Robust Multi-Variate Temporal Features of Multi-Variate Time Series

SICONG LIU, Arizona State University

SILVESTRO ROBERTO POCCIA, University of Torino

K. SELÇUK CANDAN, Arizona State University

MARIA LUISA SAPINO, University of Torino

XIAOLAN WANG, University of Massachusetts, Amherst

Many applications generate and/or consume multi-variate temporal data, and experts often lack the means to adequately and systematically search for and interpret multi-variate observations. In this article, we first observe that multi-variate time series often carry localized multi-variate temporal features that are robust against noise. We then argue that these multi-variate temporal features can be extracted by simultaneously considering, at multiple scales, temporal characteristics of the time series *along with external knowledge*, including variate relationships that are known a priori. Relying on these observations, we develop data models and algorithms to detect *robust multi-variate temporal* (RMT) features that can be indexed for efficient and accurate retrieval and can be used for supporting data exploration and analysis tasks. Experiments confirm that the proposed RMT algorithm is highly effective and efficient in identifying *robust* multi-scale temporal features of multi-variate time series.

CCS Concepts: • **Mathematics of computing** → **Time series analysis**; • **Information systems** → **Temporal data**; **Multimedia databases**; **Data mining**; *Proximity search*; *Multimedia and multimodal retrieval*; • **Computing methodologies** → **Activity recognition and understanding**; Unsupervised learning; • **Theory of computation** → Pattern matching;

Additional Key Words and Phrases: Robust multi-variate temporal features, multi-variate time series

ACM Reference format:

Sicong Liu, Silvestro Roberto Poccia, K. Selçuk Candan, Maria Luisa Sapino, and Xiaolan Wang. 2018. Robust Multi-Variate Temporal Features of Multi-Variate Time Series. *ACM Trans. Multimedia Comput. Commun. Appl.* 14, 1, Article 7 (January 2018), 24 pages.
<https://doi.org/10.1145/3152123>

This work is supported by NSF#1339835, “SI2-SSE: E-SDMS: Energy Simulation Data Management System Software”; NSF#1318788, “Data Management for Real-Time Data Driven Epidemic Spread Simulations”; NSF#1610282, “DataStorm: A Data Enabled System for End-to-End Disaster Planning and Response”; NSF#1633381, “BIGDATA: Discovering Context-Sensitive Impact in Complex Systems”; and EU-H2020 Marie Skłodowska-Curie grant agreement No 690817, “FourCModeling.”

Authors’ addresses: S. Liu and K. S. Candan, 699 S Mill Ave, CIDSE/ASU, Tempe, AZ 85281, USA; emails: {s.liu, candan}@asu.edu; S. R. Poccia and M. L. Sapino, University of Torino, via Pessinetto 12, 10149 Turin TO, Italy; emails: {poccia, mlsapino}@di.unito.it; X. Wang, Computer Science Building, 140 Governors Drive, Amherst, MA 01003, USA; email: xlwang@cs.umass.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 ACM 1551-6857/2018/01-ART7 \$15.00

<https://doi.org/10.1145/3152123>

1 INTRODUCTION

Many applications, such as motion recognition [34], generate temporal data and, in many of these applications, (a) the resulting time-series data are multi-variate, (b) relevant processes underlying these time series are of different scales [17, 29, 48], and (c) the variates (i.e., observation parameters) are dependent on each other in various ways [40].

Analysis and exploration of time series (as well as other types of data) often start with extraction of patterns and features that describe salient properties of the data. Popular approaches in the literature include extraction of global features of the time series (such as spectral properties quantified using a transformation; e.g., Discrete Cosine or Wavelet Transforms) and the use of these global features (which describe properties of the time series as a whole) for indexing [5]. Correlations, transfer functions, variate clusters, and spectral properties [47], Singular-Value Decomposition (SVD) and similar eigen-decompositions can be used for extracting global fingerprints of multi-variate time-series data [27]. The analogous analysis operation on a tensor, which can be used to represent temporal evolution of multi-modal data, is known as tensor decomposition [23]. Both matrix and tensor decomposition operations, as well as other techniques, such as probabilistic techniques (such as Dynamic Topic Modeling, DTM [3]) and AutoRegressive Integrated Moving-Average- (ARIMA) based analyses (which separate a time series into autoregressive, moving-average, and integrative components for modeling and forecasting [33]) are expensive.

Several researchers noticed that significant amount of waste in processing and exploration can be avoided if the attention is directed towards parts of a given time series that are likely to contain interesting patterns [4, 35]. One way to achieve this involves searching for frequently repeating patterns; this is commonly known as the *motif search* problem [14]. Most of the common approaches for motif search involve incrementally moving (or *shifting*) a fixed-length time window starting from the beginning of the given time series. For each window interval a temporal signature (such as Symbolic Aggregate approXimation (SAX) words [28]) is generated (to speed up the matching of subsequences) and frequent sub-sequences are discovered using different indexing and hashing algorithms and leveraging pruning techniques for eliminating non-promising subsequences [54]. Other local features of uni-variate time-series include landmarks [38], perceptually important points (PIP) [10], patterns [1], shapelets [42, 55], snippets [50], longest common subsequences (LCSS [49]), and motif-based schemes (which search for frequently repeating temporal patterns) [6]. Noting that uni-variate time series often carry localized temporal features that can be used for efficient search and analysis, in our earlier work [4] we developed an sDTW algorithm for extracting salient local features (that are robust against various types of noise) of uni-variate time series and showed that these can help align similar time series more efficiently and effectively. RPM [52] and STS3 [37] are two recent approaches that also seek informative patterns from uni-variate time series.

1.1 Contributions of this Paper: Local Features of Multi-Variate Time Series

In this article, we develop data models and algorithms to detect local, *robust multi-variate temporal* (RMT) features of multi-variate time series (Figure 1). Recently, in Reference [51], we proposed a multi-variate feature extraction technique, which considered the relationships and dependencies between the individual uni-variate time-series that make up the multi-variate series. The local, *robust multi-variate temporal* (RMT) features are extracted leveraging known correlations and dependencies among the variates. As in Reference [4], for uni-variate series, Reference [51] also (a) smoothes the data to generate different version of the input object corresponding to different scales and (b) compares neighboring points both in time (or in x and y dimensions) and scale to identify regions where the gradients are large. As in Reference [31], Scale-Invariant Feature Transform (SIFT)-like feature descriptors are extracted to support search.

Fig. 1. A multi-variate time-series dataset, where each variate is plotted as a row of grayscale pixels and sample multi-variate features identified on the dataset (each feature is marked with a different color): The figure shows 26 variates of length 150 and 5 local multi-variate features on these time series. Note that some of the features correspond to the onset of a rise in amplitude, whereas other correspond to the drop in the series amplitude. For each time series involved in a given multi-variate feature, we plot the corresponding temporal scope (i.e., duration) of that feature.

What makes the problem of extracting local features from multi-variate series difficult, however, is that the concepts of neighborhood, gradient, and smoothing are not well defined in the presence of multiple variates. In Reference [51], we argued that this difficulty can be overcome by leveraging metadata (known correlations and dependencies among the variates in the time series) to define neighborhoods, support data smoothing, and construct scale spaces in which gradients can be measured. Based on this observation, we proposed *topology-sensitive smoothing* and *topology-sensitive gradient computation* techniques to identify local features of multi-variate time series at different time/variant scales. In this article, we show that, unlike in Reference [51] (where the time and variate scales are shrunk and expanded together), more effective local-feature sets can be located if we allow for the time and variate aspects of the multi-variate time series to be considered independently from each other—leading to multi-variate features with heterogeneous time and variate scales. This is also visualized in Figure 1, where we see some features that are short in time but contain a lot of variates (such as the feature highlighted in red), whereas others are longer but contain fewer variates (such as the feature highlighted in yellow). We also provide a detailed discussion of RMT-based multi-variate time-series matching and inconsistency removal and experimentally evaluate their effectiveness in gesture and motion recognition [34].

1.2 Organization of the Paper

In the next section, we describe the related work. In Section 3.1, we introduce the metadata-enriched multi-variate time-series (MMTS) model that forms the basis of the proposed work. In Section 3.2, we present an overview of the proposed approach to locate RMT features and extract their descriptors. Sections 4 through 8 provide details of the various steps of the proposed RMT feature identification and descriptor extraction algorithm. Section 9 formally defines the RMT feature set of a multi-variate time series, and Section 10 describes how to use these features for matching multi-variate time series. We present experimental evaluations of the proposed approach in Section 11. We conclude the article in Section 12.

2 RELATED WORK

Euclidean distance and, more generally, L_p -norm measures, were among the first used to determine the similarity between two time series. They require that the time series being compared are of *same temporal length* and, since they assume strict synchrony among time series, they are

not suitable when two time series can have different speeds or are shifted in time [7, 21]. Other measures that require equal length and perfect synchrony across time series include *cosine* and (Pearson's) correlation similarity [44]. In contrast, edit distance [24] measures aim to determine the minimum sequence of *edit* operations that are required to measure similarity. In the 1970s, Sakoe [43], and then, in the mid-1990s, Berndt [2] proposed the dynamic time warping (DTW) technique to find an optimal alignment between two given (time-dependent) sequences under certain restrictions. Intuitively, DTW considers all possible warping paths that can warp (or transform) one series into the other and picks the warping path that has the lowest cost. DTW has found wide acceptance and the past two decades have seen several innovations [8, 11, 20, 21, 41]. For example, while the original DTW is not metric (does not satisfy triangular inequality), the authors of Reference [8] proposed an extended version of DTW that satisfies triangular inequality. Most of the above algorithms, including DTW, are initially designed for comparing uni-variate time series. More recently, various extensions of DTW have been proposed for *multi-dimensional* time series [39, 45]. The most prevalent of these are the *vectorized* and *independent* extensions. In *vectorized DTW*, a multi-variate time series is considered as a sequence of vectors, where the length of vector is equal to number of variates in the time series. The DTW algorithm is then applied using the distances among these vectors instead of differences in signal amplitude. In *independent DTW*, however, each variate is treated independently from the others and DTW is applied separately to each; finally, these independent DTW distances are added to compute the overall distance between the given pair of multi-variate series.

