(i, j) \in \mathcal{A}} \delta(A_i, B_j) \quad (2)$$

$$\geq \sum_{(i, j) \in \mathcal{A}} \begin{cases} \delta(A_i, B_j) & \text{if } i \leq 3 \wedge j \leq 3 & (3) \\ \delta(A_i, B_j) & \text{if } i \geq \ell - 2 \wedge j \geq \ell - 2 & (4) \\ \delta(A_i, \mathbb{L}_i^B) + \delta(B_j, \mathbb{U}_j^A) - \delta(\mathbb{U}_j^\Omega, \mathbb{U}_j^A) & \text{if } A_i < \mathbb{L}_i^B \wedge B_j > \mathbb{U}_j^\Omega > \mathbb{U}_j^A & (5) \\ \delta(A_i, \mathbb{L}_i^B) + \delta(B_j, \mathbb{U}_j^\Omega) & \text{if } A_i < \mathbb{L}_i^B \wedge B_j > \mathbb{U}_j^\Omega \leq \mathbb{U}_j^A & (6) \\ \delta(B_j, \mathbb{U}_j^A) - \delta(\mathbb{U}_j^\Omega, \mathbb{U}_j^A) & \text{if } A_i \geq \mathbb{L}_i^B \wedge B_j > \mathbb{U}_j^\Omega > \mathbb{U}_j^A & (7) \\ \delta(B_j, \mathbb{U}_j^\Omega) & \text{if } A_i \geq \mathbb{L}_i^B \wedge B_j > \mathbb{U}_j^\Omega \leq \mathbb{U}_j^A & (8) \\ \delta(A_i, \mathbb{L}_i^B) & \text{if } A_i < \mathbb{L}_i^B \wedge B_j \leq \mathbb{U}_j^\Omega & (9) \\ \delta(A_i, \mathbb{U}_i^B) + \delta(B_j, \mathbb{L}_j^A) - \delta(\mathbb{L}_j^\Omega, \mathbb{L}_j^A) & \text{if } A_i > \mathbb{U}_i^B \wedge B_j < \mathbb{L}_j^\Omega < \mathbb{L}_j^A & (10) \\ \delta(A_i, \mathbb{U}_i^B) + \delta(B_j, \mathbb{L}_j^\Omega) & \text{if } A_i > \mathbb{U}_i^B \wedge B_j < \mathbb{L}_j^\Omega \geq \mathbb{L}_j^A & (11) \\ \delta(B_j, \mathbb{L}_j^A) - \delta(\mathbb{L}_j^\Omega, \mathbb{L}_j^A) & \text{if } A_i \leq \mathbb{U}_i^B \wedge B_j < \mathbb{L}_j^\Omega < \mathbb{L}_j^A & (12) \\ \delta(B_j, \mathbb{L}_j^\Omega) & \text{if } A_i \leq \mathbb{U}_i^B \wedge B_j < \mathbb{L}_j^\Omega \geq \mathbb{L}_j^A & (13) \\ \delta(A_i, \mathbb{U}_i^B) & \text{if } A_i > \mathbb{U}_i^B \wedge B_j \geq \mathbb{L}_j^\Omega & (14) \\ 0 & \text{otherwise} & (15) \end{cases}$$

$$\geq \text{MinLRPaths}(A, B) \quad (16)$$

$$+ \sum_{i=4}^{\ell-3} \begin{cases} \delta(A_i, \mathbb{L}_i^B) & \text{if } A_i < \mathbb{L}_i^B \\ \delta(A_i, \mathbb{U}_i^B) & \text{if } A_i > \mathbb{U}_i^B \\ 0 & \text{otherwise} \end{cases} \quad (17)$$

$$\quad (18)$$

$$\quad (19)$$

$$+ \sum_{j=4}^{\ell-3} \begin{cases} \delta(B_j, \mathbb{U}_j^A) - \delta(\mathbb{U}_j^\Omega, \mathbb{U}_j^A) & \text{if } B_j > \mathbb{U}_j^\Omega > \mathbb{U}_j^A \\ \delta(B_j, \mathbb{L}_j^A) - \delta(\mathbb{L}_j^\Omega, \mathbb{L}_j^A) & \text{if } B_j < \mathbb{L}_j^\Omega < \mathbb{L}_j^A \\ \delta(B_j, \mathbb{U}_j^\Omega) & \text{if } B_j > \mathbb{U}_j^\Omega \\ \delta(B_j, \mathbb{L}_j^\Omega) & \text{if } B_j < \mathbb{L}_j^\Omega \\ 0 & \text{otherwise.} \end{cases} \quad (20)$$

$$\quad (21)$$

$$\quad (22)$$

$$\quad (23)$$

$$\quad (24)$$

□

Notes:

(2) repeats the definition of DTW as a sum over all alignments in \mathcal{A} .

(16) to (24) are the clauses of Theorem 1. (16) adds $\text{MinLRPaths}(A, B)$, a quantity no greater than the sum of the alignments between the first three elements of each series and between the last three elements of each. (17) to (19) add allowances for each $A_i : 4 \leq i \leq \ell-3$. (20) to (24) add allowances for each $B_j : 4 \leq j \leq \ell-3$.

Clauses (3) to (15) repeat the sum over all alignments $(A_i, B_j) \in \mathcal{A}$, separating them by the various possible combinations of a condition in (17) to (19) with a condition in (20) to (24). Each condition implicitly includes *and none of the above*. The key constraints that arise due to this ordering are made explicit. However, it is important to keep in mind that each of (5) to (15) includes the implicit constraint $4 \leq i \leq \ell-3 \vee 4 \leq j \leq \ell-3$. Clauses (3) to (15) do not include cases with both $A_i > \mathbb{U}_i^B$ and $B_j > \mathbb{U}_j^\Omega$ or both $A_i < \mathbb{L}_i^B$ and $B_j < \mathbb{L}_j^\Omega$ because these are mutually inconsistent, as per Observations 1 and 2.

For each clause in (3) to (15), the clause in (17) to (19) that will apply to the specific A_i and the clause in (20) to (24) that will apply to the specific B_j are uniquely determined and the addition to the sum over alignments is the sum of the values that will be added by clauses (16) to (24) for this A_i and B_j . As every A_i and B_i must appear in at least one alignment in \mathcal{A} , and as the sum of the provisions for each A_i and B_j are no greater than the corresponding $\delta(A_i, B_j)$, $\text{LB_PETITJEAN}_w(A, B)$ must be a lower bound on $\text{DTW}(A, B)$.

The following notes discuss each case in turn.

(3,4): these capture the alignments between the first three and between the last three elements of A and B . The $\text{MinLRPaths}(A, B)$ on line (16) contributes an amount not greater than the sum of these. The remaining alignments can include elements in $\{A_1, \dots, A_3, A_{\ell-2}, \dots, A_\ell\}$, but only aligned with elements outside $\{B_1, \dots, B_3, B_{\ell-2}, \dots, B_\ell\}$, and vice versa.

(5): this captures alignments (A_i, B_j) where A_i will be covered by case (17) and B_j will be covered by (20). $A_i < \mathbb{L}_i^B \vdash \Omega_i = \mathbb{L}_i^B \leq \mathbb{U}_j^\Omega$. As $B_j > \mathbb{U}_j^\Omega > \mathbb{U}_j^A$,

$$\delta(A_i, B_j) \geq \delta(A_i, \mathbb{U}_j^\Omega) + \delta(B_j, \mathbb{U}_j^A) - \delta(\mathbb{U}_j^\Omega, \mathbb{U}_j^A) \quad (25)$$

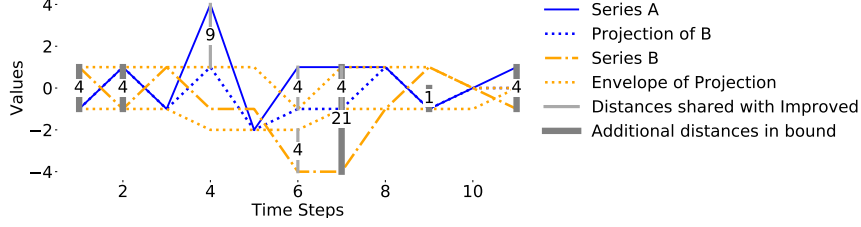


Figure 12: Illustration of LB_PETITJEAN_1 with $\delta(A_i, B_j) = (A_i - B_j)^2$. The dark gray lines represent the points where LB_PETITJEAN_1 captures greater value than LB_IMPROVED . The medium gray lines are values also captured by LB_IMPROVED .

$$\geq \delta(A_i, \mathbb{L}_i^B) + \delta(B_j, \mathbb{U}_j^A) - \delta(\mathbb{U}_j^\Omega, \mathbb{U}_j^A).$$

- (6): this captures alignments (A_i, B_j) where A_i will be covered by case (17) and B_j will be covered by (22). $A_i < \mathbb{L}_i^B \vdash \Omega_i = \mathbb{L}_i^B \leq \mathbb{U}_j^\Omega$. As $B_j > \mathbb{U}_j^\Omega$,

$$\begin{aligned} \delta(A_i, B_j) &> \delta(A_i, \mathbb{U}_j^\Omega) + \delta(B_j, \mathbb{U}_j^\Omega) \\ &\geq \delta(A_i, \mathbb{L}_i^B) + \delta(B_j, \mathbb{U}_j^\Omega). \end{aligned}$$

- (7): this captures alignments (A_i, B_j) where A_i will be covered by case (19) and B_j will be covered by (20).

$$B_j > \mathbb{U}_j^A \vdash \delta(A_i, B_j) \geq \delta(B_j, \mathbb{U}_j^A) \geq \delta(B_j, \mathbb{U}_j^A) - \delta(\mathbb{U}_j^\Omega, \mathbb{U}_j^A).$$

- (8): this captures alignments (A_i, B_j) where A_i will be covered by case (17) and B_j will be covered by (22). $A_i \geq \mathbb{L}_i^B \wedge B_j > \mathbb{U}_j^\Omega \vdash \Omega_i = B_i \leq \mathbb{U}_j^\Omega$. Hence,

$$\delta(A_i, B_j) > \delta(B_j, \mathbb{U}_j^\Omega)$$

- (9): this captures alignments (A_i, B_j) where A_i will be covered by case (19) and B_j will be covered by (24).

$$A_i \leq \mathbb{L}_i^B \vdash \delta(A_i, B_j) > \delta(A_i, \mathbb{L}_i^B).$$

- (10-14): these are equivalent to (5-9), addressing clauses (21) and (23) in place of (20) and (22) and exchanging upper envelopes for lower and vice versa.

