

Semantic Discord: Finding Unusual Local Patterns for Time Series

Li Zhang* Yifeng Gao* Jessica Lin*

Abstract

Finding anomalous subsequence in a long time series is a very important but difficult problem. Existing state-of-the-art methods have been focusing on searching for the subsequence that is the most dissimilar to the rest of the subsequences; however, they do not take into account the background patterns that contain the anomalous candidates. As a result, such approaches are likely to miss local anomalies. We introduce a new definition named *semantic discord*, which incorporates the context information from larger subsequences containing the anomaly candidates. We propose an efficient algorithm with a derived lower bound that is up to 3 orders of magnitude faster than the brute force algorithm in real world data. We demonstrate that our method significantly outperforms the state-of-the-art methods in locating anomalies by extensive experiments. We further explain the interpretability of semantic discord.

1 INTRODUCTION

Time series anomalous sequence detection is an important problem and has wide application in different domains such as medical care [8], fraud detection [3, 7] and Internet of Things (IoT) [18], as the anomalies identify unexpected and unusual items or events that are different from normal patterns. Particularly, the task involves identifying a time series subsequence that is the least similar to all other subsequences in a long time series. Such anomalous subsequence has also been referred as *discord*, which is defined as the subsequence of a given length that has the largest z -normalized Euclidean Distance to its closest match [11]. The length of the subsequence is determined by some prior knowledge on the time series.

Unfortunately, the approaches following the current discord definition, in which a subsequence is normalized using all the points in the subsequence, would highly likely miss true anomalies in at least some cases. We demonstrate the problem by a simple example generated from a cycle in dishwasher power consumption. Consider a time series shown in Figure 1 (a) with six segments labeled 1 to 6 respectively. Segments 1, 5 and 6 have the same shape, except that Segment 5 contains

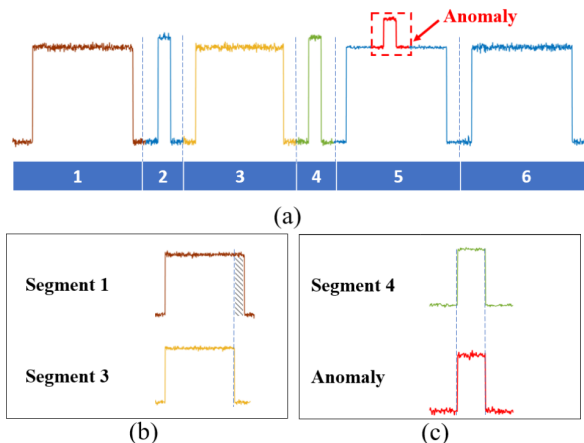


Figure 1: An illustration of a local anomalous subsequence in time series. (a) Time series with labeled segments 1 to 6. (b) Segments 1 and 3 after normalization have large distance. (c) Segment 4 and local anomaly have close to zero distance after normalization.

a small bump, which is a local anomaly. Segment 3 has a slightly different length from Segments 1 and 6. The local anomaly in Segment 5 has identical shape as Segments 2 and 4 after z -normalization.

Clearly, it is impossible to find the local anomaly with its true length via the classic discord definition, as the local anomaly and Segment 4 (or Segment 2) would be indistinguishable and resulting in a distance of zero after z -normalization. (Figure 1(c)). Moreover, if we expand the input length to some length close to the size of Segment 5, Segment 3 would be identified as the anomalous subsequence because of slight length difference to its closest match, Segment 1, as shown in Figure 1(b). Furthermore, even if some longer subsequence containing the bump could be identified given some arbitrary length, the length used to find the anomaly would be so far off from the ground truth that it loses interpretability.

To resolve this issue, in this paper, we propose a new anomaly definition, **semantic discord**. Our new discord definition incorporates the background subsequence containing the anomaly candidate, namely, a context subsequence, which provides local semantic information for detecting anomalous subsequence. Under

*Department of Computer Science, George Mason University, {lzhang18, ygao12, jessica}@gmu.edu

the new definition, we search over all possible context subsequences and use the best ones to normalize the anomalous candidates respectively (as opposed to normalizing using the candidates themselves).

One may consider finding context a trivial problem, but it is not. Bringing in context subsequences forces the search space to grow quadratic in the length of context subsequence. This is because for one anomaly subsequence candidate of length l , if the context subsequence is length L ($L > l$), there are $L - l + 1$ possible contexts to choose from. To compare two candidates, we need to search over $O((L - l + 1)^2)$ possible pairs of context subsequences. To address this issue, in this paper, we derive an effective novel lower bound to reduce the extra cost for computing the quadratic number of distances.

The contributions of this paper are summarized as follows:

- We introduce a new definition named *semantic discord*, which incorporates the context information from larger subsequences containing the anomaly candidates.
- We propose an efficient exact algorithm that is up to 3 orders of magnitude faster than the brute force algorithm in real world data.
- We demonstrate that our method significantly outperforms the state-of-the-art methods in locating anomalies by extensive experiments. We further explain the interpretability of semantic discord.

The paper is organized as follows. In Section 2, we review the related work and provide the background information. In Section 3, we provide the definitions on time series and formally define our problem. In Section 4, we provide details on our problem using smart brute force, followed by the derivation of an effective lower bound, as well as the proposed algorithm. In Section 5, we perform extensive experiments to show that our method outperforms the state-of-the-art methods. We also test the scalability with synthetic and real datasets.

2 BACKGROUND AND RELATED WORK

We describe the state-of-the-art methods in time series anomaly detection in this section. HOTSAX [11] is one of the first algorithms to investigate on the detection of anomalous time series subsequence. The authors defined the *discord* to be the subsequence that has the largest 1-nearest-neighbor distance in a single long series, and proposed an efficient algorithm by re-ordering the candidate subsequences. Recently, Matrix Profile based techniques [24, 26] have provided acceleration on computing the 1-nearest-neighbor distances.

