Discovering Multiple Subdimensional Motifs in Multivariate Time Series

Valerio Guerrini[1], Thibaut Germain[2], Charles Truong[3,4], Laurent Oudre[1]
[1]*Université Paris Saclay, Université Paris Cité, ENS Paris Saclay, CNRS, SSA, INSERM, Centre Borelli,* F-91190, Gif-sur-Yvette, France
[2] *École Polytechnique, Institut Polytechnique de Paris, CMAP,* F-91128, Palaiseau, France
[3] *The American University of Paris*
[4] *Université Paris Saclay, Université d'Evry Val d'Essonne, CNRS, LaMME,* F-91037, Evry, France

***Abstract*—Motif Discovery aims at identifying repeated patterns in time series. It is a fundamental task in time series analysis, with applications across numerous fields where recurring phenomena are observed. In multivariate time series, motifs may appear only on a subset of dimensions, making their identification more challenging. Existing approaches to multivariate Motif Discovery face limitations when many motifs are present. In particular, current methods struggle to handle ambiguities arising when motifs overlap in time but span disjoint subsets of dimensions. To address this issue, we formalize these ambiguous cases and propose a method based on a more expressive notion of overlap that accounts for both temporal alignment and dimensional involvement. Finally, we introduce a new evaluation framework that jointly assesses temporal and dimensional accuracy, and demonstrate its relevance through carefully designed experiments.**

***Index Terms*—Multivariate time series, Motif Discovery, Matrix Profile**

I. Introduction

Time series are time-ordered sequences of measurements, either univariate, when a single quantity is observed over time, or multivariate, when multiple quantities are recorded simultaneously. Many real-world systems, such as human motion, exhibit recurring phenomena, for example stereotypical movements, whose identification can help better understand the underlying processes. These phenomena are often reflected as repeated patterns in time series. The task of discovering such repeated patterns, commonly referred to as Motif Discovery, has been an active area of research for more than two decades. In practice, Motif Discovery relies on comparing subportions of time series, called subsequences. Subsequences that are sufficiently similar according to a given similarity measure are grouped together and interpreted as occurrences of the same motif. However, this task is inherently ambiguous, as the notion of being "sufficiently similar" is context-dependent and not uniquely defined. These difficulties are further amplified in the multivariate case, where motifs may only be supported by subsets of dimensions and Motif Discovery must jointly identify both temporal occurrences and active dimensions.

Related Work: Motif Discovery was first developed for univariate time series. Various distance measures have been proposed to quantify the similarity between subsequences in this setting. Among these, the Z-Normalized Euclidean [1] distance removes the mean and variance of each subsequence, enabling shape-based comparison independent of scale, whereas Dynamic Time Warping (DTW) [2] allows comparison of subsequences of different lengths. In practice, the Z-normalized Euclidean distance is most commonly used, as DTW is computationally expensive. Initial approaches relied on similarity thresholds to decide whether subsequences should be grouped together [3]. These thresholds are, however, difficult to set in practice. As a result, many works focused on the simplified problem of identifying the closest pair of matching subsequences, known as the Pair Motif problem [4], [5]. This problem was effectively addressed by the introduction of the Matrix Profile, which enables the efficient computation of nearest neighbors for all subsequences of a time series [5]. Once a Pair Motif is identified, it can then be extended into a Motif Set by searching for additional non-overlapping subsequences that lie within a relative distance threshold of the initial pair [6]. Despite numerous contributions, real-world time series still raise many challenges for Motif Discovery, and it has been shown that no existing univariate method is able to address all of them simultaneously [7].

In the multivariate setting, early approaches either reduced multivariate signals to a univariate time series via Principal Component Analysis (PCA) [8], or applied the same procedures as in the univariate, replacing the distance measures with multivariate versions, such as the Z-normalized Euclidean distance in $\mathbb{R}^d$ [9]. However, such approaches ignore the particular challenges of multivariate Motif Discovery: they neither identify the active dimensions supporting each motif nor remain robust to the presence of irrelevant or spurious dimensions. To address this limitation, the notion of subdimensional motifs [10], [11] was introduced to capture motifs defined not only by their temporal occurrences but also by the dimensions on which they consistently appear. Subdimensional motifs formalize the idea of selecting the most relevant k dimensions for a motif, where k can be specified by the user or determined using heuristics [11], [12]. The search for subdimensional motifs typically relies on the same general tools as in the univariate case, most notably the Matrix Profile, but distances between multivariate subsequences are computed using only the k most relevant dimensions, referred to as k-dimensional distances. Although subdimensional motifs are widely recognized as the most relevant formulation for multivariate time series, current approaches implement this principle only partially. In particular, the overlap conditions imposed by these methds are too strict, as they do not account for the fact that motifs occur only on subsets of dimensions. Cases where different motifs have occurrences that overlap in time but involve disjoint subsets of dimensions therefore cannot be properly handled.

