

Series2Graph: Graph-based Subsequence Anomaly Detection for Time Series

Paul Boniol
EDF R&D, LIPADE, Université de Paris
boniol.paul@gmail.com

Themis Palpanas
LIPADE, Université de Paris &
French University Institute (IUF)
themis@mi.parisdescartes.fr

ABSTRACT

Subsequence anomaly detection in long sequences is an important problem with applications in a wide range of domains. However, the approaches that have been proposed so far in the literature have severe limitations: they either require prior domain knowledge that is used to design the anomaly discovery algorithms, or become cumbersome and expensive to use in situations with recurrent anomalies of the same type. In this work, we address these problems, and propose an unsupervised method suitable for domain agnostic subsequence anomaly detection. Our method, Series2Graph, is based on a graph representation of a novel low-dimensionality embedding of subsequences. Series2Graph needs neither labeled instances (like supervised techniques), nor anomaly-free data (like zero-positive learning techniques), and identifies anomalies of varying lengths. The experimental results, on the largest set of synthetic and real datasets used to date, demonstrate that the proposed approach correctly identifies single and recurrent anomalies without any prior knowledge of their characteristics, outperforming by a large margin several competing approaches in accuracy, while being up to orders of magnitude faster.

PVLDB Reference Format:

Paul Boniol and Themis Palpanas. Series2Graph: Graph-based Subsequence Anomaly Detection for Time Series. *PVLDB*, 13(11): 1821-1834, 2020.
DOI: <https://doi.org/10.14778/3407790.3407792>

Keywords

Time series, Data series, Subsequence anomalies, Outliers.

1. INTRODUCTION

Data series¹ anomaly detection is a crucial problem with application in a wide range of domains [46, 6]. Examples of

¹A data series is an ordered sequence of real-valued points. If the dimension that imposes the ordering of the sequence is time then we talk about *time series*, but it could also be mass (e.g., mass spectrometry), angle (e.g., astronomy), or

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 13, No. 11

ISSN 2150-8097.

DOI: <https://doi.org/10.14778/3407790.3407792>

such applications can be found in manufacturing, astronomy, engineering, and other domains [44, 46], including detection of abnormal heartbeats in cardiology [27], wear and tear in bearings of rotating machines [5], machine degradation in manufacturing [41], hardware and software failures in data center monitoring [47], mechanical faults in vehicle operation monitoring [17] and identification of transient noise in gravitational wave detectors [7]. This implies a real need by relevant applications for developing methods that can accurately and efficiently achieve this goal.

[Anomaly Detection in Sequences] Anomaly detection is a well studied task [8, 53, 60, 36] that can be tackled by either examining single values, or sequences of points. In the specific context of sequences, which is the focus of this paper, we are interested in identifying anomalous subsequences [60, 51], which are not single abnormal values, but rather an abnormal *sequence* of values. In real-world applications, this distinction becomes crucial: in certain cases, even though each individual point may be normal, the trend exhibited by the sequence of these same values may be anomalous. Failing to identify such situations could lead to severe problems that may only be detected when it is too late [5].

[Limitations of Previous Approaches] Some existing techniques explicitly look for a set of pre-determined types of anomalies [27, 2]. These are techniques that have been specifically designed to operate in a particular setting, they require domain expertise, and cannot generalize.

Other techniques identify as anomalies the subsequences with the largest distances to their nearest neighbors (termed discords) [60, 51]. The assumption is that the most distant subsequence is completely isolated from the "normal" subsequences. However, this definition fails when an anomaly repeats itself (approximately the same) [55]. In this situation, anomalies will have other anomalies as close neighbors, and will not be identified as discords. In order to remedy this situation, the m^{th} discord approach has been proposed [59], which takes into account the multiplicity m of the anomalous subsequences that are similar to one another, and marks as anomalies all the subsequences in the same group. However, this approach assumes the cardinality of the anomalies to be known, which is not true in practice (otherwise, we need to try several different m values, increasing execution time). Furthermore, the majority of the previous approaches require prior knowledge of the anomaly length, and their performance deteriorates significantly when the correct length value is not used.

position (e.g., biology). In this paper, we will use the terms *time series*, *data series*, and *sequence* interchangeably.

[Proposed Approach] In this work, we address the aforementioned issues, and we propose Series2Graph, an unsupervised method suitable for domain agnostic subsequence anomaly detection. Our approach does not need labeled instances (like supervised techniques do), or clean data that do not contain anomalies (like zero-positive learning techniques require). It also allows the same model to be used for the detection of anomalies of different lengths.

Series2Graph is based on a graph representation of a novel low-dimensionality embedding of subsequences. It starts by embedding subsequences into a vector space, where information related to their shapes is preserved. This space is then used to extract overlapping trajectories that correspond to recurrent patterns in the data series. Subsequently, we construct a graph, whose nodes are derived from the overlapping trajectories, and edges represent transitions (among subsequences in different nodes) that exist in the original series.

Intuitively, this graph encodes all the subsequences of a (single, or collection of) data series, and encodes the recurring patterns in these subsequences. This allows us then to differentiate between normal behavior, i.e., frequently occurring patterns, and anomalies, i.e., subsequences that rarely occur in the data series.

Overall, the experimental results (that include comparisons to several state of the art approaches, using a superset of the publicly available datasets used in the literature for subsequence anomaly detection) demonstrate that Series2Graph dominates by a large margin the competitors in accuracy, versatility, and execution time.

[Contributions] Our contributions are the following.

- We propose a new formalization for the subsequence anomaly detection problem, which overcomes the shortcomings of existing models. Our formalization is based on the intuitive idea that anomalous are the subsequences that are not similar to the common behavior, which we call normal.
- We describe a novel low-dimensionality embedding for subsequences, and use a graph representation for these embeddings. This representation leads to a natural distinction between recurring subsequences that constitute normal behavior, and rarely occurring subsequences that correspond to anomalies.
- Based on this representation, we develop Series2Graph [11], an unsupervised method for domain agnostic subsequence anomaly detection. Series2Graph supports the identification of previously unseen single and recurring anomalies, and can be used to find anomalies of different lengths.
- Finally, we conduct an extensive evaluation using several large and diverse datasets from various domains that demonstrates the effectiveness and efficiency of Series2Graph.

2. PRELIMINARIES

[Data Series] We begin by introducing some formal notations useful for the rest of the paper.

A data series $T \in \mathbb{R}^n$ is a sequence of real-valued numbers $T_i \in \mathbb{R} [T_1, T_2, \dots, T_n]$, where $n = |T|$ is the length of T , and T_i is the i^{th} point of T . We are typically interested in local regions of the data series, known as subsequences. A subsequence $T_{i,\ell} \in \mathbb{R}^\ell$ of a data series T is a continuous subset of the values from T of length ℓ starting at position i . Formally, $T_{i,\ell} = [T_i, T_{i+1}, \dots, T_{i+\ell-1}]$.

Given two sequences, A and B , of the same length, ℓ , we can calculate their Z-normalized Euclidean dis-

tance, $dist$, as follows [16, 43, 54, 56, 57]: $dist = \sqrt{\sum_{i=1}^{\ell} (\frac{A_i - \mu_A}{\sigma_A} - \frac{B_i - \mu_B}{\sigma_B})^2}$, where μ and σ represent the mean and standard deviation, respectively, of the sequences. In the following, we will simply use the term *distance*.

Given a subsequence $T_{i,\ell}$, we say that its m^{th} Nearest Neighbor (m^{th} NN) is $T_{j,\ell}$ if $T_{j,\ell}$ has the m^{th} shortest distance to $T_{i,\ell}$ among all the subsequences of length ℓ in T , excluding trivial matches [22]; a trivial match of $T_{i,\ell}$ is a subsequence $T_{a,\ell}$, where $|i-a| < \ell/2$ (i.e., the two subsequences overlap by more than half their length).

[Discords and their Shortcomings] We now define the data series *discord*, which is the most prevalent subsequence anomaly definition in the literature.

DEFINITION 1 (DISCORD). [59, 51, 60, 36] *A subsequence $T_{i,\ell}$ is a discord if the distance between its NN, namely $T_{j,\ell}$, is the largest among all the NN distances computed between subsequences of length ℓ in T . We require that $T_{j,\ell}$ is not a trivial match of $T_{i,\ell}$.*

DEFINITION 2 (m^{th} -DISCORD). [59] *A subsequence $T_{i,\ell}$ is an m^{th} -discord if the distance between its m^{th} NN, namely $T_{j,\ell}$, is the largest among all the m^{th} NN distances computed between subsequences of length ℓ in T . We require that $T_{i,\ell}$ is not a trivial match of $T_{i,\ell}$.*

Subsequence anomaly detection based on discords has attracted a lot of interest in the past years. There exist several works that have proposed fast and scalable discord discovery algorithms in various settings [51, 32, 38, 23, 60, 15, 58, 39], including simple and m^{th} -discords (the authors of these papers define the problem as k^{th} -discord discovery), in-memory and disk-aware techniques, exact and approximate algorithms.

The strength of this definition is its mathematical simplicity. Nevertheless, we observe that it fails to address some challenges of real use cases. The reason is that large datasets may contain several anomalies that repeat themselves (approximately the same). In other words, it is likely that an anomaly has another anomaly as its close neighbor, and thus, does not correspond to a discord anymore.

Even though the m^{th} -discord definition [59] tried to address these problems, the parameter m that refers to the multiplicity of some anomaly in the dataset remains a user-defined parameter, and is not easy to set correctly. Choosing an m value that is smaller, or larger, than the correct one, will lead to both false negatives and false positives.

[Graphs] We introduce some basic definitions for graphs, which we will use in this paper.

We define a Node Set \mathcal{N} as a set of unique integers. Given a Node Set \mathcal{N} , an Edge Set \mathcal{E} is then a set composed of tuples (x_i, x_j) , where $x_i, x_j \in \mathcal{N}$. $w(x_i, x_j)$ is the weight of that edge.

Given a Node Set \mathcal{N} , an Edge Set \mathcal{E} (pairs of nodes in \mathcal{N}), a Graph G is an ordered pair $G = (\mathcal{N}, \mathcal{E})$. A directed graph or digraph G is an ordered pair $G = (\mathcal{N}, \mathcal{E})$ where \mathcal{N} is a Node Set, and \mathcal{E} is an ordered Edge Set.

In the rest of this paper, we will only use directed graphs, denoted as G .