An alternative approach to the above techniques is to extract *features* from the given time series and use these features to compute similarity/distance instead of the original series. *We provided a detailed discussion of the related work on global and local features in the previous section.* The authors of Reference [18], for example, proposed a feature-extraction algorithm that extracts minimal distinguishing subsequences that can be used as features. Morchen [36] proposed using Discrete Fourier Transform (DFT) and Discrete Wavelet Transform (DWT) for feature extraction. The *PCA-Similarity Factor* [25] and *Extended Frobenius norm* (EROS) [53], which use matrix factorization techniques, such as singular vector decomposition (SVD) and principle component analysis, have also been proposed to transform the input multi-variate time series into equal length and then apply cosine similarity over them. Applications of SVD and DTW for various multimedia tasks, such as similarity search, classification, recognition, and watermarking, include References [19, 22, 26, 32, 46].

As we discussed earlier, in Reference [51], we proposed to extract and use SIFT- [30, 31] like robust multi-variate temporal features to determine similarity between time series. In this article, we extend the approach in Reference [51] with more general scale-space construction and pruning techniques to obtain better classification performances. The general framework is named RMT for both articles, but the underlying techniques here are significant extensions of Reference [51]: In particular, Reference [51] utilizes a special case of our proposed generalized scale-space construction and pruning techniques, where only diagonal scale space is considered. In this submission, however, we argue (and experimentally show) that, in general, using a more complete scale space can be more effective. Indeed, as the experimental results reported in Section 11 show, we achieve better accuracy using a complete scale space (enabling time and variate scales to be different) rather than using a diagonal scale space (where the time and variate scales are constrained to be identical). In addition, while Reference [51] considers only one scheme for time-series matching using RMT features, this article introduces several alignment measures (including for alignment of features pairs and measuring feature significance) and provides a detailed study of the impact of these measures on the classification accuracies.

3 DATA MODEL AND OVERVIEW OF THE PROPOSED APPROACH

Before describing the process through which we extract RMT features, we first introduce the metadata-enriched, multi-variate time-series model underlying the proposed approach.

3.1 MMTS Model

In this section, we present a MMTS model which minimizes the assumptions that need to be made about the structure of the data:

Definition 3.1 (MMTS). A metadata-enriched multi-variate (MM) time series is a four-tuple $\mathbf{Y} = (\mathcal{V}, \mathcal{M}, \mathbb{Y}, \mathcal{D})$, where

- $\mathcal{V}$ is a set of variates,
- $\mathcal{M} = \{M_1, \dots, M_m\}$ is a set of metadata modalities, where each modality M_i describes how the corresponding subset $V_i \subseteq \mathcal{V}$ of variates are related to each other,
- $\mathbb{Y}$ is a $(d_1 + d_2 + \dots + d_m) \times l$ data matrix, where
 - l is the temporal length of the multi-variate time series,
 - $d_i = |V_i|$, and
 - cells of the matrix $\mathbb{Y}$ take values from the data domain $\mathcal{D}$.

As defined above, each variate is associated with a modality metadata describing how it is related with other variates. In this article, without loss of generality, we consider *graph-organized* (G) representation of variate modalities: Each modality, i , has an associated graph $G_i(V_i, E_i, W_i)$ that relates the variates V_i of the given mode. Depending on the application, the graph maybe directed or undirected and weights may have *distance* or *similarity* semantics. If the underlying graph is unweighted, then for all $e_k \in E_i$, $W_i(e_k) = 1$.

Note that graph-based description of variate relationships is a common way of modeling temporal dynamics of multi-variate time series [12, 16, 47]. Note also that, while the metadata describes the relationship between the variates, this relationship may or may not have causal impact on the observed temporal characteristics of the data:

Definition 3.2 (Metadata-Defined Variate Causality Model). Let us assume that we have metadata $\mathcal{M}$ that describe the relationship between the variates in the data. Given $\mathcal{M}$, under the variate causality model, we have $\mathbb{Y}[t] = \mathbf{R}_{\mathcal{M}}\mathbb{Y}[t-1] + \vec{E}(t)$, where $\mathbb{Y}[t]$ is a column vector extracted from $\mathbb{Y}$ and corresponds to the observations at time t , $\mathbf{R}_{\mathcal{M}}$ is a (row-normalized) matrix defining how the values of $\mathbb{Y}$ at time $t-1$ impact the values of $\mathbb{Y}$ at time t , and $\vec{E}$ is a multi-variate time series denoting independent, external inputs.

Intuitively, $\mathbf{R}_{\mathcal{M}}$ is a matrix describing how the values of one variate are impacted by the past values of the variates in the data. Alternatively, $\mathbf{R}_{\mathcal{M}}$ may be a matrix describing the relationships among simultaneous observations:

Definition 3.3 (Metadata-Defined Variate Correlation Model). Let us assume that we have metadata $\mathcal{M}$ that describe the relationship between the variates in the data. Given $\mathcal{M}$, under the variate correlation model, we have a matrix $\mathbf{R}_{\mathcal{M}}$ such that $\mathbf{R}_{\mathcal{M}}[i, j] = \Phi(\mathbb{Y}[:, i], \mathbb{Y}[:, j]) \in [0, 1]$. Here $\mathbb{Y}[:, i]$ and $\mathbb{Y}[:, j]$ are rows corresponding to observations for variates i and j , respectively, and Φ is an application specific similarity function.

Here, Φ may be computed by comparing (recent) historical data of the time series or may reflect available domain knowledge, such as the distance of the sensors recording the variates or known relationships parameters.

Fig. 2. Multi-variate features can be of different sizes (in this example, the multi-variate series represent temperature readings in a floor split into zones).

It is important to note that the algorithms presented in the article are applicable under both of the above models,¹ and we use the matrix R_M to denote both relationships.

3.2 Overview: Extracting Local, RMT Features

In this article, we propose algorithms to extract RMT features from MMTS datasets. Intuitively, as already visualized in Figure 1, an RMT feature is a fragment of a multi-variate time series that is maximally different from its immediate neighborhood (both in time and across variate relationships specified by the metadata). As in Reference [31], we rely on a four-step process to identify such RMT features and extract their feature descriptors:

Step 1: Scale-space construction: As shown in Figure 2, multi-variate temporal features of interest can be of different lengths and may cover different number of variates. In order to be able to locate such features of different sizes, the RMT features are extracted from a scale space we construct for the given multi-variate time series through iterative smoothing.² Iterative smoothing of the multi-variate data in time and variates creates different resolution versions of the input data and, thus, helps identify features with different amount of details in time and in terms of the number of variates involved. While iterative smoothing techniques are well understood for uni-variate data [4, 48], this is not the case for multi-variate time series. Therefore, in Sections 4 and 5, we describe how to construct a scale space by smoothing a multi-variate time series, leveraging available metadata that describe known relationships among variates.

Step 2: Identifying feature candidates: Next, the process identifies candidate features of interest across multiple scales of the given multi-variate time series by searching over multiple scales and variates of the given series. Each candidate RMT feature has a *temporal-scope* (a beginning and an end in time) and a *variate-scope* (a set of variates involved in the feature). These candidate features of interest are those with the largest variations with respect to their neighbors in time, variates, and scale. We describe this process in Section 6.

Step 3: Eliminating poor candidates: At the following step, those candidate features identified in the previous step that are sensitive to noise are eliminated. These include features that are poorly localized (and hence are difficult to match). This is described in Section 7.

Step 4: RMT feature descriptor creation: In the final step, for each RMT feature, a local descriptor is extracted using the information obtained in the previous steps. More specifically, the algorithm computes and samples gradients within the scope of the RMT feature. To avoid sudden changes in the descriptor with small changes in the position and to give less emphasis to gradients that are far

¹Thus, without loss of generality, we sometimes focus on the dependency model and, other times, use the correlation model.

²This is different from what is known as “multi-variate exponential smoothing,” a forecasting technique where the multi-variate models include the “smoothing parameters” and these are learned to obtain models with a better fit to the data [47].

Fig. 3. Scale-space construction through iterative smoothing in time and variates.

from the center of the descriptor, a weighing function is used to assign a weight to the magnitude of each sample point based on its distance from the center of the feature. Note that while gradient computation is well understood for uni-variate data [4, 48], this is not the case for multi-variate time series. We describe how this is achieved in Section 8.

The above approach has three key advantages:

Advantage 1: First, the identified salient features are robust against noise and common transformations, such as temporal shifts or dropped/missing variates.

Advantage 2: Scale invariance enables the extracted salient features to be robust against variations in speed and enables multi-resolution searches.

Advantage 3: The temporal and relationship scales at which a multi-variate feature is located give an indication about the *scope* (both in terms of duration and the number of variates involved) of the multi-variate feature.

4 TEMPORAL AND VARIATE SMOOTHING OF MULTI-VARIATE TIME SERIES

Let $Y = (\mathcal{V}, \mathcal{M}, \mathbb{Y}, \mathcal{D})$ be a metadata-enriched multi-variate time series, as defined in Section 3.1. The first step in identifying multi-variate features of Y is to generate a scale space representing versions of the given multi-variate series with different amounts of details. As shown in Figure 3, the scale space, $\mathbb{Y}$, of Y is obtained through iterative smoothing across both time and variate relationships, starting with an initial smoothing parameter $\Sigma_0 = \langle \sigma_{time,0}, \sigma_{var,0} \rangle$ and iteratively increasing the smoothing degree up to $\Sigma_{max} = \langle \sigma_{time,max}, \sigma_{var,max} \rangle$, obtaining differently smoothed versions of the time series.

The values of Σ_0 and Σ_{max} control the sizes of the smallest and largest features sought in the data (see Section 9). In the rest of this section, we will first describe temporal and variate smoothing techniques. We will then describe optimizations to reduce the cost of the scale-space construction step of the process. For each of the techniques, we will also discuss the relationship among Σ_0 , Σ_{max} , and the sizes of the features identified.

4.1 Temporal Smoothing

Let $Y = (\mathcal{V}, \mathcal{M}, \mathbb{Y}, \mathcal{D})$ be a metadata-enriched multi-variate time series and let $Y_v = \mathbb{Y}[* , v]$ be a uni-variate time series corresponding to one of its variates. Let $Y_v^{(\sigma)}$ indicate a version of the uni-variate time series, Y_v , smoothed through convolution with the Gaussian function, $G(t, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{t^2}{2\sigma^2}}$, with temporal smoothing parameter σ (Figure 4). Given this, $Y^{(time, \sigma)} = (\mathcal{V}, \mathcal{M}, \mathbb{Y}^{(time, \sigma)}, \mathcal{D})$ is a version of the multi-variate time series, Y , where each row of $\mathbb{Y}^{(time, \sigma)}$

Fig. 4. Gaussian smoothing of a univariate series for time instant, t .

Fig. 5. Graph smoothing for a node.

is a uni-variate time series smoothed with temporal smoothing parameter σ , independently of the other uni-variate series.