Contributions: We present a new formulation of overlaps suited to subdimensional Motif Discovery, and formalize the subproblem of co-occurring motifs. We propose `sDimMotifs`, a method based on the Matrix Profile, which accurately discovers multiple subdimensional motifs while respecting these refined non-overlapping conditions. We also introduce a novel evaluation framework for subdimensional motifs that jointly considers both temporal and dimensional aspects of motif occurrences to validate our approach. An open-source implementation of `sDimMotifs` and the code to reproduce all experiments are available on GitHub[1]

II. Background

A. Problem setting

*Definition 1 (**Multivariate Time Series**):* A multivariate real-valued time series of dimension d and length n is a time-ordered sequence $S = [s_1, \dots, s_n]$ of n coefficients in $\mathbb{R}^d$.

*Definition 2 (**Subsequence**):* The subsequence of a time series $S \in \mathbb{R}^{d \times n}$ of length ℓ and starting at index $i \in [i, \dots, n - \ell + 1]$ is the sequence $S_{i,\ell} = [s_i, \dots, s_{i+\ell-1}]$.
To avoid trivial matches between temporally adjacent subsequences, Motif Discovery typically excludes overlapping subsequences, defined as follows:

*Definition 3 (**Overlapping Subsequences**):* Two subsequences $(S_{i,\ell}, S_{j,\ell'})$ of a time series $S \in \mathbb{R}^{d \times n}$ with $i < j$ overlap if $j \leq i + \ell$.
As discussed in the introduction, Motif Discovery can be framed in two main problem settings. The first is a simplified, well-posed task:

*Problem 1 (**Pair Motif Discovery**):* Identifying the two most similar non-overlapping subsequences in a time series.
The second is a more general but ill-posed task:

*Problem 2 (**Motif Set Discovery**):* Identify sets of subsequences covering all occurrences of distinct repeated patterns in a time series.
To address these problems, the Matrix Profile was introduced as a general-purpose tool for efficiently identifying the closest non-overlapping subsequences. It operates using a distance function between subsequences, which can be chosen according to the application; in the univariate case, the Z-normalized Euclidean distance is most commonly used.

*Definition 4 (**Distance Profile**):* Given a time series S and a subsequence $S_{i,\ell}$, the Distance Profile $D \in \mathbb{R}^{n-\ell+1}$ is a vector that stores $\text{dist}(S_{i,\ell}, S_{j,\ell})\ \forall\ i, j \in [1, \dots, n - \ell + 1]$.

[1] https://github.com/grrvlr/sDimMotifs

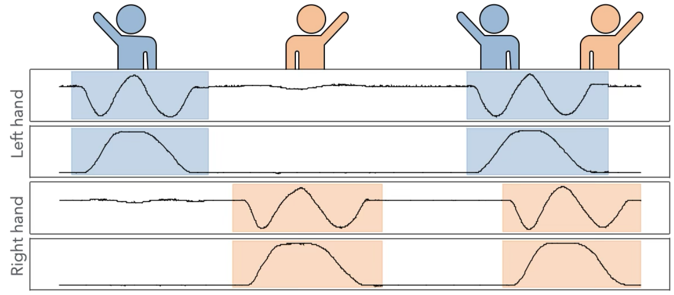

Fig. 1. Toy example from ARM-CODA showing two subdimensional motifs for left and right arm elevations. The last occurrences overlap in time but involve disjoint dimensions.

