



funBlalign: a hierarchical algorithm for functional motif discovery based on mean squared residue scores

Jacopo Di Iorio¹ · Marzia A. Cremona^{2,3} · Francesca Chiaromonte^{1,4,5}

Received: 26 October 2023 / Accepted: 11 November 2024
© The Author(s) 2024

Abstract

Motif discovery is gaining increasing attention in the domain of functional data analysis. Functional motifs are typical “shapes” or “patterns” that recur multiple times in different portions of a single curve and/or in misaligned portions of multiple curves. In this paper, we define functional motifs using an additive model and we propose *funBlalign* for their discovery and evaluation. Inspired by clustering and biclustering techniques, *funBlalign* is a multi-step procedure which uses agglomerative hierarchical clustering with complete linkage and a functional distance based on mean squared residue scores to discover functional motifs, both in a single curve (e.g., time series) and in a set of curves. We assess its performance and compare it to other recent methods through extensive simulations. Moreover, we use *funBlalign* for discovering motifs in two real-data case studies; one on food price inflation and one on temperature changes.

Keywords Functional data analysis · Functional motif discovery · Clustering · Biclustering

1 Background and motivation

The last decades have seen an increasing interest in the analysis of functional data, i.e., data that can be represented as smooth curves. Functional Data Analysis (FDA) methods (see, e.g., Ramsay and Silverman 2005; Ferraty and Vieu 2006; Kokoszka and Reimherr 2017) have been applied in a variety of scientific fields. These include biology, medicine, and genetics, where FDA has been employed to analyze, e.g., the genomic landscape of “jumping genes” and COVID-19 epidemics (Chen et al. 2020; Boschi et al. 2021); neurosciences and psychometrics, where it has been

used, e.g., to map cognitive processes analyzing response times and brain imaging data (Buckner et al. 2004; Lila et al. 2017); economics and environmental sciences, where it has been used to explore patterns and generate predictions over space or time concerning air pollution, climate indicators, stock market prices, etc (see, e.g., Das et al. 2019; Ghuman et al. 2020).

In this paper, we focus on *functional motif discovery*; that is, the identification of typical “shapes” or “patterns” that recur multiple times in different portions of a single curve and/or in misaligned portions of multiple curves. While the ability to identify such patterns offers great promise in multiple scientific fields (Cremona and Chiaromonte 2023), to the best of our knowledge, the notion of functional motif still lacks a rigorous statistical formalization. We define a functional motif Q of length l as a collection of n_Q curve portions $p_k^Q(t)$ for $k = 1, \dots, n_Q$ with $t \in [0, l]$ (the occurrences of Q) obeying the additive model

$$p_k^Q(t) = \mu^Q + \alpha_k^Q + \beta^Q(t) + \varepsilon_k(t) \quad \forall t \in [0, l]$$

where μ^Q is the mean of the motif, α_k^Q its portion-specific adjustment, $\beta^Q(t)$ its t-varying adjustment and $\varepsilon_k(t)$ an error term (see Fig. 1).

In order to discover functional motifs, both in a single curve or in a set of curves (evaluated over a grid of

✉ Jacopo Di Iorio
jacopo.di.iorio@emory.edu

¹ Department of Statistics, Penn State University, Joab L. Thomas Building, University Park 16802, PA, USA

² Department of Operations and Decision System, Université Laval, 2325 Rue de la Terrasse, Québec G1V0A6, Québec, Canada

³ CHU de Québec, Université Laval Research Center, 1050 Chemin Sainte-Foy, Québec G1S4L8, Québec, Canada

⁴ Institute of Economics and L'EMbeDS, Sant'Anna School of Advanced Studies, Piazza Martiri della Libertà 33, Pisa 56127, PI, Italy

⁵ Department of Quantitative Theory and Methods, Emory University, 36 Eagle Row, Atlanta 30302, GA, USA

equally spaced points), we develop *funBAlign*, a multi-step algorithm that requires as input, in addition to the curve(s) themselves, only the *discretized* length ℓ and the minimum cardinality n_{min} of the motifs to be discovered. *funBAlign* performs a comprehensive scan of all portions of length ℓ of all curves in the data, and arranges them in a dendrogram employing agglomerative hierarchical clustering with complete linkage and a functional generalization of the Mean Squared Residue Score (MSR) — a measure typically employed by biclustering techniques (Pontes et al. 2015). The dendrogram is dynamically cut to identify a set of candidate functional motifs, which are then post-processed to select the most interesting. While hierarchical agglomerations are commonly used for functional clustering (see Ferreira and Hitchcock 2009 for a comparison, and Jacques and Preda 2014 for a survey on functional clustering techniques), the use of means squared residues still represents a novelty in the functional framework. The MSR was originally introduced by Cheng and Church (2000), for discovery and validation of biclusters, i.e. of subsets of rows and columns of a data matrix, and has been widely used since in the multivariate setting (e.g., Liu and Wang 2007; Angiulli et al. 2008; Yang et al. 2005), but its functional generalization, fMSR, has only very recently been employed in functional clustering problems by Galvani et al. (2021) and Di Iorio and Vantini (2023).

Functional motif discovery, while gaining increasing attention, is itself an under-explored area. To the best of our knowledge, the only other approach that deals specifically with it in the FDA domain is probabilistic K -means with local alignment (*probKMA*, Cremona and Chiaromonte 2023). Following a stream of literature devoted to the simultaneous alignment and clustering of curves at the global level (see, e.g., Liu and Yang 2009; Sangalli et al. 2010), and to the simultaneous domain selection and clustering of curves (Fraiman et al. 2016; Floriello and Vitelli 2017; Vitelli 2023), *probKMA* identifies candidate functional motifs combining a probabilistic K -means algorithm and local alignment (or domain selection) techniques. A strong point of *probKMA* is the fact that, while requiring the specification of a minimum motif length, it can extend such length in a motif-specific and data-driven fashion. However, it is designed to operate on a set of curves and it requires extensive post-processing. As a counterpart to its efficacy, the need for multiple initializations of the K -means algorithm and the complexity of post-processing make *probKMA* computationally expensive. Outside the FDA domain, a problem similar to functional motif discovery has been tackled by the data mining community – seeking patterns embedded multiple times within a single time series (see, e.g. Lila et al. and Patel 2002; Mueen et al. 2009; Yeh et al. 2016). This work are based on a k -Nearest Neighbors algorithm and requires users to specify several parameters, some of which may not be intuitive and significantly affect outcomes (newer versions of these pro-

cedures have been published since 2016, but they primarily enhance computational performance on large data sets, rather than modifying required inputs or algorithmic approach).

As mentioned above, and in contrast to existing methods, *funBAlign* can naturally handle both single curves and sets of curves. In the latter case, thanks to the hierarchy it creates, it can also highlight relationships among curves harboring the same motif. On a different front, although the comprehensive scan employed by *funBAlign* can be computationally demanding, it must be performed only once – without the need for multiple runs or initializations as in *probKMA*. More generally, *funBAlign* can effectively tackle applications where functional alignment fails or may be inadequate, such as identifying motifs embedded consecutively, or appearing only in one curve. We also note that the two input parameters required by our procedure, the motif length ℓ and the minimum number of portions n_{min} , while impacting outcomes, are user-friendly and intuitive.

The remainder of the paper is organized as follows. Section 2 presents the theoretical setting of *funBAlign*, including the rigorous definition of functional motifs and of the dissimilarity measure employed. The algorithmic implementation is described in Section 3. Finally, the performance of our proposal is assessed through simulations, comparisons with other available methods, and real-data case studies in Sections 4, 5 and 6.

2 Theoretical setting

2.1 Model-based definition of functional motifs

Consider a set of real-valued curves $f_i(t), i = 1, \dots, N$, each defined on a compact interval which we assume to be $[0, T_i]$ without loss of generality. Intuitively, a functional motif is a “shape” or “pattern”, defined on a domain interval of given length l , which occurs multiple times within the set of curves – possibly with noise. A motif can occur at different positions within a curve and/or in misaligned portions of different curves in the set ($N = 1$ corresponds to the case in which motifs are sought within a single curve, e.g., a time series). For each curve $f_i(t)$, we consider all possible overlapping portions of length l , and we align them on the interval $[0, l]$. In symbols, a generic portion is $f_{i,L}(t) = f_i|_L \circ h(t)$ with $t \in [0, l]$, where L is a sub-interval of length l of $[0, T_i]$ and h is a shift transformation from $[0, l]$ to L . Let I be the collection of all portions of length l of all curves in the set, i.e. $I = \{p(t) = f_{i,L}(t) \mid i = 1, \dots, N \text{ and } L \subseteq [0, T_i] \text{ with } |L| = l\}$. We provide a rigorous definition of functional motif using an additive model as follows.