The results can be used to accelerate computation of HOTSAX. Lin et al. [15] defined anomalies as subsequences with zero or a small number of reverse nearest neighbors. Senin et al. [22] exploited grammar induction to generate rule density curve which can be used to identify anomalous subsequences corresponding to rare grammar rules. None of the existing works mentioned above considered any context information; in fact, all the existing definitions of time series anomalies would favor global anomalies and miss local unusual patterns.

There is extensive work on finding local, point-based outliers such as [1, 3, 9, 13, 14] based on local density or looking for a subspace in high dimension. Unfortunately, representing time series as high dimensional points will not solve the problem as time series subsequence anomalies are not in the sparsest region [11]. These methods are not designed for high-dimensional overlapping subsequences.

Several other existing works [5, 6, 23] on finding anomalous subsequences are worth mentioning, though the problem settings in their work are not exactly the same as ours. Wang et al. [23] developed an anomaly detection algorithm on aligned time series instances on manufacturing data. Fereman et al. [5] developed features on discrete time log covariates to guide an isolation forest on anomalous sequence discovery. In addition, with the rising interest in deep learning, many deep-learning-based unsupervised anomaly detection algorithms [2, 16, 19, 25] have been proposed. However, training such neural networks is very time-consuming, and the results lack direct interpretability comparing with traditional search-based methods [5]. Since our method does not rely on alignments or covariates, and it has the desirable property of interpretability, we do not compare with the above methods.

3 DEFINITIONS

We begin with the fundamental definitions of time series.

DEFINITION 1. A *Time series* $\mathcal{T} = [t_1, t_2, \dots, t_n]$ is an ordered list of data points, where t_i is a finite real number and n is the length of time series \mathcal{T} .

DEFINITION 2. A *time series subsequence* $S_i^L = [t_i, t_{i+1}, \dots, t_{i+L-1}]$ is a contiguous set of points in time series \mathcal{T} starting from position i with length L . Typically $L \ll n$, and $1 \leq i \leq n - L + 1$.

Subsequences can be extracted from time series \mathcal{T} by sliding a fixed-length window through the time series.

Given two subsequences of the same length, the Euclidean Distance is often used to measure their differences. To achieve scale and offset invariance, each subsequence must be properly normalized before the actual

distance computation. The normalization step in general is very critical, as noted in previous work — “*without normalization time series similarity has essentially no meaning. More concretely, very small changes in offset rapidly dwarf any information about the shape of the two time series in question.*” [10]

For discord computation, previous work [11] uses z-normalized Euclidean Distance in order to make scale-invariant subsequences comparison prior to the distance computation. We describe the z-normalized Euclidean Distance as follows:

DEFINITION 3. A **z-normalized Euclidean Distance** $d_{ED}(p, q)$ of subsequences S_p^l, S_q^l of length l is computed as

$$\sqrt{\sum_{m=1}^l \left(\frac{t_{p+m-1} - \mu_p}{\sigma_p} - \frac{t_{q+m-1} - \mu_q}{\sigma_q} \right)^2},$$

where μ_p, σ_p and μ_q, σ_q are the means and standard deviations of subsequences S_p^l and S_q^l respectively.

Under the current definition of z-normalization, the local anomalous subsequence would be normalized by its own mean and standard deviation. As a result, the local anomalies may be considered similar to some very different patterns. As shown in Figure 1, the bump of Segment 5 can be matched to Segment 2 or Segment 4, and will not be identified as discord. To overcome this problem, we propose a new definition of discord under a novel distance named *Optimal Context-Aware Distance*, to measure the dissimilarity of subsequences with contextual information. Instead of normalizing by a subsequence itself, we normalize it by using a longer context subsequence containing it, as the context subsequence would reflect the background information about the anomalous subsequence candidate. An ideal candidate of context subsequence can be in the length of some existing patterns, for example, Segment 5 in our previous example.

To be precise, we introduce the definitions related to our proposed Context-Aware Distance as follows.

DEFINITION 4. Given two subsequences S_p^l, S_q^l of length l and two context subsequences S_i^L, S_j^L of length L , where $p - L + l \leq i \leq p$ and $q - L + l \leq j \leq q$. The **Context-Aware Euclidean Distance** $d_{i,j}(p, q)$ between S_p^l and S_q^l under their contexts S_i^L and S_j^L is defined by

$$\sqrt{\sum_{m=1}^l \left(\frac{t_{p+m-1} - \mu_i}{\sigma_i} - \frac{t_{q+m-1} - \mu_j}{\sigma_j} \right)^2},$$

where μ_i, σ_i and μ_j, σ_j are means and standard deviations of subsequences S_i^L and S_j^L respectively.

To differentiate, the candidate subsequence S_p^l, S_q^l are called **target** and **reference target** respectively. Their context subsequences S_i^L and S_j^L are called the **context** and **reference context** respectively.

To simplify the notation, we denote a subsequence by T if it is a (reference) target and a subsequence by C if it is a (reference) context. We omit the length for the brevity of computation that we will describe later. For example, a target S_p^l is denoted as T_p and a context S_i^L is denoted as C_i .

As previous work [4,11,12] noted, finding discord requires excluding self-matches, which refer to the subsequence itself or those that overlap with the subsequence. In our work, we only consider non-self matches of contexts and targets. We formally define non-self match as follows:

DEFINITION 5. Given a time series \mathcal{T} , a subsequence S_i with length L is considered a **non-self match** of another subsequence S_j of length L if $|i - j| > L$.

Clearly, if C_j is a non-self match of C_i , then T_q would be a non-self match of T_p . In other words, if two contexts are not overlapping, any targets within the context pair respectively will not be overlapping. Thus, we only need to enforce non-self match between context C_i and C_j to guarantee non-self matches between any pairs of their targets, respectively.

DEFINITION 6. Given two non-overlapping (non-self match) targets T_p and T_q of length l and a context length L , the **Optimal Context-Aware Euclidean Distance** between T_p and T_q is defined as

$$d_{opt}(p, q) = \min_{(i,j) \in \Omega_{p,q}} d_{i,j}(p, q),$$

where $\Omega_{p,q} = \{(i, j) | (p - L + l \leq i \leq p) \wedge (q - L + l \leq j \leq q) \wedge d_{ED}(i, j) < \epsilon\}$.