4.1.1 Temporal Scope. Let us consider a time instant t on which we are applying Gaussian smoothing with parameter σ . Since, under Gaussian smoothing, 3 standard deviations (i.e., 3σ both directions) would cover $\sim 99.73\%$ of the contributions to the smoothed values, we can define the corresponding *temporal scope* as a time interval, centered at t , of length 6σ ; in other words, we have $\text{scope}_T(t, \sigma) = [t - 3\sigma, t + 3\sigma]$. Consequently, if the temporal length of a multi-variate time series is L , then we must have $\sigma_{\text{time}, \max} \leq L/6$. Similarly, since we expect that the smallest feature should involve a time instant and at least its two immediate neighbors, we also have $\sigma_{\text{time}, 0} \geq 2/6$.

4.1.2 Octaves of Temporal Smoothing. Let σ_1 and σ_2 be two smoothing parameters. The parameter σ_2 is said to be an *octave* larger than σ_1 if $\sigma_2 = 2\sigma_1$. Intuitively, σ_2 defines features twice as large as σ_1 by using a Gaussian smoothing parameter twice as large.

4.2 Variate Smoothing

As described above, the temporal smoothing process relies on a convolution operation that leverages the temporal ordering of the time instants in the series. The challenge is that a similar total order does not necessarily exist among the variates; therefore, the definition of variate smoothing is not as straightforward.

4.2.1 Gaussian Smoothing of Graph-Organized Variates. Let $\mathbf{Y} = (\mathcal{V}, \mathcal{M}, \mathbb{Y}, \mathcal{D})$ be a metadata-enriched multi-variate time series, where the metadata $\mathcal{M}$ is graph-structured; i.e., there is a graph $G(V, E, W)$ that relates the variates, V , of the data. Let us further define $\text{frwd}_G(v_l, \delta)$ and $\text{bkwd}_G(v_l, \delta)$, as the forward and backward neighbors of variate $v_l \in V$ at a distance $\geq (\delta - 0.5)$ and $< (\delta + 0.5)$ on G . Intuitively, $\text{frwd}_G()$ and $\text{bkwd}_G()$ functions order all the variates into a partial order relative to the variate v_l .

Given the partial order defined by the $\text{frwd}_G()$ and $\text{bkwd}_G()$ functions and a non-negative smoothing parameter σ , we then obtain the Gaussian smoothed version, $\mathbb{Y}^{(\text{var}, \sigma)}$ of the matrix $\mathbb{Y}$, as follows: Let $\mathbb{Y}[t]$ be a column vector extracted from $\mathbb{Y}$ corresponding to the observations at time t . Then, for all v_l , we have $\mathbb{Y}^{(\text{var}, \sigma)}[t, l]$ equal to

$$\left(\underset{v_h \in \underset{\text{frwd}_G(v_l, 0)}{\text{AVG}} \cup \text{bkwd}_G(v_l, 0)}{\text{AVG}} \mathbb{Y}[t, h] \right) + \sum_{\delta=1}^{\infty} G(\delta, \sigma) \left(\underset{v_h \in \text{frwd}_G(v_l, \delta)}{\text{AVG}} \mathbb{Y}[t, h] \right) + \sum_{\delta=1}^{\infty} G(\delta, \sigma) \left(\underset{v_h \in \text{bkwd}_G(v_l, \delta)}{\text{AVG}} \mathbb{Y}[t, h] \right).$$

Figure 5 shows how we apply Gaussian smoothing over a relationship graph. The lower half of the figure shows a variate a and its forward and backward k -hop neighbors in the relationship graph. As shown in the upper half of the figure, when identifying the contributions of the variate on a , Gaussian smoothing is applied along the hop distance. Since at a given hop distance there may be more than one variate, all the variates at the same distance have the same degree of contribution and the degree of contribution gets progressively smaller as we get away from the variate for which the smoothing is performed.

4.2.2 Variate Scope under Gaussian Smoothing. Similarly to the temporal scope, we define the variate scope corresponding to variate v_l at smoothing level σ as

$$scope_V(v_l, \sigma) = \{v_l\} \cup \left(\bigcup_{\delta \leq 3\sigma} frwd_G(v_l, \delta) \right) \cup \left(\bigcup_{\delta \leq 3\sigma} bkwd_G(v_l, \delta) \right).$$

The variate smoothing parameters, $\sigma_{var,0}$ and $\sigma_{var,max}$, must be selected such that for each variate v_l , $\sigma_{var,0}$ includes its immediate (one hop) graph neighbors, $forward_neighbors(v_l)$ and $backward_neighbors(v_l)$ on G , and the value of δ corresponding to $\sigma_{var,max}$ should be compatible with the diameter of the graph G .

4.2.3 Octaves of Gaussian Variate Smoothing. Let σ_1 and σ_2 be two Gaussian graph smoothing parameters. Under Gaussian smoothing, the graph smoothing parameter σ_2 is said to be an *octave* larger than σ_1 if $\sigma_2 = 2\sigma_1$.

4.3 Combined Time and Variate Smoothing

Given the above definitions of temporal and variate smoothing functions, we now define combined time and variate smoothing of metadata-enriched multi-variate time series:

Definition 4.1 (TV-Smoothing of a Multi-Variate Time Series). Let $\mathbf{Y} = (\mathcal{V}, \mathcal{M}, \mathbb{Y}, \mathcal{D})$ be a MM time series. Recall that $\mathbb{Y}$ is a $(d_1 + d_2 + \dots + d_m) \times l$ data matrix, where l is the temporal length of the multi-variate time series, $d_i = |V_i|$, and $\mathbb{Y}$ takes values from the data domain $\mathcal{D}$. For a given smoothing parameter, $\Sigma = \langle \sigma_{time}, \sigma_{var} \rangle$, the TV-smoothed version, $\mathbf{Y}(\Sigma)$, of the multi-variate time series, $\mathbf{Y}$, is defined as $\mathbf{Y}(\Sigma) = (\mathbb{Y}^{(time, \sigma_{time})})^{(var, \sigma_{var})}$, where

- $\mathbb{Y}^{(time, \sigma_{time})}$ is a version of $\mathbb{Y}$ where each row (i.e., each uni-variate time series) is temporally smoothed with smoothing parameter σ_{time} , independently from the rest and
- $\mathbb{X}^{(var, \sigma_{var})}$ is a version of $\mathbb{X}$ where each column (i.e., time instant) is smoothed with smoothing parameter σ_{var} , using the variable relationships and modalities described by the metadata $\mathcal{M}$.

5 STEP 1: SCALE-SPACE CONSTRUCTION FOR MULTI-VARIATE TIME SERIES

As we have seen in Section 4, given a metadata-enriched multi-variate time series, $\mathbf{Y} = (\mathcal{V}, \mathcal{M}, \mathbb{Y}, \mathcal{D})$, first step in identifying multi-variate features of $\mathbf{Y}$ is to generate a scale space representing versions of the multi-variate series with different amounts of details. In this article, we consider two types of scale spaces: *diagonal* and *full* scale spaces, described below.

5.1 Diagonal Scale-Spaces

Let $\Sigma_0 = \langle \sigma_{time,0}, \sigma_{var,0} \rangle$ be the user provided smallest temporal and variate smoothing parameters and let l indicate the total number of layers in the scale space. An l -layer *diagonal* state space, $\mathbb{Y}_{diag}$, is defined as a set of data matrices $\{\mathbb{Y}_0, \dots, \mathbb{Y}_l\}$, where $\mathbb{Y}_i = \mathbb{Y} \{ \langle \sigma_{time,0} \times k^i, \sigma_{var,0} \times k^i \rangle \}$, for some scaling parameter $k > 1$. Note that, in this case, we have $\sigma_{time,max} = \sigma_{time,0} \times k^l$ and

$\sigma_{var,max} = \sigma_{var,0} \times k^l$. This will generate only the diagonal entries in the scale space shown in Figure 3.

5.2 Full Scale-Spaces

In contrast, the complete scale space shown in Figure 3 is generated as follows: Let

- $\Sigma_0 = \langle \sigma_{time,0}, \sigma_{var,0} \rangle$ be the smallest temporal and variate smoothing parameters,
- $\mathcal{L} = \langle l_{time}, l_{var} \rangle$ indicate the number of temporal and variate smoothing layers, and
- $\mathcal{K} = \langle k_{time}, k_{var} \rangle$ be scaling parameters for temporal and variate smoothings.

An $\mathcal{L}$ -layer *full* state space, $\mathfrak{Y}_{full}$, is defined as a set of data matrices $\langle \mathbb{Y}_{0,0}, \dots, \mathbb{Y}_{i,j}, \dots, \mathbb{Y}_{l_{time},l_{var}} \rangle$, where $\mathbb{Y}_{i,j} = \mathbb{Y}\{\langle \sigma_{time,0} \times k_{time}^i, \sigma_{var,0} \times k_{var}^j \rangle\}$. In this case, we have $\sigma_{time,max} = \sigma_{time,0} \times k_{time}^{l_{time}}$ and $\sigma_{var,max} = \sigma_{var,0} \times k_{var}^{l_{var}}$.

5.3 Optimization: Time and Variate Subsampling

In the process described above, the multi-variate time series is incrementally smoothed both in time and relationships, halving details at each octave. We note that, once the details have been halved at an octave boundary, performing the feature extraction operation at the same level detail is going to be wasteful. To avoid such waste, we subsample the multi-variate time series at octave boundaries (Figure 3). More specifically, at temporal octave boundaries (where temporal details have been halved) we drop one out of every two consecutive temporal observations, reducing the size of the data by half. Similarly, at variate octave boundaries (where variate relationship details have been halved), we reduce the numbers of variates by half by applying a variate clustering algorithm.³

6 STEP 2: IDENTIFYING MULTI-VARIATE TEMPORAL FEATURE CANDIDATES

Building on the observation [4, 31] that robust localized features are often located where the differences between neighboring regions (possibly in different scales) are large, we seek RMT features of the given multi-variate time series at the *local extrema* of the scale space defined by the difference-of-smoothing (DoS) series. Naturally, the DoS generation and feature identification process will be slightly different depending on whether a diagonal or full scale space is used.