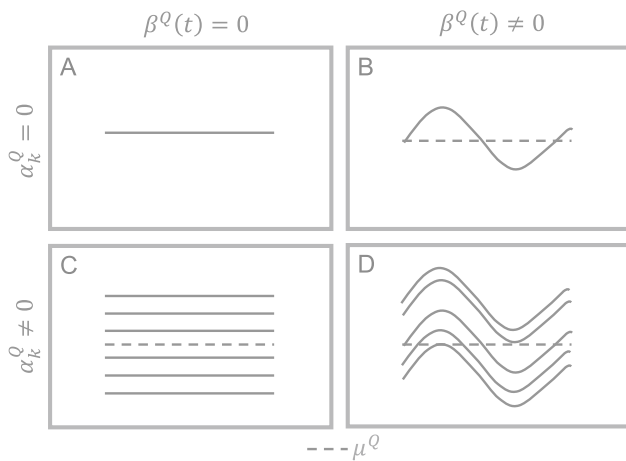


Fig. 1 Examples of ideal functional motifs, obtained by (1) with $\varepsilon_k(t) = 0$. In panel A all portions are identical and constant; $p_k^Q(t) = \mu^Q$ ($\alpha_k^Q = 0$ and $\beta^Q(t) = 0$). In panel B all portions are identical; $p_k^Q(t) = \mu^Q + \beta^Q(t)$ ($\alpha_k^Q = 0$). In panel C portions are constant and parallel; $p_k^Q(t) = \mu^Q + \alpha_k^Q$ ($\beta^Q(t) = 0$). Panel D illustrates the general case, with parallel portions sharing the same shape; $p_k^Q(t) = \mu^Q + \alpha_k^Q + \beta^Q(t)$. In all panels, the dashed line represents the motif mean μ^Q

Definition 1 A functional motif Q is a collection $I_Q \subset I$ of $n_Q > 1$ curve portions $p_k^Q(t)$ for $k = 1, \dots, n_Q$ such that

$$p_k^Q(t) = \mu^Q + \alpha_k^Q + \beta^Q(t) + \varepsilon_k(t) \quad \forall t \in [0, l] \quad (1)$$

where μ^Q is a mean level, α_k^Q a portion-specific adjustment, $\beta^Q(t)$ a t-varying adjustment, and $\varepsilon_k(t)$ an error term.

To have unique and identifiable parameters, we impose $\sum_{k=1}^{n_Q} \alpha_k^Q = 0$, $\int_0^l \beta^Q(t) dt = 0$, and $\mathbb{E}[\varepsilon_k(t)] = 0$ for $\forall k = 1, \dots, n_Q$ and $\forall t \in [0, l]$. It is important to remark that, even if the definition above does not consider explicitly the curve to which each portion belongs to, this information is crucial in the algorithm we introduce in Sect. 3.

Definition 1 is inspired by the biclustering literature, and in particular by the definition of multivariate coherent evolution biclusters (Madeira and Oliveira 2004) utilized in the seminal paper by Cheng and Church (2000). In the functional framework, Galvani et al. (2021) used a model similar to the one in (1) to discover functional biclusters in a data matrix whose cells correspond to curves, while Di Iorio and Vantini (2023) employed the model in (1) to identify local clusters in subsets of aligned curves. Following the biclustering literature, we call a functional motif “ideal” when $\varepsilon_k(t) = 0$, i.e. in the absence of noise. Such a motif is composed by perfectly parallel portions sharing the same shape (Fig. 1D). Among ideal motifs, special cases can be obtained setting α_k^Q and/or $\beta^Q(t)$ to 0. If both $\alpha_k^Q = 0$ and $\beta^Q(t) = 0$ we have a constant motif, whose portions are all $p_k^Q(t) = \mu^Q$

(Fig. 1A). Setting $\beta^Q(t) = 0$ but allowing $\alpha_k^Q \neq 0$ we have parallel constant portions $p_k^Q(t) = \mu^Q + \alpha_k^Q$ (Fig. 1C), while setting $\alpha_k^Q = 0$ but allowing $\beta^Q(t) \neq 0$ all portions will be identical to $p_k^Q(t) = \mu^Q + \beta^Q(t)$ (Fig. 1B).

2.2 Evaluating the coherence of functional motifs

To evaluate the coherence of a candidate motif, i.e. of a collection of portions, to an (unknown) additive model, we gauge the error term $\varepsilon_k(t)$ using a functional version of the Mean Squared Residue score (MSR), or H-score, first introduced by Cheng and Church (2000) to seek large biclusters among the rows (genes) and columns (experimental conditions) of a gene expression data matrix. Galvani et al. (2021) first proposed a generalization of the MSR to the functional framework to seek large biclusters in matrices of curves. Here, we define the functional Mean Squared Residue score (fMSR) as follows.

Definition 2 The functional Mean Squared Residue score (fMSR) of a functional motif Q is

$$H(Q) = \frac{1}{n_Q} \frac{1}{l} \sum_{k=1}^{n_Q} \int_0^l \left(p_k^Q(t) - (\hat{\mu}^Q + \hat{\alpha}_k^Q + \hat{\beta}^Q(t)) \right)^2 dt, \quad (2)$$

where the estimates $\hat{\mu}^Q$, $\hat{\alpha}_k^Q$ and $\hat{\beta}^Q(t)$ of the terms in (1) are

$$\begin{aligned} \hat{\mu}^Q &= \bar{p}^Q = \frac{1}{n_Q} \frac{1}{l} \sum_{k=1}^{n_Q} \int_0^l p_k^Q(t) dt, \\ \hat{\alpha}_k^Q &= \bar{p}_k^Q - \hat{\mu}^Q = \frac{1}{l} \int_0^l p_k^Q(t) dt - \hat{\mu}^Q, \\ \hat{\beta}^Q(t) &= \bar{p}^Q(t) - \hat{\mu}^Q = \frac{1}{n_Q} \sum_{k=1}^{n_Q} p_k^Q(t) - \hat{\mu}^Q. \end{aligned}$$

In forming the estimates, \bar{p}_k^Q , $\bar{p}^Q(t)$ and \bar{p}^Q represent, respectively, the mean value of each portion $p_k^Q(t)$, the functional mean of all portions $p_k^Q(t)$, $k = 1, \dots, n_Q$, and the mean value of $\bar{p}^Q(t)$ (or equivalently, the mean of the \bar{p}_k^Q 's). Using them, we can rewrite (2) as

$$H(Q) = \frac{1}{n_Q} \frac{1}{l} \sum_{i=1}^{n_Q} \int_0^l \left(p_k^Q(t) - \bar{p}_k^Q - \bar{p}^Q(t) + \bar{p}^Q \right)^2 dt \quad (3)$$

which allows us to easily implement the score calculation. We observe that the fMSR of an ideal, i.e. noiseless, functional motif Q is $H(Q) = 0$. Thus, similar to biclustering

methods which look for biclusters with low MSR, we look for functional motifs with low fMSR.

Di Iorio et al. (2020) recently proved that the MSR is biased towards small biclusters (i.e., biclusters with a small number of rows and/or columns). Here, we prove that the fMSR suffers from a similar bias towards small functional motifs; that is, motifs comprising a small number of portions. This can hinder motif discovery, as it distorts the comparisons of motifs comprising different numbers of portions. The following theorem fully characterizes this bias and suggests a way of correcting it (a proof is provided in Section S1 and the bias in the multivariate framework is extensively treated and illustrated through simulations in Di Iorio et al. 2020).

Theorem 1 *Let Q be a functional motif. For $n = 2, 3, \dots, n_Q$ let \bar{H}_n be the average fMSR of all sub-motifs of Q obtained selecting exactly n of the n_Q portions $p_k^Q(t)$ belonging to Q . Then*

$$\bar{H}_{n+1} = \bar{H}_n \frac{n^2}{n^2 - 1}. \tag{4}$$

An implication of Theorem 1 is that the ratio $r_{n,n+1} = \frac{\bar{H}_{n+1}}{\bar{H}_n}$ depends on the number of portions n included in the sub-motif. It is also straightforward to verify that

$$\bar{H}_{n+m} = \bar{H}_n \prod_{x=n}^{n+m-1} \frac{x^2}{x^2 - 1} \tag{5}$$

for every $m > 0$, allowing one to compute \bar{H}_{n+m} using \bar{H}_n alone. Hence, we have that

$$r_{n,n+m} = \frac{\bar{H}_{n+m}}{\bar{H}_n} \xrightarrow{m \rightarrow \infty} \prod_{x=n}^{\infty} \frac{x^2}{x^2 - 1}. \tag{6}$$

The infinite product in (6) converges and we have $r_{2,2+m} \rightarrow 2$. As a consequence, the bias can be at most 1, and it decreases when considering motifs comprising a large number of portions (see Di Iorio et al. 2020, for simulations and additional details in the multivariate case). However, its effects can be troublesome in comparisons involving seldom motifs, which will be non-negligibly favored. For this reason, we propose to correct the bias defining an *adjusted* version of the fMSR as follows.

Definition 3 Let Q be a functional motif comprising $n_Q \geq 2$ portions. Its adjusted functional mean squared residue score is defined as:

$$H_{adj}(Q) = \begin{cases} H(Q) & \text{if } n_Q = 2, \\ \frac{H(Q)}{\prod_{r=2}^{n_Q-1} \frac{r^2}{r^2-1}} & \text{if } n_Q > 2. \end{cases} \tag{7}$$

We note that the correction is straightforward and it requires only a multiplication of the fMSR score by a factor - with negligible additional computational burden. The adjusted fMSR is the measure we employ in the remainder of the paper (a numerical comparison of the proposed algorithm with adjusted and non-adjusted fMSR score is available in Section S2).

2.3 An fMSR-based dissimilarity measure

The adjusted fMSR can be used to construct a dissimilarity measure between two curve portions $p_1(t)$ and $p_2(t)$ in $[0, l]$ as

$$d_{fMSR}(p_1, p_2) = H_{adj}(W), \tag{8}$$

where $W = \{p_1(t), p_2(t)\}$ is the functional motif composed only by the portions $p_1(t)$ and $p_2(t)$. According to this definition, $d_{fMSR}(p_1, p_2) = 0$ if and only if W is an ideal motif.