- (15): this captures alignments (A_i, B_j) where A_i will be covered by case (19) and B_j will be covered by (24), both of which add zero to the lower bound.

This bound is illustrated in Figure 12.

When an observation holds irrespective of window size we omit the subscript from LB_PETITJEAN .

A variant of LB_PETITJEAN that omits the left and right paths,

$$\text{LB_PETITJEAN_NoLR}_w(A, B) =$$

$$\sum_{i=1}^{\ell} \begin{cases} \delta(A_i, \mathbb{U}_i^B) & \text{if } A_i > \mathbb{U}_i^B \\ \delta(A_i, \mathbb{L}_i^B) & \text{if } A_i < \mathbb{L}_i^B \\ 0 & \text{otherwise} \end{cases} + \sum_{i=1}^{\ell} \begin{cases} \delta(B_i, \mathbb{U}_i^A) - \delta(\mathbb{U}_i^\Omega, \mathbb{U}_i^A) & \text{if } B_i > \mathbb{U}_i^\Omega > \mathbb{U}_i^A \\ \delta(B_i, \mathbb{L}_i^A) - \delta(\mathbb{L}_i^\Omega, \mathbb{L}_i^A) & \text{if } B_i < \mathbb{L}_i^\Omega < \mathbb{L}_i^A \\ \delta(B_i, \mathbb{U}_i^\Omega) & \text{if } B_i > \mathbb{U}_i^\Omega \leq \mathbb{U}_i^A \\ \delta(B_i, \mathbb{L}_i^\Omega) & \text{if } B_i < \mathbb{L}_i^\Omega \geq \mathbb{L}_i^A \\ 0 & \text{otherwise} \end{cases}$$

is tighter than LB_IMPROVED because if $B_j > \mathbb{U}_j^\Omega > \mathbb{U}_j^A$ then LB_IMPROVED will add $\delta(B_j, \mathbb{U}_j^\Omega)$ to the bound whereas LB_PETITJEAN adds the greater amount $\delta(B_j, \mathbb{U}_j^A) - \delta(\mathbb{U}_j^\Omega, \mathbb{U}_j^A)$. Similarly, when $B_j < \mathbb{L}_j^\Omega < \mathbb{L}_j^A$, LB_PETITJEAN also adds a greater amount. However, it is possible, but rare in practice, for LB_PETITJEAN to be less tight than LB_PETITJEAN_NoLR and hence possible for it to be less tight than LB_IMPROVED.

We present pseudocode for calculating LB_PETITJEAN in Algorithm 1. LB_PETITJEAN requires calculation of envelopes around both series as well as around the projection. While the envelopes for the training data can be computed in advance of nearest neighbor search, and the envelope on the query need only be computed once, an envelope on the projection needs to be computed for each query-training data pair. This can be computed in $O(\ell)$ time [13], and hence LB_PETITJEAN has $O(\ell)$ complexity. Its computation is similar to LB_IMPROVED, the major additional overhead being need to compute an envelope on the query and the need to compute two distances for some elements of B rather than one. However, while the complexity is $O(\ell)$, the constants are large and the additional tightness of the bound relative to LB_KEOGH will only compensate for the additional computation in the most demanding of cases.

5. The LB_Webb lower bound

An approximation of LB_PETITJEAN can be computed without recourse to the projection or its envelopes. This more efficient variant, LB_WEBB, uses the concept that an element B_j is *free above* \mathbb{U}^A if all elements A_i within its window are within the envelope of B or cannot access above $\mathbb{L}^{\mathbb{U}^A}$.

$$\text{F}\Uparrow(j) = \forall_i (4 \leq i \leq \ell - 3 \wedge j - w \leq i \leq j + w) \rightarrow (\mathbb{L}_i^B \leq A_i \leq \mathbb{U}_i^B \vee A_i < \mathbb{L}_i^B \leq \mathbb{L}_i^{\mathbb{U}^A}).$$

Similarly, an element B_j is *free below* \mathbb{L}^A if all elements A_i within its window are within the envelope of B or cannot access below $\mathbb{U}^{\mathbb{L}^A}$.

$$\text{F}\Downarrow(j) = \forall_i (4 \leq i \leq \ell - 3 \wedge j - w \leq i \leq j + w) \rightarrow (\mathbb{L}_i^B \leq A_i \leq \mathbb{U}_i^B \vee A_i > \mathbb{U}_i^B \geq \mathbb{U}_i^{\mathbb{L}^A}).$$

Algorithm 1 Algorithm for computing LB_PETITJEAN

procedure LB_PETITJEAN(series A , series B , lower envelope of A LA , upper envelope of A UA , lower envelope of B LB , upper envelope of B UB , window w , # left-right bands k , abandon value a)
 $b \leftarrow \text{MinLRPaths}(A, B)$
 for $i \leftarrow 4$ to $\ell - 3$ **do** ▷ Compute the LB_KEOGH bridge.
 if $A_i > UB_i$ **then**
 $b \leftarrow b + \delta(A_i, UB_i)$
 $P_i \leftarrow UB_i$
 else if $A_i < LB_i$ **then**
 $b \leftarrow b + \delta(A_i, LB_i)$
 $P_i \leftarrow LB_i$
 else
 $P_i \leftarrow A_i$
 end if
 if $b > a$ **then return** b
 end for
 $(LP, UP) \leftarrow \text{compute_envelopes}(P)$ ▷ Linear time algorithm [13].
 for $i \leftarrow 4$ to $\ell - 3$ **do** ▷ Allow for B_i that LB_KEOGH could not reach.
 if $B_i > UP_i > UA_i$ **then**
 $b \leftarrow b + \delta(B_i, UA_i) - \delta(UP_i, UA_i)$
 else if $B_i < LP_i < LA_i$ **then**
 $b \leftarrow b + \delta(B_i, LA_i) - \delta(LP_i, LA_i)$
 else if $B_i > UP_i$ **then**
 $b \leftarrow b + \delta(B_i, UP_i)$
 else if $B_i < LP_i$ **then**
 $b \leftarrow b + \delta(B_i, LP_i)$
 end if
 if $b > a$ **then return** b
 end for
 return b
end procedure

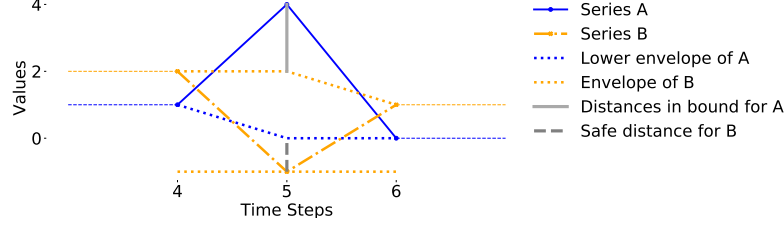


Figure 13: Illustration of free below \mathbb{L}^A . B_5 is free below \mathbb{L}^A when $w = 1$ because none of the distances included in allowance for A_i within its window extend beyond \mathbb{L}^A . A_4 and A_6 are both within the envelope of B , and so do not contribute to the bound. A_5 is above \mathbb{U}^B , so contributes $\delta(A_5, \mathbb{U}_5^B)$. However, as $\mathbb{U}_5^B > \mathbb{L}_5^A$, it does not extend beyond \mathbb{L}^A . Note, for computational efficiency, LB_WEBB uses $\mathbb{U}^{\mathbb{L}^A}$ rather than \mathbb{L}^A , as if $A_i \geq \mathbb{U}_i^B \geq \mathbb{U}_i^{\mathbb{L}^A}$ then the allowance for A_i cannot extend beyond \mathbb{L}_j^A for any j within the window of i . Hence, it is safe to include allowance from B_j to \mathbb{L}_j^A when this is true for all A_i within the window of j .

This is illustrated in Figure 13. If B_j is free above \mathbb{U}^A , then LB_KEOGH does not reach above $\mathbb{L}_j^{\mathbb{U}^A}$ within B_j 's window and hence $\delta(B_j, \mathbb{L}_j^{\mathbb{U}^A})$ can be added to LB_KEOGH. Respectively, if B_j is free below \mathbb{L}^A , then $\delta(B_j, \mathbb{U}_j^{\mathbb{L}^A})$ can be added to LB_KEOGH.

LB_WEBB uses only the envelopes of A and B , an envelope on the envelope of B and a simple record with respect to each point B_i of whether it is free above \mathbb{U}^A or below \mathbb{L}^A . The latter can be generated as a simple side effect of the calculation of the LB_KEOGH bridge.