3. PROBLEM FORMULATION

We now provide a new formulation for subsequence anomaly detection. The idea is that a data series is transformed into a sequence of abstract states (corresponding to

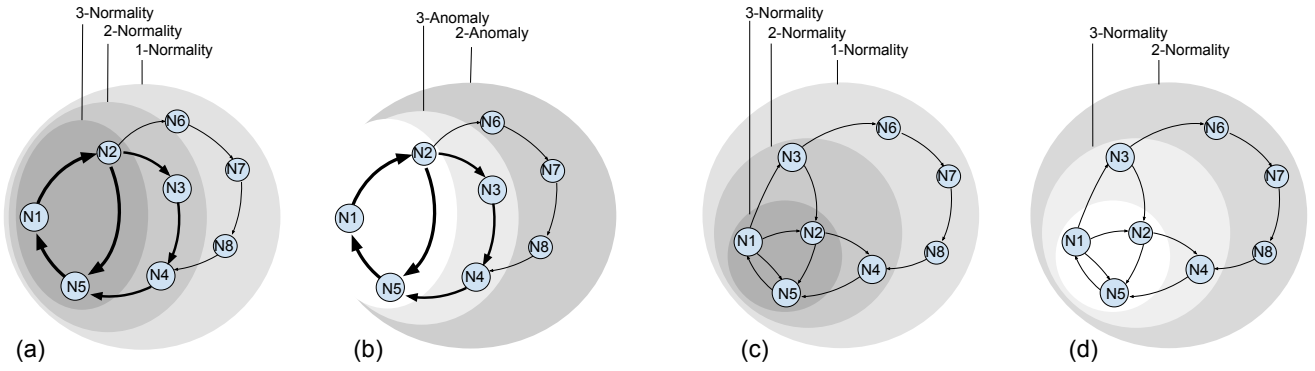


Figure 1: 3-Normality, 2-Normality, 1-Normality, and 3-Anomaly, 2-Anomaly for two given graphs ((a),(b) and (c),(d)) representing the simplified model of two data series. Edge weights and node degrees are used to define the θ -Normality and θ -Anomaly subgraphs.

different subsequence patterns), represented by nodes \mathcal{N} in a directed graph, $G(\mathcal{N}, \mathcal{E})$, where the edges \mathcal{E} encode the number of times one state occurred after another. Thus, normality can be characterized by the (i) the edge weight, which indicates the number of times two subsequences occurred one after the other in the original sequence, and (ii) the node degree, the number of edges adjacent to the node, which indicates the proximity of the subsequences in that node to other subsequences. Note that G is a connected graph (there exists a path between every pair of nodes), and thus, the degree of each node is at least equal to 1.

Under this formulation, paths in the graph composed of high-weight edges and high-degree nodes correspond to normal behavior. As a consequence, the Normality of a data series can be defined as follows.

DEFINITION 3 (θ -NORMALITY). Let a node set be defined as $\mathcal{N} = \{N_1, N_2, \dots, N_m\}$. Let also a data series T be represented as a sequence of nodes $\langle N^{(1)}, N^{(2)}, \dots, N^{(n)} \rangle$ with $\forall i \in [0, n], N^{(i)} \in \mathcal{N}$ and $m \leq n$. The θ -Normality of T is the subgraph $G_\theta^\nu(\mathcal{N}_\nu, \mathcal{E}_\nu)$ of $G(\mathcal{N}, \mathcal{E})$ with $\mathcal{E} = \{(N^{(i)}, N^{(i+1)})\}_{i \in [0, n-1]}$, such that: $\mathcal{N}_\nu \subset \mathcal{N}$ and $\forall (N^{(i)}, N^{(i+1)}) \in \mathcal{E}_\nu, w((N^{(i)}, N^{(i+1)}), (deg(N^{(i)})) - 1) \geq \theta$.

An example of θ -Normality subgraph is shown in Figures 1(a) and (c). In Figure 1(a), the subgraph composed of nodes N_1, N_2, N_5 , has edges with weights larger than 3, and a minimum node degree of 2. Therefore, it is a 3-Normality subgraph. In Figure 1(c), the subgraph composed of nodes N_1, N_2, N_5 , has edges of weight 1, but does not have any node with a degree under 4. Therefore, it is a 3-Normality subgraph. Similarly, we define an anomaly as follows.

DEFINITION 4 (θ -ANOMALY). Let a node set be defined as $\mathcal{N} = \{N_1, N_2, \dots, N_m\}$. Let a data series T be represented as a sequence of nodes $\langle N^{(1)}, N^{(2)}, \dots, N^{(n)} \rangle$ with $\forall i \in [0, n], N^{(i)} \in \mathcal{N}$ and $m \leq n$. The θ -Anomaly of T is the subgraph $G_\theta^\alpha(\mathcal{N}_\alpha, \mathcal{E}_\alpha)$ of $G(\mathcal{N}, \mathcal{E})$ with $\mathcal{E} = \{(N^{(i)}, N^{(i+1)})\}_{i \in [0, n-1]}$, such that: $G_\theta^\nu(\mathcal{N}_\nu, \mathcal{E}_\nu) \cap G_\theta^\alpha(\mathcal{N}_\alpha, \mathcal{E}_\alpha) = \emptyset$.

An example of θ -Anomaly subgraph is outlined in Figures 1(b) and (d). In Figure 1(b), the nodes that do

not belong to the 3-Normality subgraph constitute the 3-Anomaly subgraph. The 2-Anomaly subgraph is included in the 3-Anomaly subgraph, and the intersection of the 2-Anomaly and 2-Normality subgraphs is empty. Similar observations hold for Figure 1(d). We now define the membership criteria of a subsequence to a θ -Normality subgraph.

DEFINITION 5 (θ -NORMALITY MEMBERSHIP). Given a data series T represented as a sequence of abstract states $\langle N^{(1)}, N^{(2)}, \dots, N^{(n)} \rangle$, a subsequence $T_{i,\ell}$, represented by $\langle N^{(i)}, N^{(i+1)}, \dots, N^{(i+\ell)} \rangle$, belongs to the θ -Normality of T if and only if $\forall j \in [i, i+\ell], (N^{(j)}, N^{(j+1)}) \in \theta\text{-Normality}(T)$. On the contrary, $T_{i,\ell}$ belongs to the θ -Anomaly of T if and only if $\exists j \in [i, i+\ell], (N^{(j)}, N^{(j+1)}) \notin \theta\text{-Normality}(T)$.

Based on the above definitions, using θ -Normality subgraphs naturally leads to a ranking of subsequences based on their "normality". For practical reasons, this ranking can be transformed into a score, where each rank can be seen as a threshold in that score. We elaborate on this equivalence in the following section. Observe also that the subsequence length is not involved in the definition of normal/abnormal, which renders this approach more general and flexible.

Note that given the existence of graph G , the above definitions imply a way for identifying the anomalous subsequences. The problem is now how to construct this graph. Therefore, the problem we want to solve is the following.

PROBLEM 1 (PATTERN GRAPH CONSTRUCTION). Given a data series T , we want to automatically construct the graph $G(\mathcal{N}, \mathcal{E})$.

Following common practice, we assume that anomalies correspond to rare events.

Table 1 summarizes the symbols we use in this paper.

4. PROPOSED APPROACH

In this section, we describe Series2Graph, our unsupervised solution to the subsequence anomaly detection problem. For a given data series T , the overall Series2Graph process is divided into four main steps as follows (video examples of this process are available online [1]).

- 1. Subsequence Embedding:** Project all the subsequences (of a given length ℓ) of T in a two-dimensional space, where shape similarity is preserved.

Table 1: Table of symbols

Symbol	Description
T	a data series
$ T $	cardinality of T
λ	convolution size
ℓ	input subsequence length
ℓ_q	query subsequence length
ℓ_A	anomaly length
\mathcal{N}	set of nodes
\mathcal{E}	set of edges
$G(\mathcal{N}, \mathcal{E})$	directed graph corresponding to T
θ	density layer (for normality/anomaly)
θ -Normality	subgraph of G (also called G_θ^ν)
θ -Anomaly	subgraph of G (also called G_θ^α)
\mathcal{N}_ν	set of nodes of θ -Normality
\mathcal{E}_ν	set of edges of θ -Normality
\mathcal{N}_α	set of nodes of θ -Anomaly
\mathcal{E}_α	set of edges of θ -Anomaly
$w(e)$	weight of edge $e \in \mathcal{E}$
$deg(N_i)$	degree of node $N_i \in \mathcal{N}$
$Proj$	set of all embedded subsequences
$Proj_r$	reduced set $Proj$ of three dimensions
$SProj$	rotated $Proj_r$
ψ	angle
Ψ	angle set
\mathcal{I}_ψ	radius set of angle ψ
\mathcal{N}_ψ	node set in \mathcal{I}_ψ

- Node Creation:** Create a node for each one of the densest parts of the above two-dimensional space. These nodes can be seen as a summarization of all the major patterns of length ℓ that occurred in T .
- Edge Creation:** Retrieve all transitions between pairs of subsequences represented by two different nodes: each transition corresponds to a pair of subsequences, where one occurs immediately after the other in the input data series T . We represent transitions with an edge between the corresponding nodes. The weights of the edges are set to the number of times the corresponding pair of subsequences was observed in T .
- Subsequence Scoring:** Compute the normality (or anomaly) score of a subsequence of length $\ell_q \geq \ell$ (within or outside of T), based on the previously computed edges/nodes and their weights/degrees.

We note that length ℓ required in the first step of the method is user defined, but is independent of the length of the subsequences that we want to detect as anomalies, which can have different lengths. Moreover, the proposed approach is robust to the choice of ℓ , especially when the ℓ value is larger than the length of the targeted anomalies. In contrast, existing methods require prior knowledge of the anomaly length and can only discover anomalies of that length. We demonstrate these points in the experimental evaluation.

Below, we describe in detail each one of the above steps.

4.1 Subsequence Embedding

We first describe our approach for projecting a data series into a two-dimensional space. We propose a new shape-based embedding, such that two subsequences similar in shape will be geometrically close in the transformed space after the embedding. In order to achieve this, we (i) extract all the subsequences and represent them as vectors, (ii) reduce the dimensionality of these vectors to three dimensions

(that we can visualize in a three-dimensional space), (iii) rotate the space of these vectors (i.e., subsequences) such that two of the components contain the shape related characteristic, and the last one the average value. As a result, two subsequences with similar shape, but very different mean value (i.e., small Z-normalized Euclidean distance, but large Euclidean distance) will have very close values for the first two components, but very different for the third one.