The intuition behind $\Omega_{p,q}$ is simple: given targets T_p, T_q , we search over all possible combinations of context pairs that are at least similar within some large constant threshold ϵ and overlap with targets T_p and T_q , respectively.

We are ready to define our proposed definition of semantic discord:

DEFINITION 7. Given a time series \mathcal{T} , the target T_p of length l is said to be the **semantic discord** of \mathcal{T} if T_p has the largest Context-Aware Euclidean Distance to its closest non-self match reference target.

By using the Context-Aware Euclidean Distance instead of the z-normalized Euclidean Distance, the semantic discord incorporates the background information from context subsequences containing the anomaly candidates, and hence better captures local anomalies in normal patterns.

Algorithm 1 Brute Force Algorithm for Computing $d_{opt}(p, q)$

```

1: Input: Time Series  $\mathcal{T}$ , Context length  $L$ , Target
   length  $l$ ,  $\Omega_{p,q}$ 
2:  $best\_so\_far = 0$ 
3: for  $p = 1$  to  $|\mathcal{T}| - l + 1$  do
4:    $nn\_dist = Inf$ 
   /* Compute minimum  $d_{opt}(p, q)$  for each sub-
   sequence  $T_q$  */
5:   for  $q = 1$  to  $|\mathcal{T}| - l + 1$  do
6:     if  $|p - q| \leq L$  then /* Skip self-matching */
7:       continue
8:     end if
   /* Compute  $d_{i,j}(p, q)$  */
9:     for  $i = p - L + l$  to  $p$  do
10:      for  $j = q - L + l$  to  $q$  do
11:         $d_{i,j}(p, q) = GetDist_{CA}(p, q, i, j, \Omega_{p,q})$ 
12:      end for
13:    end for
14:     $dist = \min_{i,j} d_{i,j}(p, q)$ 
15:    if  $dist < nn\_dist$  then
16:       $nn\_dist = dist$ 
17:       $nn\_match\_T = (p, q)$ 
18:       $nn\_match\_C = (i, j)$ 
19:    end if
20:  end for
21: end for
22: if  $nn\_dist > best\_so\_far$  then
23:    $best\_so\_far = nn\_dist$ 
24:    $best\_match\_T = nn\_match\_T$ 
25:    $best\_match\_C = nn\_match\_C$ 
26: end if
27: return  $best\_so\_far, best\_match\_T, best\_match\_C$ 

```

4 PROPOSED METHOD

In this section, we start with a smart brute force algorithm under our new semantic discord definition. We then derive our proposed lower bound on the Optimal Context-Aware Distance. Finally, we introduce our proposed algorithm for finding Semantic Discord.

4.1 Smart Brute Force A brute force algorithm for identifying semantic discord is incredibly costly. For every pair of target subsequences T_p, T_q of length l , there are $O(L^2)$ possible combination of context C_i and C_j of length L . By computing all pairwise distances using the Optimal Context-Aware Euclidean Distance, the complexity of the brute force algorithm would reach $O(n^2L^2l)$. This is intractable for even a time series of a few thousand points. The algorithm is shown in Algorithm 1.

It is possible to use the technique introduced in

[21, 24] to come up with a smart brute force algorithm to speed up the distance computation. By expressing the Context-Aware Distance $d_{i,j}(p, q)$ between T_p and T_q as a function of means, standard deviations and inner products, $d_{i,j}(p, q)$ can be computed via:

$$\begin{aligned}
d_{i,j}^2(p, q) &= \frac{l}{\sigma_i^2}(\sigma_p^2 + (\mu_p - \mu_i)^2) \\
&- \frac{2l}{\sigma_i\sigma_j}(\frac{QT_{p,q}}{l} - \mu_i\mu_q - \mu_j\mu_p + \mu_i\mu_j) \\
&+ \frac{l}{\sigma_j^2}(\sigma_q^2 + (\mu_q - \mu_j)^2),
\end{aligned} \tag{4.1}$$

where μ_i, σ_i and μ_j, σ_j are the means and standard deviations of C_i and C_j respectively. Similarly, μ_p, σ_p and μ_q, σ_q are the means and standard deviations of T_q and T_p respectively. $QT_{p,q} = \sum_{m=1}^l t_{p+m-1}t_{q+m-1}$ is the inner product between target T_p and T_q , where l is the length of target. We can easily pre-compute the means and standard deviations for all subsequences in $O(n)$ time, thus the major cost would be computing the inner product $QT_{p,q}$. Note $QT_{p,q}$ can be computed based on previous $QT_{p-1,q-1}$ in constant complexity [24, 26]:

$$QT_{p,q} = QT_{p-1,q-1} - t_{p-1}t_{q-1} + t_{p+l-1}t_{q+l-1}.$$

Therefore, the cost to compute $d_{i,j}(p, q)$ is constant and the overall complexity of the brute force algorithm is reduced to $O(n^2L^2)$. Although it is faster than the original brute force algorithm, it is still too time-consuming and intractable for our purpose.