Theorem 2. If $\forall x, y: A_i \leq x \leq y \leq B_j \vee A_i \geq x \geq y \geq B_j$ $\delta(A_i, B_j) \geq \delta(A_i, y) + \delta(B_j, x) - \delta(x, y)$,

$$\text{LB_WEBB}_w(A, B) = \text{MinLRPaths}(A, B)$$

$$+ \sum_{i=4}^{\ell-3} \begin{cases} \delta(A_i, \mathbb{U}_i^B) & \text{if } A_i > \mathbb{U}_i^B \\ \delta(A_i, \mathbb{L}_i^B) & \text{if } A_i < \mathbb{L}_i^B \\ 0 & \text{otherwise} \end{cases}$$

$$+ \sum_{i=4}^{\ell-3} \begin{cases} \delta(B_i, \mathbb{U}_i^A) & \text{if } F\uparrow(i) \wedge B_i > \mathbb{U}_i^A \\ \delta(B_i, \mathbb{L}_i^A) & \text{if } F\downarrow(i) \wedge B_i < \mathbb{L}_i^A \\ \delta(B_i, \mathbb{U}_i^A) - \delta(\mathbb{U}_i^{\mathbb{L}^B}, \mathbb{U}_i^A) & \text{if } \neg F\uparrow(i) \wedge B_i > \mathbb{U}_i^{\mathbb{L}^B} > \mathbb{U}_i^A \\ \delta(B_i, \mathbb{L}_i^A) - \delta(\mathbb{L}_i^{\mathbb{U}^B}, \mathbb{L}_i^A) & \text{if } \neg F\downarrow(i) \wedge B_i < \mathbb{L}_i^{\mathbb{U}^B} < \mathbb{L}_i^A \\ 0 & \text{otherwise} \end{cases}$$

is a lower bound on $\text{DTW}_w(A, B)$, where $\mathbb{U}_i^{\mathbb{L}^B}$ denotes i^{th} value of the upper envelope of the lower envelope of B and $\mathbb{L}_i^{\mathbb{U}^B}$ denotes i^{th} value of the lower envelope of the upper envelope of B and k is an integer $0 \leq k \leq \ell/2$.

Proof 3.

$$\text{DTW}_w(A, B)$$

$$\begin{aligned}
&= \sum_{(i,j) \in \mathcal{A}} \delta(A_i, B_j) \\
&\geq \sum_{(i,j) \in \mathcal{A}} \begin{cases} \delta(A_i, B_j) & \text{if } i \leq 3 \wedge j \leq 3 & (26) \\ \delta(A_i, B_j) & \text{if } i \geq \ell-2 \wedge j \geq \ell-3 & (27) \\ \delta(A_i, \mathbb{U}_i^B) + \delta(B_j, \mathbb{L}_j^A) & \text{if } A_i > \mathbb{U}_i^B \wedge F\downarrow(j) & (28) \\ \delta(A_i, \mathbb{L}_i^B) + \delta(B_j, \mathbb{U}_j^A) & \text{if } A_i < \mathbb{L}_i^B \wedge F\uparrow(j) & (29) \\ \delta(B_j, A_i) + \delta(A_i, \mathbb{U}_i^B) - \delta(\mathbb{U}_i^B, A_i) & \text{if } A_i > \mathbb{U}_i^B \wedge B_j < \mathbb{L}_j^{\mathbb{U}_j^B} < \mathbb{L}_j^A & (30) \\ \delta(B_j, A_i) + \delta(A_i, \mathbb{L}_i^B) - \delta(\mathbb{L}_i^B, A_i) & \text{if } A_i < \mathbb{L}_i^B \wedge B_j > \mathbb{U}_j^{\mathbb{L}_j^B} > \mathbb{U}_j^A & (31) \\ \delta(A_i, B_j) & \text{if } A_i > \mathbb{U}_i^B & (32) \\ \delta(A_i, B_j) & \text{if } A_i < \mathbb{L}_i^B & (33) \\ \delta(A_i, B_j) & \text{otherwise} & (34) \end{cases} \\
&\geq \text{MinLRPaths}(A, B) & (35) \\
&+ \sum_{i=4}^{\ell-3} \begin{cases} \delta(A_i, \mathbb{U}_i^B) & \text{if } A_i > \mathbb{U}_i^B & (36) \\ \delta(A_i, \mathbb{L}_i^B) & \text{if } A_i < \mathbb{L}_i^B & (37) \\ 0 & \text{otherwise} & (38) \end{cases} \\
&+ \sum_{j=4}^{\ell-3} \begin{cases} \delta(B_j, \mathbb{U}_j^A) & \text{if } F\uparrow(j) \wedge B_j > \mathbb{U}_j^A & (39) \\ \delta(B_j, \mathbb{L}_j^A) & \text{if } F\downarrow(j) \wedge B_j < \mathbb{L}_j^A & (40) \\ \delta(B_j, \mathbb{L}_j^A) - \delta(\mathbb{L}_j^{\mathbb{U}_j^B}, \mathbb{L}_j^A) & \text{if } B_j < \mathbb{L}_j^{\mathbb{U}_j^B} < \mathbb{L}_j^A & (41) \\ \delta(B_j, \mathbb{U}_j^A) - \delta(\mathbb{U}_j^{\mathbb{L}_j^B}, \mathbb{U}_j^A) & \text{if } B_j > \mathbb{U}_j^{\mathbb{L}_j^B} > \mathbb{U}_j^A & (42) \\ 0 & \text{otherwise} & (43) \end{cases}
\end{aligned}$$

Notes.

- (26,27): these capture the alignments between the first three and between the last three elements of A and B . The $\text{MinLRPaths}(A, B)$ on line (35) contributes an amount not greater than the sum of these.
- (28): this captures alignments where $A_i > \mathbb{U}_i^B \geq \mathbb{L}_j^A > B_j$ and (36) will add $\delta(A_i, \mathbb{U}_i^B)$ and (40) will add $\delta(B_j, \mathbb{L}_j^A)$, which sum to less than $\delta(A_i, B_j)$.
- (29): this captures alignments where $A_i < \mathbb{L}_i^B \leq \mathbb{U}_j^A < B_j$ and (36) will add $\delta(A_i, \mathbb{L}_i^B)$ and (40) will add $\delta(B_j, \mathbb{U}_j^A)$, which sum to less than $\delta(A_i, B_j)$.
- (30): this captures alignments for which (36) and (41) will be counted.
- (31): this captures alignments for which (37) and (42) will be counted.
- (32): this captures alignments for which (36) and (43) will apply.
- (33): this captures alignments for which (37) and (43) will apply.
- (34): this captures alignments for which (38) and one of (39), (40), (41) or (42) will apply.

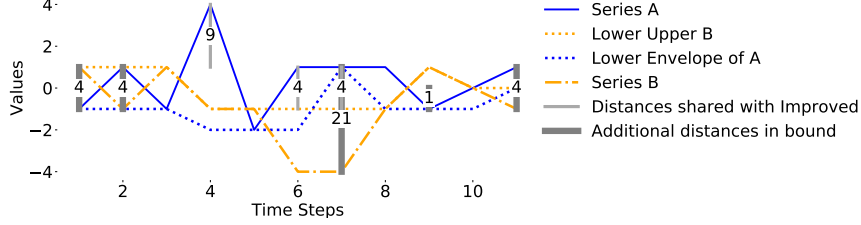


Figure 14: Illustration of LB_WEBB_1 with $\delta(A_i, B_j) = (A_i - B_j)^2$. The dark gray lines represent the points where LB_WEBB_1 captures greater value than LB_IMPROVED . The medium gray areas are those captured by all of LB_KEOGH , LB_IMPROVED and LB_WEBB_1 .

- (35): $\text{MinLRPaths}(A, B)$ equals the minimum value of a path through the first and last three elements of A and B . It cannot be greater than the value of the alignments at (a) and (28).
- (36): this adds the distance from A_i to the upper envelope of B for elements A_i that have alignments captured at (28), (30) and (32).
- (37): this adds the distance from A_i to the lower envelope of B for elements A_i that have alignments captured at (29), (31) and (33).
- (38): this adds zero for elements A_i that fall within the envelope of B , whose alignments are captured at (34).
- (39): this applies to elements B_j for which all alignments (A_i, B_j) are of type (29) or (34). If (29), (37) added $\delta(A_i, \mathbb{L}_i^B)$ and $\delta(A_i, B_j) \leq \delta(A_i, \mathbb{L}_i^B) + \delta(B_j, \mathbb{U}_j^A)$ so it is safe to add the latter term. If (34) then (36) applied, no allowance was added for A_i and hence it is also safe to add $\delta(B_j, \mathbb{U}_j^A)$.
- (40): this applies to elements B_j for which all alignments (A_i, B_j) are of type (28) or (34). If (28), (36) added $\delta(A_i, \mathbb{U}_i^B)$ and $\delta(A_i, B_j) \leq \delta(A_i, \mathbb{U}_i^B) + \delta(B_j, \mathbb{L}_j^A)$ so it is safe to add the latter term. If (34) then (36) applied, no allowance was added for A_i and hence it is also safe to add $\delta(B_j, \mathbb{L}_j^A)$.
- (41): this applies to elements B_j for which at least one alignment is of type (30). $\delta(A_i, \mathbb{U}_i^B)$ is added at (36), leaving $\delta(B_j, A_i) - \delta(\mathbb{U}_i^B, A_i) \leq \delta(B_j, \mathbb{L}_j^A) - \delta(\mathbb{L}_j^{\mathbb{U}^B}, \mathbb{L}_j^A)$.
- (42): this applies to elements B_j for which at least one alignment is of type (31). $\delta(A_i, \mathbb{L}_i^B)$ is added at (37), leaving $\delta(B_j, A_i) - \delta(\mathbb{L}_i^B, A_i) \leq \delta(B_j, \mathbb{U}_j^A) - \delta(\mathbb{U}_j^{\mathbb{L}^B}, \mathbb{U}_j^A)$.
- (43): this applies to elements of B for which all alignments are of type (34).

LB_WEBB is illustrated in Figure 14. Pseudocode is presented in Algorithm 2. It is more efficient to compute than either LB_IMPROVED or LB_PETITJEAN . It is tighter than LB_KEOGH and LB_ENHANCED and often tighter than LB_IMPROVED . It is less tight than LB_PETITJEAN .