We start by extracting subsequences using a sliding window that we slide by one point at a time², and then applying a local convolution (of size³ $\lambda = \ell/3$) to each subsequence, in order to reduce noise and highlight the important shape information. Formally, for a given subsequence length ℓ and local convolution size $\lambda = \ell/3$, we transform subsequence $T_{i,\ell}$ into a vector (of size $\ell - \lambda$):

$$\left[\sum_{k=i}^{i+\lambda} T_k, \sum_{k=i+1}^{i+1+\lambda} T_k, \dots, \sum_{k=i+\ell-\lambda}^{i+\ell} T_k \right]$$

We insert the vectors corresponding to all subsequences $T_{i,\ell}$ in matrix $Proj(T, \ell, \lambda) \in \mathbb{M}_{|T|, \ell-\lambda}(\mathbb{R})$, where \mathbb{M} is the set of real-valued matrices with $|T|$ rows and $\ell - \lambda$ columns.

In order to reduce the dimensionality of matrix $Proj(T, \ell, \lambda)$, we apply a Principal Component Analysis (PCA) transform. For the sake of simplicity, we keep only the first three components (PCA_3), and denote the reduced three-column matrix as $Proj_r(T, \ell, \lambda)$.

We note that using the first three components was sufficient for our study. Consider that for the 25 datasets used in our experimental evaluation, the three most important components explain on average 95% of the total variance. Generalizing our solution to a larger number of important components is part of our current work.

Since we are interested in subsequence anomalies, which correspond to anomalous shapes (trends), we need to emphasize (out of the three components obtained by the aforementioned reduced projection) the components that explain the most the shape of the subsequences. Let $min(T)$ and $max(T)$ be the minimum and maximum values of the data series T . We extract the vector $\vec{v}_{ref} = \overrightarrow{O_{mn}O_{mx}}$, where $O_{mn} = PCA_3(min(T) * \lambda * \mathbf{1}_{\ell-\lambda})$ and $O_{mx} = PCA_3(max(T) * \lambda * \mathbf{1}_{\ell-\lambda})$ (PCA_3 returns the three most important components using the trained PCA applied on $Proj(T, \ell, \lambda)$). Intuitively, the vector \vec{v}_{ref} describes the time dimension along which the values change (bounded by $\lambda * min(T)$ and $\lambda * max(T)$, where the multiplication with λ corresponds to a local convolution). The other dimensions (orthogonal vectors of \vec{v}_{ref}) describe how the values change. Thus, overlapping points/sequences in these other dimensions indicate recurrent behaviors and isolated points/sequences indicate possible anomalies. Given the unit vectors $(\vec{u}_x, \vec{u}_y, \vec{u}_z)$ that represent the axes of the cartesian coordinate system of the PCA, the angle $\phi_x = \angle \vec{u}_x \vec{v}_{ref}$, $\phi_y = \angle \vec{u}_y \vec{v}_{ref}$, $\phi_z = \angle \vec{u}_z \vec{v}_{ref}$ and their corresponding rotation matrices $R_{u_x}(\phi_x), R_{u_y}(\phi_y)$ and $R_{u_z}(\phi_z)$, we define $SProj(T, \ell, \lambda)$ as follows: $SProj(T, \ell, \lambda) = R_{u_x}(\phi_x)R_{u_y}(\phi_y)R_{u_z}(\phi_z)Proj_r(T, \ell, \lambda)^T$. The matrix

²This is equivalent to using a *time delay embedding* [30] with a delay $\tau = 1$.

³We use $\lambda = \ell/3$ in all our experiments. Varying λ between $\ell/10$ and $\ell/2$ does not affect the results (we omit this graph for brevity).

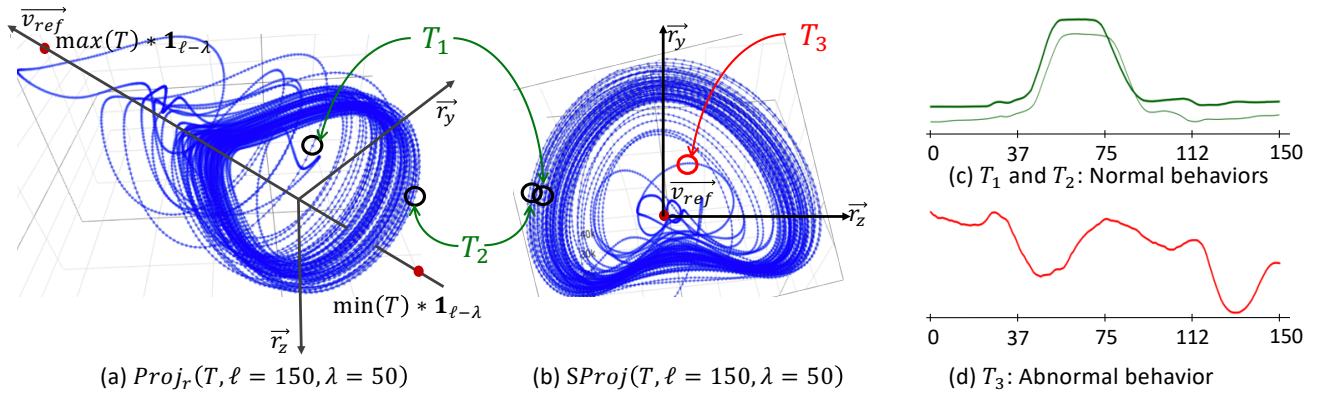


Figure 2: (a) $Proj_r(T, \ell, \lambda)$ and (b) $SProj(T, \ell, \lambda)$ of a data series T corresponding to the movement of an actor’s hand that (c) takes a gun out of the holster and points to a target (normal behavior); (d) the anomaly (red subsequence) corresponds to a moment when the actor missed the holster [31]. We rotate (a) into (b) such that \vec{v}_{ref} is invariant in two dimensions.

$SProj(T, \ell, \lambda)$ is the reduced projection $Proj_r(T, \ell, \lambda)$ rotated in order to have the unit vector \vec{u}_x aligned to the offset vector \vec{v}_{ref} .

Figure 2 depicts the rotation procedure to transform $Proj_r$ into $SProj$ for an example data series T that corresponds to the movement of an actor’s hand that takes a gun out of the holster and points to a target (normal behavior). This rotation is using vector \vec{v}_{ref} , defined by the minimal and maximal constant sequences mentioned earlier (marked with the red dots in Figure 2(a)). The unit vectors of the rotated space are $(\frac{\vec{v}_{ref}}{\|\vec{v}_{ref}\|}, \vec{r}_y, \vec{r}_z)$, where \vec{r}_y and \vec{r}_z are the rotated vectors \vec{u}_y and \vec{u}_z .

What this rotation achieves is that (similarly to Z-normalization) subsequences with a different mean but the same shape in the space before the rotation (e.g., subsequences T_1 and T_2 in Figure 2(c)) will have very close \vec{r}_y and \vec{r}_z components in the new coordinate system (as shown in Figures 2(a) and (b)). Therefore, subsequences with similar shapes will appear close together, shapes that repeat often in the dataset will form dense clusters in the space (like subsequences T_1 and T_2), and rare shapes (anomalies) will appear relatively isolated (like subsequence T_3). Figures 2(c) and (d) depict the normal (T_1 and T_2) and abnormal (T_3) subsequences. The anomaly (T_3) corresponds to a case when the actor missed the holster [31].

We observe that in the rotated space (see Figure 2(b)), the shape differences are easy to distinguish, and the identification of normal behavior (dense clusters of repeated patterns) and anomalies (isolated patterns) is visually obvious.

In the rest of the paper, $SProj(T, \ell, \lambda)$ will refer to the 2-dimensions matrix keeping only the r_y and r_z components.

Algorithm 1 describes the computation of the pattern embeddings. A naive solution is to compute all the convolutions for all the subsequences of T , which leads to a complexity of magnitude $O(|T|\ell\lambda)$. Nevertheless, by using the previously computed convolutions (Line 4), the complexity is reduced to $O(|T|\lambda)$. The PCA step of Algorithm 1 is implemented with a randomized truncated Singular Value Decomposition (SVD), using the method of Halko et al. [28] with a time complexity of $O(|T|(\ell - \lambda)|component|)$. The last step consists of matrix multiplications, and therefore

Algorithm 1: Pattern Embedding

```

input : Data series  $T$ , input length  $\ell$ ,  $\lambda$ 
output: 3-dimension points sequence  $SProj$ 
// Transform first subsequence
1  $P \leftarrow \left( \sum_{k=j}^{j+\lambda} T_k \right)_{j \in [0, \ell-\lambda]}$ ;
2 add  $P$  in  $Proj$ ;
// Transform every other subsequences in  $T$ 
3 foreach  $i \in [1, |T| - \ell]$  do
4    $P[0 : \ell - \lambda - 1] \leftarrow P[1 : \ell - \lambda]$ ;
5    $P[\ell - \lambda] \leftarrow \sum_{k=i+\ell-\lambda}^{i+\ell} T_k$ ;
6   add  $P$  in  $Proj$ ;
7 end
// Reduce to three dimensions
8  $pca \leftarrow PCA_3.fit(Proj)$ ;
9  $Proj \leftarrow pca.transform(Proj)$ ;
// Get rotation characteristics
10  $v_{ref} \leftarrow pca.transform((\max(T) - \min(T)) * \lambda * \mathbf{1}_{\ell-\lambda})$ ;
11  $\phi_x, \phi_y, \phi_z \leftarrow getAngle((u_x, u_y, u_z), v_{ref})$ ;
12  $R_{u_x}, R_{u_y}, R_{u_z} \leftarrow GetRotationMatrices(\phi_x, \phi_y, \phi_z)$ ;
// Rotate  $SProj$ 
13  $SProj \leftarrow R_{u_x}.R_{u_y}.R_{u_z}.Proj^T$ 

```

has a complexity of $O(|T||R_{u_x}|^2)$. In our case, the size of the rotation matrices are much smaller than λ , which leads to a global complexity of $O(3|T|(\ell - \lambda))$.

4.2 Node Creation

At this point, we are ready to extract shape related information, as in Figure 2, where recurrent and isolated trajectories can be distinguished. The idea is to extract the most *crossed* sections of the 2-dimensional space defined by the unit vector (\vec{r}_y, \vec{r}_z) . These sections will be the nodes in the graph we want to construct.

First, we define the *radius subset*.

DEFINITION 6 (RADIUS SET). *Given a data series T and its projection matrix $P = SProj(T, \ell, \lambda)$, the radius set \mathcal{I}_ψ is the set of intersection points between the vector $\vec{u}_\psi = \cos(\psi)\vec{r}_y + \sin(\psi)\vec{r}_z$ and every segment $[x_{i-1}, x_i]$, where x_{i-1}, x_i are two consecutive rows of P : $\mathcal{I}_\psi = \{x | (\vec{u}_\psi \times$*




Figure 3: Node extraction from $SProj(T, \ell, \lambda)$ by measuring the density of the intersected trajectories to a given vector. The densest points are added to the Node Set \mathcal{N} .