4.2 Lower Bound In this subsection, we introduce an effective lower bound to reduce the number of times to compute the actual optimal Context-Aware Distance $d_{opt}(p, q)$. We would like to use the lower bound to determine the order of comparisons and hope to stop the search at the earliest time. We will explain this idea in detail in Section 4.3. Inspired by [17], instead of enumerating all possible combinations of context subsequences, we consider it as an optimization problem. To simplify the optimization process, we relax all the constraints and aim to find a pair of values μ' and σ' which minimizes the value of $d_{opt}(p, q)$. Thus, the lower bound obtained from μ' and σ' will be a global minimum value independent from data and smaller than the actual distance. The objective function is relaxed as shown as follows:

$$\begin{aligned}
d_{opt}(p, q) &= \min_{\Omega_{p,q}} d_{i,j}(p, q) \\
&= \min_{\mu_i, \mu_j, \sigma_i, \sigma_j} \sqrt{\sum_{m=1}^l \left(\frac{t_{p+m-1} - \mu_i}{\sigma_i} - \frac{t_{q+m-1} - \mu_j}{\sigma_j} \right)^2} \\
&\geq \min_{\mu', \sigma'} \left(\min_{\mu_j, \sigma_j} \sqrt{\sum_{m=1}^l \left(\frac{t_{p+m-1} - \mu'}{\sigma'} - \frac{t_{q+m-1} - \mu_j}{\sigma_j} \right)^2} \right) \\
&= \min_{\mu_j, \sigma_j} \left(\min_{\mu', \sigma'} D(\mu', \sigma', \mu_j, \sigma_j) \right) = LB(p, q),
\end{aligned} \tag{4.2}$$

where $LB(p, q)$ is the lower bound which we would like to achieve.

We split the lower bound computation to two steps to optimize the equation. We first optimize the inner part $D^* = \min_{\mu', \sigma'} D(\mu', \sigma', \mu_j, \sigma_j)$, and then solve the outer minimization problem $\min_{\mu_j, \sigma_j} D^*$ for $LB(p, q)$.

4.2.1 Solving the inner problem To optimize the inner part $D^* = \min_{\mu', \sigma'} D(\mu', \sigma', \mu_j, \sigma_j)$, we can simply get minimum value of $D(\mu', \sigma', \mu_j, \sigma_j)$ in terms of μ' and σ' by solving $\frac{\partial D}{\partial \mu'} = 0$ and $\frac{\partial D}{\partial \sigma'} = 0$. We optimize on D^2 instead of D for the simplicity of computation. Specifically,

$$(4.3) \quad D^2 = \sum_{m=1}^l \left(\frac{t_{p+m-1} - \mu'}{\sigma'} - \frac{t_{q+m-1} - \mu_j}{\sigma_j} \right)^2.$$

Starting by solving μ' , we have

$$(4.4) \quad \frac{\partial D^2}{\partial \mu'} = \sum_{m=1}^l \frac{-2}{\sigma'} \left(\frac{t_{p+m-1} - \mu'}{\sigma'} - \frac{t_{q+m-1} - \mu_j}{\sigma_j} \right) = 0$$

By solving equation (4.4), we can obtain μ^* , the optimal μ' value, as:

$$(4.5) \quad \mu^* = \mu_p + \frac{\sigma'}{\sigma_j} (\mu_j - \mu_q).$$

Substituting μ' with μ^* in equation (4.3), we get

$$(4.6) \quad \begin{aligned} D^2 &= \sum_{m=1}^l \left(\frac{t_{p+m-1} - \mu_p}{\sigma'} - \frac{\mu_j - \mu_q}{\sigma_j} - \frac{t_{m+q-1} - \mu_j}{\sigma_j} \right)^2 \\ &= \sum_{m=1}^l \left(\frac{t_{p+m-1} - \mu_p}{\sigma'} - \frac{t_{m+q-1} - \mu_q}{\sigma_j} \right)^2. \end{aligned}$$

Intuitively, μ^* adjusts T_p to the same offset of T_q under context C_j to minimize D^2 . From equation 4.3, we see that T_p and T_q after simplification, have the same mean value.

We next take derivative with regard to σ' on equation 4.6 to solve $\frac{\partial D(\mu', \sigma', \mu_j, \sigma_j)}{\partial \sigma'} = 0$:

$$(4.7) \quad \begin{aligned} \frac{\partial D^2}{\partial \sigma'} &= \frac{\partial}{\partial \sigma'} \sum_{m=1}^l \left(\frac{t_{p+m-1} - \mu_p}{\sigma'} - \frac{t_{m+q-1} - \mu_q}{\sigma_j} \right)^2 \\ &= \frac{-2l\sigma_p}{(\sigma')^2} \left[\frac{\sigma_p}{\sigma'} - \frac{\sigma_q}{\sigma_j} \delta_{p,q} \right] = 0, \end{aligned}$$

where $\delta_{p,q}$ is the correlation similarity between target subsequences T_p and T_q :

$$(4.8) \quad \delta_{p,q} = \frac{\sum_{m=1}^l (t_{p+m-1} t_{q+m-1}) - l\mu_p \mu_q}{l\sigma_p \sigma_q}$$

By solving equation (4.7), we get σ^* , the optimal value of σ' as:

$$\sigma^* = \begin{cases} \frac{\sigma_p \sigma_j}{\delta_{p,q} \sigma_q} & \delta_{p,q} > 0 \\ \infty & \delta_{p,q} \leq 0 \end{cases}$$

Now we are ready to compute the square of the lower bound in equation (4.6) by the optimal μ^* and σ^* ,

$$(4.9) \quad \begin{aligned} (D^*)^2 &= \overbrace{\frac{l\sigma_p^2}{(\sigma')^2}}^{D1} + \overbrace{\left(-\frac{2l}{\sigma' \sigma_j} (QT_{p,q} - \mu_p \mu_q) \right)}^{D2} + \overbrace{\frac{l\sigma_q^2}{\sigma_j^2}}^{D3} \\ &= D1 + D2 + D3, \end{aligned}$$

where

$$\begin{aligned} D1 &= \frac{l\sigma_p^2}{(\sigma')^2} = \frac{l\delta_{p,q}^2 \sigma_q^2}{\sigma_j^2}, \\ D2 &= -\frac{2l\delta_{p,q} \sigma_q}{\sigma_p \sigma_j^2} \left(\frac{QT_{p,q}}{l} - \mu_p \mu_q \right) = -\frac{2l\delta_{p,q}^2 \sigma_q^2}{\sigma_j^2}, \text{ and} \\ D3 &= \frac{l\sigma_q^2}{\sigma_j^2}. \end{aligned}$$

Summing up $D1$, $D2$ and $D3$, we have

$$(D^*)^2 = \frac{l\sigma_q^2(1 - \delta_{p,q}^2)}{\sigma_j^2}.$$

After taking the square root, we obtain optimal D^* ,

$$(4.10) \quad D^* = \frac{\sigma_q}{\sigma_j} \sqrt{l(1 - \delta_{p,q}^2)}.$$

Surprisingly, the optimal value of D^* is very simple – it only depends on the standard deviation of the reference context σ_j and the reference target σ_q , rather than the mean values μ_j or μ_q . This makes sense intuitively, as the ratio of standard deviations of target and its context reflects to some extent on the scale variant between the two targets under normalization of their respective contexts.

