

Temporal Association Rules from Motifs

Mauro Milella 

Department of Mathematics and Computer Science, University of Ferrara, Italy

Giovanni Pagliarini 

Department of Mathematics and Computer Science, University of Ferrara, Italy

Guido Sciavicco 

Department of Mathematics and Computer Science, University of Ferrara, Italy

Ionel Eduard Stan 

Department of Informatics, Systems, and Communications, University of Milano-Bicocca, Italy

Abstract

A motif is defined as a frequently occurring pattern within a (multivariate) time series. In recent years, various techniques have been developed to mine time series data. However, only a few studies have explored the idea of using motif discovery in temporal association rule mining. Interval-based temporal association rules have been recently defined and studied, along with the temporal version of the known frequent patterns, and therefore, association rule extraction algorithms (such as APRIORI and FP-Growth). In this work, we define a vocabulary of propositional letters wrapping motifs, and show how to extract temporal association rules starting from such a vocabulary. We apply our methodology to time series datasets in the fields of hand signs execution and gait recognition, and we discuss how they capture curious insights within data, keeping a high level of interpretability.

2012 ACM Subject Classification Theory of computation → Modal and temporal logics; Theory of computation → Theory and algorithms for application domains

Keywords and phrases Motifs, Interval Temporal Logic, Association Rules

Digital Object Identifier 10.4230/LIPIcs.TIME.2025.19

Category Short Paper

Supplementary Material *Software (Source Code)*: <https://github.com/aclai-lab/Sole.jl> [11]
archived at `swb:1:dir:dd723aee72578208606649ff12168e891cdae221`

Software (Source Code): <https://github.com/aclai-lab/ModalAssociationRules.jl> [10]
archived at `swb:1:dir:697da0b30a22cd23450ab445a887ebf1a602db8f`

Funding We acknowledge the support of the FIRD project Methodological Developments in Modal Symbolic Geometric Learning, funded by the University of Ferrara. Moreover, this research has also been funded by the Italian Ministry of University and Research through PNRR – M4C2 – Investimento 1.3 (Decreto Direttoriale MUR n. 341 del 15/03/2022), Partenariato Esteso PE00000013 – “FAIR – Future Artificial Intelligence Research” – Spoke 8 “Pervasive AI”, funded by the European Union under the NextGeneration EU programme.

1 Introduction

In machine learning, we distinguish between *functional* and *symbolic* learning. The former encompasses strategies for representing the theory underlying a certain phenomenon as *functions*, while the latter derives *logical descriptions* of that phenomenon. Traditional symbolic methods rely on propositional logic and assume static data, in which each instance is described by n attributes. Temporal data, however, cannot be successfully dealt within the same schema in a native way, and it is commonly pre-processed (e.g., via averaging attributes along all dimensions) to appear static, enabling standard symbolic techniques.



© Mauro Milella, Giovanni Pagliarini, Guido Sciavicco, and Ionel Eduard Stan;
licensed under Creative Commons License CC-BY 4.0

32nd International Symposium on Temporal Representation and Reasoning (TIME 2025).

Editors: Thierry Vidal and Przemysław Andrzej Wałęga; Article No. 19; pp. 19:1–19:7

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Modal, and in particular temporal symbolic learning [8, 12] takes a different point of view: modal symbolic methods are characterized by being based on modal logic (e.g., temporal logic) so that non-static data can be dealt with natively, and that the extracted knowledge takes the form of interpretable modal logic formulas. Association rule extraction techniques can benefit from a similar approach [9, 14], with the introduction of the modal adaptation of the known frequent set extraction algorithms, namely APRIORI [1] and FP-Growth [5]. In particular, the temporal case is considered a representative example to illustrate the qualities and characteristics of modal association rules. Nevertheless, the original approach presents some important limitations, originating from the basic definition of temporal alphabet, which may cause difficult-to-interpret association rules. One way to overcome the limits in alphabet definition towards temporal association rule extraction is to consider time series *motifs*, that is, patterns that are considered interesting because they frequently occur.

In this paper, we consider the problem of temporal rule extraction from time series using a motif-based alphabet. We adapt the original definition of *local support*, introduced in [9, 14], to accommodate motifs suitably, and apply our methodology to two temporal datasets, showing how the obtained rules have an immediate natural language translation.

2 Background

Motif discovery for time series was introduced in 2003 [4], generating quite a body of research. Virtually every time series data mining technique has been applied to the motif discovery problem, including indexing, data discretization, triangular-inequality pruning, hashing, early abandoning.