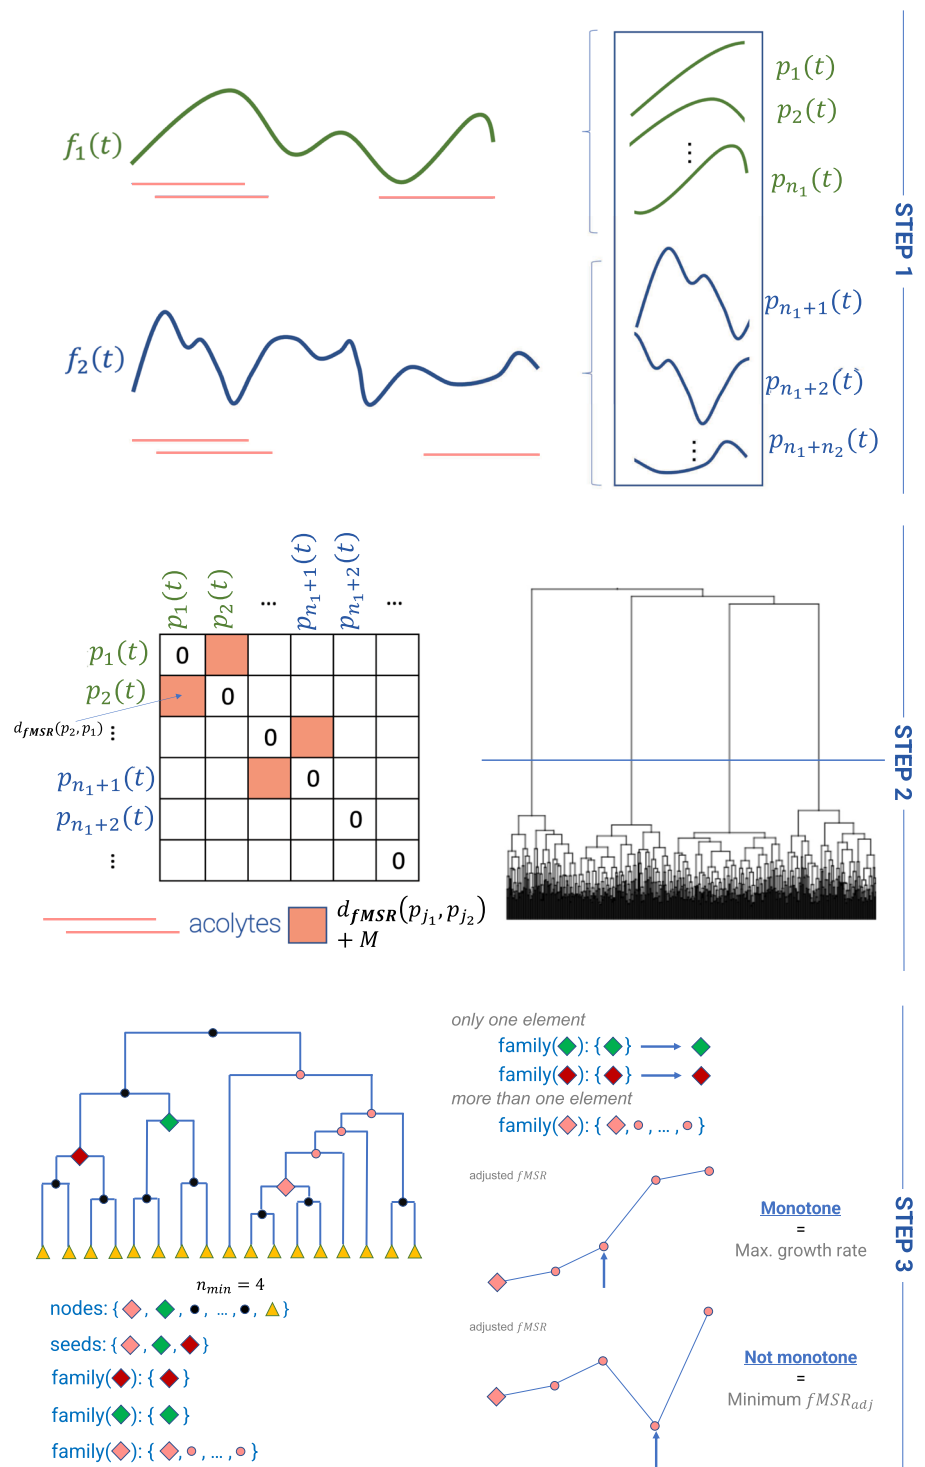
3 The funBlalign algorithm

Given a set of $N \geq 1$ real-valued curves $f_i(t)$ defined on $[0, T_i]$, $i = 1, \dots, N$, *funBlalign* discovers recurrent and coherent motifs as defined by (1). The algorithm considers the evaluation of the curves over a grid of equally spaced points $t_0 = 0, t_1 = \delta t, t_2 = 2\delta t, \dots$ and requires as input the *discretized* motif length ℓ (the number of grid points corresponding to the length l in (1)) and the minimum number of portions n_{min} . It comprises four steps, as described below and presented in the schematic of Fig. 2).

Step 1 — Portion creation and alignment. For every curve $f_i(t)$, $i = 1, \dots, N$, we create all n_i portions of length ℓ (starting at t_0, t_1, t_2 , etc.), and align them so that their domains all start at $t = 0$. We indicate the resulting overall set of aligned portions with $p_j(t)$, $j = 1, \dots, n$, where $n = \sum_{i=1}^N n_i$. For each portion, we keep track of the originating curve $i(j)$ and of the ℓ grid points occupied.

Step 2 — Hierarchical clustering based on the fMSR dissimilarity. We compute the fMSR-based dissimilarity $d_{fMSR}(p_{j_1}, p_{j_2})$ of every portion pair (p_{j_1}, p_{j_2}) , $1 \leq j_1, j_2 \leq n$, and calculate $M = \max_{1 \leq j_1, j_2 \leq n} d_{fMSR}(p_{j_1}, p_{j_2})$, i.e. the dissimilarity of the most dissimilar portion pair. Note that since all portions $p_k(t)$ are evaluated over a grid, the integrals involved in computing dissimilarities are approximated by sums. Whenever two portions originate from the same curve and share at least 50% of their grid points, we name them “acolytes” and artificially increase their dissimilarity by M , to prevent obvious similarities among curve portions

Fig. 2 A summary schematics of the three main steps of the *funBlalign* algorithm



with large overlaps from dominating the agglomeration. Our dissimilarities are thus

$$d_{j_1, j_2} = \begin{cases} d_{fMSR}(p_{j_1}, p_{j_2}) + M & \text{if } p_{j_1} \text{ and } p_{j_2} \text{ are acolytes,} \\ d_{fMSR}(p_{j_1}, p_{j_2}) & \text{otherwise} \end{cases} \tag{9}$$

and we use their matrix $D_{fMSR} = [d_{j_1, j_2}]_{1 \leq j_1 \leq n, 1 \leq j_2 \leq n}$ to perform agglomerative hierarchical clustering with complete linkage (see, e.g., Murtagh and Contreras 2012). Let *Tree* indicate the resulting dendrogram, whose nodes represent clusters of portions. Due to the use of complete linkage and to the addition of M , the longest dendrogram branches occur

when two nodes comprising acolytes are merged (see Fig. 2 Step 2). We use such branches to cut the dendrogram, generating N_{tree} sub-trees $tree_s, s = 1, \dots, N_{tree}$, which do not contain acolytes.

Step 3 — Collection of candidate functional motifs. For every sub-tree $tree_s, s = 1, \dots, N_{tree}$, we consider the sets of all nodes $nodes_s$ and all leaf nodes $leaves_s$ ($leaves_s \subseteq nodes_s$), and for every $x \in nodes_s$ we consider the sets of all descendants $des(x)$ and ascendants $asc(x)$. We define the seeds of $tree_s$ as the set of nodes with at least n_{min} leaf descendants, and with no descendant meeting the same criterion; in symbols

$$seed_s = \{x \in nodes_s : |des(x) \cap leaves_s| \geq n_{min} \wedge \forall y \in des(x), y \notin seed_s\}. \tag{10}$$

Note that a sub-tree can have zero, one or multiple seeds. Finally, we define the family of $v \in seed_s$ as the union of v and all its ascendants that are not shared with other seeds:

$$family(v) = \{v\} \cup \{x \in asc(v) : \nexists y \in seed_s, y \neq v \text{ s.t. } x \in asc(y)\}. \tag{11}$$

Next, we select a “recommended representative” for the family. If $|family(v)| = 1$, this is trivially the one family member. If $|family(v)| > 1$, we sort the nodes $x \in family(v)$ in order of increasing cardinality $|x|$ (i.e. number of portions) and consider $H_{adj}(x)$ (adjusted $fMSR$ value computed on the portions). If $H_{adj}(x)$ increases with $rank(|x|)$, we use an “elbow” approach and select the node just before the maximum increase. Otherwise, we select the node with minimum $H_{adj}(x)$. Our collection of candidate motifs \mathcal{C} is composed of recommended representatives of all families from all sub-trees.