*Definition 5 (**Matrix Profile**):* Given a time series S and a subsequence length ℓ, the Matrix Profile is a meta time series $P \in \mathbb{R}^{n-\ell+1}$ such that

$$P[i] = \min_{j \in \mathcal{N}(i)} \text{dist}\big(S_{i,\ell}, S_{j,\ell}\big),$$

where $\mathcal{N}(i)$ is the set of subsequences not overlapping $S_{i,\ell}$
To illustrate the challenges of multivariate Motif Discovery, we consider a toy example extracted from the ARM-CODA dataset, a 3D motion capture dataset recording upper-body movements [13]. The signal in Fig 1 shows the (y, z) coordinates of the left and right hand sensors for repeated left and right arm elevations. Each movement corresponds to a motif visible only on a subset of dimensions: for example, a right arm movement primarily affects sensors on the right arm, while other dimensions carry little relevant information. Considering all dimensions jointly can distort motif detection by inflating distances between occurrences. Identifying the active dimensions is thus crucial for robust Motif Discovery, motivating the concept of *subdimensional motifs*.

*Definition 6 (**Subdimensional Subsequence**):* A subdimensional subsequence $S^{\mathcal{X}}_{i,\ell} \in \mathbb{R}^{k \times m}$ is a subsequence restricted to a subset of dimensions $\mathcal{X} \subseteq \{1, \dots, d\}$, where $k = \text{Card}(\mathcal{X})$ denotes the number of selected dimensions.

*Definition 7 (**k-dimensional distance**):* The k-dimensional distance function, denoted $\text{dist}^{(k)}$, measures the distance between two multivariate subsequences by considering only the k most similar dimensions. Formally,

$$\text{dist}^{(k)}(S_{i,\ell}, S_{j,\ell}) = \min_{\substack{\mathcal{X} \subseteq \{1,\dots,d\} \\ \text{Card}(\mathcal{X}) = k}} \text{dist}\big(S^{\mathcal{X}}_{i,\ell},\, S^{\mathcal{X}}_{j,\ell}\big)\,.$$

Under this formulation, the Subdimensional Pair Motif problem consists in finding the two closest non-overlapping subdimensional subsequences according to the k-distance. It can be efficiently addressed by extending the Matrix Profile framework: when computed with the k-dimensional distance, the resulting profiles are called the k-dimensional Distance Profile and k-dimensional Matrix Profile.

B. Subdimensional Motif Discovery in practice

a) Computation of the k-dimensional Matrix Profiles: The `mSTAMP` algorithm [11] was proposed to efficiently compute the k-dimensional Distance Profiles and Matrix Profiles

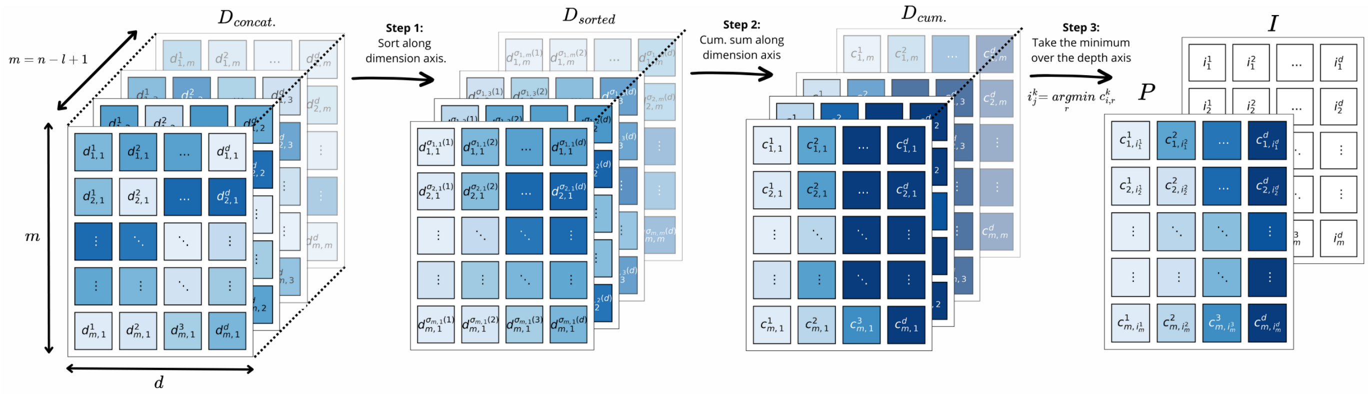

Fig. 2. Illustration of the mSTAMP algorithm. Let $D_{\text{concat}} \in \mathbb{R}^{m \times d \times m}$ be the concatenation of the univariate distance profiles across all dimensions. **Step 1:** For each (i, j), D_{concat} is sorted along the dimension axis, inducing a permutation $\sigma_{i,j}$. **Step 2:** A cumulative sum is computed along this axis, yielding the k-dimensional distance profile D_{cum}. **Final step:** For each (i, j), the minimum over the depth axis is taken to produce the k-dimensional Matrix Profiles. The output consists of two matrices, P and I, storing the k-dimensional Matrix Profiles and their Index Profiles.