4.2.2 Solving the outer problem With the optimal value of D^* for equation (4.2), we are now left with the outer problem $\min_{\mu_j, \sigma_j} D^*(\mu_j, \sigma_j)$, which only depends on σ_j . We can obtain the minimum value by simply picking the maximum value of σ_j , that is $\frac{\sigma_q}{\max(\sigma_j)} \sqrt{l(1 - \delta_{p,q}^2)}$.

Moreover, as our Optimal Context-Aware Euclidean Distance is symmetric, i.e. $d_{opt}(p, q) = d_{opt}(q, p)$, by symmetry, it is effortless to obtain the other version of lower bound, $\frac{\sigma_p}{\max(\sigma_i)} \sqrt{l(1 - \delta_{q,p}^2)}$. To obtain a tighter lower bound, we choose the larger LB of the two versions. Therefore, our final lower bound $LB(p, q)$ is the following:

$$(4.11) \quad LB(p, q) = \begin{cases} \gamma_{i,j}(p, q) \sqrt{l(1 - \delta_{p,q}^2)} & \delta_{p,q} > 0, \\ \gamma_{i,j}(p, q) \sqrt{l} & \delta_{p,q} \leq 0. \end{cases}$$

Where $\delta_{p,q}$ is the correlation as shown in equation (4.8), and $\gamma_{i,j}(p, q) = \max\left(\frac{\sigma_q}{\max(\sigma_j)}, \frac{\sigma_p}{\max(\sigma_i)}\right)$. We can compute $\gamma_{i,j}(p, q)$ with moving standard deviation in $O(1)$ complexity. Thus, we are able to compute the lower bound $LB(p, q)$ in $O(1)$, which is much less costly than computing the actual Optimal Context-Aware Distance as the latter requires $O(L^2)$ complexity. In the next section, we will utilize the lower bound $LB(p, q)$ to prune the unnecessary distance computations.

Algorithm 2 Proposed Algorithm

```
1: Input: Time Series  $\mathcal{T}$ , Context Length  $L$ , Target  
   Length  $l$ ,  $\Omega_{p,q}$   
2: Output:  $best\_so\_far$ ,  $best\_match\_T$ ,  $best\_match\_C$   
3:  $best\_so\_far = 0$   
4: for  $p = 1$  to  $|\mathcal{T}| - l + 1$  do  
5:    $nn\_dist = Inf$   
   /* Compute Lower Bound based on equation  
   (4.11) for target subsequences in  $\mathcal{T}^*$  */  
6:    $LB = \text{ComputeAllLB}(T_p)$   
7:    $LB\_Sort, LB\_Index = \text{Sort}(LB)$   
8:   for  $k = 1$  to  $|\mathcal{T}| - l + 1$  do  
9:      $q = LB\_Index[k]$   
10:    if  $IsSelfMatch(p, q)$  then  
11:      continue /* Skip Self-Matching */  
12:    end if  
13:    if  $nn\_dist \leq LB\_Sort[k]$  then  
14:      break  
15:    end if  
16:    for  $i = p - L + l$  to  $p$  do  
17:      for  $j = q - L + l$  to  $q$  do  
18:         $d_{i,j}(p, q) = \text{GetDist}_{CA}(p, q, i, j, \Omega_{p,q})$   
19:      end for  
20:    end for  
21:     $dist = \min_{i,j} d_{i,j}(p, q)$   
22:    if  $dist < nn\_dist$  then  
23:       $nn\_dist = dist$   
24:       $nn\_match\_T = (p, q)$   
25:       $nn\_match\_C = (i, j)$   
26:    end if  
27:  end for  
28: end for  
29: if  $nn\_dist > best\_so\_far$  then  
30:    $best\_so\_far = nn\_dist$   
31:    $best\_match\_T = nn\_match\_T$   
32:    $best\_match\_C = nn\_match\_C$   
33: end if  
34: return  $best\_so\_far, best\_match\_T, best\_match\_C$ 
```

4.3 Proposed Pruning Algorithm Our algorithm is based on the proposed lower bound and Algorithm 1. For each target, we reorder all the reference targets according to their lower bound value in ascending order. If the target is compared with reference targets in this order, we can simply stop the search once the best-so-far distance is smaller than the current lower bound value, as the remaining unchecked actual distance must be greater than the best-so-far distance, hence significantly reduce the computation cost.

The proposed algorithm is shown in Algorithm 2. The largest closest match of the Optimal Context-Aware Distance is denoted as $best_so_far$ and is initialized to zero (Line 3). For every target subsequence T_p , the closest match distance of current target is initialized to infinity (Line 5).

We then compute $LB(p, q)$ and sort all the reference targets in ascending order (Lines 6-7) with the cost of $O(n \log(n))$ and T_q is checked in this order. Lines 8-28 are the inner loop which compares a target with all possible reference targets. Lines 10-12 check that T_p and T_q are from non-self match contexts to avoid trivial solution. Line 13 compares the nn_dist with the lower bound of the target T_p with the current reference target T_q . If nn_dist is smaller or equal to the lower bound, it means that we have already found the closest match for the target T_p , so we can skip the rest of reference targets (Line 14) for current target T_p and proceed to check the next target subsequence T_{p+1} ; otherwise, the actual distance $dist$ is computed (Lines 16-21) and the nn_dist distance and the corresponding indices are updated accordingly (Lines 22-26). Finally, the global $best_so_far$ is updated based on nn_dist and is returned along with the indices of its corresponding contexts and targets (Lines 29-34).