► **Definition 1.** A time series is a sequence $T : t_1, \dots, t_N$ of N real-valued observations. A set of $\mathcal{T}$ time series $T_1, \dots, T_n$ is called *multivariate time series*. A region $T_{ij} : t_i, \dots, t_j$ of $j - i + 1$ consecutive observations in a time series is called a *subsequence*. Given two subsequences T_i and T_k , their distance is the Euclidean distance between their Z-normalized forms.

Given a subsequence, we can compute its distance to all subsequences in the same time series.

► **Definition 2.** Given an N -observations time series T and a subsequence T_{ij} , the distance profile D_i is the vector of all distances between T_{ij} and $T_{k(k+j-i+1)}$, for every $1 \leq k \leq N$.

Distance profiles are collected in a distance matrix, from which a matrix profile is built.

► **Definition 3.** Given a time series T and a length l , the full distance matrix $M_{T,l}$ is the squared matrix of dimension $N - l + 1$ whose i -th row is the distance profile D_i . Given an integer s , the matrix profile P is the vector that at position i contains the minimum value in D_i , ignoring all the values in $D_{i(i-s)}$ and $D_{i(i+s)}$, as these distances are considered trivialities.




For each position i in a matrix profile P , if the value P_i is lower than a threshold α , we say that $T_{i(i+l)}$ is an instance of a *motif*.

Note that both the literature and the maintained packages on motif discovery are relatively extensive; in this work, we enrich the existing software ecosystem with the framework *Sole.jl*¹, an open source solution for symbolic learning with non-tabular data, leveraging the *MatrixProfile.jl*² library, which implements the STAMP algorithm [15] for computing the matrix profile of a time series, allowing for motifs discovery.

¹ See <https://github.com/aclai-lab/Sole.jl> and <https://github.com/aclai-lab/ModalAssociationRules.jl>.

² <https://github.com/baggepinnen/MatrixProfile.jl>

■ **Table 1** Allen’s relation and their notation within HS; equality and inverse relations are omitted.

HS modality	Definition w.r.t. the interval structure		Example
$\langle A \rangle$ (after)	$[i, j]R_A[i', j']$	$\Leftrightarrow j = i'$	
$\langle L \rangle$ (later)	$[i, j]R_L[i', j']$	$\Leftrightarrow j < i'$	
$\langle B \rangle$ (begins)	$[i, j]R_B[i', j']$	$\Leftrightarrow i = i' \wedge j' < j$	
$\langle E \rangle$ (ends)	$[i, j]R_E[i', j']$	$\Leftrightarrow j = j' \wedge i < i'$	
$\langle D \rangle$ (during)	$[i, j]R_D[i', j']$	$\Leftrightarrow i < i' \wedge j' < j$	
$\langle O \rangle$ (overlaps)	$[i, j]R_O[i', j']$	$\Leftrightarrow i < i' < j < j'$	

In this paper, we are interested in learning association rules from data. Fixed an alphabet of propositional literals $\mathcal{P} = \{p_1, \dots, p_k\}$, a *propositional rule* is an object of the type $\rho : X \Rightarrow Y$, where $X \subset \mathcal{P}$ is called *antecedent*, $Y \subset \mathcal{P}$ is called *consequent*, and $X \cap Y = \emptyset$; following the classical jargon, we refer to each literal as *item* and we call a set of items as *itemset*. In this scenario, $X \cup Y$ is considered interesting when it is *frequent*, that is, if it occurs more often than a predetermined threshold referred to as *minimum support*, and a rule $X \Rightarrow Y$ is extracted if the ratio between the support of X and that of $X \cup Y$ is higher than another predetermined threshold known as *minimum confidence*.

While classical association rules are designed for propositional patterns to emerge, temporal association rules are designed to generalize this idea to patterns with a temporal component. The natural choice to describe temporalized co-occurrence of events or patterns with a duration is *Halpern and Shoham modal logic of time intervals* (HS), defined over Allen’s relations, as show in Tab. 1.

Interval temporal logic gives us a way to naturally describe temporal association rules, as time series can be naturally seen as interval models. Let $\mathfrak{T} = \{\mathcal{T}_1, \dots, \mathcal{T}_m\}$ be a set of (multivariate) time series, or *temporal dataset*, and fix a propositional alphabet $\mathcal{P}$; let us also assume that each time series in $\mathfrak{T}$ is based on the same temporal domain D . Each single multivariate time series $\mathcal{T}$ is a collection of n time series $T_1, \dots, T_n$; elements of $\mathcal{P}$ are naturally associated to a specific time series. In this way, if $\mathcal{T}$ is a multivariate time series (i.e., an interval model), $[i, j]$ is an interval, and φ an interval formula, the notion of $\mathcal{T}, [i, j] \models \varphi$ can be interpreted as the notion of $\mathcal{T}$ *satisfies* φ at $[i, j]$.