$\vec{x} = \vec{0}) \wedge (\vec{x}_{i-1} \times \vec{x}_{i-1} \times \vec{x}_i = \vec{0})\}$, where \times operator is the cross product.

Figure 3 (a) displays two radius subsets (marked with the red points). We can now define the *Pattern Node Set* as follows.

DEFINITION 7 (PATTERN NODE SET). *Given a data series T , its projection $P = SProj(T, \ell, \lambda)$ and a set of \mathcal{I}_ψ ($\psi \in \Psi$), the Pattern Node Set of T is:*

$$\mathcal{N} = \cup_{\psi \in \Psi} \mathcal{N}_\psi$$

$$\mathcal{N}_\psi = \{x | \exists \epsilon, \forall y, |x - y| > \epsilon \implies f_h(x, \mathcal{I}_\psi) > f_h(y, \mathcal{I}_\psi)\}$$

$$\text{with } f_h(x, \mathcal{I}_\psi) = \frac{1}{nh\sqrt{2\pi\sigma(\mathcal{I}_\psi)^2}} \sum_{x_i \in \mathcal{I}_\psi} e^{-\frac{(x-x_i-h\mu(\mathcal{I}_\psi))^2}{2h\sigma(\mathcal{I}_\psi)^2}}$$

In the above definition, f is a kernel density estimation function, applied on a radius subset using a Gaussian distribution. Then, nodes become the areas in the 2-dimensional space, where the trajectories of the patterns are the most likely to pass through. In other words, each node corresponds to a set of very similar patterns. The bandwidth parameter h affects the granularity of the extraction: the smaller the h value is, the more local maxima, and therefore the more nodes we will end up with; the larger the h value is, the less nodes the graph will have, and therefore the more general it will be. We define parameter $r = |\Psi|$ as the number of angles ψ that we use in order to extract the pattern node set. In other words, this parameter is sampling the space (refer to Algorithm 2, Line 1). Once again, a big number of angles will lead to high precision, but at the cost of increased computation time.

In practice, we observed that parameter r is not critical, and we thus set $r = 50$ for the rest of this paper. Regarding the bandwidth parameter of the density estimation, we set it following the Scott's rule [50]: $h_{scott} = \sigma(\mathcal{I}_\psi) \cdot |\mathcal{I}_\psi|^{-\frac{1}{5}}$.

Algorithm 2 outlines the above process for extracting the node set from $SProj$. In Lines 6-8, we compute for each radius subset \mathcal{I}_ψ all intersection points between a radius vector and the possible segments composed of two consecutive points in $SProj$. The complexity of this operation is bounded by $O(|SProj|r) \simeq O(|T|r)$. The time complexity of the kernel density estimation is $O(|\mathcal{I}_\psi|)$ (since

Algorithm 2: Node Extraction

input : 2-dimensional point sequence $SProj$, rate r , bandwidth h
output: Node Set \mathcal{N}

// Set the number of radius

- 1 $\Psi \leftarrow (i \frac{2\pi}{r})_{i \in [0, r]}$;
- 2 $\mathcal{N} \leftarrow \{\}$;
- 3 **foreach** $\psi \in \Psi$ **do**
- 4 $\mathcal{I}_\psi \leftarrow \{\}$;
- 5 **foreach** $i \in [0, |SProj| - 1]$ **do**
- 6 // Find intersected points
- 6 $radius \leftarrow \max_{x,y}(SProj_i, SProj_{i+1})$;
- 7 $P_\psi \leftarrow (radius_x \cdot \cos(\psi), radius_y \cdot \sin(\psi))$;
- 8 add $Intersect((\Omega, P_\psi), (SProj_i, SProj_{i+1}))$ in \mathcal{I}_ψ
- 9 **end**
- 10 // Extract Nodes
- 10 $\mathcal{N}_\psi \leftarrow \operatorname{argmax}_{x \in \mathcal{I}_\psi} f_h(x, \mathcal{I}_\psi)$;
- 11 add \mathcal{N}_ψ in \mathcal{N} ;
- 12 **end**

$|\mathcal{I}_\psi| \leq |SProj|$). Actually, we experimentally observed that $|\mathcal{I}_\psi| \ll |SProj|$. Therefore, the overall time complexity is bounded by $O(|T|r)$. We can improve this complexity using the following observation. Instead of checking the intersection with every possible radius, we can select those that bound the position of the points i and $i+1$ in $SProj$ (only the radius with ψ between $\psi_i = \angle \vec{u}_x \cdot SProj_i$ and $\psi_{i+1} = \angle \vec{u}_x \cdot SProj_{i+1}$). Therefore, the worst case complexity becomes $O(|T|r)$, and the best case complexity is reduced to $O(|T|)$.

4.3 Edge Creation

Once we identify the nodes, we need to compute the edges among them. Recall that the set of extracted nodes corresponds to all the possible states, where subsequences of the data series T can be. In order to compute the transitions among these nodes, we loop through the entire projection $SProj(T, \ell, \lambda)$ and we extract the sequence $\langle N^{(0)}, N^{(1)}, \dots, N^{(n)} \rangle$ of the nodes N_i in \mathcal{N} that the embedded subsequences $(SProj(T, \ell, \lambda)_0, SProj(T, \ell, \lambda)_1, \dots, SProj(T, \ell, \lambda)_n)$ belong to. Intuitively, the above node sequence involves all the nodes in \mathcal{N} (some of them more than once) and represents the entire input data series. We use this sequence to identify the edges of the graph G_ℓ we want to construct. In practice, we extract the edges (all the pairs of successive nodes in the above sequence) and set their weights to the number of times the edge is observed in the sequence. Formally, the edges set \mathcal{E} is defined as follows.

DEFINITION 8 (PATTERN EDGES SET). *Given a data series T , its projection $P = SProj(T, \ell, \lambda)$ and its Pattern Node Set \mathcal{N} , the edges set \mathcal{E} is equal to: $\mathcal{E} = \{(S(P_i), S(P_{i+1}))\}_{i \in [1, |P| - 1]}$, where function S finds for a given projection point, the closest node in \mathcal{N} . Formally: $S(x) = \operatorname{argmin}_{n \in \mathcal{N}} d(x, n)$, where $x \in P$ and d is the geometrical Euclidean distance.*

Since the weight of each edge is equal to the cardinality of this edge in the edge set \mathcal{E} , this weight is proportional to the number of times two subsequences follow each other in the input data series. For efficiency, $S(x)$ is computed as follows:

Algorithm 3: Edge Extraction

input : 2-dimension points sequence $SProj$, and a node set \mathcal{N}
output: a Edge Set \mathcal{E}

```
1  $\Psi \leftarrow (i \frac{2\pi}{r})_{i \in [0, r]}$ ;  
2  $NodeSeq \leftarrow []$ ;  
3 foreach  $i \in [0, |SProj| - 1]$  do  
    // Get the radius that bound  $SProj_i$  and  $SProj_{i+1}$   
4  $\psi_i \leftarrow getAngle(u_x, SProj_i)$ ;  
5  $\psi_{i+1} \leftarrow getAngle(u_x, SProj_{i+1})$ ;  
6 foreach  $(\psi \in \Psi) \wedge (\psi \in [\psi_i, \psi_{i+1}])$  do  
    // Fill the sequence of node  $NodeSeq$   
7  $radius \leftarrow max_{x,y}(SProj_i, SProj_{i+1})$ ;  
8  $P_\psi \leftarrow (radius_x \cdot \cos(\psi), radius_y \cdot \sin(\psi))$ ;  
9  $x_{int} \leftarrow Intersect((\Omega, P_\psi), (SProj_i, SProj_{i+1}))$ ;  
10  $n_{int} \leftarrow argmin_{n \in \mathcal{N}_\psi} (|x_{int} - n|)$ ;  
11 add  $n_{int}$  in  $NodeSeq$   
12 end  
    // Extract edges from  $NodeSeq$   
13  $\mathcal{E} \leftarrow \{(NodeSeq_i, NodeSeq_{i+1})\}_{i \in [0, |fullPath|]}$ ;  
14 end
```

for a given projection point, we first find the node subset \mathcal{N}_ψ of \mathcal{N} (with $\psi \in \Psi$), such that $|\angle \vec{x} \vec{u}_\psi|$ is minimal. We then compute $S(x)$ such as $S(x) = argmin_{n \in \mathcal{N}_\psi} |x \cdot \vec{u}_\psi - n|$, where $\vec{x} \cdot \vec{u}_\psi$ is the scalar product between \vec{x} and \vec{u}_ψ . As depicted in Figures 3(a) and (b), a total of $n_1 + n_2$ subsequences are intersected by \mathcal{I}_ψ and represented by node N_ψ^0 . At $\mathcal{I}_{\psi+1}$, these subsequences are divided between nodes $N_{\psi+1}^0$ (n_1 subsequences) and $N_{\psi+1}^1$ (n_2 subsequences). Therefore, we have $w(N_\psi^0, N_{\psi+1}^0) = n_1$ and $w(N_\psi^0, N_{\psi+1}^1) = n_2$.

Algorithm 3 outlines the steps we follow to extract the edges among the nodes in \mathcal{N} . For each point in the input data series T , we identify the radius it belongs to and we choose the closest node. Therefore, the complexity is bounded by $O(|T|)$ and varies based on the number of radius we have to check and the number of nodes in each \mathcal{N}_ψ . The former is bounded by parameter r : on average, we have no more than $|T|/r$ points per \mathcal{N}_ψ . The overall complexity is in the worst case $O(|T|^2)$, and in the best case $O(|T|)$. We note that this worst case corresponds to the situation where each subsequence in T belongs to a different node. This is not what we observe in practice: for all our datasets, the overall complexity is close to the best case.

4.4 Subsequence Scoring

We now describe how we can use the information in the graph to identify the normal and anomalous behaviors.

[Subsequence to Path Conversion] We start with the conversion of a subsequence to a path in a given graph. For an already computed graph $G_\ell(\mathcal{N}, \mathcal{E})$, we define function $Time2Path(G_\ell, T_{i, \ell_q})$ that converts a subsequence T_{i, ℓ_q} into a path (i.e., a sequence of nodes) in G_ℓ , by (i) computing the pattern embedding SP of T_{i, ℓ_q} (using the PCA transformation and rotation matrices, computed in Lines 8 and 12, respectively, of Algorithm 1), and (ii) extracting the edges using $EdgeExtraction(SP, \mathcal{N})$ (output of Algorithm 3) on the node set \mathcal{N} of graph G_ℓ .