Overall, in the best case, for each target T_p , we only need to sort all lower bounds with the cost of $O(n \log(n))$ and only compute actual distance once, which has a complexity of $O(L^2)$. Therefore, in the best case, the total cost for detecting semantic discord is $O(n^2 \log(n))$, which is a great improvement compared with smart brute force $O(n^2 L^2)$ as $\log n \ll L^2$ in most real world applications in time series [24]. We will show that the pruning rate of our algorithm reaches above 99.9% in the experimental section.

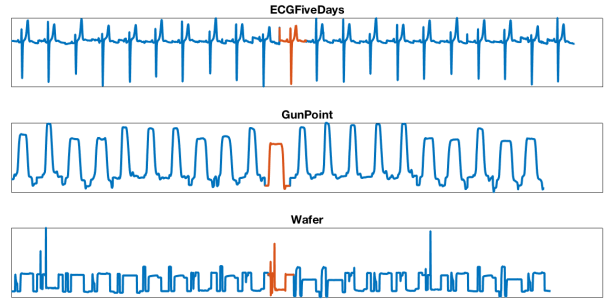


Figure 2: Time series generated from concatenation on different classes of UCR datasets. The subsequence in red is the ground truth of anomaly.

5 EXPERIMENTS

5.1 Experiment Setup We first describe the experiment setup on the proposed method and the baseline. Due to the unsupervised nature of discord, the ground truth is often not available. We use the widely-used UCR time series classification benchmark¹ to generate time series. The datasets under the category of synthetic or frequency domain with high frequency noises and those with missing values are excluded as we are interested in real world data with interpretation in time domain. We select all the remaining data with lengths between 80 and 500. For each dataset, the

¹https://www.cs.ucr.edu/~eamonn/time_series_data/

time series is generated by *concatenating* twenty instances from one class and one instance from a different class. We randomly generate twenty time series for each dataset and report the average result. Figure 2 shows three synthetic datasets that we generated from *ECGFiveDays*, *GunPoint* and *Wafer* datasets.

5.1.1 Baseline Methods We compare the proposed approach with the following state-of-the-art discord detection methods :

- **HOTSAX** [11] identifies the anomaly by the largest 1-nearest-neighbor distance. Various papers [24, 26] have proposed acceleration methods on computing the 1-nearest-neighbor distances. We use *STOMP* [26] as the accelerated version of HOTSAX to find the discord. We will use **Discord context** and **Discord target** to denote the lengths of discord context and target, respectively.
- **RKNN** [15] identifies the discord as the subsequence that has the lowest reverse nearest neighbor count. We use the exact version without approximating heuristic to achieve the best performance. The subsequence lengths of context and target, respectively, for the RKNN approach are denoted as **RKNN context** and **RKNN target**.
- **Rule Density Curve** [22] uses Sequitur grammar rules density curve to identify the anomalous subsequences corresponding to rare grammar rules. We use **RD target** and **RD context** to denote the target length and context length, respectively, for the rule density curve method.

The actual lengths of time series instances in the UCR datasets are used as the window sizes across all methods. Our proposed method uses the actual instance length as the context length, and we fix the target length arbitrary as 40% of the context length across all datasets. The ϵ value is obtained by top 40% percentile from 2,000 randomly-sampled contexts.

To ensure a fair comparison, we run both the target and context lengths as window size for all baseline methods. For the additional parameters of the baselines, we follow the original papers and employ $k = 3$ for RKNN, and use $w = 4$ and $a = 4$ for rule density curve.

5.2 Performance measurement We use the overlapping rate of detected discord with the ground truth as the metric of evaluation. The overlapping rate is defined as

$$\text{Overlapping rate} = \frac{|Discord \cap GroundTruth|}{|Discord|}.$$

The range is between 0 and 1, where a higher overlapping rate indicates a better performance. We pick overlapping rate as the evaluation criterion as it provides more direct quantitative measure on the quality of detection instead of simply setting a threshold on hit or miss. Moreover, overlapping rate allows effective measurement on anomaly length difference from the ground truth.

The average overlapping rate of twenty synthetic time series generated from each dataset is reported.

5.3 Results Table 1 shows the overlapping rate for all the methods on the generated datasets from the UCR database. Our proposed approach has the highest overlapping rate of 0.65 with the ground truth, and three results are higher than 0.5. In contrast, other methods mostly report overlapping rates that are lower than 0.4 with one exception of Discord Context on the GunPoint dataset. Discord with context length is the second best method which achieves 4 wins. Comparing with the three baseline methods, each with two sets of length options, Semantic Discord achieves the best overlapping rates on 11 datasets. From the result, we conclude that our method is better in terms of the overall wins, has the best performance compared to the state-of-the-art methods, and is able to identify anomalous subsequences that other methods could not find.

We perform the Wilcoxon signed rank test on the result of overlapping rate. The p -value of Wilcoxon test between our proposed approach and all other methods are 0.0392, 0.00386, 0.0012, 0.03, 0.0019, 0.0078 respectively. All the p -values are less than 0.05 and we conclude that the proposed method significantly outperforms Discord, RKNN and Rule density curve. The results show that our proposed approach has statistically significant 'wins' over the other methods.

Our method performs generally better on the datasets with some pattern structure with local differences such as ECG data, but it does not perform well in the cases where there is no existing semantic context information.