Step 4 — Post-processing. We sort candidate motifs $x \in \mathcal{C}$ based either on adjusted $fMSR$, i.e. $rank(H_{adj}(x))$, or on a combination of adjusted $fMSR$ and inverse cardinality, i.e. $rank(H_{adj}(x)) + rank(-|x|)$ (alternative ranking criteria can be used, see Section 6). Starting from the top, we compare each motif to those with higher rank. If all portions of x are acolytes to portions of a higher ranking motif, we filter it out; otherwise we retain it. This produces a final collection of discovered motifs $\mathcal{D} \subseteq \mathcal{C}$.

We remark that Step 2 relies on M and the definition of acolytes. While larger values of M do not affect the solution, the way we define acolytes (e.g., the percent overlap) can have an impact on the results; this definition may be changed depending on applications and user needs. In addition, the dynamic cut in Step 3 selects candidate motifs controlling a trade off between cohesiveness (small adjusted $fMSR$) and

prevalence (large number of occurrences). In contrast, the post-processing in Step 4 eliminates overlapping candidates which might have been selected in different sub-trees. We note that Step 4 actually provides an “importance ranking” which can be used, more generally, to further explore and select among candidates. This is particularly useful in applications where the algorithm identifies a very large number of candidate motifs.

4 Simulations

We assess the performance of *funBlalign* through an extensive simulation study. To simulate a smooth curve embedding multiple occurrences of functional motifs we use the flexible B-spline-based model proposed by Cremona and Chiaromonte (2023). Briefly, this model generates a smooth curve as $f(t) = \sum_{\gamma=1}^J c_{\gamma} \phi_{\gamma}(t)$ where $\{\phi_{\gamma}\}_{\gamma=1}^{\Gamma}$ is a B-spline basis of order ν with equally spaced knots $t_1, \dots, t_{\Gamma-\nu+2}$ and $c_{\gamma}, \gamma = 1, \dots, \Gamma$ are real coefficients. In every simulation conducted in our study, we generate a single smooth curve of length 7000 using order $\nu = 3$ and knots at distance $T_{knots} = 10$. The curve is then evaluated across a grid of 7001 equally spaced points ($t_0 = 0, t_1 = 1, \dots$). We then randomly embed in the curve the same number (8 or 10) of occurrences of 4 distinct motifs of length $4T_{knot} = 40$ (i.e., $\ell = 41$ points). Coefficients defining both the curve and the motifs are randomly generated from a $Beta(0.45, 0.45)$ and then rescaled to $[-15, 15]$. We incorporate into every motif occurrence a vertical shift drawn uniformly from $[-10, 10]$ as well as noise – adding to the coefficients independent draws from $\mathcal{N}(0, \sigma^2)$. We note that this way of generating motifs within a curve does *not* match the additive model in (1); thus, in our simulations we are challenging *funBlalign* with motifs that may be harder for it to identify, showcasing the flexibility of the algorithm. We also note that the curve background can, by chance, comprise segments that resemble one of the motifs; that is, extra portions that were not intentionally embedded but do follow the motif pattern. Moreover, the background may reveal entirely different and distinguishable motifs; that is, patterns that, while not intentionally inserted in multiple occurrences, happen to repeat themselves along the curve. Extra portions or additional motifs emerging from the background introduce an added complexity in interpreting simulation results (see below). Because of this complexity, we prefer to refer to motifs and occurrences as “embedded” (vs. not embedded) instead of “true” (vs. false).

We construct a total of 100 simulations. We consider 10 alternative motifs sets, each one with 4 distinct motifs; 2 alternative numbers of occurrences (8 or 10) which, for simplicity, are the same for all motifs; and 4 alternative levels of noise, expressed by $\sigma = 0.1, 0.5, 1$ or 2 . In a first batch of $10 \times 2 \times 4 = 80$ simulations, all motifs share the same σ , and

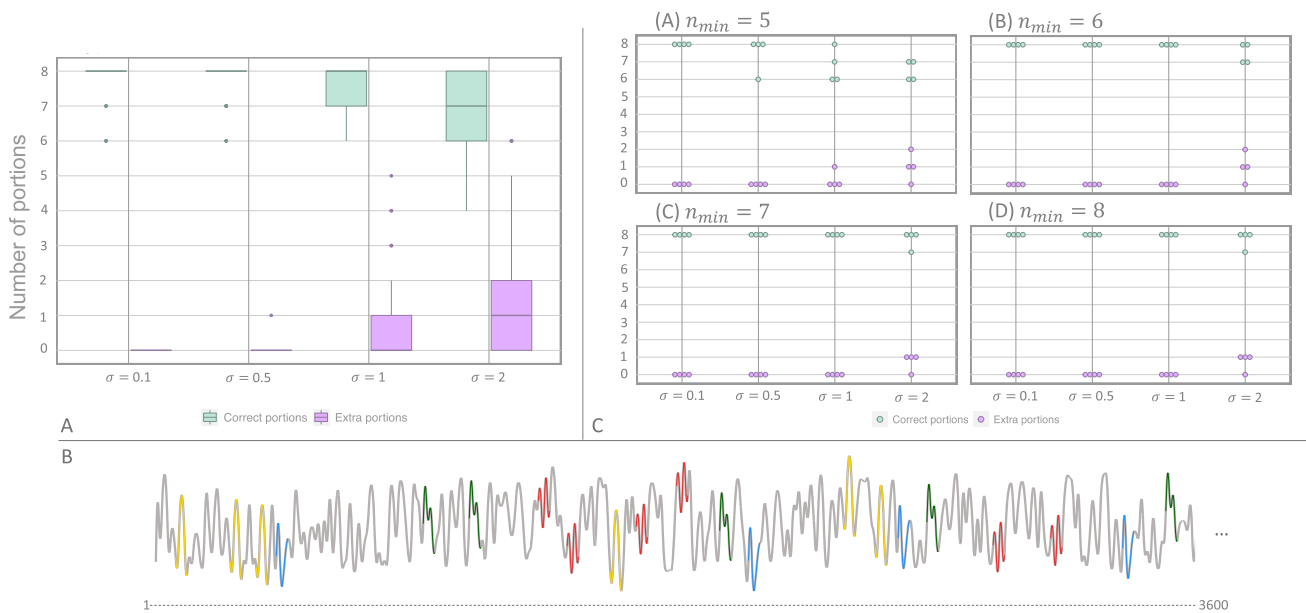


Fig. 3 **A** Performance of *funBAlign* for all simulations where motifs have 8 occurrences and shared noise levels. The algorithm is run with minimum cardinalities $n_{min} = 5, 6, 7, 8$ and results are pooled. For each σ , the green and pink boxplots (left and right) represent correctly identified portions (the ones which were embedded in the simulation, and correctly identified by *funBAlign*) and extra portions (the ones which were not embedded in the simulation, but still identified by *funBAlign*),

respectively. **B** First half of the curve for a simulation employing motif set n. 7 and $\sigma = 0.1$. Occurrences of the motifs are color-coded. The three wired dots (...) express the fact that the curve continues. **C** Performance of *funBAlign* for simulations employing motif set n. 7, shown separately for runs with varying n_{min} . Green and pink jittered dots represent correctly identified portions and extra portions, respectively

we use all σ values in turn. In a second batch of $10 \times 2 = 20$ simulations, each motif is attributed a different σ . We run *funBAlign* setting $\ell = 41$ (the true length of the motifs) and using different minimum cardinalities $n_{min} = 5, 6, 7, 8$ (progressively closer to the true number of occurrences, 8 or 10). We post-process the candidate motifs produced in each run, ranking them according to two criteria – the adjusted fMSR and the rank sum – and identifying results which are most similar to the intentionally embedded motifs (the ones we use as targets). Here we focus on those; that is, on the detection of occurrences of the motifs we embedded in the curve, plus potential *extra* portions of the background that the algorithm associates with such motifs. *funBAlign* may “discover” in the background motifs other than those we created, but we will ignore them in the discussion to follow. We also restrict the main text presentation to the 40 simulations where motifs have 8 occurrences and shared noise levels. Results for simulations where motifs have 10 occurrences and shared noise levels, or where motifs have 8 or 10 occurrences and different noise levels, are entirely consistent with those presented here and are provided in Sections S4 and S5.