6.1 Local Extrema in Diagonal Scale-Spaces

An l -layer *diagonal* state space of a metadata-enriched multi-variate time series, $\mathbf{Y} = (\mathcal{V}, \mathcal{M}, \mathbb{Y}, \mathcal{D})$, is defined as a set of data matrices $\{\mathbb{Y}_0, \dots, \mathbb{Y}_l\}$, where $\mathbb{Y}_i = \mathbb{Y}\{\langle \sigma_{time,0} \times k^i, \sigma_{var,0} \times k^i \rangle\}$, for some scaling parameter $k > 1$. Given this, we create the corresponding DoS by considering a sequence of difference matrices $\{\mathbb{D}_0, \dots, \mathbb{D}_{l-1}\}$, where $\mathbb{D}_i = |\mathbb{Y}_{i+1} - \mathbb{Y}_i|$. We detect RMT feature candidates by seeking the local maxima and minima of the resulting DoS: each variate-time-scale (VTS) triple, $\langle v, t, s \rangle$, is compared to its neighbors (both in time and variate relationships) in the same scale as well as the scales above and below, and the triple is selected as a candidate only if it is close to being an extremum; i.e., each $\langle v, t, s \rangle$ is compared against its 26 ($= 3^3 - 1$) neighbors in time, scale, and variate relationships.⁴ More specifically, for each $\langle v, t, s \rangle$, we compare $\mathbb{D}_s[v, t]$ against

³In the experiments reported in Section 11, we use a k -means algorithm, where k is equal to the half of the number of variates, based on the distances among sensors on the underlying sensor-distance graph.

⁴The number of neighboring triples may be less than 26 if the triple is at the boundary in terms of time, scale, or variate relationship graph.

$$\max \left\{ \begin{array}{ccc} \mathbb{D}_{s-1}[v, t-1] & \mathbb{D}_s[v, t-1] & \mathbb{D}_{s+1}[v, t-1] \\ \mathbb{D}_{s-1}[v, t] & \mathbb{D}_s[v, t] & \mathbb{D}_{s+1}[v, t] \\ \mathbb{D}_{s-1}[v, t+1] & \mathbb{D}_s[v, t+1] & \mathbb{D}_{s+1}[v, t+1] \\ \mathbb{FD}_{s-1}[v, t-1] & \mathbb{FD}_s[v, t-1] & \mathbb{FD}_{s+1}[v, t-1] \\ \mathbb{FD}_{s-1}[v, t] & \mathbb{FD}_s[v, t] & \mathbb{FD}_{s+1}[v, t] \\ \mathbb{FD}_{s-1}[v, t+1] & \mathbb{FD}_s[v, t+1] & \mathbb{FD}_{s+1}[v, t+1] \\ \mathbb{BD}_{s-1}[v, t-1] & \mathbb{BD}_s[v, t-1] & \mathbb{BD}_{s+1}[v, t-1] \\ \mathbb{BD}_{s-1}[v, t] & \mathbb{BD}_s[v, t] & \mathbb{BD}_{s+1}[v, t] \\ \mathbb{BD}_{s-1}[v, t+1] & \mathbb{BD}_s[v, t+1] & \mathbb{BD}_{s+1}[v, t+1] \end{array} \right\},$$

where $\mathbb{FD}_s[v, t] = (\mathbb{FD}_s)[v, t]$, $\mathbb{BD}_s[v, t] = (\mathbb{BD}_s)[v, t]$, and F and B are two matrices describing forward and backward relationships among variates. Intuitively, $\mathbb{FD}$ accounts for the combined DoS values of the forward neighbors and $\mathbb{BD}$ accounts for the combined DoS values of the backward neighbors of v . We declare the triple, $\langle v, t, s \rangle$, a candidate if the corresponding DoS value, $\mathbb{D}_s[v, t]$, is greater than $\Theta\%$ of the maximum of its 26 scale-neighbors in DoS, for some user provided $\Theta \sim 100$.

6.2 Local Extrema in Full Scale-Spaces

As seen earlier, an $\mathcal{L}$ -layer *full* state space, $\mathfrak{Y}_{full}$, is defined as a set of data matrices $\langle \mathbb{Y}_{0,0}, \dots, \mathbb{Y}_{i,j}, \dots, \mathbb{Y}_{l_{time}, l_{var}} \rangle$, where $\mathbb{Y}_{i,j} = \mathbb{Y}\{\langle \sigma_{time,0} \times k_{time}^i, \sigma_{var,0} \times k_{var}^j \rangle\}$. Given this, for each $s = \langle i, j \rangle$ pair, we can define three differences:

$$\mathbb{D}_{i,j}^t = |\mathbb{Y}_{i+1,j} - \mathbb{Y}_{i,j}|, \quad \mathbb{D}_{i,j}^v = |\mathbb{Y}_{i,j+1} - \mathbb{Y}_{i,j}|, \quad \text{and} \quad \mathbb{D}_{i,j}^{t,v} = |\mathbb{Y}_{i+1,j+1} - \mathbb{Y}_{i,j}|.$$

Local extrema are then identified by considering each variate-time-scale (VTS) triple, $\langle v, t, s \rangle$, and comparing $\max(\mathbb{D}_{i,j}^t, \mathbb{D}_{i,j}^v, \mathbb{D}_{i,j}^{t,v})$ to 78 ($= 3 \times 26$) neighboring triples⁵ of $\langle v, t, s \rangle$ in time, scale, and relationships for each of the $\mathbb{D}^t$, $\mathbb{D}^v$, and $\mathbb{D}^{t,v}$. We finally declare the triple, $\langle v, t, s \rangle$, a candidate if the corresponding DoS value, $\mathbb{D}_s[v, t]$, is greater than $\Theta\%$ of the maximum of its 78 neighbors in DoS, for some user provided $\Theta \sim 100$.

7 STEP 3: ELIMINATING POOR RMT FEATURE CANDIDATES

Local extrema of DoS can include candidate triples that are poorly localized. In order to identify whether a triple $\langle v, t, s \rangle$ is well or poorly localized in the scale space, we can consider the principal curvatures at the point $\langle v, t, s \rangle$ of the scale space generated earlier: a poorly defined peak in the difference-in-smoothing will have a large principal curvature in the scale space in one direction, but a small one in the perpendicular direction. Consequently, as was observed in References [15, 31], we can search for well-localized candidates by considering the ratio of the eigenvalues of the 2×2 Hessian matrix, which describes the local curvature of the scale space in terms of the second-order partial derivatives.

Given the above observation, the major challenge, in this case, is to define and compute the partial derivatives for metadata enhanced multi-variate time series to obtain the Hessian matrix we seek. More specifically, for each VTS triple, $\langle v, t, s \rangle$, we need to construct a 2×2 time/variates Hessian matrix, $\mathfrak{D}_{v,t,s}^{TV} = \begin{bmatrix} D_{T,T} & D_{T,V} \\ D_{V,T} & D_{V,V} \end{bmatrix}$, where

- $D_{T,T} = D_T D_T$ is the second derivative along time for the triple $\langle v, t, s \rangle$,
- $D_{V,V} = D_V D_V$ is the second derivative along “variate relationships” for $\langle v, t, s \rangle$,

⁵The number of neighboring triples may be less than 78 if the triple is at the boundary in terms of time, scale, or variate relationship graph.

- $D_{T,V} = D_T D_V$ is the partial derivative along time of the partial derivative along variate relationships of the triple $\langle v, t, s \rangle$, and
- $D_{V,T} = D_V D_T$ is the partial derivative along variate relationships of the partial derivative along time of the triple $\langle v, t, s \rangle$.

In this article, we propose to estimate the derivatives along time and variate relationships by taking differences of neighboring sample points:

$$D_T(v, t, s) = \mathbb{Y}_s[v, t+1] - \mathbb{Y}_s[v, t-1],$$

$$D_V(v, t, s) = \begin{cases} (\mathbb{F}\mathbb{Y}_s[v, t] - \mathbb{B}\mathbb{Y}_s[v, t]) & \text{for directed relationships} \\ (\mathbb{F}\mathbb{Y}_s[v, t] - \mathbb{Y}_s[v, t]) & \text{for undirected relationships.} \end{cases}$$

Here, $\mathbb{F}\mathbb{Y}_s$ and $\mathbb{B}\mathbb{Y}_s$ account for the (weighted) averages of the forward and backward variate neighbors at the corresponding scale: i.e., $\mathbb{F}\mathbb{Y}_s[v, t] = (F\mathbb{Y}_s)[v, t]$ and $\mathbb{B}\mathbb{Y}_s[v, t] = (B\mathbb{Y}_s)[v, t]$, where (as was discussed in the previous section) F and B are two matrices describing forward and backward relationships among variates.

Once the Hessian matrix, $\mathfrak{D}_{v,t,s}^{TV}$, is constructed for the triple $\langle v, t, s \rangle$, whether the triple is poorly localized can be checked using eigenvalue-based techniques [15, 31]. Note that derivatives (with respect to time) will be high at the boundaries of time (i.e., the beginning and end of the time series). Similarly, in directed variate relationship graphs, source and sink nodes are likely to have large derivatives with respect to the relationship space. Since many of these triples at the boundary of time and relationship do not correspond to real features of the data, but are essentially boundary noises, such candidate triples are removed even if they are well localized.

8 STEP 4: RMT FEATURE DESCRIPTOR CREATION

For data objects that can be represented as two-dimensional matrices (such as images), Reference [31] proposed that a gradient histogram-based descriptor around the given point $\langle x, y \rangle$ on the matrix could be constructed by computing a gradient for each element in the neighborhood of the point [31]. The resulting gradients are then quantized into c orientations. Finally, a $2a \times 2b$ grid is superimposed on the neighborhood region centered around the point and the gradients for the elements that fall into each cell are aggregated into a c -bin gradient histogram. This process leads to a feature descriptor vector of length $2a \times 2b \times c$. In Reference [4], we have shown that gradient histograms (created from data vectors instead of data matrices) are also effective in describing temporal features of uni-variate time series. In the case of multi-variate time series, however, we cannot directly apply the above techniques. Instead, we first need to construct an *extractor matrix* to enable the gradient extraction process.

8.1 Extractor Matrix

Let $\mathfrak{V}$ be a scale space defined over the given metadata-enriched multi-variate time series and the VTS triple, $\langle v, t, s \rangle$, be an RMT feature identified from $\mathfrak{V}$. The multi-variate feature defined by a variate-time-scale triple, $\langle v, t, s \rangle$, has an associated scope, defined by the scale, s , in which it is identified. The pair, $\langle v, t \rangle$, forms the center of the feature in time and variates. Given this feature center, under scale, s , which corresponds to temporal and variate smoothing parameter pair, $\Sigma = \langle \sigma_{time}, \sigma_{var} \rangle$, the temporal and variate scopes of the feature are computed as described in Sections 4.1.1 and 4.2.2, respectively.