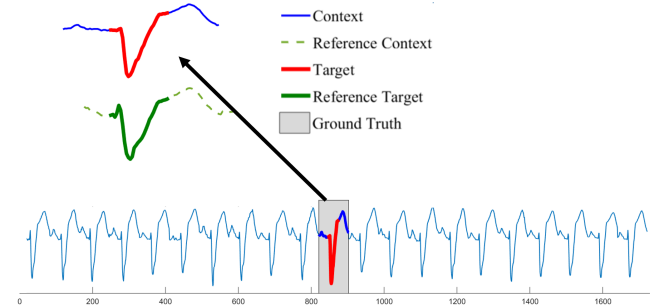


Figure 3: Target and context pair found agrees with ground truth on ECGFiveDays Synthetic data

5.3.1 Interpretability Besides the superior performance, the Semantic Discord provides additional interpretability. Figure 3 shows one of the synthetic data generated from ECGFiveDays. The context subsequence successfully locates the anomaly at the exact location of the ground truth anomaly shown in the grey shaded area with starting position of 821. The normalized context and target pair are shown in red and blue, and the normalized reference target and reference context pair are drawn in light dotted green line and dark green line respectively. We can see that the beginning of the target subsequence in blue does not

Table 1: Datasets description and performance comparison

Datasets	Type	Length	Discord context	Discord target	RKNN context	RKNN target	RD context	RD target	Proposed method
Coffee	Spectro	286	0.1	0.394	0.082	0.197	0	0.05	0.55
DistalPhalanxOutlineCorrect	Image	80	0.283	0.234	0	0.05	0.05	0	0.3
ECG200	ECG	96	0.2	0.154	0.15	0.05	0.05	0	0.283
ECGFiveDays	ECG	136	0.172	0.4	0	0	0	0.066	0.65
GunPoint	Motion	150	0.461	0.215	0	0.008	0.032	0	0.44
GunPointMaleVersusFemale	Motion	150	0.1	0.103	0	0	0	0.05	0.528
GunPointOldVersusYoung	Motion	150	0.001	0.095	0	0.006	0.05	0.017	0.468
Ham	Spectro	431	0.209	0.05	0	0	0.034	0	0
MiddlePhalanxOutlineCorrect	Image	80	0.197	0.15	0	0.034	0.05	0.2	0.198
MoteStrain	Sensor	84	0.014	0.117	0.047	0	0	0	0.011
PhalangesOutlinesCorrect	Image	80	0.052	0.153	0	0	0	0.081	0.12
PowerCons	Power	144	0	0.109	0.043	0	0	0	0.249
ProximalPhalanxOutlineCorrect	Image	80	0.284	0.05	0	0	0.05	0.203	0.059
SonyAIBORobotSurface1	Sensor	70	0	0.095	0.248	0.068	0.093	0.05	0.02
SonyAIBORobotSurface2	Sensor	65	0.017	0.05	0.019	0.05	0.05	0	0.194
Strawberry	Spectro	235	0.341	0.418	0	0.02	0.009	0.093	0.452
ToeSegmentation1	Motion	277	0.008	0	0.077	0.029	0	0.188	0
ToeSegmentation2	Motion	343	0	0.132	0.05	0.05	0.1	0.108	0.031
TwoLeadECG	ECG	82	0.103	0.25	0	0.056	0	0.078	0.608
Wafer	Sensor	152	0.05	0.109	0	0.022	0.05	0.1	0.129
Wine	Spectro	234	0.2	0.1	0	0.039	0.1	0.116	0.053
Yoga	Image	426	0.115	0.116	0	0	0.144	0	0.057
Number of Win			4	3	1	3	1	2	11
(Wilcoxon Test) Proposed vs Other			0.0392	0.0386	0.0012	0.0003	0.0019	0.0078	

have the small sharp edge, and the ascend is straighter compared to the reference target. Our method provides insight to the analyst on the discovered semantic discord, as well as improved interpretability.

5.4 Scalability As we propose a new definition in this work, we compare the scalability against the smart brute force of our definition on both synthetic and real data. The data we use include Random Walk (a synthetic random walk data), ECG (an ECG trace of 530,000 samples), Dishwasher data (an electrical consumption data of 180,000 samples in Watts for 20 households at aggregate and appliance level [20]). We test up to the first 128,000 samples and compare the total number of distance calls to the Euclidean distance function on the search for semantic discord. The length of the target discord is set to 160, and the context length is set to 400. All the experiments are performed on a computer with i5-7400HQ CPU of 2.80 Ghz with 32.0 GB memory on Matlab.

Figure 4 shows the numbers of distance calls by our method and the brute force algorithm, respectively. Under log scale, it is easy to observe that the difference increases as the dataset gets larger. At the length of 16,000, the pruning rate of our method is higher than 99.3% for all three datasets. At the size of 128,000, our speed-up technique is about 2,000 times faster than brute force method for the ECG data, and the average pruning rate for the distance call is 99.95%.

6 CONCLUSION

Finding time series anomalous subsequence is a critical problem and has broad applications. In this work, we introduce a new definition named *Semantic Discord*, which incorporates the context information from larger subsequences containing the anomaly candidates. We propose an efficient pruning al-

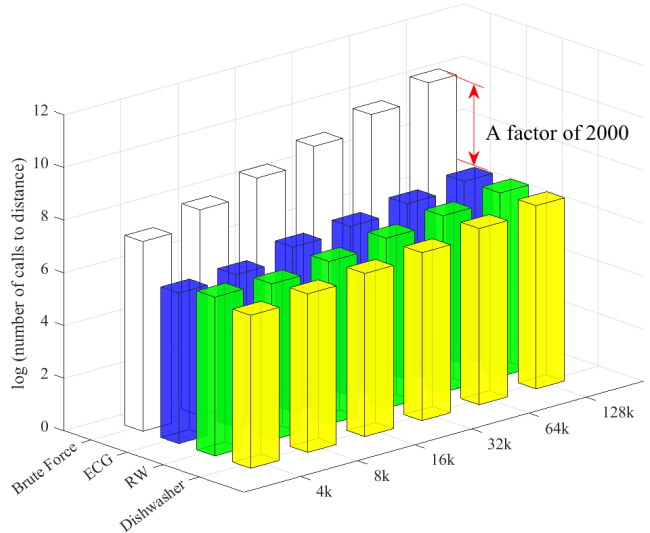


Figure 4: A comparison of our algorithm and brute force on number of distance calls (log scale) on three datasets. For brute force, the number of distance calls is the same for all datasets.

gorithm with a derived lower bound that is up to 3 orders of magnitude faster compared to the smart brute force algorithm. Through the experiments, we demonstrate that our method outperforms the state-of-the-art methods, and is well suited for applications in different domains.