for all values of $k \in \{1, \ldots, d\}$. Its key idea is to avoid explicitly computing Distance Profiles for all $\binom{d}{k}$ possible combinations of dimensions. Instead, `mSTAMP` starts from the univariate Distance Profiles computed independently for each dimension and exploits their ordering to construct the subdimensional profiles. The procedure is illustrated in Fig. 2.

b) Retrieval of the best Subdimensional Pair Motif: For each value of k, the best k-dimensional Pair Motif is obtained by selecting the minimum value in column k of P, which corresponds to the k-dimensional Matrix Profile, and retrieving the associated index from the Index Profile I. Two cases can then be distinguished. When k is specified by the user, the Pair Motif corresponding to this value of k is directly selected. When k is unknown, which is the setting considered throughout the remainder of this work, the optimal value of k is determined using the Minimum Description Length (MDL) principle, which selects the model that provides the most compact description of the data. For each candidate Pair Motif, the set of active dimensions is recomputed by selecting the k dimensions yielding the smallest distances between the two occurrences. The Pair Motif candidate minimizing the MDL criterion is then selected as the final solution.

c) Extension of Pair Motifs to Motif Sets: A classical approach to construct a Motif Set from a Pair Motif consists in identifying all non-overlapping subsequences whose distance to one of the two occurrences of the Pair Motif is below a specified threshold, where distances are computed over the active dimensions of the selected Pair Motif. [6]

d) Retrieving Multiple Motifs: The problem of retrieving multiple motifs has often been overlooked or treated as a trivial extension in the literature. For instance, in the context of [11], it is stated that: *"In the case where multiple semantically meaningful k-dimensional motifs are present in the multidimensional time series (...) There are two steps in each iteration: (1) apply the MDL-based method to find the k-dimensional motif with the minimum bit size, and (2) remove the found k-dimensional motif by replacing the Matrix Profile values of the found motif (and its trivial match) with infinity."* The implementation provided in the `stumpy` library [14], referred to as `mMotifs`, follows this strategy by combining such iterative masking of the Matrix Profile with the aforementioned extension from Pair Motifs to Motif Sets, thereby producing multiple Motif Sets in a greedy and sequential manner.

III. Method

In this section, we formalize configurations in which the presence of multiple motifs poses difficulties for existing approaches, and introduce a method designed to effectively handle settings involving multiple, potentially overlapping (in the sense of Def. 3) subdimensional motifs.

A. Co-occurring subdimensional motifs

Revisiting the toy example in Fig. 1, we see that the last occurrences of the right and left arm elevation patterns overlap in time, with the left arm movement starting before the right arm finishes. Under the classical temporal overlap definition (Def. 3), these occurrences would be considered overlapping and therefore mutually exclusive, preventing current methods from detecting both simultaneously. Yet, the movements clearly co-exist, each supported by disjoint subsets of dimensions. This motivates a refined notion of overlap that considers both temporal and dimensional aspects, which we call *subdimensional overlap*.

*Definition 8 (**Subdimensional overlap**):* Let $S \in \mathbb{R}^{d \times n}$ be a multivariate time series. Two subdimensional subsequences $(S_{i,\ell}^{\mathcal{X}}, S_{j,\ell'}^{\mathcal{X}'})$, with $i < j$, are *subdimensionally overlapping* if

$$j \leq i + \ell \quad \text{and} \quad \mathcal{X} \cap \mathcal{X}' \neq \emptyset.$$

That is, they overlap in time and share at least one active dimension.

We formally distinguish between two types of overlap: (i) *temporal overlap*, which occurs when two subsequences overlap in time, and (ii) *subdimensional overlap*, which additionally requires the overlap to occur on at least one shared active dimension.

*Definition 9 (**Co-occurring subdimensional motifs**):* Two subdimensional motifs are said to *co-occur* if there exists at least one occurrence of each motif, $S^{\mathcal{X}}_{i,\ell}$ and $S^{\mathcal{X}'}_{j,\ell'}$, such that they temporally overlap but do not subdimensionally overlap, i.e.,

$$j \leq i + \ell \quad \text{and} \quad \mathcal{X} \cap \mathcal{X}' = \emptyset.$$

B. sDimMotifs

We propose `sDimMotifs`, a method that builds upon the k-dimensional Matrix Profile to iteratively discover motifs satisfying the subdimensional overlap constraints introduced above. In the following, we describe the proposed approach in detail, with particular emphasis on how these constraints are explicitly enforced throughout the Motif Discovery process.