A *temporal itemset* is a set of temporal items, that is, items enriched with a temporal relation, and temporal rules are such that antecedents and consequents are temporal itemsets. The notion of support is generalized to the case of a temporal dataset $\mathfrak{T}$ by distinguishing a *local support*, computing the relative frequency of a temporal itemset occurring in some instance $\mathcal{T} \in \mathfrak{T}$, and a *global support*, counting the number of instances such that their local support is higher than a minimum local support threshold. An itemset X is said to be *frequent* if its global support is greater than a minimum global support threshold. In this scenario, ModalAPRIORI [14] and ModalFP-Growth [9] can be used to extract temporal patterns as a particular case of modal patterns.

3 Extracting Temporal Association Rule from Motifs

Let us consider a set of *feature extraction functions* $\mathcal{F} = \{F_1, \dots, F_k\}$, where each function F is defined as $F : \mathbb{R}^d \rightarrow \mathbb{R}$ for some natural value $d \leq N$. Theoretically speaking, given a multivariate time series $\mathcal{T}$, it is possible to define an alphabet of items based on $\mathcal{F}$, that is, $\mathcal{P} = \{\alpha \leq F(T) \leq \beta \mid F \in \mathcal{F}, T \in \mathcal{T}, \alpha \in \mathbb{R} \cup \{-\infty\}, \beta \in \mathbb{R} \cup \{+\infty\}\}$, which allows for mining temporal association rules, as items can be immediately interpreted over intervals, obtaining a scalar value which can be compared with lower and upper bounds α, β .

Unfortunately, it can be shown that this approach may introduce strong bias during (local) support computation, leading to promising association rules which, however, encode trivialities. An intuition about this is that the scalar value obtained by applying a feature extraction function to an interval $[i, j]$, could be redundant with many other identical values obtained by applying the same function on sub-intervals or super intervals of $[i, j]$ (e.g., considering *max*, *min*, *average* functions, but even more refined functions such as *catch22* [7]).

To avoid flattening intervals with feature extraction functions, we modify the definition of the alphabet. We consider a set $\mathcal{A}_{\mathfrak{T}}$ of the *most representative motifs* for the temporal dataset $\mathfrak{T}$, that is, the distinct motifs approximating the largest fraction of data in $\mathfrak{T}$ [6]. Fixed a distance function δ and a constant α , we define a motif-based temporal alphabet as:

$$\mathcal{P} = \{\delta(T, \mu) \leq \alpha \mid T \in \mathfrak{T}, \mu \in \mathcal{A}_{\mathfrak{T}}, \alpha \in \mathbb{R}\}.$$

This is a first step towards dealing with intervals natively, as a given non-temporal item $p \in \mathcal{P}$ is true on T_{ij} if and only if the behaviour of T within the segment is close enough to the motif μ encapsulated by p . However, the definition above is insufficient to guarantee that the computation of local support is unbiased. For example, we can consider a temporal item enriched with Allen's $\langle B \rangle$ relation: if it is true on T_{ij} , then it automatically inflates the support computation, as it is trivially true on all the interval $T_{i(j+1)} \dots T_{iN}$.

To ensure a fair support computation, we must “anchor” temporal items to a specific temporal frame, that is, the temporal relations must be applied to the subset of intervals having at least one non-temporal item true on them, instead of possibly any interval. This consideration is also crucial when translating an itemset to natural language, as it naturally filters out two tricky scenarios: firstly, the case in which a conjunction of only temporal items is considered, as it is no clear the set of intervals they refer to, and, secondly, it ensures that two non-temporal items wrapping motifs of different lengths are not mixed, as the number of intervals the two items can be true at the same time is 0.

► **Definition 4.** *An itemset X is said to be anchored iff contains at least one non-temporal item, and all non-temporal items $\Omega \subseteq X$, called anchor of X , are based on motifs of the same length. Let $l(\Omega)$ denote the length of the interval on which an anchor may hold.*

In this way, the length of intervals that may potentially satisfy the entire itemset is fixed, allowing us to define the frequency of that itemset in a truly representative way.