[Normality Score] We are now ready to measure normality. As mentioned earlier, modeling a data series using a cyclic graph results in the graph encoding information con-

cerning the recurrence of subsequences. Then, the path normality score function can be defined as follows.

DEFINITION 9 (PATH NORMALITY SCORE). *Given a data series T and its graph $G_\ell(\mathcal{N}, \mathcal{E})$, and a subsequence T_{i, ℓ_q} of length $\ell_q \geq \ell$, the normality of a path $P_{th} = Time2Path(G_\ell(\mathcal{N}, \mathcal{E}), T_{i, \ell_q}) = \langle N^{(i)}, N^{(i+1)}, \dots, N^{(i+\ell_q)} \rangle$ is equal to: $Norm(P_{th}) = \sum_{j=i}^{i+\ell_q-1} \frac{w(N^{(j)}, N^{(j+1)}) \cdot (deg(N^{(j)}) - 1)}{\ell_q}$*

We can thus infer a normality score for subsequences in T using the $Time2Path$ function defined earlier (the opposite of this normality score is the anomaly score). Formally, the normality score is defined as follows.

DEFINITION 10 (SUBSEQUENCE NORMALITY SCORE). *Given a data series T , its graph $G_\ell(\mathcal{N}, \mathcal{E})$ and a subsequence T_{i, ℓ_q} of length $\ell_q \geq \ell$, the Normality score T_{i, ℓ_q} is equal to: $Normality(T_{i, \ell_q}) = Norm(Time2Path(G_\ell(\mathcal{N}, \mathcal{E}), T_{i, \ell_q}))$*

Observe that the two previous definitions are consistent with the definition of θ -Normality, such that every P_{th} in θ -Normal subgraph will have $N(P_{th}) \geq \theta$, and every P_{th} that is exclusively in a lower normality level will have $N(P_{th}) \leq \theta$. As a matter of fact, the rank generated by the normality score is similar to the θ -Normality rank, and in both rankings, the anomalies are found at the bottom of the ranking. The following lemma formalizes this statement.

LEMMA 1. *Given a data series T , its graph $G_\ell(\mathcal{N}, \mathcal{E})$, a subsequence T_{i, ℓ_q} , and its path $P_{th} = Time2Path(G_\ell(\mathcal{N}, \mathcal{E}), T_{i, \ell_q})$, we have: $\forall \theta \in \mathbb{N}_{>0}, N(P_{th}) < \theta \implies P_{th} \in \theta$ -Anomaly(T)*

PROOF. Consider a subsequence T_{i, ℓ_q} corresponding to $P_{th} = \langle N^{(i)}, N^{(i+1)}, \dots, N^{(i+\ell_q)} \rangle$. If $P_{th} \in \theta$ -Normality(T), then according to Definition 5, we have:

$$\begin{aligned} & \forall j \in [i, i + \ell_q - 1], (N^{(j)}, N^{(j+1)}) \in \theta\text{-Normality}(T) \\ \implies & \forall j \in [i, i + \ell_q - 1], w(N^{(j)}, N^{(j+1)}) \cdot (deg(N^{(j)}) - 1) \geq \theta \\ \implies & \sum_{j=i}^{i+\ell_q-1} \frac{w(N^{(j)}, N^{(j+1)}) \cdot (deg(N^{(j)}) - 1)}{\ell_q} \geq \theta \end{aligned}$$

As a consequence, $P_{th} \notin \theta$ -Normality(T) and according to Definition 5, $P_{th} \in \theta$ -Anomaly(T). \square

Therefore, the subsequences of T with a low score are those that compose the θ -Anomaly subgraph, where the value of θ is low (close to one for the discords). We note that this process identifies both single anomalies (discords) and recurrent anomalies.

4.5 Series2Graph Summary

Algorithm 4 summarizes all the steps of our approach. In Lines 1-3, we compute the subsequence embedding, the node set \mathcal{N} and then the edge set \mathcal{E} in order to build the graph G_ℓ . Line 7 computes the *Normality* score for all subsequences of the input data series: we use a sliding window over the input data series to extract all subsequences, we score each one of them and store the result in the vector $NormalityScore = [Normality(T_{0, \ell_q}), \dots, Normality(T_{n-\ell_q, \ell_q})]$, initialized in Line 5. Finally, we apply a moving average filter on the *NormalityScore* vector (Line 9). This filter tries to rectify possible small inaccuracies of the scoring function, by ensuring that two highly overlapping subsequences will have similar *Normality* scores (as we would normally expect).

Algorithm 4: Series2Graph

```
input : data series  $T$ , input length  $\ell$ , query length  $\ell_q$ 
output: a data series  $AnomScore$ 
1  $SProj \leftarrow PatternEmbedding(T, \ell, \lambda)$ ;
2  $\mathcal{N} \leftarrow NodeExtraction(SProj, r = 50, h = h_{opt})$ ;
3  $\mathcal{E} \leftarrow EdgeExtraction(SProj, \mathcal{N})$ ;
4  $G_\ell \leftarrow Graph(\mathcal{N}, \mathcal{E})$ ;
   // vector of Normality scores
5  $NormalityScore \leftarrow [0]_{0, |T| - \ell_q}$ ;
   // compute Normality score for all subsequences of
   // length  $\ell_q$  in  $T$ 
6 foreach  $i \in [1, |T| - \ell_q]$  do
7    $NormalityScore[i] \leftarrow Norm(Time2Path(G_\ell, T_i, \ell_q))$ ;
8 end
9  $NormalityScore \leftarrow movingAverage(NormalityScore, \ell)$ ;
```

5. EXPERIMENTAL EVALUATION

We now present the results of the experimental evaluation with several real datasets from different domains, including all the annotated datasets that have been used in the discord discovery literature. In order to ensure the reproducibility of our experiments, we created a web page [1] with the source code, datasets, and other supporting material.

5.1 Experimental setup and Datasets

We have implemented our algorithms [11] in C (compiled with gcc 5.4.0). The evaluation was conducted on a server with Intel Xeon CPU E5-2650 v4 2.20GHz, and 250GB of RAM.

We benchmark our approach using different annotated real and synthetic datasets, listed in Table 2. Following previous work [52], we use several synthetic datasets that contain sinusoid patterns at fixed frequency, on top of a random walk trend. We then inject different numbers of anomalies, in the form of sinusoid waveforms with different phases and higher than normal frequencies, and add various levels of Gaussian noise on top. We refer to those datasets using the label SRW-[# of anomalies]-[% of noise]-[length of anomaly] and use them in order to test the performance of the algorithms under different, controlled conditions.

Our real datasets are: (i) Simulated engine disks data (SED) collected by the Rotary Dynamics Laboratory at NASA [3, 4]. This data series represents disk revolutions recorded over several runs. (ii) MIT-BIH Supraventricular Arrhythmia Database (MBA) [21, 42], which are electrocardiogram recordings from 5 different patients, containing multiple instances of two different kinds of anomalies. (iii) Five additional real datasets from various domains that have been studied in earlier works [33, 51, 51], and their anomalies are simple *discords* (usually only 1): aerospace engineering (Space Shuttle Marotta Valve [33]), gesture recognition (Ann’s Gun dataset [51]) and medicine (patient’s respiration measured by the thorax extension [33], and the record 15 of the BIDMC Congestive Heart Failure Database [33]).

We measure *Top-k* accuracy (i.e., the correctly identified anomalies among the k retrieved by the algorithm, divided by k) and execution time.

5.2 Length Flexibility

We first evaluate the influence of the subsequence length parameters ℓ and ℓ_q . Ideally, we would like the method to be robust to variations in these parameters. This will

Table 2: List of dataset characteristics: data series length, anomaly length (ℓ_A), number of annotated anomalies (N_A) and domain.

Datasets	Length	ℓ_A	N_A	Domain
SED	100K	75	50	Electronic
MBA (803)	100K	75	62	Cardiology
MBA (804)	100K	75	30	Cardiology
MBA (805)	100K	75	133	Cardiology
MBA (806)	100K	75	27	Cardiology
MBA (820)	100K	75	76	Cardiology
MBA (14046)	100K	75	142	Cardiology
Marotta Valve	20K	1K	1	Aerospace engineering
Ann Gun	11K	800	1	Gesture recognition
Patient Respiration	24K	800	1	Medicine
BIDMC CHF	15K	256	1	Cardiology
SRW-[20-100]-[0%]-[200]	100K	200	var.	Synthetic
SRW-[60]-[5%-25%]-[200]	100K	200	60	Synthetic
SRW-[60]-[0%]-[100-1600]	100K	var.	60	Synthetic




Figure 4: (a) MBA ECG recording (4000 points snippet from patient 803), with one annotated Supraventricular contraction (V). (b) Euclidean distances of each subsequence (for the entire MBA ECG recording) of length 80 (green) and 90 (red). In this two cases, the subsequences with the highest distances are not the same (length 90: Normal Beat, length 80: Anomaly Type V).

ensure that the constructed model can accurately identify anomalies, even if these anomalies are of several lengths, significantly different than what was originally expected by the users. We stress that, in contrast, all previous techniques [51, 60, 58, 14, 40] require knowledge of the *exact* anomaly length, and are very brittle otherwise.

Previous approaches developed for discord discovery, like *STOMP*[60], are very sensitive to even slight variations of the subsequence length parameter (for the identification of the anomalous subsequences). Figure 4 displays the Euclidean distances of each subsequence of the MBA ECG records 803 to its nearest neighbor (computed using *STOMP*), using lengths 80 and 90, given the length of the anomaly $\ell_A = 80$. The results demonstrate that a small variation in the input length can lead to very different outcomes (using a length equal to 90, the discord is a normal heart beat, and therefore a false positive).

On the other hand, *Series2Graph* is robust to variations in the query subsequence length. Figure 5 depicts the $G_\ell(\mathcal{N}, \mathcal{E})$ graphs for ℓ equals to 80, 100, 120, while the anomalies length is 75 for Type S anomalies, and 120 for type V anomalies. The results show that in all cases, irrespective of the




Figure 5: $G_\ell(\mathcal{N}, \mathcal{E})$ of the MBA(820) electrocardiogram data series for ℓ of 80, 100 and 120. In the three cases, the different kinds of anomalies (S: blue, V: red) are well separable with lower edges weights.

length ℓ used to construct the graph, the anomaly trajectories (Type V highlighted in red and Type S highlighted in blue) are distinct from the highly-weighted trajectories (thick black) that correspond to normal behavior.