Fig. 3A summarizes performance for the 40 simulations, pooling results across algorithm runs with different n_{min} ’s. For every level of σ we display two boxplots, each comprising $4 \times 10 \times 4 = 160$ values: for each of the 4 motifs in each of the 10 sets and across the 4 runs, we count the number of

portions correctly identified (left boxplot) and the number of extra portions (right boxplot). We can see that *funBAlign* is quite effective in identifying embedded motif occurrences. However, as expected, as the level of noise increases some embedded portions are missed, and some extra portions are found – though these are usually very similar to the embedded portions (see Fig. 5). Similar results, again pooling across runs with different n_{min} ’s, but separating each of the 10 alternative motif sets, are provided in Fig. S4. Fig. 3C shows results (identification of true occurrences and of extra portions) for one motif set, n. 7, but separately for the 4 runs with varying n_{min} . The minimum cardinality used in the algorithm can indeed impact performance, especially through the ranking of the candidate motifs (see Figs. S5-S6). When n_{min} is much lower than the true number of occurrences, rank sum is preferable to adjusted fMSR as a ranking criterion, because it tends to privilege results more similar to the embedded motifs. On the other hand, when n_{min} is close to the true number of occurrences, best results do not necessarily have highest rank sums. This fact can hinder their identification. In addition, as expected, when n_{min} is too low we miss some occurrences and, as n_{min} gets higher, we identify more extra portions. Fig. 3B shows the first half of the curve for the simulation using motif set n. 7, with noise level $\sigma = 0.1$, color-coding motif occurrences interspersed across the curve. Fig. 4 provides details on the motifs identified in

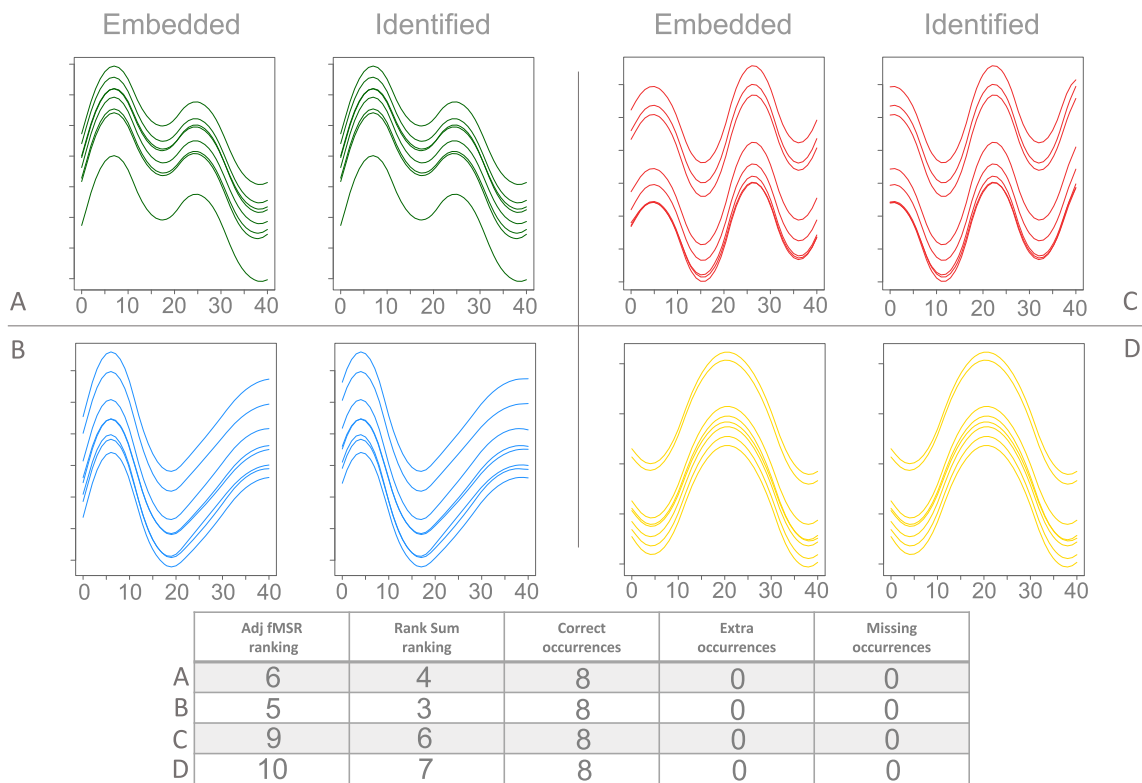


Fig. 4 Motif identification for the simulation employing motif set n. 7 and $\sigma = 0.5$. *funBlalign* is run with $n_{min} = 6$. For each of the 4 motifs, color-coded in green, red, blue and yellow, left and right panels show the 8 occurrences embedded in the curve and the most similar portions

identified by the algorithm, respectively. The table provides rankings and numbers of correctly identified, extra and missing occurrences for each motif. With low noise, no embedded occurrences are missed, and no extra occurrences are identified

the simulation with noise level $\sigma = 0.5$, and Fig. 5 on those corresponding to $\sigma = 2$, running *funBlalign* with $n_{min} = 6$ in both cases. Performance is excellent, though noisier motifs, as expected, cause a slight deterioration. We also see how, at the same level of noise, motifs may be easier or harder to identify depending on their shapes. Importantly, we note that when *funBlalign* identifies extra portions, i.e. segments of the background that resemble a motif by chance, these are indeed very similar to the occurrences intentionally embedded in the curve (Fig. 5). In effect, they should be thought of as “unplanned” true positives, not as false positives.

In our simulation study we treat ℓ as a known quantity, using its true value. However, in many real data applications a reasonable value for ℓ may not be known in advance. In such settings, provided the user can at least identify a reasonable range for ℓ , we suggest to run the algorithm for each ℓ in such range, and compute the average adjusted fMSR score of the top motifs. An elbow method strategy can then be employed to select a satisfactory ℓ value within the explored range. Figure S2 shows the results of this procedure on the simulations employing motif set n. 7.

We studied how the choice of n_{min} could impact the correct and complete identification of motifs. We simulated four

versions of a 7001 point long curve embedding 20 instances of the same motif, using four levels of noise; $\sigma = 0.1, 0.5, 1,$ and 2 . We then verified the correct identification of the motif with values of n_{min} that differ markedly from the actual number of embedded portions; namely, $n_{min} = 6, 8, \dots, 32, 34$ 4. As shown in Figure S3, a too low n_{min} can result in incomplete identification, especially for higher levels of noise. However, in general, the identification is complete or almost complete even if $n_{min} < 20$. Conversely, a too high n_{min} forces the algorithm to identify extra portions – in addition to the embedded 20, which are all identified. This highlights the importance of a reasonable choice of n_{min} , but also shows that the proposed method has some tolerance towards non-trivial misspecifications.

5 Comparisons with related methods

We compare *funBlalign* with two different methods; namely, *probKMA* (Cremona and Chiaromonte 2023), a functional motif discovery algorithm employing probabilistic K -mean with local alignment, and a motif discovery method based on *SCRIMP-MP* (Zhu et al. 2018), one of the most recent

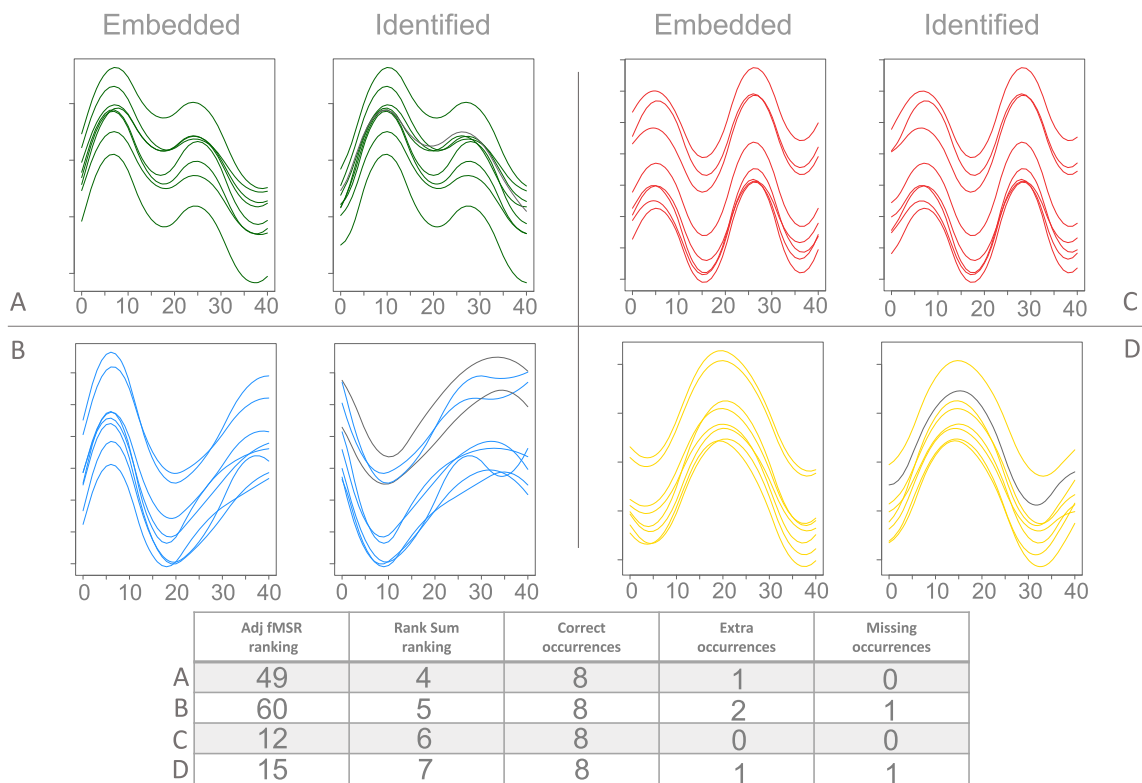


Fig. 5 Motif identification for the simulation employing motif set n. 7 and $\sigma = 2$. *funBlalign* is run with $n_{min} = 6$. See legend for Fig. 4. With high noise, 2 embedded occurrences are missed, and 4 extra occurrences are identified (represented in gray among the identified portions panels)

Matrix Profile (MP) algorithms for motif discovery in univariate time series. We focus on the ability of the three methods to correctly identify embedded motifs. Concerning computational burden, we restrict ourselves to a targeted comparison between *funBlalign* and *probKMA*, the other functional-native method (see below; a rigorous, more comprehensive comparison is hard to perform due to differences in coding languages, degree of code optimization, and overall structure of the pipelines – e.g., number and nature of algorithmic and post-processing parameters to be fixed or tuned).