Algorithm 1 sDimMotifs

Require: S multivariate time series, ℓ window length, r radius ratio, N_m number of motifs
Ensure: Set of Subdimensional Motifs

$D_{\text{concat}} \leftarrow$ ConcatenateUnivariateProfiles(S, ℓ)
$M \leftarrow \mathbf{0}_{m \times d \times m}$ Global spatio-temporal mask
$(M, D_{\text{concat}}) \leftarrow$ RemoveTrivialMatches(M, D_{concat})
for $K = 1$ to N_m **do**
  $(P, I) \leftarrow$ mSTAMP(D_{concat})
  $(S^{\mathcal{X}}_{i,\ell}, S^{\mathcal{X}}_{j,\ell}) \leftarrow$ FindBestPair(P, I)
  $\mathcal{M}_K \leftarrow$ ExtendToMotifSet$(S^{\mathcal{X}}_{i,\ell}, S^{\mathcal{X}}_{j,\ell}, D_{\text{concat}}, r)$
  $M \leftarrow$ UpdateMask$(M, \mathcal{M})$
  $D_{\text{concat}} \leftarrow$ ApplyMask(D_{concat}, M)
end for
return $\{\mathcal{M}_1, \ldots, \mathcal{M}_{N_m}\}$

a) Overview of the algorithm: Algorithm 1 presents the proposed procedure. At each iteration, the best subdimensional Pair Motif is identified and extended into a Motif Set. A global spatio-temporal mask, which acts simultaneously on dimensions and temporal indices, is updated to prevent invalid overlaps in later iterations.

b) Initialization: Univariate Distance Profiles are first computed for each dimension and subsequently concatenated to form the tensor D_{concat}. To prevent trivial matches, i.e., pairs of subsequences overlapping in time, all entries (i, j) such that $j < i + \ell$ or $i < j + \ell$ are treated as invalid. These entries are then set to $+\infty$ in D_{concat} and marked in the global spatio-temporal mask M.

c) Selection of the Motif Set: At each iteration, the k-dimensional Matrix Profile is computed using the `mSTAMP` procedure. The candidate Pair Motif is then selected using the MDL criterion. This step yields both the temporal locations of the two motif occurrences and the subset of active dimensions supporting the pattern. The Pair Motif is extended into a Motif Set by considering all subsequences over the active dimensions. A subsequence is included if its distance to one of the two motif occurrences is below a threshold R. This threshold is defined as a fraction of the distance between the two occurrences of the Pair Motif: $R := r \cdot \text{dist}(S^{\mathcal{X}}_{i,\ell}, S^{\mathcal{X}}_{j,\ell})$ where r is a user-defined parameter controlling the allowed variation. The validity of each subsequence, i.e., its non-overlap with previously selected subsequences, is ensured by consulting the global spatio-temporal mask.

d) Update of M and D_{concat}: Once a Motif Set is identified, the global spatio-temporal mask M and the concatenated distance tensor D_{concat} are updated. For each occurrence, all temporally overlapping subsequences on the motif's active dimensions are masked in D_{concat}, excluding them from future distance computations. Because this operation changes the distance ordering across dimensions, the k-dimensional Distance and Matrix Profiles must be recomputed at the next iteration. The recomputation of these profiles ensures the absence of subdimensional overlap and preserves the correctness of the Matrix and Index Profiles, a property that is not guaranteed by approaches that mask the Matrix Profiles alone.

With these mechanisms in place, our method provides an accurate and efficient framework for detecting multiple, potentially co-occurring, subdimensional motifs which we now evaluate in the following experimental section.