In order to confirm this observation, we conduct several experiments. First, we measure the *Top-k* accuracy as the input length ℓ and the query length ℓ_q vary (using a query length $2\ell_q/3 = \ell$, with the anomaly length $\ell_A = 80$). Figure 6(a) demonstrates the stable behavior of Series2Graph. Even though the *Top-k* accuracy varies for small lengths, the performance remains relatively stable when the input lengths ℓ we use to construct the graph are larger than the anomaly length ℓ_A . This means that simply selecting an ℓ value larger than the expected anomaly length ℓ_A will lead to good performance.

In contrast, as Figure 6(b) demonstrates, the performance of STOMP (a discord-based approach) varies widely. Thus, such approaches need careful tuning, requiring domain expertise and good knowledge of the possible anomaly lengths. Furthermore, even though STOMP accuracy seems to converge to a stable value as the length is increasing, the Series2Graph accuracy stays significantly higher and much more stable in average, as shown in Figure 6(c).

5.3 Optimal Bandwidth

We now evaluate the effectiveness of the kernel bandwidth h_{scott} in $f_h(x, \mathcal{I})$ in the node extraction step. We set a constant value for ℓ and ℓ_q ($\ell = 80$, $\ell_q = 160$), and we measure the accuracy for different bandwidths. Figure 7(a) displays the *Top-k* accuracy on the MBA and SED datasets as a function of $h/\sigma(\mathcal{I}_\theta)$ (logarithmic scale). As expected, a small bandwidth ratio breaks down too much the normal pattern, and therefore reduces its Normality score, while a large bandwidth ratio (above 0.7) hinders some key nodes to detect anomalies in two of the six datasets, namely MBA(806) and MBA(820). The anomalies in these two datasets are close to the normal behavior, thus the abnormal trajectories can be easily missed. In contrast, using the Scott bandwidth ratio h_{scott} (marked with the dotted line) leads to very good accuracy for all the datasets we tested (We used the datasets with the same anomaly and pattern lengths, so that we can compare Scott bandwidth ratios).

5.4 Convergence of Edge Set

In this section, we evaluate Series2Graph accuracy on previously unseen data series (i.e., on different data series than

the one used to build the graph G , with potentially different anomalies). Note that the Normality score of a non-existing pattern in G will be 0 (see Definitions 9 and 10) and therefore close to the score of the anomalies in G .

In the experiment of Figure 7(b), we build the graph using only a prefix of the input series. We then vary the size of this prefix and measure the *Top-k* accuracy for the entire series. As we can see, the *Top-k* accuracy usually reaches its maximum value without having to build the graph using all the available data. We observe that on average *Top-k* accuracy already reaches more than 85% of its maximum value, when we use as little as 40% of the input data series. Nevertheless, one can also see that the *Top-k* accuracy of MBA(820) and MBA(806) converge slower than the other datasets. These two datasets contain anomalies of Type S, which means they are very similar to a normal heartbeat. Therefore, Series2Graph requires more subsequences in order to build a model that effectively separates the anomalous behavior from the normal behavior.

Finally, Figure 7(c) demonstrates the stability of Series2Graph in the anomaly discovery task as we vary the query subsequence length ℓ_q using a fixed input length ℓ . The results show that we can identify anomalies with very high accuracy, by scoring candidate subsequences of a wide range of lengths, provided they are larger than the length of the anomalies ($\ell_q \geq \ell_A$).

5.5 Discord Identification

The next experiments evaluate the capability of Series2Graph to identify discords. We used the datasets that have appeared in the discord discovery literature (mentioned above). Figure 8 shows the graphs obtained for Ann Gun, Marotta Valve, Patient’s respiration, and Record 15 of the BIDMC Congestive Heart Failure Database. The thickness of the lines correspond to the weights of the edges. This figure shows that the discords of these datasets (red subsequences) always correspond to trajectories with low weights (θ -Anomaly, for a small θ value), whereas the normal subsequences (green subsequences) correspond to trajectories with high weights (θ -Normality, for a large θ). Therefore, the anomaly scores of the discords are in every case the largest, and hence are correctly identified.

5.6 Accuracy Evaluation

In this section, we report the anomaly detection accuracy results. We compare Series2Graph to the current state-of-the-art data series subsequence anomaly detection algorithms, using $\ell_q = \ell_A$. For Series2Graph, we use the same $\ell = 50$ and latent $\lambda = 16$ for all datasets, even though different values would be optimal in every case, thus demonstrating the robustness of the approach. We also evaluate Series2Graph’s learning ability by comparing the accuracy obtained by building the graph using only the first half of the data series ($S2G_{|T|/2}$), compared to using the full data series ($S2G_{|T|}$). We consider two techniques that enumerate *Top-k* 1^{st} discords, namely GrammarViz (GV) [51] and STOMP [60]. Moreover, we compare Series2Graph against the Disk Aware Discord Discovery algorithm (DAD) [58], which finds m^{th} discords. We also compare to Local Outlier Factor (LOF) [14] and Isolation Forest [37]. Note that the last two methods are not specific to subsequence anomaly discovery. Finally, we include in our comparison LSTM-AD [40], a *supervised* deep learning technique. We stress




Figure 6: On the MBA and SED datasets: (a) $Top-k$ accuracy of Series2Graph varying the input and query length ($2\ell_q/3 = \ell$) to build the graph G_ℓ . (b) STOMP $Top-k$ accuracy varying the input length ℓ , (c) STOMP and Series2Graph in $Top-k$ accuracy average compared to the input length ℓ .




Figure 7: $Top-k$ accuracy of Series2Graph on the MBA and SED datasets while: (a) varying the bandwidth ratio $h/\sigma(\mathcal{I}_\psi)$ in $f_h(x, \mathcal{I})$ (logarithmic scale), (b) varying the length of the prefix snippet used to build the graph, (c) varying the query length ℓ_q (starting from the input length ℓ) to compute the normality score.

that the comparison to LSTM-AD is not fair to all the other techniques: LSTM-AD has to first train on *labeled* data, which gives it an unfair advantage; all the other techniques are *unsupervised*. We include it to get an indication as to how the unsupervised techniques compare to a state-of-the-art supervised anomaly detection algorithm.

In Table 3, we show the $Top-k$ accuracy. We set k equal to the number of anomalies. These experiments test the capability of each method to correctly retrieve the k anomalous subsequences in each dataset. For Series2Graph, we simply have to report the $Top-k$ anomalies that Algorithm 4 produces. For the *discord* based techniques, we have to consider the $Top-k$ 1st *discord* and the m th *discord* (with $m = k$). Finally, LSTM-AD marks as anomalies the subsequences that have the largest errors (distances) to the sequences that the algorithm predicts; we compute accuracy considering the subsequences with the k largest errors.

In the first section of Table 3, we report the results of all techniques on the annotated real datasets with multiple (diverse and similar) anomalies. Series2Graph (both built on half and full dataset) is clearly the winner, with nearly perfect accuracy. As expected, $Top-k$ 1st *discord* techniques (GV and STOMP) have in most of the cases lower accuracy than Series2Graph, since anomalies do not correspond to rare subsequences (i.e., isolated discords). We also observe that the m th *discord* technique (DAD), which is able to detect groups of m similar anomalous subsequences, does not perform well. This is due to the many false positives produced by the algorithm.

In the rest of Table 3, we report the accuracy of the evaluated methods on all the synthetic datasets (where we vary

the number of anomalies, the % of Gaussian noise and the anomaly subsequence length ℓ). We note that the accuracy of the *discord* discovery techniques substantially improves since in this case, most anomalies correspond to rare and isolated subsequences (i.e., different from one another). Even in these cases though, Series2Graph is on average significantly more accurate than the competitors. Moreover, in contrast to GV, STOMP and DAD, we observe that Series2Graph’s performance is stable as the noise increases between 0%-25%.

Regarding LSTM-AD, we note that, in general, it is more accurate than the *discord* based algorithms. However, LSTM-AD cannot match the performance of Series2Graph, and in some cases it completely misses the anomalies (i.e., for the SED, MBA(806) and MBA(820) datasets). Regarding LOF, we observe that it does not perform well. Isolation Forest on the other hand achieves a surprisingly good accuracy, which makes it a strong competitor.

Overall, we observe that regular Series2Graph ($S2G_{|T|}$) is considerably more accurate than all competitors (with rare exceptions, for which its performance is still very close to the best one), in all the settings we used in our evaluation.

5.7 Scalability Evaluation

We now present scalability tests (we do not consider LSTM-AD, since supervised methods have a completely different way of operation and associated costs, e.g., data labeling and model training). In Figures 9(a,b,c), we report the execution time of the algorithms (log scale) versus the size of the dataset. We use several prefix snippets (50K, 100K, 500K, 1M, 2M points) of the MBA(14046) dataset, a

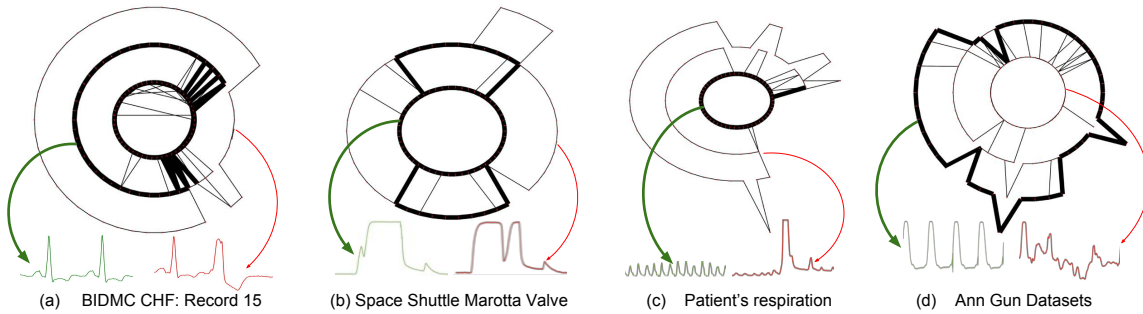


Figure 8: (a) G_{80} for BIDMC Congestive Heart Failure Database: Record 15, (b) G_{200} for Space Shuttle Marotta Valve (TEK 16), (c) G_{50} Patient’s respiration, (d) G_{150} Ann Gun Datasets. Green subsequences belong to θ -Normality subgraph with large θ . Red subsequences belong to θ -Anomaly subgraph with small θ .

Table 3: $Top-k$ accuracy for DAD, STOMP, GrammarViz, LSTM-AD, $S2G_{|T|/2}$ (Series2Graph built using half of the dataset) and $S2G_{|T|}$ (Series2Graph built using the entire dataset) with k equal to number of anomalies.