funBlalign and *probKMA* differ in various respects. *funBlalign* employs agglomerative clustering with adjusted fMSR, which generates a complete hierarchy of all curve portions. In contrast, *probKMA* relies on local functional *K*-means with Sobolev distance. Moreover, while *probKMA* can extend the length of motifs endogenously and discover motifs of varying and unknown sizes, starting from a user-defined set of minimum motif lengths, *funBlalign* requires the length of all motifs to be the same and fixed beforehand. However, being based on *K*-means, *probKMA* must be run several times with different initialization, which can add to the computational cost. In addition, *probKMA* implements a more complex post-processing, involving several tuning parameters – and this could make the method less attractive

for non-specialized user. Finally, *funBlalign* has the benefit of being equally applicable to a single curve or to sets of curves, whereas *probKMA* is designed to operate on sets of curves; to use it in applications involving a single curve, this must be split at the outset – which requires further arbitrary choices, and may potentially lead to the loss of interesting motifs. We compare the performance of the two algorithms using two simulation settings for scenario (1) introduced in Section 4.2 of Cremona and Chiaromonte (2023). In both, two motifs of length 61, say A and B, occur each 12 times across 20 curves. In particular, 12 curves contain 1 occurrence of a single motif (of A for 6, and of B for 6 curves), 4 curves contain 2 occurrences of a single motif (of A for 2, and of B for 2 curves), 2 curves contain 1 occurrence of both motifs, and 2 curves contain no motif. The two settings differ in terms of length of the curves *L* and of noise incorporated in the motifs; one has short curves and low noise ($L = 200$ and $\sigma = 0.1$), and the other long curves and high noise ($L = 500$ and $\sigma = 2$).

We run *funBlalign* with $\ell = 61$ (the true length of the motifs) and $n_{min} = 8, 10$, and post-process candidate motifs with the rank sum criterion. *probKMA* is run using $K = 2, 3$, minimum motif length $v = 41, 51, 61$, and 20 random initializations for each (*K*, *v*) pair; results from the 120 runs with different parameters/initializations are pooled following the

Table 1 Comparison between *funBlalign* and *probKMA*

Setting	Motif	Portion	<i>probKMA</i>	<i>funBlalign</i> - n_{min}	
				= 8	= 10
$L = 200 \sigma = 0.1$	Motif A	Correct	12	12	12
		Extra	0	0	0
	Motif B	Correct	12	12	12
		Extra	0	0	0
$L = 500 \sigma = 2$	Motif A	Correct	11	12	12
		Extra	2	1	1
	Motif B	Correct	12	10	12
		Extra	1	0	3

For *probKMA*, median results across 10 runs are reported

Table 2 Comparison between *funBlalign* and *SCRIMP-MP*

	Motif	<i>funBlalign</i> - n_{min}				MP <i>kneighbor</i>	<i>R</i>			
		5	6	7	8		3	10	25	50
$\sigma = 0.5$	Motif 1	(8,0)	(8,0)	(8,0)	(8,0)	4	(6,0)	(6,0)	(6,0)	(6,0)
		(8,0)	(8,0)	(8,0)	(8,0)	6	(6,0)	(8,0)	(8,0)	(8,0)
		(8,0)	(8,0)	(8,0)	(8,0)	8	(7,3)	(8,0)	(8,2)	(8,2)
	Motif 2	(8,0)	(8,0)	(8,0)	(8,0)	4	(6,0)	(6,0)	(6,0)	(6,0)
		(8,0)	(8,0)	(8,0)	(8,0)	6	(8,0)	(8,0)	(8,0)	(8,0)
		(8,0)	(8,0)	(8,0)	(8,0)	8	(8,0)	(8,0)	(8,2)	(8,2)
	Motif 3	(6,0)	(8,0)	(8,0)	(8,0)	4	(6,0)	(6,0)	(6,0)	(6,0)
		(8,0)	(8,0)	(8,0)	(8,0)	6	(8,0)	(8,0)	(8,0)	(8,0)
		(8,0)	(8,0)	(8,0)	(8,0)	8	(8,0)	(8,0)	(8,2)	(8,2)
	Motif 4	(8,0)	(8,0)	(8,0)	(8,0)	4	(5,0)	(6,0)	(6,0)	(6,0)
		(8,0)	(8,0)	(8,0)	(8,0)	6	(5,0)	(8,0)	(8,0)	(8,0)
		(8,0)	(8,0)	(8,0)	(8,0)	8	(5,0)	(8,1)	(8,2)	(8,2)
$\sigma = 2$	Motif 1	(7,0)	(8,1)	(8,1)	(8,1)	4	(3,0)	(6,0)	(6,0)	(6,0)
		(8,0)	(8,0)	(8,0)	(8,0)	6	(3,0)	(8,0)	(8,0)	(8,0)
		(8,0)	(8,0)	(8,0)	(8,0)	8	(3,0)	(8,2)	(8,2)	(8,2)
	Motif 2	(7,2)	(7,2)	(8,1)	(8,1)	4	(5,1)	(4,2)	(4,2)	(4,2)
		(7,2)	(7,2)	(8,1)	(8,1)	6	(6,2)	(5,3)	(5,3)	(5,3)
		(7,2)	(7,2)	(8,1)	(8,1)	8	(7,3)	(6,4)	(6,4)	(6,4)
	Motif 3	(6,1)	(8,0)	(8,0)	(8,0)	4	(5,1)	(5,1)	(5,1)	(5,1)
		(8,0)	(8,0)	(8,0)	(8,0)	6	(5,1)	(6,2)	(6,2)	(6,2)
		(8,0)	(8,0)	(8,0)	(8,0)	8	(5,1)	(7,3)	(7,3)	(7,3)
	Motif 4	(6,1)	(7,1)	(7,1)	(7,1)	4	(5,1)	(5,1)	(5,1)	(5,1)
		(7,1)	(7,1)	(7,1)	(7,1)	6	(5,3)	(5,3)	(5,3)	(5,3)
		(7,1)	(7,1)	(7,1)	(7,1)	8	(6,4)	(7,3)	(7,3)	(7,3)

Cases in which all occurrences are correctly identified without the addition of any extra portions are in bold

motif discovery post-processing recommended in Cremona and Chiaromonte (2023). The whole procedure is repeated 10 times, and medians are taken over these 10 repetitions when evaluating performance. Results are summarized in Table 1. As expected, the choice of n_{min} has an impact on the performance of *funBlalign*: smaller values can lead to missed portions, and larger values to extra portions. However, with an appropriate n_{min} , *funBlalign* can match and even exceed the performance of *probKMA*. In terms of com-

putational burden, we compare the current implementations of *funBlalign* and *probKMA* in the aforementioned simulated setting with $L = 200$ and $\sigma = 0.1$. Specifically, we run *funBlalign* with $\ell = 61$ and $n_{min} = 10$, and post-process candidate motif length with the rank sum criterion. Conversely, we run *probKMA* using $K = 2, 3$, minimum motif length $v = 61$, no motifs elongation, and 10 random initializations for each (K, v) pairs. Both methods are run 10 times on a local machine (64GB of memory, 8 performance cores

and 2 efficiency cores). On average, the current implementation of *funBlalign* outperforms *probKMA*: while our method requires 1.69 seconds to run sequentially and 1.29 seconds to run in parallel on 10 cores, *probKMA* runs in 69.96 and 24.25 seconds, respectively.

We next consider the motif discovery method based on *SCRIMP-MP* as implemented in the R library *tsmp*. This method employs a k Nearest Neighbours (kNN) routine and a similarity measure based on the z -normalized Euclidean distance, requiring the user to fix several parameters *ex ante*; namely, the number of motifs to be discovered u , their length ℓ , a motif radius R (i.e., the distance within which two portions of the time series are identified as belonging to the same motif), and a maximum number of neighbors to consider $k_{neighbor}$. Specifically, the algorithm will start from pairs of most similar portions, to which other portions of the curve are added only if they are within a distance R from the starting pair, and only up to a maximum of $k_{neighbor}$. In contrast to *probKMA*, this method is designed to operate with a single time series; to use it in applications involving a set of curves, those need to be joined – which can be problematic. Another MP-based method called *Ostinato* (Kamgar et al. 2019) can find “consensus” motifs shared by some curves in a set; however, the definition of “consensus” motif differs from that of functional motif since it considers only the best matching occurrence in each curve. We compare the performance of *funBlalign* and the *SCRIMP-MP* based method using two curves from our own simulation study in Section 4, namely the ones employing motifs set n. 7 with 8 occurrences each, with shared noise level $\sigma = 0.5$ and $\sigma = 2$, respectively (see Figures 4 and 5). We run *funBlalign* with $\ell = 41$ (the true length of the motifs) and $n_{min} = 5, 6, 7, 8$. We use the *SCRIMP-MP* method, as implemented in the R library *tsmp*, to identify a maximum of $u = 50$ motifs with $\ell = 41$. All possible combinations of $R = 3, 10, 25, 50$ and $k_{neighbor} = 4, 6, 8$ are tested. Table 2 summarizes the results most similar to the embedded motifs. Also here the choice of n_{min} has an impact on the performance of *funBlalign*, but R and $k_{neighbor}$ have a yet stronger impact on the performance of *SCRIMP-MP*, which struggles more – especially with higher noise. In addition, tuning these two parameters is not trivial, because intuition and knowledge about them is unlikely to be available *a priori*. A similar comparison using motifs with different noise levels is presented in Section S6.