IV. Experimental Results

A. Experimental Setup

a) Baselines: We compare our method against multiple baselines. First, we consider naive adaptations of univariate motif discovery methods that do not explicitly model subdimensional subsequences: (i) PCA followed by univariate `STOMP` [5], and (ii) `STOMP` using multivariate distances. These approaches capture temporal similarity but ignore the fact that motifs may be supported by only a subset of dimensions. Second, we include (iii) `mMotifs`, which explicitly addresses Subdimensional Motif Discovery.

b) Metrics: Until now, evaluation has focused exclusively on temporal accuracy. In the multivariate setting, however, correctly identifying the relevant dimensions is as important as localizing motifs in time, and both aspects should therefore be jointly assessed. Temporal accuracy is evaluated as in the univariate setting [7], using range-based precision, recall, and F1-score. We propose to assess dimensional accuracy using metrics based on the same principles. Temporal and dimensional scores are jointly used to determine the best matching between predicted and ground-truth motif occurrences, but are reported separately in the final results.

c) Experiments: Qualitative results are shown on a toy example derived from the arm-coda dataset [13], which measures repetitions of 15 movements among 16 subjects using 34 sensors. We only consider the (y, z) dimensions of the left and right hand sensors, selecting two occurrences of two movements (left and right arm elevations) from the first subject. For quantitative experiments, we generate synthetic multivariate time series with 5 dimensions, where each motif spans 3 dimensions and occurs 4 times. We vary the number of motifs from 1 to 50. Data and code are available on GitHub.

B. Results

a) Qualitative results: The results in Fig. 3 align with the expected behavior of the methods. When all dimensions

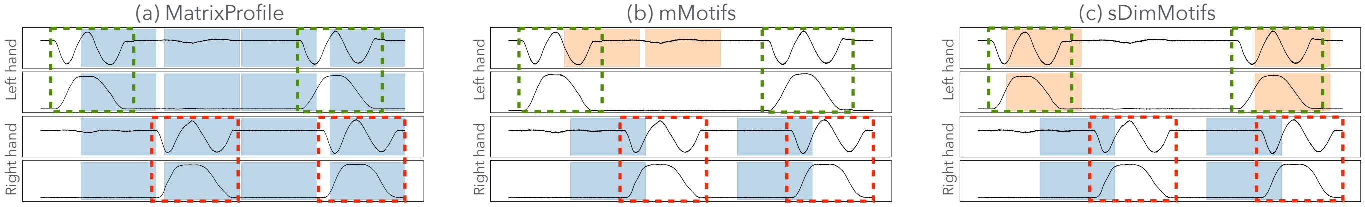

Fig. 3. Qualitative results on the toy example shown in Fig. 1. Dashed lines indicate ground-truth occurrences, each color corresponds to a different motif.

are jointly considered, STOMP with a $\mathbb{R}^d$ distance fails to distinguish between the different motifs. Both mMotifs and sDimMotifs identify the same first submotif (shown in blue, corresponding to a right arm elevation). The detected occurrences are slightly shifted with respect to the ground truth, which is a common behavior in Motif Discovery: methods often align detections with abrupt signal variations at the beginning or end of occurrences. While sDimMotifs successfully retrieves the second submotif, mMotifs fails to do so because the corresponding time indices have already been masked after extraction of the first motif.

b) Quantitative results: Figure 4 shows that, even for a single motif, methods using subdimensional formulations outperform the others in temporal detection (top, temporal F-score). Dimensional scores are not shown for STOMP ($\mathbb{R}^d$) and PCA+STOMP as they are undefined for these methods. While mMotifs' temporal and dimensional performance drops quickly when the number of motifs exceeds two, sDimMotifs remains stable up to 50 motifs.

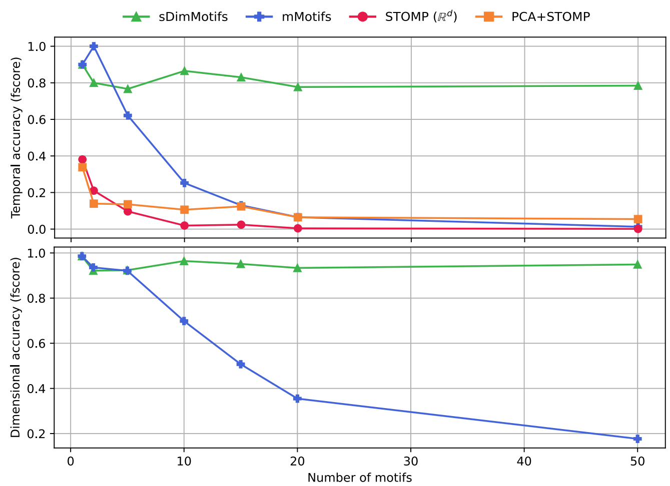

Fig. 4. Evolution of temporal F-score (top) and dimensional F-score (bottom) as the number of motifs in the multivariate time series increases.