Algorithm 2 Algorithm for computing LB_WEBB

procedure LB_WEBB(series A , series B , lower envelope of A LA , upper envelope of A UA , lower envelope of B LB , upper envelope of B UB , lower envelope of UB LUB , upper envelope of LB ULB , window w , # left-right bands k , abandon value a)

```

 $b \leftarrow \text{MinLRPaths}(A, B)$ 
 $c\uparrow \leftarrow w$   $\triangleright$  Count of the number of  $F\uparrow$  elements to the left of  $i$ .
 $c\downarrow \leftarrow w$   $\triangleright$  Count of the number of  $F\downarrow$  elements to the left of  $i$ .
 $F\uparrow \leftarrow \langle \text{false} \rangle^w$   $\triangleright$  true if  $F\uparrow(i)$ . Initialize all elements as false.
 $F\downarrow \leftarrow \langle \text{false} \rangle^w$   $\triangleright$  true if  $F\downarrow(i)$ . Initialize all elements as false.
for  $i \leftarrow k + 1$  to  $\ell - k$  do  $\triangleright$  Compute the LB_KEOGH bridge.
    if  $A_i > UB_i$  then
         $b \leftarrow b + \delta(A_i, UB_i)$ 
         $c\uparrow \leftarrow 0; c\downarrow \leftarrow c\downarrow + 1$ 
    else if  $A_i < LB_i$  then
         $b \leftarrow b + \delta(A_i, LB_i)$ 
         $c\downarrow \leftarrow 0; c\uparrow \leftarrow c\uparrow + 1$ 
    else
         $c\uparrow \leftarrow c\uparrow + 1; c\downarrow \leftarrow c\downarrow + 1$ 
    end if
    if  $c\uparrow > 2 \times w$  then  $F\uparrow(i) \leftarrow \text{true}$ 
    if  $c\downarrow > 2 \times w$  then  $F\downarrow(i) \leftarrow \text{true}$ 
    if  $b > a$  then return  $b$ 
end for
for  $i \leftarrow \max(1, \ell - k - c\uparrow + w)$  to  $\ell$  do  $\triangleright$  Remaining free elements.
     $F\uparrow(i) \leftarrow \text{true}$ 
end for
for  $i = \max(1, \ell - k - c\downarrow + w)$  to  $\ell$  do  $\triangleright$  Remaining free elements.
     $F\downarrow(i) \leftarrow \text{true}$ 
end for
for  $i \leftarrow k + 1$  to  $\ell - k$  do  $\triangleright$  Allow for  $B_i$  that LB_KEOGH could not reach.
    if  $F\uparrow(i) \wedge B_i > UA_i$  then
         $b \leftarrow b + \delta(B_i, UA_i)$ 
    else if  $F\downarrow(i) \wedge B_i < LA_i$  then
         $b \leftarrow b + \delta(B_i, LA_i)$ 
    else if  $B_i > ULB_i \geq UA_i$  then
         $b \leftarrow b + \delta(B_i, UA_i) - \delta(ULB_i, UA_i)$ 
    else if  $B_i < LUB_i \leq LA_i$  then
         $b \leftarrow b + \delta(B_i, LA_i) - \delta(LUB_i, LA_i)$ 
    end if
    if  $b > a$  then return  $b$ 
end for
return  $b$ 

```

end procedure

5.1. LB_WEBB*

In some cases a simplified variant of LB_WEBB can be deployed. Where $\delta(A_i, B_j) = |A_i - B_j|$, $\text{LB_WEBB}_w(A, B) =$

$$\begin{aligned} \text{LB_WEBB}_w^*(A, B) = & \text{MinLRPaths}(A, B) \\ & + \sum_{i=4}^{\ell-3} \begin{cases} \delta(A_i, \mathbb{U}_i^B) & \text{if } A_i > \mathbb{U}_i^B \\ \delta(A_i, \mathbb{L}_i^B) & \text{if } A_i < \mathbb{L}_i^B \\ 0 & \text{otherwise} \end{cases} \\ & + \sum_{i=4}^{\ell-3} \begin{cases} \delta(B_i, \mathbb{U}_i^A) & \text{if } F\uparrow(i) \wedge B_i > \mathbb{U}_i^A \\ \delta(B_i, \mathbb{L}_i^A) & \text{if } F\downarrow(i) \wedge B_i < \mathbb{L}_i^A \\ \delta(B_i, \mathbb{U}_i^{\mathbb{L}^B}) & \text{if } \neg F\uparrow(i) \wedge B_i > \mathbb{U}_i^{\mathbb{L}^B} > \mathbb{U}_i^A \\ \delta(B_i, \mathbb{L}_i^{\mathbb{U}^B}) & \text{if } \neg F\downarrow(i) \wedge B_i < \mathbb{L}_i^{\mathbb{U}^B} < \mathbb{L}_i^A \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

$\text{LB_WEBB}^*(A, B)$ does not require that $\forall x, y: A_i \leq x \leq y \leq B_j \vee A_i \geq x \geq y \geq B_j$, $\delta(A_i, B_j) \geq \delta(A_i, y) + \delta(B_j, x) - \delta(x, y)$ and is a lower bound for DTW where $\delta(A_i, B_j)$ increases monotonically with $|A_i - B_j|$. This is the class of δ for which LB_KEOGH, LB_IMPROVED and LB_ENHANCED are lower bounds of DTW.

5.2. LB_WEBB_ENHANCED

As the DTW window increases in size, lower bounds based on envelopes, such as LB_KEOGH, LB_IMPROVED, LB_PETITJEAN and LB_WEBB, are likely to decline in tightness due to each point in the envelope representing a maximum or minimum over an ever increasing proportion of a whole series. In this case, the method underlying LB_ENHANCED is likely to excel, as it does not use envelopes. To this end, a parameterized variant of LB_WEBB that employs the left and right bands of LB_ENHANCED is likely to come to the fore.

$$\begin{aligned} \text{LB_WEBB_ENHANCED}_w^k(A, B) = & \sum_{i=1}^k [\min(\mathcal{L}_i^w) + \min(\mathcal{R}_{\ell-i+1}^w)] \\ & + \sum_{i=k+1}^{\ell-k} \begin{cases} \delta(A_i, \mathbb{U}_i^B) & \text{if } A_i > \mathbb{U}_i^B \\ \delta(A_i, \mathbb{L}_i^B) & \text{if } A_i < \mathbb{L}_i^B \\ 0 & \text{otherwise} \end{cases} \\ & + \sum_{i=k+1}^{\ell-k} \begin{cases} \delta(B_i, \mathbb{U}_i^A) & \text{if } F\uparrow(i) \wedge B_i > \mathbb{U}_i^A \\ \delta(B_i, \mathbb{L}_i^A) & \text{if } F\downarrow(i) \wedge B_i < \mathbb{L}_i^A \\ \delta(B_i, \mathbb{U}_i^A) - \delta(\mathbb{U}_i^{\mathbb{L}^B}, \mathbb{U}_i^A) & \text{if } \neg F\uparrow(i) \wedge B_i > \mathbb{U}_i^{\mathbb{L}^B} > \mathbb{U}_i^A \\ \delta(B_i, \mathbb{L}_i^A) - \delta(\mathbb{L}_i^{\mathbb{U}^B}, \mathbb{L}_i^A) & \text{if } \neg F\downarrow(i) \wedge B_i < \mathbb{L}_i^{\mathbb{U}^B} < \mathbb{L}_i^A \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

6. Experiments

In order to assess the practical merits of the new lower bounds relative to the prior state of the art, we compare their performance on the 85 dataset Bakeoff Paper [17] version of the widely utilized UCR benchmark time series datasets [18]. We use this version because of the availability of comparative results of many techniques for it.

All experiments are performed using single threaded implementations of the respective algorithms in Java and executed on Intel Xeon CPU E5-2680 v3 2.50GHz CPUs. In the interests of reproducible science the source code is made available at <https://github.com/GIWebb/DTWBounds>.

The experiments are run in a heterogeneous grid environment. While the amount of RAM may differ from experiment to experiment, all comparative runs for a single dataset were performed on a single machine, ensuring that comparative results for each dataset are commensurable even though results between datasets may not be.

We use $\delta = (A_i - B_j)^2$.

6.1. Tightness

We first seek to quantify the relative tightness of our new bounds relative to LB_KEOGH and LB_IMPROVED. Relative tightness will vary greatly depending on window size. To obtain an evaluation that is relevant to real world practice, we use for each dataset the window size recommended by the archive. These recommended window sizes are those that provide most accurate nearest neighbor classification using leave-one-out cross-validation on the training set. Some recommended window sizes are 0. There is no value in computing a linear time lower bound for a window size of zero, as it is quicker to simply compute the full distance. Hence we do not include these datasets in this evaluation and use only the 60 datasets with recommended window sizes of one or more.

We compute tightness for every pair of a training (T) and a test (Q) series. We calculate the tightness of a lower bound $\lambda_w(Q, T)$ on $DTW_w(Q, T)$ as $\lambda_w(Q, T)/DTW_w(Q, T)$. We exclude pairs (Q, T) for which $DTW_w(Q, T) = 0.0$. We compare the average tightness on each dataset for each of LB_PETITJEAN and LB_WEBB against each of LB_IMPROVED and LB_KEOGH in Figures 1, 2 and 15 to 18. LB_PETITJEAN is always tighter than LB_IMPROVED and often substantially tighter. In the most extreme case, for ShapeletSim, LB_PETITJEAN has average tightness of 0.038 while LB_IMPROVED has average tightness of only 0.009. The advantage of LB_PETITJEAN relative to LB_KEOGH is even greater.

LB_WEBB is also necessarily tighter than LB_KEOGH. Figure 1 shows that the advantage is often substantial. It is tighter on average than LB_IMPROVED for 47 datasets and less tight for 13. Figure 2 shows that it is never substantially less tight than LB_IMPROVED and is often substantially tighter.