As we have also seen in Section 4, observations closer in time and relationships to the triple will have significantly larger contributions to the feature than the points closer to the boundaries of the scope. Therefore, to identify gradients across time and variate relationships, we first construct an N -step aggregation series:

Definition 8.1 (*N*-Step Aggregation Series). For directed variate relationships, we define the *N*-step aggregation series corresponding to scale *s* as follows: For $-N < a \leq N$,

$$W_s[a] = \begin{cases} \text{if } a > 0 & (F^a \mathbb{Y}_s) \\ \text{if } a = 0 & \mathbb{Y}_s \\ \text{if } a < 0 & (B^a \mathbb{Y}_s), \end{cases}$$

where, as before, *F* and *B* are two matrices describing forward and backward relationships among variates. Similarly, in the case of undirected variate relationships, we define the *N*-step aggregation series, such that for $0 \leq a \leq N$ we have

$$W_s[a] = \begin{cases} \text{if } a > 0 & (F^a \mathbb{Y}_s) \\ \text{if } a = 0 & \mathbb{Y}_s. \end{cases}$$

Once the *N*-step aggregation series are obtained, we can then construct the extractor matrices from which the feature descriptors will be obtained.

Definition 8.2 (*Extractor Matrix*). Let $\langle v, t, s \rangle$ be a VTS triple on the scale space. In the case of directed variate relationships, we define the corresponding extractor matrix as a $2N \times 2M$ matrix, $X_{v,t,s}$, such that for $-N < a \leq N$ and $-M < b \leq M$, we have $X_{v,t,s}[a, b] = (W_s[a])[v, t + b]$. In the case of undirected variate relationships, we define the extractor matrix as a $(N + 1) \times 2M$ matrix, $X_{v,t,s}$, such that for $-N < a \leq N$ and $0 \leq b \leq M$, we have $X_{v,t,s}[a, b] = (W_s[a])[v, t + b]$.

The values of *N* and *M* should be selected to roughly cover the scope of the feature.

8.2 Descriptor Extraction

Given this extractor matrix, $X_{v,t,s}$, the feature descriptor is created as a *c*-directional gradient histogram of this matrix, sampling the gradient magnitudes around the salient point using a $2a \times 2b$ grid (or $2a \times b$ grid for undirected relationship graphs) superimposed on the matrix, $X_{v,t,s}$. To give less emphasis to gradients that are far from the point $\langle v, t \rangle$, a Gaussian weighting function is used to reduce the magnitude of elements further from $\langle v, t \rangle$.

This process leads to a feature descriptor vector of length $2a \times 2b \times c$ (or $2a \times b \times c$ for undirected graphs). The descriptor size must be selected in a way that reflects the temporal characteristics of the time series; if a multi-variate time series contains many similar features, then it might be more advantageous to use large descriptors that can better discriminate: these large descriptors would not only include information that describe the corresponding features, but would also describe the temporal contexts in which these features are located.

9 RMT FEATURE SET OF A MULTI-VARIATE TIME SERIES

Given the above, the RMT features of a MM time series, $Y = (\mathcal{V}, \mathcal{M}, \mathbb{Y}, \mathcal{D})$, with respect to the parameters

- $\Sigma_0 = \langle \sigma_{time,0}, \sigma_{var,0} \rangle$; i.e., the smallest temporal and variate smoothing parameters,
- $\mathcal{L} = \langle l_{time}, l_{var} \rangle$; i.e., the number of temporal and variate smoothing layers, and
- $\mathcal{K} = \langle k_{time}, k_{var} \rangle$; i.e., the scaling parameters for temporal and variate smoothings,

is defined as a set, $\mathcal{F}$, where each feature, $f \in \mathcal{F}$, extracted from *Y*, is a pair of the form, $f = \langle pos, \vec{d} \rangle$:

- $pos = \langle v, t, s \rangle$ is a VTS triple denoting the position of the feature in the scale space of the multi-variate time series, where *v* is the index of the variate at which the feature is *centered*,

t is the time instant around which the duration of the feature is *centered*, and s is the temporal/variante smoothing scale in which the feature is identified. Note that this triple also defines the *temporal and variates scopes* of the RMT feature.

- $\vec{d}$ is a vector of length $2a \times 2b \times c$ for directed relationship graphs and $2a \times b \times c$ for undirected graphs, as described in the previous section.

Note that this set contains RMT features of potentially different sizes. In particular, we have $\sigma_{time,min} = \sigma_{time,0}$, $\sigma_{var,min} = \sigma_{var,0}$, $\sigma_{time,max} = \sigma_{time,0} \times k_{time}^{l_{time}}$, and $\sigma_{var,max} = \sigma_{var,0} \times k_{var}^{l_{var}}$, and these define the minimum and maximum temporal and variate scopes of the features identified from the given multi-variate time series.

10 TIME SERIES MATCHING USING RMT FEATURES

This feature set can be used for various applications, including alignment, indexing, and classification of multi-variate series. Let us be given two MM time series, Y_1 and Y_2 and their feature sets $\mathcal{F}_1$ and $\mathcal{F}_2$. We rely on the alignments of the feature pairs in $\mathcal{F}_1$ and $\mathcal{F}_2$ to measure how well these two series match each other.

10.1 Alignment of Feature Pairs

Let $f_1 = \langle \langle v_1, t_1, s_1 \rangle, \vec{d}_1 \rangle$ and $f_2 = \langle \langle v_2, t_2, s_2 \rangle, \vec{d}_2 \rangle$ be two RMT features in $\mathcal{F}_1$ and $\mathcal{F}_2$, respectively. When matching f_1 and f_2 , we consider how well aligned as well as how important these two features are.

10.1.1 Temporal Alignment of a Pair of Features. Two features are said to be temporally aligned if their temporal scopes overlap significantly and temporal centers are close.

Definition 10.1 (Temporal Overlap). Let $[t_{s_1}, t_{e_1})$ denote the temporal scope of the first feature defined by t_1 and the temporal smoothing parameter corresponding to the feature scale s_1 . Similarly, let $[t_{s_2}, t_{e_2})$ denote the temporal scope of the first feature defined by t_2 and the feature scale s_2 . We define the temporal overlap score of the two features as $Overlap_T(f_1, f_2) = \frac{\min(t_{e_1}, t_{e_2}) - \max(t_{s_1}, t_{s_2})}{\max(t_{e_1}, t_{e_2}) - \min(t_{s_1}, t_{s_2})}$.

Definition 10.2 (Temporal Center Proximity). We define the temporal proximity score as $Prox_T(f_1, f_2) = 1 - \frac{|t_1 - t_2|}{maxLength}$, where $maxLength$ is the length of the time series.

Given these, we define temporal alignment score as follows:

Definition 10.3 (Temporal Alignment). We define the temporal alignment score of the two features as $Align_T(f_1, f_2) = \frac{Overlap_T(f_1, f_2) + Prox_T(f_1, f_2)}{2}$.

10.1.2 Variate Alignment of a Pair of Features. Two features are said to be variate aligned if their variate scopes overlap significantly:

Definition 10.4 (Variate Alignment). Let $scope(v_1, \sigma_{var,1})$ denote the variate scope of the first feature defined by parameter $\sigma_{var,1}$ corresponding to feature scale s_1 . Similarly, let $scope(v_2, \sigma_{var,2})$ be the variate scope of the second feature. We define the variate alignment score of the two features as $Align_V(f_1, f_2) = \frac{scope(v_1, \sigma_{var,1}) \cap scope(v_2, \sigma_{var,2})}{scope(v_1, \sigma_{var,1}) \cup scope(v_2, \sigma_{var,2})}$.

10.1.3 Descriptor Alignment of a Pair of Features. Two features are said to be descriptor aligned if their descriptor vectors are similar to each other:

Definition 10.5 (Descriptor Alignment). We define the descriptor alignment score of the two features as $Align_D(f_1, f_2) = sim(\vec{d}_1, \vec{d}_2)$ or as $Align_D(f_1, f_2) = (1 + \Delta(\vec{d}_1, \vec{d}_2))^{-1}$ for a given similarity, $sim()$, or distance, $\Delta()$, function.

10.1.4 Amplitude Alignment of a Pair of Features. Two features are said to be amplitude aligned if the average amplitudes of the time series within the corresponding feature scopes are similar to each other.

Definition 10.6 (Amplitude Alignment). We define the amplitude alignment score as $Align_A(f_1, f_2) = (1 + |ampl_1 - ampl_2|)^{-1}$, where $ampl_1$ and $ampl_2$ are the average amplitudes of the time series, Y_1 and Y_2 , within the scopes of f_1 and f_2 .

10.2 Feature Significance

10.2.1 Scope Significance of a Given Pair of Features. The size of the temporal and/or variate scopes may impact the significance of a feature.

Definition 10.7 (Temporal Scope Significance). The combined temporal scope significance of f_1 and f_2 is defined as $Sig_T(f_1, f_2) = \frac{\sigma_{time,1} + \sigma_{time,2}}{2}$, where $\sigma_{time,1}$ and $\sigma_{time,2}$ are the two temporal smoothing parameters corresponding to temporal scales, s_1 and s_2 .

Definition 10.8 (Variate Scope Significance). The combined variate scope significance of f_1 and f_2 is defined as $Sig_V(f_1, f_2) = \frac{\sigma_{var,1} + \sigma_{var,2}}{2}$, where $\sigma_{var,1}$ and $\sigma_{var,2}$ are the two variate smoothing parameters corresponding to feature scales, s_1 and s_2 .

10.2.2 Contextual Significance of Features. In many applications, we also need to consider how discriminating or contextually important a feature is. For example, in a classification scenario, when comparing features f_1 and f_2 from these time series, we may also consider how representative (frequent and discriminating) f_1 and f_2 are as the contextual importance measure, $Sig_C(f_1, f_2)$.

10.3 Overall Feature Matching Score

Given the above, the overall matching score of two features is a combination of the individual measures of alignment and importance:

Definition 10.9 (Overall Feature Matching Score). We define the overall matching score, $match(f_1, f_2)$, of the two features as

$$\mu \left(\begin{matrix} Align_D(f_1, f_2), Align_T(f_1, f_2), Align_V(f_1, f_2), Align_A(f_1, f_2), \\ Sig_T(f_1, f_2), Sig_V(f_1, f_2), Sig_C(f_1, f_2) \end{matrix} \right),$$

where μ is a merge function that combines the individual scores.