► **Definition 5.** *Let $\mathfrak{T}$ be a temporal dataset, $\Lambda_{\mathcal{P}}$ be the set of temporal items built on the motif-based alphabet $\mathcal{P}$, let $X \subseteq \Lambda_{\mathcal{P}}$ be an anchored itemset, and let $\Omega \subseteq X$ be its anchor. The motif-based local and global supports of X on some instance $\mathcal{T} \in \mathfrak{T}$ are defined respectively as:*

$$mbls_{\mathcal{T}}(X) = \frac{|\{[i, j] \in I(D) \mid T, [i, j] \models X\}|}{|\{[i, j] \in I(D) \mid j - i + 1 = l(\Omega)\}|}, \quad mbgs_{\mathfrak{T}}^{s_l}(X) = \frac{|\{\mathcal{T} \in \mathfrak{T} \mid mbls_{\mathcal{T}}(X) \geq s_l\}|}{|\mathfrak{T}|}.$$

Using the motif-based support, we proceed to comment our experiments.

4 Experiments

We consider two well-known public datasets concerning human gestures and gait recognition, namely NATOPS [13], from which we consider an arm gesture called “I have command”, and HuGaDB [3], from which we consider the “Walk” movement, aiming to describe insightful common patterns and to express them in natural language.

In both cases, we mine frequent itemsets by leveraging ModalAPRIORI. First of all, we extract the top 5 motifs with length 10 and the top 3 with length 20, in order to capture qualitatively appreciable patterns; the two class of motifs encodes a behaviour expressed respectively in nearly half a second and one second: shorter subsequences would bring little informativeness, while longer one would be too coarse. We establish two thresholds for minimum local and global support, respectively s_l and s_g , to 0.1, which is relatively low: this is not a problem, as the vast majority of the association rules generated after the itemset extraction phase (14233 for “I have command”, 253 for “Walking”) are filtered out by leveraging confidence and *lift* [2] meaningfulness measures. In particular, the higher the lift, the more the antecedent and the consequent of a rule are positively correlated.

The motif-based alphabet is generated considering the Z-normalized Euclidean distance δ . We choose the threshold α associated to a specific motif μ to be the tenth percentile of the values in the distance profile of μ ; in this way, we ensure each single non-temporal item to be frequent. We enrich the propositional alphabet with the set of temporal items $\Lambda_{\mathcal{P}}$ obtained by considering every Allen’s relation.

Even if we did not graphically shown an example of “I have command” gesture, we try to describe a chunk of it with the following rule we extracted, which exhibits a perfect confidence (of 1.0) and an high lift (of 6.7). The rule is expressed in a compact format, fixing a coordinate, leveraging superscripts to provide an intuition about the movement captured by the motif underlying each item, and subscripts to indicate the length of the motif and which body part it refers to (e.g., *re* for right elbow, *rh* for right hand).

$$y_{20, re}^{up} \wedge z_{20, rh}^{front} \Rightarrow \langle O \rangle x_{10, rh}^{left\&invert} \wedge \langle B \rangle y_{10, re}^{rest\&up}$$

The rule above can be read as *whenever the right hand of the operator is completely stretching in front of him/her and their elbow goes all the way up on the y-axis, the same elbow started the movement range in a rest position and, near the end of the movement range, the operator’s right hand is moving to the left, but will soon change direction.*

The following pair of rules describes a non-trivial behaviour typical of the walking gait; note that *rf* (*lf*) stands for right (left) foot, while *rt* (*lt*) stands for right (left) thigh.

$$x_{10, rf}^{front} \Rightarrow \langle A \rangle x_{10, lt}^{up}, \quad x_{10, lf}^{back} \Rightarrow \langle A \rangle x_{10, rt}^{up}$$

The rules states that *when the right (left) foot accelerates forward (backwards) for approximately half of a second, then the left (right) thigh accelerates upward immediately after.*

5 Conclusions

In this paper we introduced a motif-based approach to temporal alphabet definition and rule extraction. By leveraging motifs-based frequently recurring patterns in time series, we obtained a more interpretable and structurally robust framework for mining temporal association rules, solving the biases introduced by naive alphabet definitions and enhancing semantic clarity. We formalized the concept of anchored itemsets and introduced a novel definition of motif-based local and global support, ensuring that patterns are both meaningful and computationally tractable. Experimental validation on temporal datasets demonstrates the expressiveness and interpretability of the extracted rules, showing promise for applications in explainable temporal data analysis.