The tightness of LB_ENHANCED varies with k , tending to, but not always, growing with k . Tan *et. al.* [16] identify $k = 8$ as providing an effective trade-off

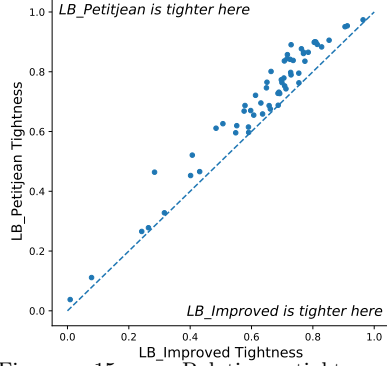


Figure 15: Relative tightness of LB_PETITJEAN and LB_IMPROVED

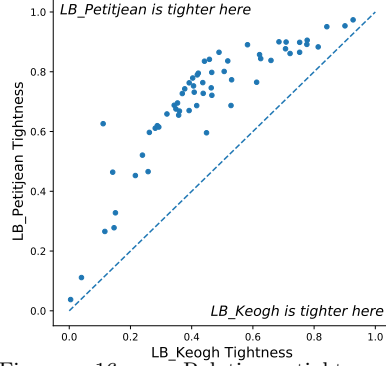


Figure 16: Relative tightness of LB_PETITJEAN and LB_KEOGH

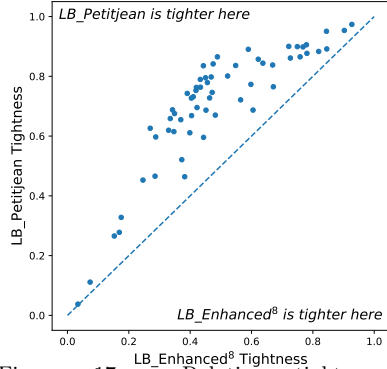


Figure 17: Relative tightness of LB_PETITJEAN and LB_ENHANCED⁸

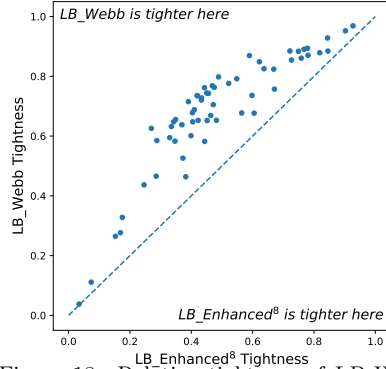


Figure 18: Relative tightness of LB_WEBB and LB_ENHANCED⁸

between tightness and computation. Both LB_PETITJEAN and LB_WEBB are always tighter than LB_ENHANCED⁸. See Figures 17 and 18.

6.2. Classification times with optimal windows

The utility of a lower bound for a given application is determined by the trade-off it provides between tightness and speed. LB_PETITJEAN requires slightly more computation than LB_IMPROVED, while LB_WEBB requires substantially less. To assess the relative utilities of these trade-offs, we next test the efficiency of nearest neighbor search utilizing these bounds, again employing optimal window sizes and hence limiting the evaluation to the 60 datasets for which the optimal window size is greater than zero.

We conduct two types of nearest neighbor search. Each finds for each test series Q , $\arg \min_{T \in \mathcal{T}} \text{DTW}_w(Q, T)$, the training series T that is the nearest neighbor to Q using DTW with the optimal window, w . The first approach, described in Algorithm 3, tests a test series Q against each training series T in random order, first applying the relevant lower bound and then only computing

Algorithm 3 Experimental procedure for nearest neighbor search with random order

```

procedure RANDEXP(set of query series  $\mathcal{Q}$ , set of training series  $\mathcal{T}$ , lower
bound  $\lambda$ )
  for  $Q \in \mathcal{Q}$  do
    if  $\lambda$  requires  $\mathbb{U}^Q$  and  $\mathbb{L}^Q$  then
      Calculate and save  $\mathbb{U}^Q$  and  $\mathbb{L}^Q$ 
    end if
     $b \leftarrow \emptyset$ 
    for  $T \in \mathcal{T}$  do
      if  $b = \emptyset$  then
         $d \leftarrow \text{DTW}(Q, T)$ 
         $b \leftarrow T$ 
      else
        if  $\lambda(Q, T, b) < b$  then
           $d' \leftarrow \text{DTW}(Q, T)$ 
          if  $d' < d$  then
             $b \leftarrow T$ 
             $d \leftarrow d'$ 
          end if
        end if
      end if
    end for
  end for
end procedure

```

the full distance if the lower bound is less than the best distance so far. The second approach, described in Algorithm 4, for each query Q , first computes the lower bound for every training series T , then sorts the training series in ascending order and finally computes the full distances on successive training series until the minimum distance found is less than the next lower bound.

Each approach is repeated ten times for each dataset and average results are presented in order to smooth out variations in time due to extraneous factors and in performance due to randomization for the random order approach. The envelopes for the training series, \mathbb{L}^T , \mathbb{U}^T , $\mathbb{U}^{\mathbb{L}^T}$ and $\mathbb{L}^{\mathbb{U}^T}$, are precalculated and the time for calculating these envelopes is not included in the experimental timings. The calculation of all other envelopes is included in timings. Calculation of \mathbb{L}^Ω and \mathbb{U}^Ω is considered part of a the calculation of the bound and must be done once for each bound calculation. In contrast, calculation of \mathbb{U}^Q , \mathbb{L}^Q , $\mathbb{U}^{\mathbb{U}^Q}$, $\mathbb{L}^{\mathbb{L}^Q}$ need only be done once per query series.

Note that early abandoning is used for the random order search, whereby the lower bound calculation is abandoned as soon as the cumulative calculation of the lower bound exceeds the distance to the nearest neighbor found so far. This is not possible for the sorted approach, as the lower bounds are computed before any of the full distances.

Figures 19 to 22 present the comparisons of LB_WEBB and LB_PETITJEAN

Algorithm 4 Experimental procedure for nearest neighbor search with sorted series

```

procedure SORTEDEXP(set of query series  $Q$ , set of training series  $\mathcal{T}$ , lower
bound  $\lambda$ )
  for  $Q \in \mathcal{Q}$  do
    if  $\lambda$  requires  $\mathbb{U}^Q$  and  $\mathbb{L}^Q$  then
      Calculate and save  $\mathbb{U}^Q$  and  $\mathbb{L}^Q$ 
    end if
    for  $T \in \mathcal{T}$  do
       $D[T] \leftarrow \lambda(Q, T)$ 
    end for
     $d \leftarrow \infty$ 
    for  $T \in \mathcal{T}$  in ascending order on  $D[T]$  until  $D[T] \geq d$  do
      if  $d = \infty$  then
         $d \leftarrow \text{DTW}(Q, T)$ 
         $b \leftarrow T$ 
      else
        if  $D[T] < b$  then
           $d' \leftarrow \text{DTW}(Q, T)$ 
          if  $d' < d$  then
             $b \leftarrow T$ 
             $d \leftarrow d'$ 
          end if
        end if
      end if
    end for
  end for
end procedure

```

against each of LB_KEOGH and LB_IMPROVED. These and all subsequent relative compute-time scatter plots plot the mean of ten runs together with error bars that show one standard deviation either side of the mean in each dimension. As the plots use log-log scale, these plots extend further to the left than right and further below than above the point. In most cases the error bars are not visible, as they do not extend beyond the dot centered on the mean.

LB_WEBB delivers faster nearest neighbor DTW search than either LB_KEOGH or LB_IMPROVED for the majority of datasets under both approaches. When the training examples are processed in random order, LB_WEBB delivers faster nearest neighbor DTW search than LB_KEOGH for 59 out of 60 datasets. The greatest difference is for the FordB dataset for which LB_KEOGH takes on average 8 minutes and 4 seconds and LB_WEBB takes 1 minute and 12 seconds. When the training series are sorted on the lower bounds, LB_WEBB is faster 52 times and LB_KEOGH 8. The greatest difference is again for the FordB dataset for which LB_KEOGH takes on average 6 minutes and 54 seconds compared with 42 seconds for LB_WEBB.

LB_PETITJEAN supports faster DTW nearest neighbor search than

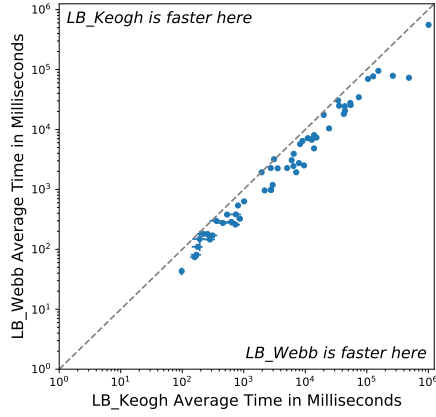


Figure 19: Relative compute time for nearest neighbor search in random order using LB_WEBB and LB_KEOGH.

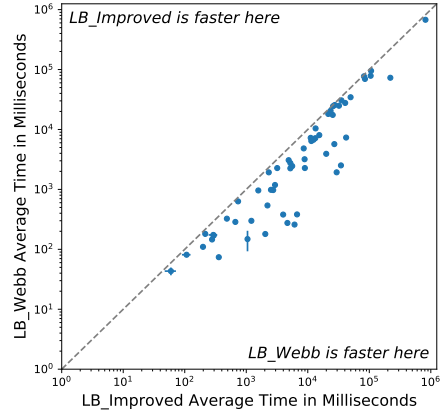


Figure 20: Relative compute time for nearest neighbor search in random order using LB_WEBB and LB_IMPROVED.