While there exist different merge functions (such as *min*, *max*, *sum*, *avg*, *product*), in the experiments reported in Section 11 we use *product*, which approximates the Boolean operator *and* when individual scores are zeros and ones [5].

10.4 Identifying Candidate Matching Pairs

Given a query time series, Y_q , and a data series, Y_d , and their feature sets $\mathcal{F}_q$ and $\mathcal{F}_d$, the next step is to identify a set $\mathcal{P} \subseteq \mathcal{F}_q \times \mathcal{F}_d$ of candidate feature pairs, such that

- $\forall f_{q,i} \in \mathcal{F}_q \exists f_{d,j} \in \mathcal{F}_d$ s.t. $\langle f_{q,i}, f_{d,j} \rangle \in \mathcal{P}$ (i.e., for each query RMT feature on the query object, at least one matching feature on the data object is located),
- $\forall \langle f_{q,i}, f_{d,j} \rangle, \langle f_{q,h}, f_{d,k} \rangle \in \mathcal{P} \ (f_{q,i} = f_{q,h}) \rightarrow (f_{d,j} = f_{d,k})$ (i.e., for each query RMT feature on the query object, at most one matching feature on the data object is located), and
- $\sum_{\langle f_{q,i}, f_{d,j} \rangle \in \mathcal{P}} match(f_{q,i}, f_{d,j})$ is maximized.

It is easy to see that, since, for each query feature, $\mathcal{P}$ contains one and only one matching data feature and since we aim to maximize the overall matching score, the set $\mathcal{P}$ can be obtained by considering each feature $f_{q,i} \in \mathcal{F}_q$ and selecting the feature $f_{d,j} \in \mathcal{F}_d$ with the maximum

Fig. 6. Example scope boundary conflicts: Blue lines mark corresponding starting points of the matching scopes, whereas red lines mark the corresponding end points.

Fig. 7. (a) Candidate RMT feature pairs for two multi-variate time series and (b) the remaining subset of matching RMT feature pairs after inconsistency pruning.

$match(f_{q,i}, f_{d,j})$ value. However, the feature pairs in $\mathcal{P}$ obtained this way may not be mutually consistent and such inconsistencies need to be eliminated.

10.5 Inconsistency Pruning of Candidate Pairs

Intuitively, we call a set of feature matchings *temporally consistent* if the corresponding features are similarly ordered in both time series. Figure 6 shows several temporal inconsistencies, where temporal scope boundaries of matching features are not similarly ordered in two time series. Formally, we define temporal consistency as follows:

Definition 10.10 (Temporal Consistency). Let us be given two MM time series, Y_1 and Y_2 , and their feature sets $\mathcal{F}_1$ and $\mathcal{F}_2$. Let $p_1 = \langle f_1^1, f_2^1 \rangle, p_2 = \langle f_1^2, f_2^2 \rangle \in \mathcal{P}$ be two candidate feature pairs and $bounds_b^a = \{ts_b^a, te_b^a\}$ be the start and end points of the temporal scope of feature f_b^a for $a, b \in \{1, 2\}$. We call p_1 and p_2 *temporally consistent* if and only if

- for all pairs of end points $t_i^1, t_j^1 \in bounds_1^1 \cup bounds_2^1$ in the first pair, we have

$$\left((t_i^1 > t_j^1) \rightarrow (t_i^2 \neq t_j^2) \right) \wedge \left((t_i^1 < t_j^1) \rightarrow (t_i^2 \neq t_j^2) \right),$$

where $t_i^2, t_j^2 \in bounds_1^2 \cup bounds_2^2$ are the two end points in the second pair corresponding to t_i^1 and t_j^1 ; and

- for all pairs of end points $t_i^2, t_j^2 \in bounds_1^2 \cup bounds_2^2$ in the second pair, we have

$$\left((t_i^2 > t_j^2) \rightarrow (t_i^1 \neq t_j^1) \right) \wedge \left((t_i^2 < t_j^2) \rightarrow (t_i^1 \neq t_j^1) \right),$$

where $t_i^1, t_j^1 \in bounds_1^1 \cup bounds_2^1$ are the two end points in the first pair corresponding to t_i^2 and t_j^2 .

Figure 7 provides an example with inconsistent matches: Here we see that the matching process identified some very distant pairs of RMT features as matches. Note also that there are many matching pairs that cross each other in time, implying temporal features that are differently ordered in time in two time series. To improve the accuracy of the matching process, we need to eliminate such inconsistencies. The outline of the process to eliminate inconsistencies is as follows

- (1) For each pair, $\langle f_1, f_2 \rangle \in \mathcal{P}$ of matching features, we compute a dominance score, $dom(f_1, f_2)$, as

$$\rho \left(\begin{matrix} Align_D(f_1, f_2), Align_T(f_1, f_2), Align_V(f_1, f_2), Align_A(f_1, f_2), \\ Sig_T(f_1, f_2), Sig_V(f_1, f_2), Sig_C(f_1, f_2) \end{matrix} \right).$$

Note that this dominance score may, but is not required to, be the same as the overall matching score discussed in Section 10.3.

- (2) We next initialize an empty set ($\mathcal{R}$) to collect the committed consistent feature pairs and two empty lists ($list_1$ and $list_2$) to keep track of their temporal scopes: i.e., we set $\mathcal{R} = \emptyset$, $list_1 = \perp$, and $list_2 = \perp$.
- (3) Next, we consider all pairs of matching features in $\mathcal{P}$ in descending order of their dominance scores. Let $\langle f_1, f_2 \rangle \in \mathcal{P}$ be the pair we are currently considering.
- (a) *Temporal consistency verification:* Let $\langle ts_1, te_1 \rangle$ and $\langle ts_2, te_2 \rangle$ be the temporal scopes of f_1 and f_2 , respectively
- (i) We *attempt* to insert the ts_1 and te_1 into $list_1$ ordered in increasing order of time; similarly, we *attempt* to insert ts_2 and te_2 into the list, $list_2$, also ordered in increasing order of time.
 - (ii) Let $rank(ts_1)$, $rank(ts_2)$, $rank(te_1)$, and $rank(te_2)$ be the corresponding ranks of the time points in their respective time ordered lists.
 - (iii) If $rank(ts_1) = rank(ts_2)$ and $rank(te_1) = rank(te_2)$, then we confirm the insertion and we keep the pair.⁶
 - (iv) Else, we drop the pair $\langle f_1, f_2 \rangle$ and eliminate the corresponding scope boundaries from the lists $list_1$ and $list_2$.
- (b) If the candidate pair $\langle f_1, f_2 \rangle$ has not been dropped due to temporal inconsistency, then insert the pair in $\mathcal{R}$: i.e., $\mathcal{R} \rightarrow \mathcal{R} \cup \{\langle f_1, f_2 \rangle\}$.

Note that the reason why the feature pairs are considered in descending order of dominance scores is that, when an inconsistency is identified, the most recently considered pair—which is less dominant (relatively less aligned, smaller, and less similar)—can be eliminated without affecting the already-committed boundaries.

10.6 RMT-Based Multi-variate Time Series Matching Score

Given two MM time series, Y_1 and Y_2 , and their feature sets $\mathcal{F}_1$ and $\mathcal{F}_2$, the above process results in a set $\mathcal{R}(\mathcal{F}_1, \mathcal{F}_2) = \{\langle f_{1,i}, f_{2,i} \rangle\}$, where $f_{1,i} \in \mathcal{F}_1$ and $f_{2,i} \in \mathcal{F}_2$, respectively. We define the overall matching score, $score(Y_1, Y_2)$, of the two multi-variate series, using this set of matching feature pairs:

$$\sum_{\langle f_1, f_2 \rangle \in \mathcal{R}(\mathcal{F}_1, \mathcal{F}_2)} \phi \left(\begin{matrix} Align_D(f_1, f_2), Align_T(f_1, f_2), Align_V(f_1, f_2), Align_A(f_1, f_2), \\ Sig_T(f_1, f_2), Sig_V(f_1, f_2), Sig_C(f_1, f_2) \end{matrix} \right),$$

where ϕ is a combined scoring function.

While there exist different merge functions (such as *min*, *max*, *sum*, *avg*, *product*), in the experiments reported in Section 11 we use *product*, which approximates the Boolean operator *and* when individual scores are zeros and ones [5].

⁶The process is slightly more complex in that there can be exceptions where the ranks are different, but time values are the same. We also confirm the insertion in these special cases.

11 EVALUATION

In this section, we present experiment results that assess the efficiency and effectiveness of the RMT⁷ feature extraction algorithms. In our preliminary work [51], we had shown that the *diagonal* scale-space-based RMT features (Section 5.1) are more effective in partial time-series search and classification tasks than alternative techniques, including SVD, where we created a single fingerprint for each multi-variate time series using the SVD transformation, and DTW, where distances were computed directly using dynamic time warping [9]. In the appendix, we also consider SAX [28] DTW, which provides time savings over DTW, possibly at the expense of accuracy. Therefore, instead of replicating the experiments reported in Reference [51], we focus on the impact of *full* scale-space-based RMT (Section 5.2) features with respect to the use of *diagonal* scale-space-based RMT (Section 5.1) and also investigate the impacts of the alternative matching and inconsistency removal strategies described in Section 10, within the context of a motion recognition task.

11.1 Settings

11.1.1 Hardware/Software. In order to ensure results are comparable to those reported in Reference [51], all experiments were run on the identical setup, including four-core Intel Core i5-2400 3.10GHz machines with 8GB memory, running 64-bit Windows 7 Enterprise, using Matlab.

11.1.2 Dataset. For the experiments in this section, we use the Mocap time series dataset [34]: The dataset consists of movement records from markers placed on subjects' bodies as they perform eight types of tasks. We use the Acclaim Skeleton File (ASF) and Acclaim Motion Capture (AMC) format where the original coordinate readings are converted into 62 joint angles data. We treat readings for each joint angle as a different uni-variate time series. The hierarchical spatial distribution (e.g. left foot, right foot, left leg, etc.) of the joint angles on the body is used to create the underlying correlation matrix used as metadata.⁸

We consider additional datasets in the online appendix.

11.1.3 Evaluation Metrics. For evaluating accuracy, we use take-one-out methodology with the following criteria: (a) *top-5 precision*: The number of series, among the nearest five results, that are of the same class of movement as the query series, and (b) *top- $\|c\|$ precision*: The number of series, among the nearest $\|c\|$ results (where $\|c\|$ is the size of the movement class containing the query series), that are of the same class as the query series. The first measure reflects how effective a particular approach is for nearest-neighbor classification, whereas the second measure reflects how well defined the classes.