6 Case studies

In this section, we assess the ability of *funBlalign* to discover functional motifs in real data sets through two case studies. The first concerns food price inflation over a period of around 60 years, and the second temperature changes over

a period of around 30 years. In both cases we utilize monthly measurements across the world provided by FAO.¹

6.1 Case study 1: food price inflation

The FAOSTAT data used in this case study comprise monthly food price inflation measurements from January 2001 to June 2022 (a total of 258 measurements) for different countries and geographical regions (details are available in the repository metadata section, and geographical regions, as defined by the United Nations, are in Section S7). We first seek motifs in a single, world-wide food price inflation curve, and then seek motifs in the curves for 19 distinct geographical regions. Before running *funBlalign*, we smooth the data using local polynomials with Gaussian kernels (*locpoly* function of the R package *KernSmooth*, Wand and Ripley 2006); a bandwidth parameter equal to 1.5 seems appropriate to avoid over-smoothing possibly interesting peaks in the data.

For the world-wide curve, we seek annual patterns setting $\ell = 12$ (months), and fix the minimum cardinality to $n_{min} = 4$; a small n_{min} is appropriate since we are considering a total of only 258 measurements. *funBlalign* identifies 30 motifs capturing various types of shapes. Two, depicting “valleys” and “peaks” are displayed in Fig. 6A. Notably, the “peak” motif has occurrences corresponding to well-known economic crises; the 2001 recession, the global financial crisis of 2007–2008, and the 2020 COVID-19 recession. For the 19 regional curves, which comprise overall many more measurements, we seek longer, biannual patterns setting $\ell = 24$ (months) and fix a larger minimum cardinality $n_{min} = 6$. Here *funBlalign* identifies 415 motifs. To sieve through such a large output, we suggest that users post-process and rank results based on various approaches. For instance, if one is interested in cohesive results regardless of how frequently a motif recurs, the adjusted fMSR criterion is the best choice. In contrast, if cardinality is important, it may be preferable to utilize the rank sum criterion. We note that, due to the definition of fMSR, motifs with lower variance – which look constant – tend to rank higher, a fact that could overshadow some patterns. For instance, among the 415 motifs identified in the 19 regional curves, some rather interesting high variance motifs are at the bottom of the fMSR ranking. Two motifs are shown in Fig. 6B. One, occurring 6 times, is the top-ranking in fMSR; it depicts a “mild ascent”, with low variance. The other, occurring 9 times, is the top-ranking in terms of variance; it depicts a “peak followed by a valley” and ranks only 401st in fMSR.

We end this section remarking on the very steep ascent in food price inflation at the end of the time domain covered by the data (2021 and first half of 2022). This can be noticed in

¹ Data used in these cases studies can be downloaded at <https://www.fao.org/faostat/en/#data/ET>.

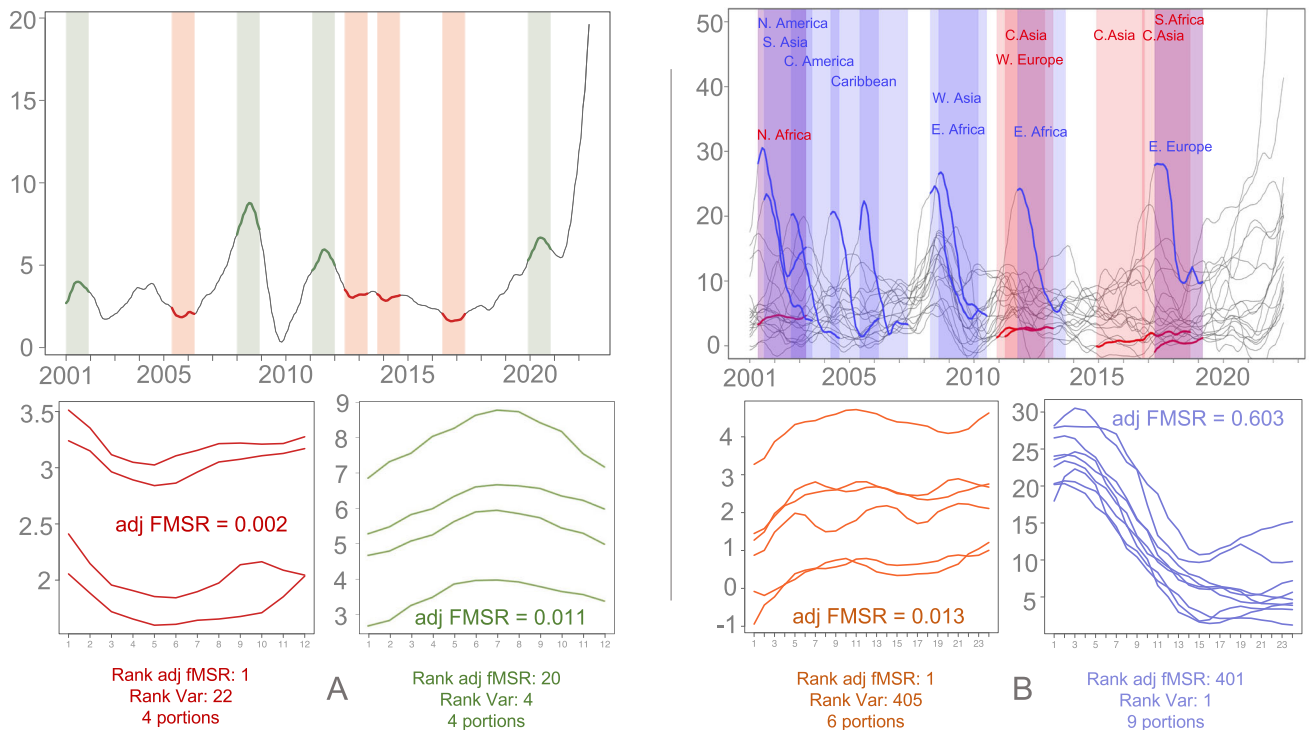


Fig. 6 Two of the motifs identified by *funBlalign* in the world-wide food price inflation curve (panel A) and two of the motifs identified in the curves for 19 geographical regions (panel B). In both panels, we show the top-ranking adjusted fMSR motif (bottom left plot) along with

the world-wide curve (Fig. 6A), and it appears also in some of the 19 regional curves (Fig. 6B) – in particular that of Western Asia, which raises to ≈ 100 and is cut in the plot. Digging deeper, the rise in the Western Asia curve seems itself driven by countries such as Lebanon and the Syrian Arab Republic, which were experiencing conflicts and thus likely additional inflation pressures on top of world-wide COVID-19 related trends.

6.2 Case study 2: Temperature changes

Data used in this case study comprise monthly measurements of temperature changes with respect to a baseline climatology corresponding to the average temperature in the period 1951–1980 for different countries and geographical regions (further details can be found in the repository metadata section). We focus again on the 19 geographical regions defined by the United Nations, and on a period from 1961 to 2021, for a total of 732 monthly measurements (the measurements for a region are obtained by averaging those for the countries belonging to it). Also in this case study, before running *funBlalign*, we smooth the data – but this time we use cubic smoothing B-splines with knots at each month and roughness penalty on the curve second derivative. The smoothing parameter is selected minimizing the average generalized cross-validation error across curves.

a high-ranking motif in terms of variance (bottom right plot) – which, due to the criterion definition, may rank relatively low in terms of fMSR. In the top plot, N, S, E, W, and C stand for Northern, Southern, Eastern, Western, and Central, respectively

In this analysis we seek long, 10-year patterns setting $\ell = 120$ (months). Considering the relative size of motifs to time period covered by the data, we fix again a small minimum cardinality $n_{min} = 3$. Also here *funBlalign* identifies a very large number of motifs; 2367. We post-process and rank them using three different approaches: adjusted fMSR, rank sum, and variance. Top-ranking motifs based on each and the 19 curves are shown in Fig. 7. Notably, a motif can occur simultaneously in multiple regions, and then separately in others – e.g., the top variance motif, with peaks as much as 4°C above baseline, characterizes East Europe and North Africa from the mid '80s to the mid '90s, and then Central Asia almost 30 years later. Also notably, occurrences of different motifs can overlap – e.g., occurrences of the top fMSR and the top rank sum motifs in East Africa, North Africa and South America in the '70s present a large overlap, almost extending one another along the time domain. We end pointing out the increasing upward departures from baseline climatology shown by temperatures in all 19 regions over the 60 years covered by the data.

7 Conclusions

We contribute to the recent literature on functional motif discovery a definition of functional motifs based on an explicit




Fig. 7 Three motifs identified by *funBalign* in the temperature change curves for 19 geographical regions. We show the top-ranking variance (top right – 3 portions), adjusted fMSR (middle right – 3 portions) and

rank sum (bottom right – 4 portions) motifs. In the left plot, N,S,E,W,C and M stand for Northern, Southern, Eastern, Western, Central, and Middle respectively

additive model, and an algorithm designed to discover such motifs; *funBalign*. Related to our additive model, we also introduce the adjusted functional Mean Squared Residue (fMSR) score. The fMSR is a functional extension of the MSR score widely used in the multivariate biclustering literature. Building upon our own past work (Di Iorio et al. 2020), we prove it to be biased towards motifs that occur less often in the data, and formulate a de-biasing adjustment. *funBalign* is a very flexible multi-step algorithm which requires only two, easy to interpret input parameters: the length ℓ and the minimum number of occurrences n_{min} of the motifs to be discovered. It uses agglomerative clustering to produce a hierarchy capturing the relationships among curve portions, followed by a dynamic cutting procedure to identify the most interesting candidate motifs based on such hierarchy, and by a post-processing step to eliminate redundant results.