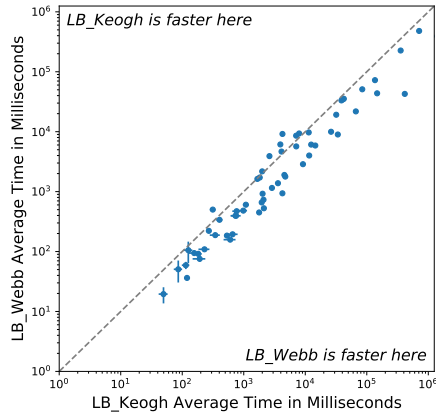


Figure 21: Relative compute time for nearest neighbor search in sorted order using LB_WEBB and LB_KEOGH.

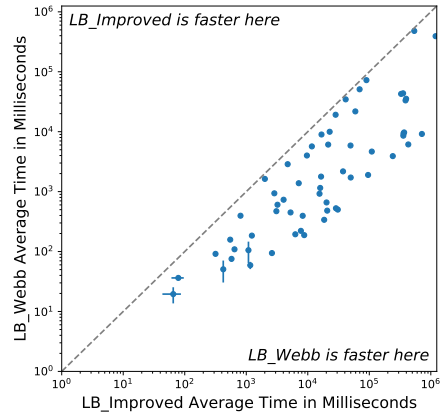


Figure 22: Relative compute time for nearest neighbor search in sorted order using LB_WEBB and LB_IMPROVED.

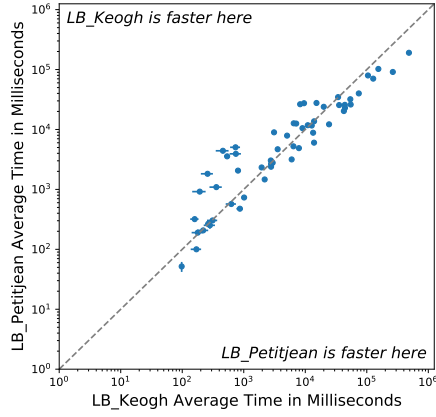


Figure 23: Relative compute time for nearest neighbor search in random order using LB_PETITJEAN and LB_KEOGH.

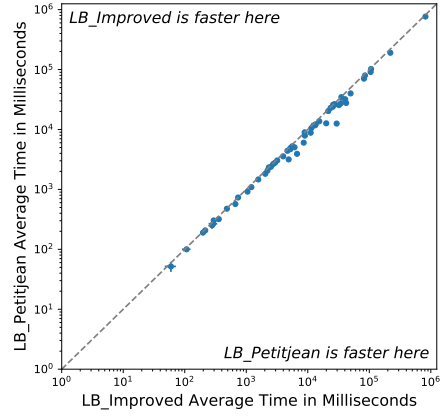


Figure 24: Relative compute time for nearest neighbor search in random order using LB_PETITJEAN and LB_IMPROVED.

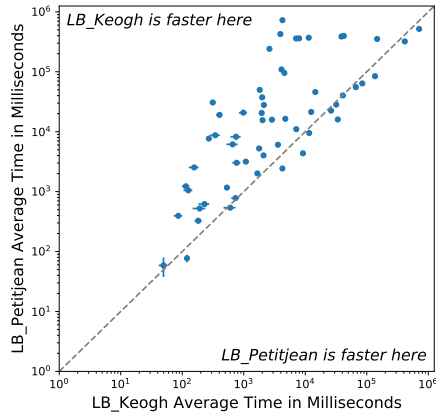


Figure 25: Relative compute time for nearest neighbor search in sorted order using LB_PETITJEAN and LB_KEOGH.

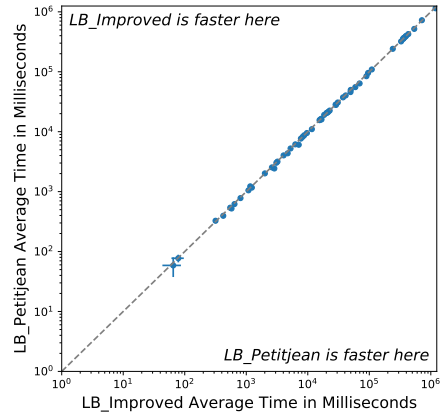


Figure 26: Relative compute time for nearest neighbor search in sorted order using LB_PETITJEAN and LB_IMPROVED.

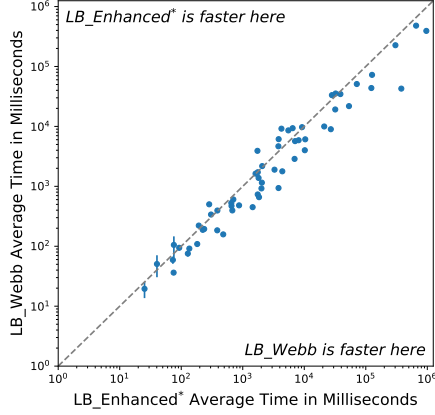


Figure 27: Relative compute time for nearest neighbor search in sorted order using LB_WEBB and LB_ENHANCED with the best performing value of k .

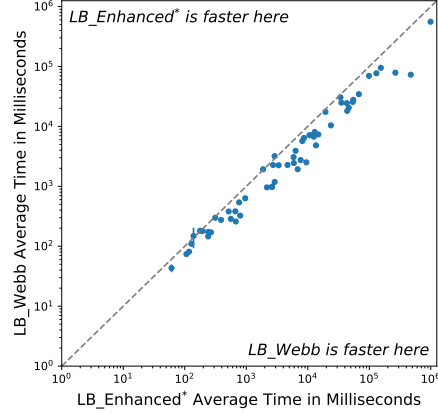


Figure 28: Relative compute time for nearest neighbor search in random order using LB_WEBB and LB_ENHANCED with the best performing value of k .

LB_IMPROVED for the majority of datasets when the series are not sorted, due to its improved tightness and similar compute time. When the series are sorted, however, the slight increase in compute time sometimes outweighs the increase in tightness. It tends to be more efficient than LB_KEOGH when the datasets are unsorted, because its more expensive computations can be abandoned as soon as the lower bound is sufficiently tight to allow a candidate to be pruned. It tends to be less efficient than LB_KEOGH when the candidates are sorted on the bound, as the bound must be calculated in full for every candidate.

We next compare LB_WEBB to LB_ENHANCED on the 60 datasets with optimal window size greater than zero. We first look at the case where the datasets are first sorted by lower bound. As the performance of LB_ENHANCED varies with parameter k , we test all values of k up to 16, at which point LB_ENHANCED is clearly beyond its optimal setting. The total average (again over ten runs) compute time for LB_WEBB is 26 minutes. LB_ENHANCED using the fastest setting of k for each dataset requires on average 49 minutes to obtain the same results. Note that this assessment does not take account of the issue of how the optimal value for k might be predetermined. Figure 27 shows the scatter plot of times per dataset for LB_WEBB relative to the best performance of LB_ENHANCED for any setting of k .

The results for processing the datasets in random order are presented in Figure 28. LB_WEBB is faster for 56 datasets and slower for 4. LB_WEBB requires on average 54 minutes to classify all 60 test sets while LB_ENHANCED with optimal k requires 1 hour and 50 minutes.

For both these tasks, LB_WEBB is faster and unlike LB_ENHANCED does not require any parameter tuning.

6.3. Classification times with varying window sizes

To explore how the bounds interact with varying window sizes, and to assess whether the advantage to LB_WEBB is specific to the 60 datasets whose optimal

Comparison	win/loss	Total time ratio
LB_WEBB vs LB_KEOGH	62 / 23	0:09:13/0:24:39= 0.37
LB_WEBB vs LB_IMPROVED	85 / 0	0:09:13/3:32:25= 0.04
LB_WEBB vs LB_PETITJEAN	85 / 0	0:09:13/3:32:05= 0.04
LB_WEBB vs LB_ENHANCED*	30 / 55	0:09:13/0:22:00= 0.42
LB_PETITJEAN vs LB_KEOGH	4 / 81	3:32:05/0:24:39= 8.60
LB_PETITJEAN vs LB_IMPROVED	56 / 29	3:32:05/3:32:25= 1.00
LB_PETITJEAN vs LB_WEBB	0 / 85	3:32:05/0:09:13=22.97
LB_PETITJEAN vs LB_ENHANCED*	4 / 81	3:32:05/0:22:00= 9.64

Table 1: Results on all UCR datasets, $w = 0.01 \cdot \ell$

Comparison	win/loss	Total time ratio
LB_WEBB vs LB_KEOGH	84 / 1	1:21:45/2:58:00=0.46
LB_WEBB vs LB_IMPROVED	85 / 0	1:21:45/4:53:11=0.28
LB_WEBB vs LB_PETITJEAN	85 / 0	1:21:45/4:43:24=0.29
LB_WEBB vs LB_ENHANCED*	79 / 6	1:21:45/2:06:49=0.64
LB_PETITJEAN vs LB_KEOGH	22 / 63	4:43:24/2:58:00=1.59
LB_PETITJEAN vs LB_IMPROVED	66 / 19	4:43:24/4:53:11=0.97
LB_PETITJEAN vs LB_WEBB	0 / 85	4:43:24/1:21:45=3.47
LB_PETITJEAN vs LB_ENHANCED*	11 / 74	4:43:24/2:06:49=2.23

Table 2: Results on all UCR datasets, $w = 0.10 \cdot \ell$

window sizes are greater than zero, we here assess classification time when the training data are sorted on the respective lower bound and the window size is a specified percentage of series length. We use three window sizes, 1% (Table 1), 10% (Table 2) and 20% (Table 3). In each case we round fractional values up in order to avoid windows of size zero. In each of the three tables of results, for each pairwise comparison, we present first a win/loss summary and then the total time taken, on average, in hours, minutes and seconds to classify the entire 85 test sets in the repository, followed by the ratio of the two times. The win/loss summary states the number of datasets for which the first algorithm required less time to classify the test set (wins) and the number for which the second algorithm required less time. There are no draws. Thus, when the window size is 1% of the total time series length (Table 1), LB_WEBB requires less time than LB_KEOGH on 62 datasets and more on 23 and requires 9 minutes and 13 seconds to classify the entire repository, which is just 37% of the 24 minutes and 39 seconds required by LB_KEOGH.