In addition, we also report pairwise matching times for the alternative approaches. Note that since we are using top-5 and top- $\|c\|$ classification, the classification time is a function of the value of $\|c\|$, the number of labeled data in the training dataset, and the pairwise matching time. To ensure that the efficiency of different algorithms can be compared independently of the value of $\|c\|$ and the training dataset, in the article, we report the pairwise matching time as an indicator of the classification cost.

11.1.4 Alternative RMT Features. We consider different types of RMT features:

- *Diagonal scale-space-based RMT (DIA)*: This is the version of the RMT features studied in our prior work. As described in Section 5.1, these features are extracted only by considering the

⁷RMT source code is available upon request see Reference [13].

⁸Note that this provides an *intentionally rough* metadata, enabling us to observe accuracy of RMT features under imperfect domain knowledge.

Table 1. (a) Default Configuration and (b) Alternative Matching/Pruning Strategies

(a) Default configuration		(b) Matching/pruning strategies	
RMT		TO	Temporal overlap (Definition 10.1)
# iterations, L	6	TP	Temporal proximity (Definition 10.2)
# of octaves, o	3	TA	Temporal alignment (Definition 10.3)
initial smoothing for time, $\sigma_{time,0}$	2.8	VA	Variate alignment (Definition 10.4)
initial smoothing for relationships, $\sigma_{var,0}$	0.5	DA	Descriptor alignment (Definition 10.5)
candidate pruning threshold, ω_T	10	AA	Amplitude alignment (Definition 10.6)
descriptor size, $2a \times 2b \times c$	$(4 \times 4 \times 8 =) 128$	TS	Temporal scope significance (Definition 10.7)
relationship reduction algorithm	k-means	VS	Variate scope significance (Definition 10.7)
SVD			
degree of energy preservation	95%		

diagonal scales of the scale space; in other words, the features' temporal and variate scopes grow in synch to each other.

- *Full scale-space-based RMT (FULL)*: This is the version of the RMT features proposed in this article. As described in Section 5.2, these are extracted by considering all scales of the scale space; consequently, the features' temporal and variate scopes grow independently from each other, enabling heterogeneously shaped features.
- *Hybrid RMT (HYB)*: We also consider hybrid feature sets, where diagonal scale-space features and full scale-space features are combined. Note that, due to the feature candidate elimination strategy described in Section 7, feature set obtained using the full scale-space is not necessarily the superset of the features obtained using the diagonal scale-space. This hybrid strategy re-introduces the diagonal scale-space features that may have been eliminated due to some features in the non-diagonal scales of the space.
- *Diagonal scale-space-based RMT-alt. 2 (DIA2)*: Note that diagonal scale-space-based RMT features can be obtained either by using only the diagonal scales of the scale-space as described in Section 5.1 or can be obtained by selecting the subset of the full scale-space-based RMT features such that the temporal and variate scales are the same. We refer to this second alternative as DIA2.

Table 1(a) provides the outline of the default parameter configuration and describes how these parameters are varied in the experiments.

11.1.5 Alternative Alignment Strategies. In this section, we experiment with the various temporal and variate alignment metrics presented in Section 10.1 and listed in Table 1(b). When needed, for combining these measures in Table 1(b), we use multiplication as the merge function. In addition, we consider two alignment strategies as follows: (a) *all octaves alignment (AoA)*: Under this strategy, any two pairs of features can be considered for alignment irrespective of their scales. (b) *same octave alignment (SoA)*: Under this strategy, only those pairs of features that have the same time and variate octaves are considered for alignment.

11.1.6 Alternative Inconsistency Elimination (Pruning) Strategies. In this section, we also consider the impact of the measures presented in Table 1 on the inconsistency elimination process (Section 10.5). In addition, we consider two pruning strategies: (a) *all octaves pruning (AoP)*: Under this strategy, any two pairs of features can be considered inconsistent irrespective of their scales. (b) *same octave pruning (SoP)*: Under this strategy, only those pairs of features that have the same time and variate octaves can be considered inconsistent.

Fig. 8. Top-five matching accuracy and matching time—default configuration: descriptor alignment– (DA) based feature matching and DA-based inconsistency pruning and overall score computation.

11.2 Discussion of the Results

11.2.1 Overview. Figure 8(a) compares the classification accuracy of RMT using the eight classes with 184 motions in the Mocap dataset against the alternative approaches reported in [51]:

- *Variate-paired* alignment: This is the default configuration where we assume that the pairing of the variates in the query and in the database are known in advance. DTW requires that this pairing is known. In the case of RMT, we leverage the pairing information by ignoring feature matches during the feature alignment phase unless at least 50% of the variates are common. As we see in Figure 8, paired RMT provides the best overall accuracy.
- *Non-variate-paired* alignment: Both SVD and RMT can operate without requiring pairing of the variates. Given two multi-variate time series, SVD uses the decomposed series rather than the series themselves, and thus it does not require the series to be variate paired. Similarly, RMT can be implemented in such a way that variate alignments are completely ignored during the matching phase. As we see in Figure 8, non-paired RMT works better than SVD and thus is applicable when pairing information is not available. While SVD supports fast matching, the accuracy is significantly lower to render it a feasible approach.

Note that Figure 8(b) also includes confidence intervals for the accuracies of various techniques. As we see here, RMT's confidence intervals do not overlap with the other techniques' accuracy confidence intervals, providing additional evidence for the advantage of using RMT features. Moreover, the confidence intervals of RMT are significantly tighter than the confidence intervals of other techniques, again providing evidence that RMT is more robust than the other approaches.

Figure 8(a) shows that, as expected, we obtain highest accuracy when we consider the full scale space. It is also important to note that in these experiments we do not leverage RMT feature significance (FS) to boost matching accuracy. Unlike DTW, RMT can further boost accuracy through relevance feedback and other (semi-)supervised learning techniques.

11.2.2 Impact of the Alternative Feature Matching and Inconsistency Pruning Strategies. As we have seen in Section 10 and Table 1, one can use several strategies to match features across multi-variate time series and prune inconsistencies. For the results above, as the default configuration, we considered descriptor alignment– (DA) based feature matching, inconsistency, pruning, and overall score computation. While the best strategy is application dependent, as we see in Figure 9, in this application, RMT is able to achieve high accuracy by using DA for feature matching,

Fig. 11. Impact of feature discovery and octave management using DA for feature matching, inconsistency pruning, and overall scoring.

This shows that the RMT feature descriptors are highly informative. This is further confirmed by Figure 10, where we see the average number of (post-pruning) matching feature pairs for the alternative strategies considered in Figure 9. As we see in this figure, a higher number of matching feature pairs does not translate into more accurate matches. This indicates that the resulting RMT features are highly informative and a small number of feature pairs at different scales and shapes are sufficient to characterize different types of motion.

Impact of the Feature Discovery and Octave Management Strategies. For the default results presented above, we leveraged full (FULL) feature set with same octave alignment (AoA) and same octave pruning (SoP) strategies. In Figure 11, we study the impact of these strategies in

Table 2. # Feature Pairs Before and After Inconsistency Pruning for the Optimal Configuration in Figure 11

	Before inconsistency pruning	After inconsistency pruning
DIA	952.2	21.5
FULL	1252.1	48.2
HYB	2204.3	56.3

further detail. As we see in this figure, the default configuration indeed leads to highest accuracy: First, as expected, feature matches and inconsistencies need to be considered at each octave scale separately. Second, the figure shows that the full scale space provide more information than the diagonal scale space—in fact, extending the FULL feature set with diagonal features (i.e., using the HYB strategy) does not lead to any better results than just using the FULL or DIA feature sets.

This is further studied in Table 2, which shows the average matching # feature pairs, before and after inconsistency pruning, for the optimal configuration in Figure 11. As we see here, inconsistency pruning eliminates a large number of feature pairs. We also see that the hybrid option (HYB), has more feature pairs than both DIA and FULL, but, as we have seen Figure 11, these additional feature pairs do not contribute to the accuracy.

12 CONCLUSIONS

Many time-series datasets are (a) multi-variate, (b) interrelated, and (c) multi-resolution. These include motion and gesture data, as described in this article, as well as data from other domains: In this article, we presented a metadata-enriched multi-variate time series model, in which a dependency/correlation model relates the individual variates to each other. Recognizing that multi-variate temporal features can be extracted more effectively by simultaneously considering, at multiple scales, differences among individual variates along with the dependency/correlation model that relates them, we further developed algorithms to detect RMT features that are multi-resolution, local, and invariant against various types of noise. Experiments using human motion data, where labeled ground truth is available, confirmed that the RMT features are highly effective in multi-variate series search and classification.

REFERENCES

- [1] Iyad Batal, Dmitriy Fradkin, James Harrison, Fabian Moerchen, and Milos Hauskrecht. 2012. Mining recent temporal patterns for event detection in multivariate time series data. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'12)*. 280–288.
- [2] Donald J. Bemdt and James Clifford. 1994. *Using Dynamic Time Warping to Find Patterns in Time Series*.
- [3] David M. Blei and John D. Lafferty. 2006. Dynamic topic models. In *Proceedings of the 33rd International Conference on Machine Learning (ICML'06)*. 113–120.
- [4] K. Selçuk Candan, Rosaria Rossini, Maria Luisa Sapino, and Xiaolan Wang. 2012. sDTW: Computing DTW distances using locally relevant constraints based on salient feature alignments. *Proc. VLDB* 5, 11 (Jul. 2012), 1519–1530.
- [5] K. Selçuk Candan and Maria Luisa Sapino. 2010. *Data Management for Multimedia Retrieval*. Cambridge University Press, New York, NY.
- [6] N. Castro and P. Azevedo. 2010. Multiresolution motif discovery in time series. In *Proceedings of the SIAM International Conference on Data Mining (SDM'10)*. SIAM, Columbus, OH, 665–676.
- [7] Lei Chen. 2005. *Similarity Search over Time Series and Trajectory Data*. Ph.D. Dissertation. University of Waterloo.
- [8] Lei Chen and Raymond Ng. 2004. On the marriage of Lp-norms and edit distance. In *Proceedings of the Thirtieth international conference on Very large data bases (VLDB'04)* 30 (2004), 792–803.
- [9] Yanping Chen, Eamonn Keogh, Bing Hu, Nurjahan Begum, Anthony Bagnall, Abdullah Mueen, and Gustavo Batista. 2015. The UCR Time Series Classification Archive. Retrieved from www.cs.ucr.edu/~eamonn/time_series_data/.
- [10] F. L. Chung, T. C. Fu, R. Luk, and V. Ng. 2001. Flexible time series pattern matching based on perceptually important points. In *Proceedings of the IJCAI Workshop on Learning from Temporal and Spatial Data* (2001), 1–7.