References

- 1 R. Agrawal and R. Srikant. Fast Algorithms for Mining Association Rules in Large Databases. In *Proceedings of 20th International Conference on Very Large Data Bases (VLDB)*, pages 487–499, 1994.
- 2 Sergey Brin, Rajeev Motwani, Jeffrey D. Ullman, and Shalom Tsur. Dynamic itemset counting and implication rules for market basket data. In *SIGMOD 1997, Proceedings ACM SIGMOD International Conference on Management of Data, May 13-15, 1997, Tucson, Arizona, USA*, pages 255–264, 1997. doi:10.1145/253260.253325.
- 3 Roman Chereshevnev and Attila Kertész-Farkas. Hugadb: Human gait database for activity recognition from wearable inertial sensor networks. In *Analysis of Images, Social Networks and Texts - 6th International Conference, AIST 2017, Moscow, Russia, July 27-29, 2017, Revised Selected Papers*, pages 131–141. Springer, 2017. doi:10.1007/978-3-319-73013-4_12.
- 4 B.Y. Chiu, E.J. Keogh, and S. Lonardi. Probabilistic discovery of time series motifs. In *Proc. of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 493–498. ACM, 2003. doi:10.1145/956750.956808.
- 5 J. Han, J. Pei, and Y. Yin. Mining Frequent Patterns without Candidate Generation. In *Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data*, pages 1–12, 2000.
- 6 S. Imani, F. Madrid, W. Ding, S. Crouter, and E. Keogh. Matrix profile xiii: Time series snippets: A new primitive for time series data mining. In *Proc. of the IEEE international conference on big knowledge (ICBK)*, pages 382–389. IEEE, 2018.
- 7 Carl H Lubba, Sarab S Sethi, Philip Knaute, Simon R Schultz, Ben D Fulcher, and Nick S Jones. catch22: Canonical time-series characteristics: Selected through highly comparative time-series analysis. *Data mining and knowledge discovery*, 33(6):1821–1852, 2019. doi:10.1007/S10618-019-00647-X.
- 8 Federico Manzella, Giovanni Pagliarini, Guido Sciavico, and Ionel Eduard Stan. Interval temporal random forests with an application to COVID-19 diagnosis. In Carlo Combi, Johann Eder, and Mark Reynolds, editors, *28th International Symposium on Temporal Representation and Reasoning, TIME 2021, September 27-29, 2021, Klagenfurt, Austria*, volume 206 of *LIPICs*, pages 7:1–7:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.TIME.2021.7.
- 9 M. Milella, G. Pagliarini, G. Sciavico, and I.E Stan. Modalfp-growth: Efficient extraction of modal association rules from non-tabular data. In *Proc. of the 25th Italian Conference on Theoretical Computer Science (ICTCS)*, volume 3811 of *CEUR*, pages 241–254. CEUR-WS.org, 2024. URL: <https://ceur-ws.org/Vol-3811/paper150.pdf>.
- 10 Mauro Milella, Giovanni Pagliarini, Guido Sciavico, and Ionel Eduard Stan. ModalAssociationRules.jl. Software, version 0.1.0., swhId: `swh:1:dir:697da0b30a22cd23450ab445a887ebf1a602db8f` (visited on 2025-09-18). URL: <https://github.com/aclai-lab/ModalAssociationRules.jl>, doi:10.4230/artifacts.24783.
- 11 Mauro Milella, Giovanni Pagliarini, Guido Sciavico, and Ionel Eduard Stan. Sole.jl. Software, version 0.6.2., swhId: `swh:1:dir:dd723aee72578208606649ff12168e891cdae221` (visited on 2025-09-18). URL: <https://github.com/aclai-lab/Sole.jl>, doi:10.4230/artifacts.24782.
- 12 Guido Sciavico and Ionel Eduard Stan. Knowledge extraction with interval temporal logic decision trees. In Emilio Muñoz-Velasco, Ana Ozaki, and Martin Theobald, editors, *27th International Symposium on Temporal Representation and Reasoning, TIME 2020, September 23-25, 2020, Bozen-Bolzano, Italy*, volume 178 of *LIPICs*, pages 9:1–9:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPICs.TIME.2020.9.
- 13 Y. Song, D. Demirdjian, and R. Davis. Tracking body and hands for gesture recognition: NATOPS aircraft handling signals database. In *Proc. of the 9th IEEE International Conference on Automatic Face and Gesture Recognition (FG)*, pages 500–506, 2011.

- 14 I.E Stan, G. Sciavicco, E. Muñoz-Velasco, G. Pagliarini, M. Milella, and A. Paradiso. On modal logic association rule mining. In *Proc. of the 23rd Italian Conference on Theoretical Computer Science (ICTCS)*, volume 3284 of *CEUR*, pages 53–65. CEUR-WS.org, 2022. URL: <https://ceur-ws.org/Vol-3284/492.pdf>.
- 15 Y. Zhu, Z. Zimmerman, N.S. Senobari, C-C.M. Yeh, G.J. Funning, A. Mueen, P. Brisk, and E.J. Keogh. Matrix profile II: exploiting a novel algorithm and GPUs to break the one hundred million barrier for time series motifs and joins. In *Proc. of the 16th International Conference on Data Mining (ICDM)*, pages 739–748. IEEE, 2016.