V. Conclusion

We introduced a refined formulation of overlaps that jointly considers temporal and dimensional aspects for subdimensional Motif Discovery, and proposed a method based on this notion. Our experiments, both qualitative and quantitative, confirm the effectiveness of the approach. In particular, the method remains robust in the presence of numerous motifs. A limitation of the approach can be its computational cost: recomputing k-dimensional Distance Profiles and Matrix Profiles for each new motif increases time complexity as motifs grow, especially for high-dimensional or long time series.

References

[1] J. Paparrizos, P. Boniol, T. Palpanas, R. S. Tsay, A. Elmore, and M. J. Franklin, "Volume under the surface: a new accuracy evaluation measure for time-series anomaly detection," *Proc. VLDB Endow.*, vol. 15, no. 11, p. 2774–2787, Jul. 2022. [Online]. Available: https://doi.org/10.14778/3551793.3551830

[2] H. Sakoe and S. Chiba, "Dynamic programming algorithm optimization for spoken word recognition," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 26, no. 1, pp. 43–49, 1978.

[3] J. Lonardi and P. Patel, "Finding motifs in time series," in *Proc. of the 2nd Workshop on Temporal Data Mining*, 2002, pp. 53–68.

[4] A. Mueen, E. Keogh, Q. Zhu, S. Cash, and B. Westover, "Exact discovery of time series motifs," in *Proceedings of the 2009 SIAM International Conference on Data Mining*. SIAM, 2009, pp. 473–484.

[5] Y. Zhu, Z. Zimmerman, N. S. Senobari, C.-C. M. Yeh, G. Funning, A. Mueen, P. Brisk, and E. Keogh, "Matrix Profile II: Exploiting a novel algorithm and GPUs to break the one hundred million barrier for time series motifs and joins," in *2016 IEEE 16th International Conference on Data Mining (ICDM)*. IEEE, 2016, pp. 739–748.

[6] M. Linardi, Y. Zhu, T. Palpanas, and E. Keogh, "Matrix Profile X: Valmod-scalable discovery of variable-length motifs in data series," in *Proceedings of the 2018 International Conference on Management of Data*, 2018, pp. 1053–1066.

[7] V. Guerrini, T. Germain, C. Truong, L. Oudre, and P. Boniol, "Time Series Motif Discovery: A Comprehensive Evaluation," *Proceedings of the VLDB Endowment*, vol. 18, no. 7, 2025.

[8] Y. Tanaka, K. Iwamoto, and K. Uehara, "Discovery of time-series motif from multi-dimensional data based on MDL principle," *Machine Learning*, vol. 58, pp. 269–300, 2005.

[9] D. Minnen, T. Starner, I. Essa, and C. Isbell, "Discovering characteristic actions from on-body sensor data," in *2006 10th IEEE International Symposium on Wearable Computers*. IEEE, 2006, pp. 11–18.

[10] D. Minnen, C. Isbell, I. Essa, and T. Starner, "Detecting Subdimensional Motifs: An Efficient Algorithm for Generalized Multivariate Pattern Discovery," in *Seventh IEEE International Conference on Data Mining (ICDM 2007)*, 2007, pp. 601–606.

[11] C.-C. M. Yeh, N. Kavantzas, and E. Keogh, "Matrix Profile VI: Meaningful Multidimensional Motif Discovery," in *2017 IEEE International Conference on Data Mining (ICDM)*, 2017, pp. 565–574.

[12] P. Schäfer and U. Leser, "Discovering leitmotifs in multidimensional time series," *Proc. VLDB Endow.*, vol. 18, no. 2, p. 377–389, Oct. 2024. [Online]. Available: https://doi.org/10.14778/3705829.3705852

[13] S. W. Combettes, P. Boniol, A. Mazarguil, D. Wang, D. Vaquero-Ramos, M. Chauveau, L. Oudre, N. Vayatis, P.-P. Vidal, A. Roren, and M.-M. Lefèvre-Colau, "Arm-CODA: A Data Set of Upper-limb Human Movement During Routine Examination," *Image Processing On Line*, vol. 14, pp. 1–13, 2024.

[14] S. M. Law, "STUMPY: A Powerful and Scalable Python Library for Time Series Data Mining," *The Journal of Open Source Software*, vol. 4, no. 39, p. 1504, 2019.