- [11] Hui Ding, Goce Trajcevski, Peter Scheuermann, Xiaoyue Wang, and Eamonn Keogh. 2008. Querying and mining of time series data: Experimental comparison of representations and distance measures. In *Proceedings of the VLDB Endowment* 1, 2 (2008), 1542–1552.
- [12] Michael Eichler. 2006. Granger causality and path diagrams for multivariate time series. *J. Econometr.* (2006).
- [13] EmitLab-ASU. 2017. RMT Code. (2017). Available upon request.
- [14] Philippe Esling and Carlos Agon. 2012. Time-series data mining. *ACM Comput. Surv.* 45, 1 (Dec. 2012), 12:1–12:34.
- [15] C. Harris and M. Stephens. 1988. A combined corner and edge detector. In *Proceedings of the 4th Alvey Vision Conference* (1988), 147–151.
- [16] A. Harvey and S. Koopman. 1997. Multivariate structural time series model. In *System Dynamics in Economic and Financial Models*. John Wiley & Sons, 269–296.
- [17] A. S. Hopkins, A. Lekov, J. Lutz, G. Rosenquist, and L. Gu. 2011. Simulating a nationally representative housing sample using energyplus. (2011), 55.
- [18] Xiaonan Ji, James Bailey, and Guozhu Dong. 2007. Mining minimal distinguishing subsequence patterns with gap constraints. In *Knowledge and Information Systems* 11, 3 (2007), 259–286.
- [19] Yohan Jin and Balakrishnan Prabhakaran. 2011. Knowledge discovery from 3D human motion streams through semantic dimensional reduction. *ACM Trans. Multimedia Comput. Commun. Appl.* 7, 2 (2 2011). DOI : <http://dx.doi.org/10.1145/1925101.1925104>
- [20] Eamonn Keogh. 2002. Exact indexing of dynamic time warping. In *Proceedings of the 28th International Conference on Very Large Data Bases (VLDB'02)*. Hong Kong, China, 406–417.
- [21] Eamonn Keogh and Chotirat Ann Ratanamahatana. 2005. Exact indexing of dynamic time warping. *Knowledge and Information Systems* 7, 3 (2005), 358–386.
- [22] Duk-Jin Kim and B. Prabhakaran. 2011. Motion fault detection and isolation in body sensor networks. *Perv. Mobile Comput.* 7, 6 (2011), 727–745. DOI : <http://dx.doi.org/10.1016/j.pmcj.2011.09.006>
- [23] Tamara G. Kolda and Brett W. Bader. 2009. Tensor decompositions and applications. *SIAM Rev.* 51, 3 (2009), 455–500.
- [24] Joseph B. Kruskal. 1983. An overview of sequence comparison: Time warps, string edits, and macromolecules. *SIAM Rev.* 25, 2 (1983), 201–237.
- [25] W. Krzanowski. 1979. Between-groups comparison of principal components. *Journal of the American Statistical Association* 74, 367 (1979), 703–707.
- [26] Chuanjun Li, S. Q. Zheng, and B. Prabhakaran. 2007. Segmentation and recognition of motion streams by similarity search. *ACM Trans. Multimedia Comput. Commun. Appl.* 3, 3, Article 16 (Aug. 2007). DOI : <http://dx.doi.org/10.1145/1236471.1236475>
- [27] Lei Li, B. Aditya Prakash, and Christos Faloutsos. 2010. Parsimonious linear fingerprinting for time series. *Proceedings of the VLDB Endowment* 3, 1–2 (2010), 385–396.
- [28] Jessica Lin, Eamonn Keogh, Stefano Lonardi, and Bill Chiu. 2003. A symbolic representation of time series, with implications for streaming algorithms. In *Proceedings of the 8th ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery (DMKD'03)*. ACM, 2–11.
- [29] Sicong Liu, Yash Garg, K Selcuk Candan, Maria Luisa Sapino, and Gerardo Chowell. 2015. NOTES2: Networks-of-traces for epidemic spread simulations. In *Proceedings of the AAAI International Workshop on Computational Sustainability (co-located with AAAI'15)*.
- [30] D. G. Lowe. 1999. Object recognition from local scale-invariant features. In *Proceedings of the International Conference on Computer Vision (ICCV'99)*.
- [31] D. G. Lowe. 2004. Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vis.* 60, 2 (2004), 91–110.
- [32] Sachin Mehta, Rajarathnam Nallusamy, Ranjeet Vinayak Marawar, and Balakrishnan Prabhakaran. 2013. A study of DWT and SVD based watermarking algorithms for patient privacy in medical images. In *Proceedings of the IEEE International Conference on Healthcare Informatics (ICHI'13)*. 287–296. DOI : <http://dx.doi.org/10.1109/ICHI.2013.41>
- [33] Terence C. Mills. 1990. *Time Series Techniques for Economists*. Cambridge University Press.
- [34] Mocap. 2001. CMU Mocap data set. Retrieved from <http://mocap.cs.cmu.edu/>.
- [35] Yasser Mohammad and Toyoaki Nishida. 2009. Constrained motif discovery in time series. *New Generat. Comput.* 27, 4 (2009), 319–346.
- [36] Fabian Mörchen. 2003. Time series feature extraction for data mining using DWT and DFT. Philipps-University Marburg.
- [37] Jinglin Peng, Hongzhi Wang, Jianzhong Li, and Hong Gao. 2016. Set-based similarity search for time series. In *Proceedings of the 2016 International Conference on Management of Data, SIGMOD Conference 2016*, Fatma Özcan, Georgia Koutrika, and Sam Madden (Eds.). ACM, 2039–2052. DOI : <http://dx.doi.org/10.1145/2882903.2882963>
- [38] C. S. Perng, H. Wang, S. R. Zhang, and D. S. Parker Jr. 2000. Landmarks: A new model for similarity-based pattern querying in time series databases. In *Proceedings of the IEEE International Conference on Data Engineering (ICDE'00)*. 33–42.

- [39] Silvestro Roberto Poccia and Yash Garg. 2017. On the effectiveness of distance measures for similarity search in multi-variate sensory data: Effectiveness of distance measures for similarity search. In *Proceedings of the 8th International Conference on Management Research (ICMR'17)*. 489–493.
- [40] Silvestro Roberto Poccia, Maria Luisa Sapino, Xilun Chen Sicong Liu, Yash Garg, Shengyu Huang, Jung Hyun Kim, Xinsheng Li, Parth Nagarkar, and K. Selçuk Candan. 2017. SIMDMS: Data management and analysis to support decision making through large simulation ensembles. In *Proceedings of the International Conference on Extending Database Technology (EDBT'17)*. 582–585.
- [41] Thanawin Rakthanmanon, Bilson Campana, Abdullah Mueen, Gustavo Batista, Brandon Westover, Qiang Zhu, Jesin Zakaria, and Eamonn Keogh. 2012. Searching and mining trillions of time series subsequences under dynamic time warping. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'12)*.
- [42] T. Rakthanmanon and E. Keogh. 2013. Fast-shapelets: A scalable algorithm for discovering time series shapelets. In *Proceedings of the SIAM International Conference on Data Mining (SDM'13)*.
- [43] Hiroaki Sakoe and Seibi Chiba. 1978. Dynamic programming algorithm optimization for spoken word recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 26, 1 (1978), 43–49.
- [44] Gerard Salton and Michael J. McGill. 1983. *Introduction to Modern Information Retrieval*.
- [45] Parinya Sanguansat. 2012. Multiple multidimensional sequence alignment using generalized dynamic time warping. *WSEAS Transactions on Mathematics* 11, 8 (2012), 668–678.
- [46] Liang Shuai, Chao Li, Xiaohu Guo, Balakrishnan Prabhakaran, and Jinxiang Chai. 2017. Motion capture with ellipsoidal skeleton using multiple depth cameras. *IEEE Trans. Vis. Comput. Graph.* 23, 2 (2017), 1085–1098. DOI: <http://dx.doi.org/10.1109/TVCG.2016.2520926>
- [47] A. de Silva, R. J. Hyndman, and R. Snyder. 2010. The vector innovations structural time series framework: A simple approach to multivariate forecasting. *Statist. Model.* 10, 4 (2010), 353–374.
- [48] Ugo Vespier, Arno Knobbe, Siegfried Nijssen, and Joaquin Vanschoren. 2012. *MDL-Based Analysis of Time Series at Multiple Time-Scales*. 371–386.
- [49] Michail Vlachos, Marios Hadjieleftheriou, Dimitrios Gunopulos, and Eamonn Keogh. 2006. Indexing multidimensional time-series. *VLDB J.* 15, 1 (Jan. 2006), 1–20. DOI: <http://dx.doi.org/10.1007/s00778-004-0144-2>
- [50] Xinxin Wang and K. Selçuk Candan. 2010. Relevant shape contour snippet extraction with metadata supported hidden markov models. In *Proceedings of the Conference on Image and Video Retrieval (CIVR'10)*. 430–437.
- [51] X. Wang, K. S. Candan, and M. L. Sapino. 2014. Leveraging metadata for identifying local, robust multi-variate temporal (RMT) features. In *Proceedings of the 2014 IEEE 30th International Conference on Data Engineering (ICDE'14)*. IEEE, 388–399.
- [52] Xing Wang, Jessica Lin, Pavel Senin, Tim Oates, Sunil Gandhi, Arnold P. Boedihardjo, Crystal Chen, and Susan Frankenstein. 2016. RPM: Representative pattern mining for efficient time series classification. In *Proceedings of the 19th International Conference on Extending Database Technology (EDBT'16)*. 185–196. DOI: <http://dx.doi.org/10.5441/002/edbt.2016.19>
- [53] Kiyoungh Yang and Cyrus Shahabi. 2004. A PCA-based similarity measure for multivariate time series. In *Proceedings of the ACM International Workshop on Multimedia Databases (MMDB'04)*. ACM, 65–74.
- [54] Dragomir Yankov, Eamonn Keogh, Jose Medina, Bill Chiu, and Victor Zordan. 2007. Detecting time series motifs under uniform scaling. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'07)*. ACM, 844–853.
- [55] Lexiang Ye and Eamonn Keogh. 2009. Time series shapelets: A new primitive for data mining. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'09)*. 947–956.

Received February 2017; revised June 2017; accepted August 2017