Comparison	win/loss	Total time ratio
LB_WEBB vs LB_KEOGH	85 / 0	3:23:09/5:25:55=0.62
LB_WEBB vs LB_IMPROVED	85 / 0	3:23:09/7:04:20=0.48
LB_WEBB vs LB_PETITJEAN	85 / 0	3:23:09/6:45:17=0.50
LB_WEBB vs LB_ENHANCED*	76 / 9	3:23:09/3:51:42=0.88
LB_PETITJEAN vs LB_KEOGH	29 / 56	6:45:17/5:25:55=1.24
LB_PETITJEAN vs LB_IMPROVED	76 / 9	6:45:17/7:04:20=0.96
LB_PETITJEAN vs LB_WEBB	0 / 85	6:45:17/3:23:09=1.99
LB_PETITJEAN vs LB_ENHANCED*	15 / 70	6:45:17/3:51:42=1.75

Table 3: Results on all UCR datasets, $w = 0.20 \cdot \ell$

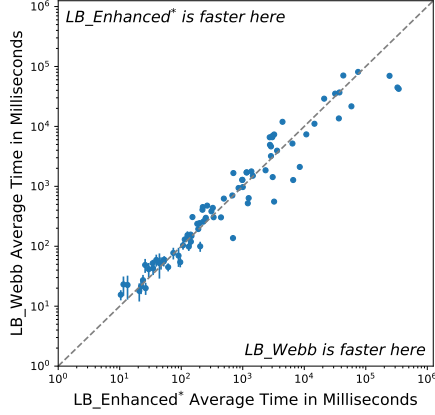


Figure 29: Relative compute time for nearest neighbor search in sorted order with window size set to 1% of series length. LB.WEBB vs LB.ENHANCED with the most effective value of k for each dataset. LB.WEBB is faster for 30 datasets and slower for 55. LB.WEBB requires on average under 9 minutes to classify all 85 test sets while LB.ENHANCED requires 22 minutes.

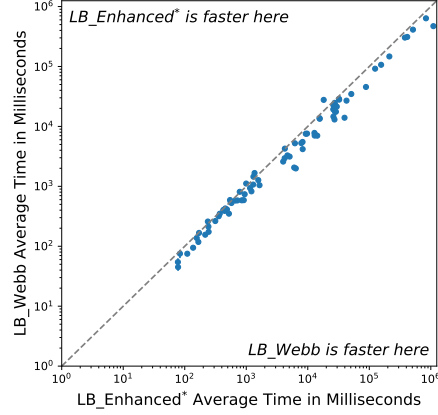


Figure 30: Relative compute time for nearest neighbor search in sorted order with window size set to 20% of series length. LB.WEBB vs LB.ENHANCED with the most effective value of k for each dataset. LB.WEBB is faster for 76 datasets and slower for 9. LB.WEBB requires on average 3 hours and 23 minutes to classify all 85 test sets while LB.ENHANCED requires 3 hours and 51 minutes.

As the window size increases from 1 to 10 to 20%, LB.WEBB consistently provides an advantage relative to LB.KEOGH, but the magnitude of that advantage decreases. Thus, at a window size of 20% of total series length, LB.WEBB is faster for all datasets. At this window size, LB.WEBB requires 3 hours and 23 minutes on average to classify all 85 datasets in the repository while LB.KEOGH requires 5 hours and 25 minutes.

LB.WEBB is faster than LB.IMPROVED for all datasets at all three window sizes. The relative magnitude of the improvement decreases as window size increases. Nonetheless, LB.WEBB requires only 3 hours and 23 minutes to classify all 85 datasets at a window size of 20% of total series length compared to 7 hours and 4 minutes for LB.IMPROVED.

LB.WEBB is faster than LB.ENHANCED at the best setting for k for only 30 out of the 85 datasets when the window size is set to 1% of series length. Nonetheless it requires less than half the time to classify the full repository. As illustrated in Figure 29, this is due to the losses being for datasets that require less computation and the wins being predominantly for datasets that require more. As the window size increases, LB.WEBB wins more often relative to LB.ENHANCED with the optimal setting of k , but the magnitudes of the wins and losses shrink, as illustrated in Figure 30.

LB.PETITJEAN delivers faster nearest neighbor search than LB.IMPROVED for the majority of datasets at all window sizes. However, the magnitudes of the wins and losses are extremely small. Neither LB.PETITJEAN nor LB.IMPROVED is competitive with LB.KEOGH, LB.ENHANCED or LB.WEBB on these tasks. This is because presorting does not offer any chance to early abandon lower

bound computation. As a result, these bounds are often computed to greater precision than required by the task. As shown in Figure 23, LB_PETITJEAN is more likely to excel in contexts where early abandon can be deployed.

7. On the effect of the left and right paths

In this section we investigate the role of the left and right paths that are incorporated in the two new bounds. To this end we compare LB_WEBB to a variant, LB_WEBB_NOLR without the left and right paths and a variant, LB_WEBB_ENHANCED, that replaces the left and right paths with left and right bands of the form used by LB_ENHANCED, presented in Section 5.2.

$$\begin{aligned} \text{LB_WEBB_NOLR}_w(A, B) = & \sum_{i=1}^{\ell} \begin{cases} \delta(A_i, \mathbb{U}_i^B) & \text{if } A_i > \mathbb{U}_i^B \\ \delta(A_i, \mathbb{L}_i^B) & \text{if } A_i < \mathbb{L}_i^B \\ 0 & \text{otherwise} \end{cases} \\ & + \sum_{i=1}^{\ell} \begin{cases} \delta(B_i, \mathbb{U}_i^A) & \text{if } F\uparrow(i) \wedge B_i > \mathbb{U}_i^A \\ \delta(B_i, \mathbb{L}_i^A) & \text{if } F\downarrow(i) \wedge B_i < \mathbb{L}_i^A \\ \delta(B_i, \mathbb{U}_i^A) - \delta(\mathbb{U}_i^B, \mathbb{U}_i^A) & \text{if } \neg F\uparrow(i) \wedge B_i > \mathbb{U}_i^B > \mathbb{U}_i^A \\ \delta(B_i, \mathbb{L}_i^A) - \delta(\mathbb{L}_i^B, \mathbb{L}_i^A) & \text{if } \neg F\downarrow(i) \wedge B_i < \mathbb{L}_i^B < \mathbb{L}_i^A \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

Figure 31 shows the relative tightness of LB_WEBB and LB_WEBB_NOLR using the optimal window size on all 60 datasets for which the optimal window size is greater than zero. LB_WEBB provides a tighter lower bound for every dataset except wafer, for which its tightness is 0.96891 versus 0.96904 a difference of just 0.00007. For many datasets the difference is small, but for a few datasets, where there is considerable variation in the start and end of the series, the difference is substantial. The largest difference is for FacesUCR for which LB_WEBB has tightness of 0.4639 relative to 0.2839 for LB_WEBB_NOLR. The average difference between the tightness of the two variants is 0.0124.

Figure 32 shows the relative tightness of LB_WEBB and LB_WEBB_ENHANCED³ using the optimal window size on all 60 datasets for which the optimal window size is greater than zero. LB_WEBB provides a tighter lower bound for every dataset. However, the difference is always small. The largest difference is for ECG5000 for which LB_WEBB has tightness of 0.8845 relative to 0.8724 for LB_WEBB_ENHANCED³. The average difference between the tightness of the two variants is 0.0008.

Figure 33 shows the relative time with LB_WEBB and LB_WEBB_NOLR for nearest neighbor search using the optimal window size on all 60 datasets for which the optimal window size is greater than zero. LB_WEBB is faster for all but 6 datasets. However, the relative differences are mainly small. The biggest difference is for ElectricDevices for which LB_WEBB requires 5 minutes and 39 seconds and LB_WEBB_NOLR requires 6 minutes and 17 seconds. On average

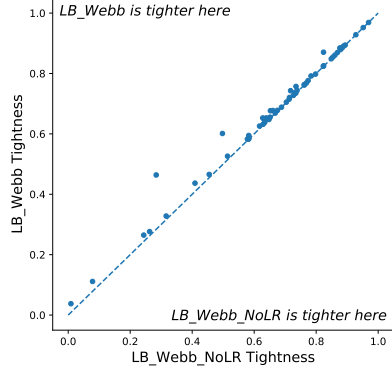


Figure 31: Relative tightness of LB_WEBB and LB_WEBB_NoLR

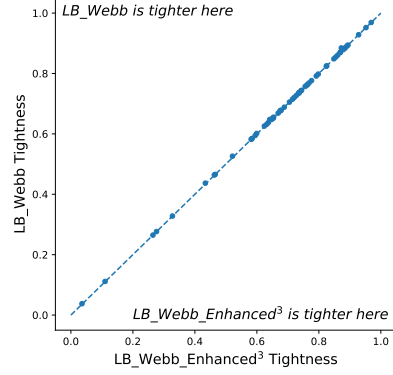


Figure 32: Relative tightness of LB_WEBB and LB_WEBB_Enhanced³

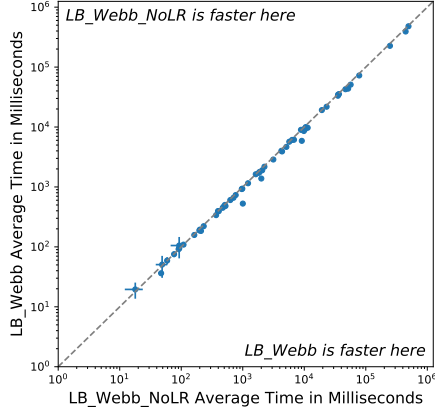


Figure 33: Relative compute time for nearest neighbor search in sorted order using LB_WEBB and LB_WEBB_NoLR. LB_WEBB is faster for 54 datasets and slower for 6. LB_WEBB requires on average 35 minutes and 21 seconds to classify all 60 test sets while LB_WEBB_NoLR requires 37 minutes and 20 seconds.