Notwithstanding the simplistic additive model that underlies the fMSR score, *funBalign* shows very good performance both in extensive simulations and in two real-data case studies. However, the selection of ℓ could be problematic in some applications, requiring the user to run the algorithm with a range of alternative values for this tuning parameter. On a different front, as noted discussing our simulation study, *funBalign* can identify extra instances of a given motif, or even entirely different motifs, that were not intentionally embedded in the simulated curves. While this is an explainable bi-product of the procedure used to create our simulated curves (certain patterns may occur and recur by chance in the curve background), it points to the need for a more rigorous statistical treatment of motif discovery. In particular, in future work, we plan to develop a measure of significance to be used in conjunction with the adjusted

fMSR for motif discovery. In addition, we are planning to extend the algorithm as to allow the identification of motifs invariant to vertical scaling (this would refer to an underlying multiplicative model, instead of the current additive one).

Parsing through all curve portions of a given length allows *funBalign* to tackle applications where classical functional alignment fails or may be inadequate, such as identifying motifs embedded consecutively, or appearing only in one curve. However, this comes at a cost; storing the dissimilarity matrix D_{fMSR} may necessitate a massive amount of memory when analyzing large functional data sets. Nonetheless, we note that this matrix is only calculated once, and in our simulations (each involving one curve encompassing 7001 measurements) and case studies (involving a maximum of 19 curves of 732 measurements) memory did not pose a challenge;

in the examples we considered, *funBalign* runs in under 1 minute (or even less as shown in the comparison with *probKMA*) without RAM problems on a local machine (64GB of memory, 8 performance cores and 2 efficiency cores). We do believe that running time could be further improved leveraging parallelization even more. In future work, we plan to optimize *funBalign* to allow it to scale efficiently also on much larger datasets.

Supplementary information

Supplement material is available online. Current R code for the algorithm is available in the GitHub repository <https://>

github.com/JacopoDior/funBIalign and an R package is in preparation.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s11222-024-10537-y>.

Acknowledgements J. Di Iorio acknowledges the support of the Penn State Eberly College of Science. M.A. Cremona acknowledges the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), of the Fonds de recherche du Québec Health (FRQS), and of FSA, Université Laval. F. Chiaromonte acknowledges support from the Huck Institutes of the Life Sciences, Penn State.

Author Contributions All authors conceived ideas and analysis approaches. J.Di I., implemented the methodology and performed analyses. All authors interpreted findings and participated to the writing of the manuscript.

Funding The authors have no relevant financial or non-financial interests to disclose.

Declarations

Competing interests The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

References

- Angiulli, F., Cesario, E., Pizzuti, C.: Random walk biclustering for microarray data. *Inf. Sci.* **178**(6), 1479–1497 (2008)
- Boschi, T., Di Iorio, J., Testa, L., Cremona, M.A., Chiaromonte, F.: Functional data analysis characterizes the shapes of the first COVID-19 epidemic wave in Italy. *Sci. Rep.* **11**, 17054 (2021). <https://doi.org/10.1038/s41598-021-95866-y>
- Buckner, R.L., Head, D., Parker, J., Fotenos, A.F., Marcus, D., Morris, J.C., Snyder, A.Z.: A unified approach for morphometric and functional data analysis in young, old, and demented adults using automated atlas-based head size normalization: reliability and validation against manual measurement of total intracranial volume. *Neuroimage* **23**(2), 724–738 (2004)
- Chen, D., Cremona, M.A., Qi, Z., Mitra, R.D., Chiaromonte, F., Makova, K.D.: Human L1 transposition dynamics unraveled with functional data analysis. *Mol. Biol. Evol.* **37**, 3576–3600 (2020). <https://doi.org/10.1093/molbev/msaa194>
- Cheng, Y., Church, G.M.: Biclustering of expression data. In *Proceedings of the 8th International Conference on Intelligent Systems for Molecular Biology, La Jolla, CA*, pp. 93–103 (2000)
- Cremona, M.A., Chiaromonte, F.: Probabilistic k -means with local alignment for clustering and motif discovery in functional data. *J. Comput. Graph. Stat.* (2023). <https://doi.org/10.1080/10618600.2022.2156522>
- Das, S., Demirer, R., Gupta, R., Mangisa, S.: The effect of global crises on stock market correlations: evidence from scalar regressions via functional data analysis. *Struct. Chang. Econ. Dyn.* **50**, 132–147 (2019)
- Di Iorio, J., Chiaromonte, F., Cremona, M.A.: On the bias of h-scores for comparing biclusters, and how to correct it. *Bioinformatics* **36**(9), 2955–2957 (2020)
- Di Iorio, J., Vantini, S.: funloci: a local clustering algorithm for functional data. *arXiv:2305.12991* (2023)
- Ferraty, F., Vieu, P.: Nonparametric functional data analysis: theory and practice (2006)
- Ferreira, L., Hitchcock, D.B.: A comparison of hierarchical methods for clustering functional data. *Commun. Stat.-Simul. Comput.* **38**(9), 1925–1949 (2009)
- Floriello, D., Vitelli, V.: Sparse clustering of functional data. *J. Multivar. Anal.* **154**, 1–18 (2017)
- Fraiman, R., Gimenez, Y., Svarc, M.: Feature selection for functional data. *J. Multivar. Anal.* **146**, 191–208 (2016)
- Galvani, M., Torti, A., Menafoglio, A., Vantini, S.: Funcc: a new bi-clustering algorithm for functional data with misalignment. *Comput. Stat. Data Anal.* **160**, 107219 (2021)
- Ghumman, A.R., Haider, H., Shafiqzaman, M.: Functional data analysis of models for predicting temperature and precipitation under climate change scenarios. *J. Water Clim. Chang.* **11**(4), 1748–1765 (2020)
- Jacques, J., Preda, C.: Functional data clustering: a survey. *Adv. Data Anal. Classif.* **8**(3), 231–255 (2014)
- Kamgar, K., Gharghabi, S., Keogh, E.: Matrix profile xv: Exploiting time series consensus motifs to find structure in time series sets. In *2019 IEEE International Conference on Data Mining (ICDM)*, pp. 1156–1161. IEEE (2019)
- Kokoszka, P., & Reimherr, M. (2017). *Introduction to Functional Data Analysis* (1st ed.). Chapman and Hall/CRC. <https://doi.org/10.1201/9781315117416>
- Lila, E., Aston, J.A., Sangalli, L.M.: Functional data analysis of neuroimaging signals associated with cerebral activity in the brain cortex. In *Functional Statistics and Related Fields*, pp. 169–172. Springer (2017)
- Liu, X., Wang, L.: Computing the maximum similarity bi-clusters of gene expression data. *Bioinformatics* **23**(1), 50–56 (2007)
- Liu, X., Yang, M.C.: Simultaneous curve registration and clustering for functional data. *Comput. Stat. Data Anal.* **53**(4), 1361–1376 (2009)
- Lonardi, J., Patel, P.: Finding motifs in time series. In *Proc. of the 2nd Workshop on Temporal Data Mining*, pp. 53–68 (2002)
- Madeira, S.C., Oliveira, A.L.: Biclustering algorithms for biological data analysis: a survey. *IEEE/ACM Trans. Comput. Biol. Bioinf.* **1**(1), 24–45 (2004)
- Mueen, A., Keogh, E., Zhu, Q., Cash, S., Westover, B.: Exact discovery of time series motifs. In *Proceedings of the 2009 SIAM international conference on data mining*, pp. 473–484. SIAM(2009)
- Murtagh, F., Contreras, P.: Algorithms for hierarchical clustering: an overview. *Wiley Interdiscip. Rev. Data Min. Knowl. Discov.* **2**(1), 86–97 (2012)
- Pontes, B., Giráldez, R., Aguilar-Ruiz, J.S.: Biclustering on expression data: a review. *J. Biomed. Inform.* **57**, 163–180 (2015)
- Ramsay, J., Silverman, B.W.: *Functional data analysis* (2005). Springer
- Sangalli, L.M., Secchi, P., Vantini, S., Vitelli, V.: K-mean alignment for curve clustering. *Comput. Stat. Data Anal.* **54**(5), 1219–1233 (2010)
- Vitelli, V.: A novel framework for joint sparse clustering and alignment of functional data. *J. Nonparametr. Stat.* **36**(1), 182–211 (2024)

- Wand, M., Ripley, B.: Kernsmooth: Functions for kernel smoothing for wand & jones (1995). R package version 2, 22–19 (2006)
- Yang, J., Wang, H., Wang, W., Yu, P.S.: An improved biclustering method for analyzing gene expression profiles. *Int. J. Artif. Intell. Tools* **14**(05), 771–789 (2005)
- Yeh, CCM., Zhu, Y., Ulanova, L., Begum, N., Ding, Y., Dau, HA., Silva, DF., Mueen, A., Keogh, E.: Matrix profile i: all pairs similarity joins for time series: a unifying view that includes motifs, discords and shapelets. In *2016 IEEE 16th international conference on data mining (ICDM)*, pp. 1317–1322. Ieee (2016)
- Zhu, Y., Yeh, CCM., Zimmerman, Z., Kamgar, K., Keogh, E.: (2018). Matrix profile xi: Scrimp++: time series motif discovery at interactive speeds. In *2018 IEEE International Conference on Data Mining (ICDM)*, pp. 837–846. IEEE

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.