Datasets	GV	STOMP	DAD	LOF	Isolation Forest (std)	LSTM-AD	$S2G_{ T /2}$	$S2G_{ T }$
SED	0.46	0.73	0.44	0.65	0.65 (0.02)	0.10	1.00	1.00
MBA (803)	0.15	0.60	0.01	0.08	1.00 (0.00)	0.35	1.00	1.00
MBA (805)	0.09	0.10	0.03	0.42	0.99 (0.01)	0.85	0.99	0.99
MBA (806)	0.01	0.59	0.66	0.92	0.75 (0.06)	0.10	0.96	1.00
MBA (820)	0.05	0.92	0.04	0.42	0.92 (0.03)	0.09	0.76	0.96
MBA (14046)	0.09	0.54	0.71	0.64	0.99 (0.01)	1.00	0.95	0.95
SRW-[20]-[0%]-[200]	1.0	0.77	0.55	0.74	0.75 (0.05)	0.94	0.95	1.00
SRW-[40]-[0%]-[200]	0.975	1.0	0.05	0.89	0.92 (0.02)	1.00	1.00	1.00
SRW-[60]-[0%]-[200]	0.96	0.88	0.10	0.76	0.87 (0.02)	0.92	1.00	1.00
SRW-[80]-[0%]-[200]	0.96	0.43	0.14	0.82	0.86 (0.01)	0.95	0.98	1.00
SRW-[100]-[0%]-[200]	0.95	0.99	0.11	0.75	0.92 (0.02)	1.00	1.00	1.00
SRW-[60]-[5%]-[200]	1.0	0.73	0.21	0.88	0.89 (0.01)	0.96	1.00	1.00
SRW-[60]-[10%]-[200]	0.83	0.98	0.01	0.70	0.80 (0.01)	0.94	0.96	0.98
SRW-[60]-[15%]-[200]	0.76	0.62	0.17	0.66	0.82 (0.01)	0.94	0.98	0.98
SRW-[60]-[20%]-[200]	0.73	1.0	0.01	0.73	0.85 (0.02)	0.96	1.00	1.00
SRW-[60]-[25%]-[200]	0.63	0.64	0.09	0.67	0.80 (0.01)	0.83	0.98	0.98
SRW-[60]-[0%]-[100]	0.98	1.0	0.23	0.74	0.88 (0.02)	1.00	0.96	0.96
SRW-[60]-[0%]-[200]	0.96	0.60	0.19	0.85	0.83 (0.01)	1.00	0.98	0.98
SRW-[60]-[0%]-[400]	0.98	1.0	0.63	0.76	0.88 (0.01)	0.88	0.96	0.96
SRW-[60]-[0%]-[800]	0.91	0.86	-	0.69	0.87 (0.01)	0.76	0.95	0.98
SRW-[60]-[0%]-[1600]	1.0	1.0	-	0.52	0.64 (0.02)	0.90	0.91	0.94
Average	0.62	0.73	0.24	0.68	0.85	0.78	0.96	0.98

2 million points concatenated version of the Marotta Valve dataset, and a 2 million points concatenated version of SED dataset. For all three datasets, k is set to be equal to the number of anomalies in each snippet. We observe that Series2Graph is faster than the competitors, especially when both the input series length ℓ and anomaly subsequence length ℓ_A take large values, as in Figure 9(b), and gracefully scales with the dataset size.

We also measure the execution time of the algorithms (log scale) as we vary the number of anomalies. We use the MBA(14046) dataset, as well as the synthetic datasets SRW-[20-100]-[0%]-[200] (Figure 9(e)). As expected, we observe that the performance of Series2Graph is not influenced by the number of anomalies. Similarly, STOMP and IF are not affected either, but GrammarViz, LOF and DAD are negatively impacted by the number of anomalies.

Finally, Figure 9(f) depicts the time performance results as we vary the length of the anomalies between 100-1600 points in the synthetic data series SRW-[60]-[0%]-[100-1600]. The performance of STOMP is constant, because its complexity is not affected by the (anomaly) subsequence length. Moreover, we note that GV, IF, LOF and DAD perform poorly as the length of the anomalies is increasing. Observe

that the execution time of Series2Graph increases slightly for larger subsequence lengths. This is due to the scoring function (last step of the algorithm). This function sums up all the edge weights of the subsequences we are interested in. Therefore, if the subsequence is large, the number of relevant edges is large as well, which slightly affects computation time. Nevertheless, Series2Graph remains the fastest algorithm among all competitors.

6. RELATED WORK

[Discord Discovery] The problem of subsequence anomaly discovery has been studied by several works that use the *discord* definition [60, 51, 32, 38, 23, 15, 39, 36]. In these studies, anomalies are termed the isolated subsequences, i.e., the ones that have the highest Euclidean distances to their nearest neighbors. The proposed solutions either operate directly on the raw values [60, 38, 23, 39], or on discrete representations of the data, e.g., Haar wavelets [23, 15], or SAX [32, 51]. Even though the information that the discords carry is interesting and useful for some applications, these approaches (that are based on the *discord* definition)

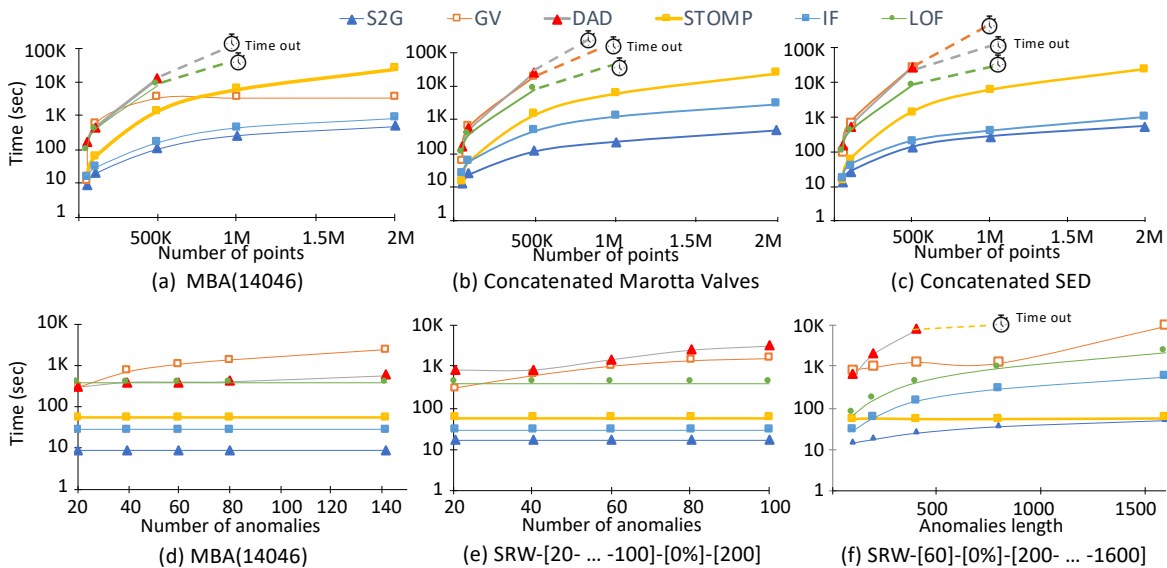


Figure 9: Execution time vs data size (a-c), number of anomalies (d,e), anomaly length (f). Time out: 8h.

fail when the dataset contains multiple anomalies that are similar to one another.

[Multiple *discord* Discovery] The notion of m^{th} *discord* has been proposed in order to solve the issue of multiple similar anomalies [59]. The approach described in this study finds the sequence that has the furthest m^{th} nearest neighbor in Euclidean space. During the search, a space pruning strategy based on the intermediate results of the simple *discord* discovery is applied. As we have already discussed, the m^{th} -*discord* definition fixes the main problem of simple *discord*, but is very sensitive to the m parameter and can lead to false positives.

[Outlier Detection] Local Outlier Factor [14] is a degree of being an outlier assigned to a multidimensional data point. This degree depends on how much the data point is isolated (in terms of distance) to the surrounding neighborhood. Similarly, Isolation Forest [37] is a classical machine learning technique that isolates anomalies instead of modeling normality. It first builds binary trees with random splitting nodes to partition the dataset. The anomaly score is defined as a function of the averaged path lengths between a particular sample and the roots of the trees. The above two methods are not specifically targeted to data series subsequences anomaly detection, which is reflected in the low accuracy they achieve in several of the datasets we tested.

[Deep Learning Approaches] Subsequence anomaly detection has also been studied in the context of supervised deep learning techniques, with the use of Long Short Term Memory (LSTM) architectures [29]. The studies that use this recurrent neural network are based on a forecasting model [40, 12]. First, the LSTM network is trained using the data segments that do *not* contain anomalies. Then, the sequence is examined and the LSTM network is used to forecast the subsequent values: when the error between the forecast and the real value is above some threshold, the subsequence is classified as an anomaly. The system learns the threshold using the validation set, picking the value that maximizes the F1-score of the classification results. While the aforementioned approach originally used

the annotated anomalies to learn the threshold, the LSTM model has also been used in a zero positive learning framework, where the annotated anomalies are not necessary for the training phase [35].

[Phase Space Reconstruction] Phase space reconstruction is a technique that has been used for pattern embedding and for non-linear data series analysis [30, 13]. This technique transforms the data series into a set of vectors and has been used to visualize the evolution of the data series. Previous studies have also proposed the construction of a complex network based on the phase space reconstruction of a data series [24, 25], or on its visibility graph [34], which can then be used to identify different patterns of interest. Series2Graph shares the same goal of converting the data series into a graph in order to reveal significant features. However, differently from the above methods that convert each point of the series into a separate node, Series2Graph uses a single node to represent several subsequences.

7. CONCLUSIONS AND FUTURE WORK

Even though subsequence anomaly detection in data series has attracted a lot of attention, existing techniques have several shortcomings. In this work, we describe a novel approach, based on a graph representation of the data series, which enables us to detect both single and recurrent anomalies (as well as the normal subsequences), in an unsupervised and domain-agnostic way. Experiments with several real datasets demonstrate the benefits of our approach in terms of both efficiency and accuracy. As future work, we plan to use modern data series indices [19, 20, 49, 48, 45, 26, 18] for accelerating the operation of Series2Graph, extend our approach to operate on streaming and multivariate data, and compare to the recently proposed NorM approach [9, 10].

Acknowledgments: We thank M. Meftah and E. Remy, from EDF R&D, for their support and advice, and the reviewers for the numerous constructive comments. Work partially supported by EDF R&D and ANRT French program.