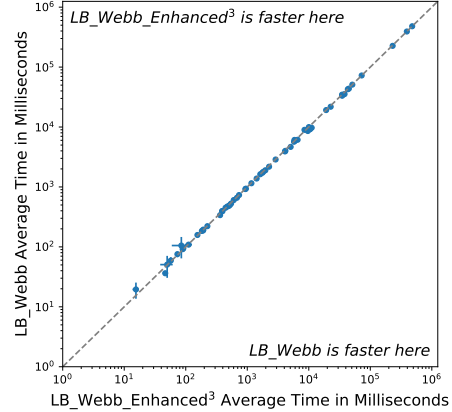


Figure 34: Relative compute time for nearest neighbor search in sorted order using LB_WEBB and LB_WEBB_Enhanced³. LB_WEBB is faster for 41 datasets and slower for 19. LB_WEBB requires on average 35 minutes and 21 seconds to classify all 60 test sets while LB_WEBB_Enhanced³ requires 35 minutes and 16 seconds.

LB_WEBB requires 35 minutes and 21 seconds to classify the entire 60 datasets while LB_WEBB_NO_LR requires 37 minutes and 20 seconds.

Figure 34 shows the relative time with LB_WEBB and LB_WEBB_ENHANCED³ for nearest neighbor search using the optimal window size on all 60 datasets for which the optimal window size is greater than zero. LB_WEBB is faster for all but 19 datasets. However, the relative differences are all small. The biggest difference is for UWaveGestureLibraryAll for which LB_WEBB requires 57 seconds and LB_WEBB_ENHANCED³ requires 1 minute and 4 seconds. On average LB_WEBB requires 35 minutes and 21 seconds to classify the entire 60 datasets while LB_WEBB_ENHANCED³ requires 35 minutes and 16 seconds.

In summary, the addition of the left and right paths to LB_WEBB almost invariably increases tightness. When there is substantial variation in the beginnings and endings of the series it can increase tightness substantially. It always increases tightness relative to using left and right bands, but in practice this increase in tightness appears to have little impact. These results suggest that in some applications it might be advantageous to use LB_WEBB_ENHANCED rather than LB_WEBB, if an appropriate value for k can be determined.

8. Conclusions

We have derived four new DTW lower bounds. To the best of our knowledge, LB_PETITJEAN is the tightest bound that has linear time complexity with respect to series length and is invariant to window size. LB_WEBB shares the same complexity, but provides a trade-off between efficiency and tightness that is more effective in many applications.

Both these bounds lend themselves to early abandoning. LB_PETITJEAN is likely to be most useful in contexts where this can be deployed, such as in a form of nearest neighbor search where the bound is calculated immediately before calculating DTW and thus can be abandoned if a closer candidate has already been encountered.

Both bounds also lend themselves to cascading. This is a process by which successive bounds providing successive trade-offs between compute time and tightness are employed in succession. For example, Rakthanmanon and Keogh [19] employ first the constant time LB_KIM [12], followed by LB_KEOGH, followed by a second evaluation of LB_KEOGH with the order of the two series reversed. Reversing the order of the two series in LB_KEOGH will obtain a tighter bound than applying LB_KEOGH in the original order in approximately 50% of cases, as the order of the series affects the bound, but neither order is a priori superior. Both LB_PETITJEAN and LB_WEBB can be deployed in a similar manner, by first computing the constant time left and right paths, then computing the bridging LB_KEOGH, before finally computing the additional final pass. This cascade provides intermediate lower bounds of successive strength that build upon one another, using the value calculated for the looser bound as a starting point for the tighter one and ending with a bound that is

likely to be substantially tighter than the best of LB_KEOGH under both orders. This is a promising direction for further research.

LB_WEBB_ENHANCED is a parameterized variant of LB_WEBB that employs the left and right bands of LB_ENHANCED in place of LB_WEBB’s left and right paths. This variant, with suitably large values of parameter k , may be useful in contexts where bounds based on distance to the envelope are less effective, such as when window sizes are large.

LB_PETITJEAN and LB_WEBB require that $\forall_{x,y \in \mathcal{R}, \mathcal{R}: A_i \leq x \leq y \leq B_j \vee A_i \geq x \geq y \geq B_j} \delta(A_i, B_j) \geq \delta(A_i, y) + \delta(B_j, x) - \delta(x, y)$, a condition satisfied by the two common pairwise distance measures, $\delta(A_i, B_j) = |A_i - B_j|$ and $\delta(A_i, B_j) = (A_i - B_j)^2$. A further variant, LB_WEBB*, supports faster computation of LB_WEBB when $\delta(A_i, B_j) = |A_i - B_j|$, and provides a tight lower bound for DTW so long as $\delta(A_i, B_j)$ increases monotonically with $|A_i - B_j|$, the same class of pairwise distance functions as for which LB_KEOGH, LB_IMPROVED and LB_ENHANCED are DTW lower bounds.

LB_WEBB has similar tightness to LB_IMPROVED, but requires substantially less computation. Our experiments show that it provides a highly effective trade-off between speed and tightness in a wide variety of contexts.

9. Acknowledgment

This research has been supported by the Australian Research Council under award DP210100072. The authors would like to also thank Prof Eamonn Keogh and all the contributors to the UCR time series classification archive.

References

- [1] H. Sakoe, S. Chiba, A dynamic programming approach to continuous speech recognition, in: International Congress on Acoustics, Vol. 3, 1971, pp. 65–69.
- [2] H. Cheng, Z. Dai, Z. Liu, Y. Zhao, An image-to-class dynamic time warping approach for both 3d static and trajectory hand gesture recognition, Pattern Recognition 55 (2016) 137–147.
- [3] M. Okawa, Online signature verification using single-template matching with time-series averaging and gradient boosting, Pattern Recognition 102 (2020) 107227.
- [4] Z. Yaseen, A. Verroust-Blondet, A. Nasri, Shape matching by part alignment using extended chordal axis transform, Pattern Recognition 57 (2016) 115–135.
- [5] G. Singh, D. Bansal, S. Sofat, N. Aggarwal, Smart patrolling: An efficient road surface monitoring using smartphone sensors and crowdsourcing, Pervasive and Mobile Computing 40 (2017) 71–88.

- [6] Y. Cao, N. Rakhilin, P. H. Gordon, X. Shen, E. C. Kan, A real-time spike classification method based on dynamic time warping for extracellular enteric neural recording with large waveform variability, *Journal of Neuroscience Methods* 261 (2016) 97–109.
- [7] R. Varatharajan, G. Manogaran, M. K. Priyan, R. Sundarasekar, Wearable sensor devices for early detection of alzheimer disease using dynamic time warping algorithm, *Cluster Computing* 21 (1) (2018) 681–690.
- [8] E. Chávez, G. Navarro, R. Baeza-Yates, J. L. Marroquín, Searching in metric spaces, *ACM Computing Surveys* 33 (3) (2001) 273–321.
- [9] J. Z. Lai, Y.-C. Liaw, J. Liu, Fast k-nearest-neighbor search based on projection and triangular inequality, *Pattern Recognition* 40 (2) (2007) 351–359.
- [10] C. A. Ratanamahatana, E. Keogh, Three myths about dynamic time warping data mining, in: *Proceedings of the 2005 SIAM International Conference on Data Mining*, SIAM, 2005, pp. 506–510.
- [11] E. Keogh, C. A. Ratanamahatana, Exact indexing of dynamic time warping, *Knowledge and Information Systems* 7 (3) (2005) 358–386.
- [12] S.-W. Kim, S. Park, W. W. Chu, An index-based approach for similarity search supporting time warping in large sequence databases, in: *17th International Conference on Data Engineering.*, IEEE, 2001, pp. 607–614.
- [13] D. Lemire, Faster retrieval with a two-pass dynamic-time-warping lower bound, *Pattern Recognition* 42 (9) (2009) 2169–2180.
- [14] Y. Shen, Y. Chen, E. Keogh, H. Jin, Accelerating time series searching with large uniform scaling, in: *Proceedings of the 2018 SIAM International Conference on Data Mining*, SIAM, 2018, pp. 234–242.
- [15] B.-K. Yi, H. Jagadish, C. Faloutsos, Efficient retrieval of similar time sequences under time warping, in: *Data Engineering, 1998. Proceedings., 14th International Conference on*, IEEE, 1998, pp. 201–208.
- [16] C. W. Tan, F. Petitjean, G. I. Webb, Elastic bands across the path: A new framework and methods to lower bound dtw, in: *Proceedings of the 2019 SIAM International Conference on Data Mining*, 2019, pp. 522–530.
- [17] A. Bagnall, J. Lines, A. Bostrom, J. Large, E. Keogh, The great time series classification bake off: a review and experimental evaluation of recent algorithmic advances, *Data Mining and Knowledge Discovery* 31 (3) (2017) 606–660.
- [18] H. A. Dau, A. Bagnall, K. Kamgar, C.-C. M. Yeh, Y. Zhu, S. Gharghabi, C. A. Ratanamahatana, E. Keogh, The UCR time series archive, *IEEE/CAA Journal of Automatica Sinica* 6 (6) (2019) 1293–1305.

- [19] T. Rakthanmanon, E. Keogh, Data mining a trillion time series subsequences under dynamic time warping, in: Twenty-Third International Joint Conference on Artificial Intelligence, 2013, pp. 3047–3051.