References

- [1] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander. Lof: identifying density-based local outliers. In *ACM sigmod record*, volume 29, pages 93–104. ACM, 2000.

- [2] R. Chalapathy and S. Chawla. Deep learning for anomaly detection: A survey. *arXiv preprint arXiv:1901.03407*, 2019.
- [3] V. Chandola, A. Banerjee, and V. Kumar. Anomaly detection: A survey. *ACM computing surveys (CSUR)*, 41(3):15, 2009.
- [4] B. Chiu and et al. Probabilistic discovery of time series motifs. In *Proceedings of the 9th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 493–498, 2003.
- [5] L. Feremans, V. Verduyck, B. Cule, W. Meert, and B. Goethals. Pattern-based anomaly detection in mixed-type time series. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 2019.
- [6] Y. Gao, Q. Li, X. Li, J. Lin, and H. Rangwala. Trajviz: a tool for visualizing patterns and anomalies in trajectory. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 428–431. Springer, 2017.
- [7] S. Ghosh and D. L. Reilly. Credit card fraud detection with a neural-network. In *System Sciences, 1994. Proceedings of the Twenty-Seventh Hawaii International Conference on*, volume 3, pages 621–630. IEEE, 1994.
- [8] C. Hong and M. Hauskrecht. Multivariate conditional anomaly detection and its clinical application. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- [9] F. Keller, E. Muller, and K. Bohm. Hics: High contrast subspaces for density-based outlier ranking. In *2012 IEEE 28th international conference on data engineering*, pages 1037–1048. IEEE, 2012.
- [10] E. Keogh and S. Kasetty. On the need for time series data mining benchmarks: a survey and empirical demonstration. *Data Mining and knowledge discovery*, 7(4):349–371, 2003.
- [11] E. Keogh, J. Lin, and A. Fu. Hot sax: Efficiently finding the most unusual time series subsequence. In *Fifth IEEE International Conference on Data Mining (ICDM'05)*, pages 8–pp. IEEE, 2005.
- [12] S. Kitaguchi. Extracting feature based on motif from a chronic hepatitis dataset. In *Proceedings of the 18th annual conference of the Japanese society for artificial intelligence (JSAI)*, 2004.
- [13] H.-P. Kriegel, P. Kröger, E. Schubert, and A. Zimek. Outlier detection in axis-parallel subspaces of high dimensional data. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 831–838. Springer, 2009.
- [14] H.-P. Kriegel, M. Schubert, and A. Zimek. Angle-based outlier detection in high-dimensional data. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 444–452. ACM, 2008.
- [15] J. Lin, D. Etter, and D. DeBarr. Exact and approximate reverse nearest neighbor search for multimedia data. In *Proceedings of the 2008 SIAM International Conference on Data Mining*, pages 656–667. SIAM, 2008.
- [16] P. Malhotra, A. Ramakrishnan, G. Anand, L. Vig, P. Agarwal, and G. Shroff. Lstm-based encoder-decoder for multi-sensor anomaly detection. *arXiv preprint arXiv:1607.00148*, 2016.
- [17] T. P. Michele Linardi, Yan Zhu and E. Keogh. Matrix profile x: Valmod - scalable discovery of variable-length motifs in data series. In *SIGMOD*. ACM, 2018.
- [18] M. Munir, S. Erkel, A. Dengel, and S. Ahmed. Pattern-based contextual anomaly detection in hvac systems. In *2017 IEEE International Conference on Data Mining Workshops (ICDMW)*, pages 1066–1073. IEEE, 2017.
- [19] M. Munir, S. A. Siddiqui, A. Dengel, and S. Ahmed. Deepant: A deep learning approach for unsupervised anomaly detection in time series. *IEEE Access*, 7:1991–2005, 2018.
- [20] D. Murray, L. Stankovic, and V. Stankovic. An electrical load measurements dataset of united kingdom households from a two-year longitudinal study. *Scientific data*, 4:160122, 2017.
- [21] T. Rakhmanan, B. Campana, A. Mueen, G. Batista, B. Westover, Q. Zhu, J. Zakaria, and E. Keogh. 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*, pages 262–270. ACM, 2012.
- [22] P. Senin, J. Lin, X. Wang, T. Oates, S. Gandhi, A. P. Boedihardjo, C. Chen, and S. Frankenstein. Time series anomaly discovery with grammar-based compression. In *EDBT*, pages 481–492, 2015.
- [23] X. Wang, J. Lin, N. Patel, and M. Braun. A self-learning and online algorithm for time series anomaly detection, with application in cpu manufacturing. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, pages 1823–1832. ACM, 2016.
- [24] C.-C. M. Yeh, Y. Zhu, L. Ulanova, N. Begum, Y. Ding, H. A. Dau, D. F. Silva, A. Mueen, and E. Keogh. Matrix profile i: all pairs similarity joins for time series: a unifying view that includes motifs, discords and shapelets. In *Data Mining (ICDM), 2016 IEEE 16th International Conference on*, pages 1317–1322. IEEE, 2016.
- [25] Y. Yuan, G. Xun, F. Ma, Y. Wang, N. Du, K. Jia, L. Su, and A. Zhang. Muvan: A multi-view attention network for multivariate temporal data. In *2018 IEEE International Conference on Data Mining (ICDM)*, pages 717–726. IEEE, 2018.
- [26] Y. Zhu, Z. Zimmerman, N. S. Senobari, C.-C. M. Yeh, G. Funning, A. Mueen, P. Brisk, and E. Keogh. Matrix profile ii: Exploiting a novel algorithm and gpus to break the one hundred million barrier for time series motifs and joins. In *Data Mining (ICDM), 2016 IEEE 16th International Conference on*, pages 739–748. IEEE, 2016.