References

- [1] Series2Graph Webpage. <http://helios.mi.parisdescartes.fr/~themisp/series2graph/>, 2020.
- [2] D. Abboud, M. Elbadaoui, W. Smith, and R. Randall. Advanced bearing diagnostics: A comparative study of two powerful approaches. *MSSP*, 114, 2019.
- [3] A. Abdul-Aziz, M. R. Woike, N. C. Oza, B. L. Matthews, and J. D. Iekki. Rotor health monitoring combining spin tests and data-driven anomaly detection methods. *Structural Health Monitoring*, 2012.
- [4] M. Ali Abdul-Aziz, N. Woike, B. Oza, Matthews, and G. Baakilini. Propulsion health monitoring of a turbine engine disk using spin test data, 2010.
- [5] J. Antoni and P. Borghesani. A statistical methodology for the design of condition indicators. *Mechanical Systems and Signal Processing*, 2019.
- [6] A. J. Bagnall, R. L. Cole, T. Palpanas, and K. Zoumpatianos. Data series management (dagstuhl seminar 19282). *Dagstuhl Reports*, 9(7), 2019.
- [7] S. Bahaadini, V. Noroozi, N. Rohani, S. Coughlin, M. Zevin, J. Smith, V. Kalogera, and A. Katsagelos. Machine learning for gravity spy: Glitch classification and dataset. *Information Sciences*, 444:172–186, 5 2018.
- [8] V. Barnett and T. Lewis. *Outliers in Statistical Data*. John Wiley and Sons, Inc., 1994.
- [9] P. Boniol, M. Linardi, F. Roncallo, and T. Palpanas. Automated anomaly detection in large sequences. In *ICDE*, 2020.
- [10] P. Boniol, M. Linardi, F. Roncallo, and T. Palpanas. SAD: an unsupervised system for subsequence anomaly detection. In *ICDE*, 2020.
- [11] P. Boniol, T. Palpanas, M. Meftah, and E. Remy. Graphan: Graph-based subsequence anomaly detection. *PVLDB*, 13(11), 2020.
- [12] L. Bontemps, V. L. Cao, J. McDermott, and N. LeKhac. Collective anomaly detection based on long short term memory recurrent neural network. *CoRR*, abs/1703.09752, 2017.
- [13] E. Bradley and H. Kantz. Nonlinear time-series analysis revisited. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 2015.
- [14] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander. Lof: Identifying density-based local outliers. In *SIGMOD*, 2000.
- [15] Y. Bu, O. T. Leung, A. W. Fu, E. J. Keogh, J. Pei, and S. Meshkin. WAT: finding top-k discords in time series database. In *SIAM*, 2007.
- [16] B. Y. Chiu, E. J. Keogh, and S. Lonardi. Probabilistic discovery of time series motifs. In *SIGKDD 2003*, pages 493–498, 2003.
- [17] N. Daouayry, A. Mechouche, P.-L. Maisonneuve, V.-M. Scuturici, and J.-M. Petit. Data-centric helicopter failure anticipation: The mgb oil pressure virtual sensor case. *IEEE BigData*, 2019.
- [18] K. Echihabi, K. Zoumpatianos, and T. Palpanas. Scalable Machine Learning on High-Dimensional Vectors: From Data Series to Deep Network Embeddings. In *WIMS*, 2020.
- [19] K. Echihabi, K. Zoumpatianos, T. Palpanas, and H. Benbrahim. The lernaean hydra of data series similarity search: An experimental evaluation of the state of the art. *PVLDB*, 2019.
- [20] K. Echihabi, K. Zoumpatianos, T. Palpanas, and H. Benbrahim. Return of the Lernaean Hydra: Experimental Evaluation of Data Series Approximate Similarity Search. *PVLDB*, 2019.
- [21] G. et al. Physiobank, physiotookit, and physionet. *Circulation*.
- [22] Y. Z. et al. Matrix profile II: exploiting a novel algorithm and gpus to break the one hundred million barrier for time series motifs and joins. In *ICDM 2016*.
- [23] A. W. Fu, O. T. Leung, E. J. Keogh, and J. Lin. Finding time series discords based on haar transform. In *ADMA*, 2006.
- [24] Z.-K. Gao and N. Jin. Complex network from time series based on phase space reconstruction. *Chaos (Woodbury, N. Y.)*, 19:033137, 09 2009.
- [25] Z.-K. Gao, M. Small, and J. Kurths. Complex network analysis of time series. *EPL (Europhysics Letters)*, 116(5):50001, dec 2016.
- [26] A. Gogolou, T. Tsandilas, K. Echihabi, A. Bezerianos, and T. Palpanas. Data Series Progressive Similarity Search with Probabilistic Quality Guarantees. In *SIGMOD*, 2020.
- [27] M. Hadjem, F. Naït-Abdesselam, and A. A. Khokhar. St-segment and t-wave anomalies prediction in an ECG data using rusboost. In *Healthcom*, 2016.
- [28] N. Halko, P. G. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 2011.
- [29] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, Nov. 1997.
- [30] H. Kantz and T. Schreiber. *Nonlinear Time Series Analysis*. Cambridge University Press, New York, NY, USA, 2003.
- [31] E. Keogh, S. Lonardi, and C. A. Ratanamahatana. Towards parameter-free data mining. In *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '04, pages 206–215, New York, NY, USA, 2004. ACM.
- [32] E. Keogh, S. Lonardi, C. A. Ratanamahatana, L. Wei, S.-H. Lee, and J. Handley. Compression-based data mining of sequential data. *Data Mining and Knowledge Discovery*, 2007.
- [33] E. J. Keogh, J. Lin, and A. W. Fu. HOT SAX: efficiently finding the most unusual time series subsequence. In *ICDM*, 2005.
- [34] L. Lacasa, B. Luque, F. Ballesteros, J. Luque, and J. C. Nuno. From time series to complex networks: The visibility graph. *Proceedings of the National Academy of Sciences*, 105(13):4972–4975, 2008.
- [35] T. Lee, J. Gottschlich, N. Tatbul, E. Metcalf, and S. Zdonik. Greenhouse: A zero-positive machine learning system for time-series anomaly detection. *CoRR*, abs/1801.03168, 2018.
- [36] M. Linardi, Y. Zhu, T. Palpanas, and E. J. Keogh. Matrix Profile Goes MAD: Variable-Length Motif And Discord Discovery in Data Series. In *DAMI*, 2020.
- [37] F. T. Liu, K. M. Ting, and Z.-H. Zhou. Isolation forest. In *ICDM*, ICDM, 2008.
- [38] Y. Liu, X. Chen, and F. Wang. Efficient Detection of Discords for Time Series Stream. *Advances in Data and Web Management*, pages 629–634, 2009.

- [39] W. Luo and M. Gallagher. Faster and parameter-free discord search in quasi-periodic time series. In J. Z. Huang, L. Cao, and J. Srivastava, editors, *Advances in Knowledge Discovery and Data Mining*, 2011.
- [40] P. Malhotra, L. Vig, G. Shroff, and P. Agarwal. Long short term memory networks for anomaly detection in time series. 2015.
- [41] K. Mirylenka, A. Marascu, T. Palpanas, M. Fehr, S. Jank, G. Welde, and D. Groeber. Envelope-based anomaly detection for high-speed manufacturing processes. *European Advanced Process Control and Manufacturing Conference*, 2013.
- [42] G. B. Moody and R. G. Mark. The impact of the mit-bih arrhythmia database. *IEEE Engineering in Medicine and Biology Magazine*, 2001.
- [43] A. Mueen, E. J. Keogh, Q. Zhu, S. Cash, and M. B. Westover. Exact discovery of time series motifs. In *SDM 2009*.
- [44] T. Palpanas. Data series management: The road to big sequence analytics. *SIGMOD Rec.*, 44(2):47–52, Aug. 2015.
- [45] T. Palpanas. Evolution of a Data Series Index. *CCIS*, 1197, 2020.
- [46] T. Palpanas and V. Beckmann. Report on the First and Second Interdisciplinary Time Series Analysis Workshop (ITISA). *ACM SIGMOD Record*, 48(3), 2019.
- [47] T. Pelkonen, S. Franklin, P. Cavallaro, Q. Huang, J. Meza, J. Teller, and K. Veeraraghavan. Gorilla: A fast, scalable, in-memory time series database. *PVLDB*, 8(12):1816–1827, 2015.
- [48] B. Peng, P. Fatourou, and T. Palpanas. MESSI: In-Memory Data Series Indexing. In *ICDE*, 2020.
- [49] B. Peng, T. Palpanas, and P. Fatourou. Paris+: Data series indexing on multi-core architectures. *TKDE*, 2020.
- [50] D. W. Scott. *Multivariate Density Estimation. Theory, Practice, and Visualization*. Wiley, 1992.
- [51] 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*, 2015.
- [52] P. Senin, J. Lin, X. Wang, T. Oates, S. Gandhi, A. P. Boedihardjo, C. Chen, and S. Frankenstein. Grammarviz 3.0: Interactive discovery of variable-length time series patterns. *TKDD*, 2018.
- [53] S. Subramaniam, T. Palpanas, D. Papadopoulos, V. Kalogeraki, and D. Gunopulos. Online outlier detection in sensor data using non-parametric models. In *VLDB 2006*, pages 187–198, 2006.
- [54] J. Wang, A. Balasubramanian, L. M. de la Vega, J. Green, A. Samal, and B. Prabhakaran. Word recognition from continuous articulatory movement time-series data using symbolic representations. In *SLPAT*.
- [55] L. Wei, E. J. Keogh, and X. Xi. Saxually explicit images: Finding unusual shapes. In *Proceedings of the 6th IEEE International Conference on Data Mining (ICDM 2006), 18-22 December 2006, Hong Kong, China*, pages 711–720, 2006.
- [56] C. Whitney, D. Gottlieb, S. Redline, R. Norman, R. Dodge, E. Shahar, S. Surovec, and F. Nieto. Reliability of scoring respiratory disturbance indices and sleep staging. *Sleep*, November 1998.
- [57] D. Yankov, E. J. Keogh, J. Medina, B. Y. Chiu, and V. B. Zordan. Detecting time series motifs under uniform scaling. In *ACM*.
- [58] D. Yankov, E. J. Keogh, and U. Rebbapragada. Disk aware discord discovery: Finding unusual time series in terabyte sized datasets. In *ICDM*, 2007.
- [59] D. Yankov, E. J. Keogh, and U. Rebbapragada. Disk aware discord discovery: finding unusual time series in terabyte sized datasets. *Knowl. Inf. Syst.*, 17(2):241–262, 2008.
- [60] C. M. Yeh, Y. Zhu, L. Ulanova, N. Begum, Y. Ding, H. A. Dau, D. F. Silva, A. Mueen, and E. J. Keogh. Matrix profile I: all pairs similarity joins for time series: A unifying view that includes motifs, discords and shapelets. In *ICDM*, pages 1317–1322, 2016.