Rule Transform: A Shapelet-based Temporal Association Rule Miner for Multivariate Time Series Classification

Omar Bahri*, Pouya Hosseinzadeh*, Soukaina Filali Boubrahimi*, Shah Muhammad Hamdi †
Department of Computer Science, *Utah State University and † New Mexico State University

Shapelet-based Classification

- A shapelet is defined as a characteristic, phase-independent subsequence that happens frequently in a time series.
- Shapelet-based time series classification algorithms are inherently interpretable.
- Shapelet Transform (ST) has proved to be one of the most powerful time series classification algorithms, with accuracies close to state-of-the-art deep learning, domain transformation-based, and ensemble learning algorithms.
- ST separates the shapelet discovery step from the classification. Therefore, different classifiers can be trained independently in the temporary feature space generated by ST.

Temporal Rule Mining

- Inspired by association rule mining and rule discovery algorithms, temporal association rule mining approaches aim to extract meaningful temporal rules from time series datasets.
- Temporal relationships can then be interpreted by domain experts to achieve a better understanding of the problem at hand, in a purely qualitative way.
- In its simplest form, the relationship between the two is *precedes*, i.e. $A \rightarrow B$, meaning that A happens before B.

Allen's Interval Algebra

precedes	meets	overlaps	finished by	contains	starts	equals
p	m	o	F	D	s	
preceded by	met by	overlapped by	finishes	contains	started by	
P	M	O	f	d	S	e

Table 1: Allen's Interval relationships

- Allen's Interval Algebra is a system for reasoning about temporal relations.
- It is built on a set of 13 relationships between time intervals. These relationships are distinct, exhaustive, and qualitative.

Abstract

The amount of time series data being collected has exploded in parallel with the usage of sensing devices and the improvement of storage capacities. As a result, massive time series datasets have been produced in recent years. In order to analyze these datasets, substantial efforts are being invested in the implementation of temporal rule mining algorithms. Coupled with the knowledge of domain experts, these approaches are extremely helpful for interpreting time series data and reaching important conclusions on the underlying systems. On the other hand, numerous algorithms have been proposed for the time series classification task. In this work, we aim to increase the interpretability of current time series classification methods, by filling the gap between rule discovery and classification in time series mining. We propose rule transform (RT), an algorithm that uses temporal algebra to transform a time series dataset into a new feature space consisting of the support of the most prominent temporal rules. Thus, the generated feature space can be 1) qualitatively interpreted by domain experts and 2) used for classification. We evaluate our algorithm on the UEA archive, and prove that it produces accuracies superior to state-of-the-art time series classification algorithms with the additional interpretability edge. Then, we apply it to a real-life solar flare dataset and discuss our findings.

Rule Transform

RT aims to fill the gap between rule mining and classification in time series mining. It uses temporal algebra to extract the most prominent temporal rules from a given dataset and relies on shapelets as building blocks. The intuition is that if temporal rules can provide domain experts with qualitative insights into the data, they can also help a machine learning classifier distinguish between different classes.

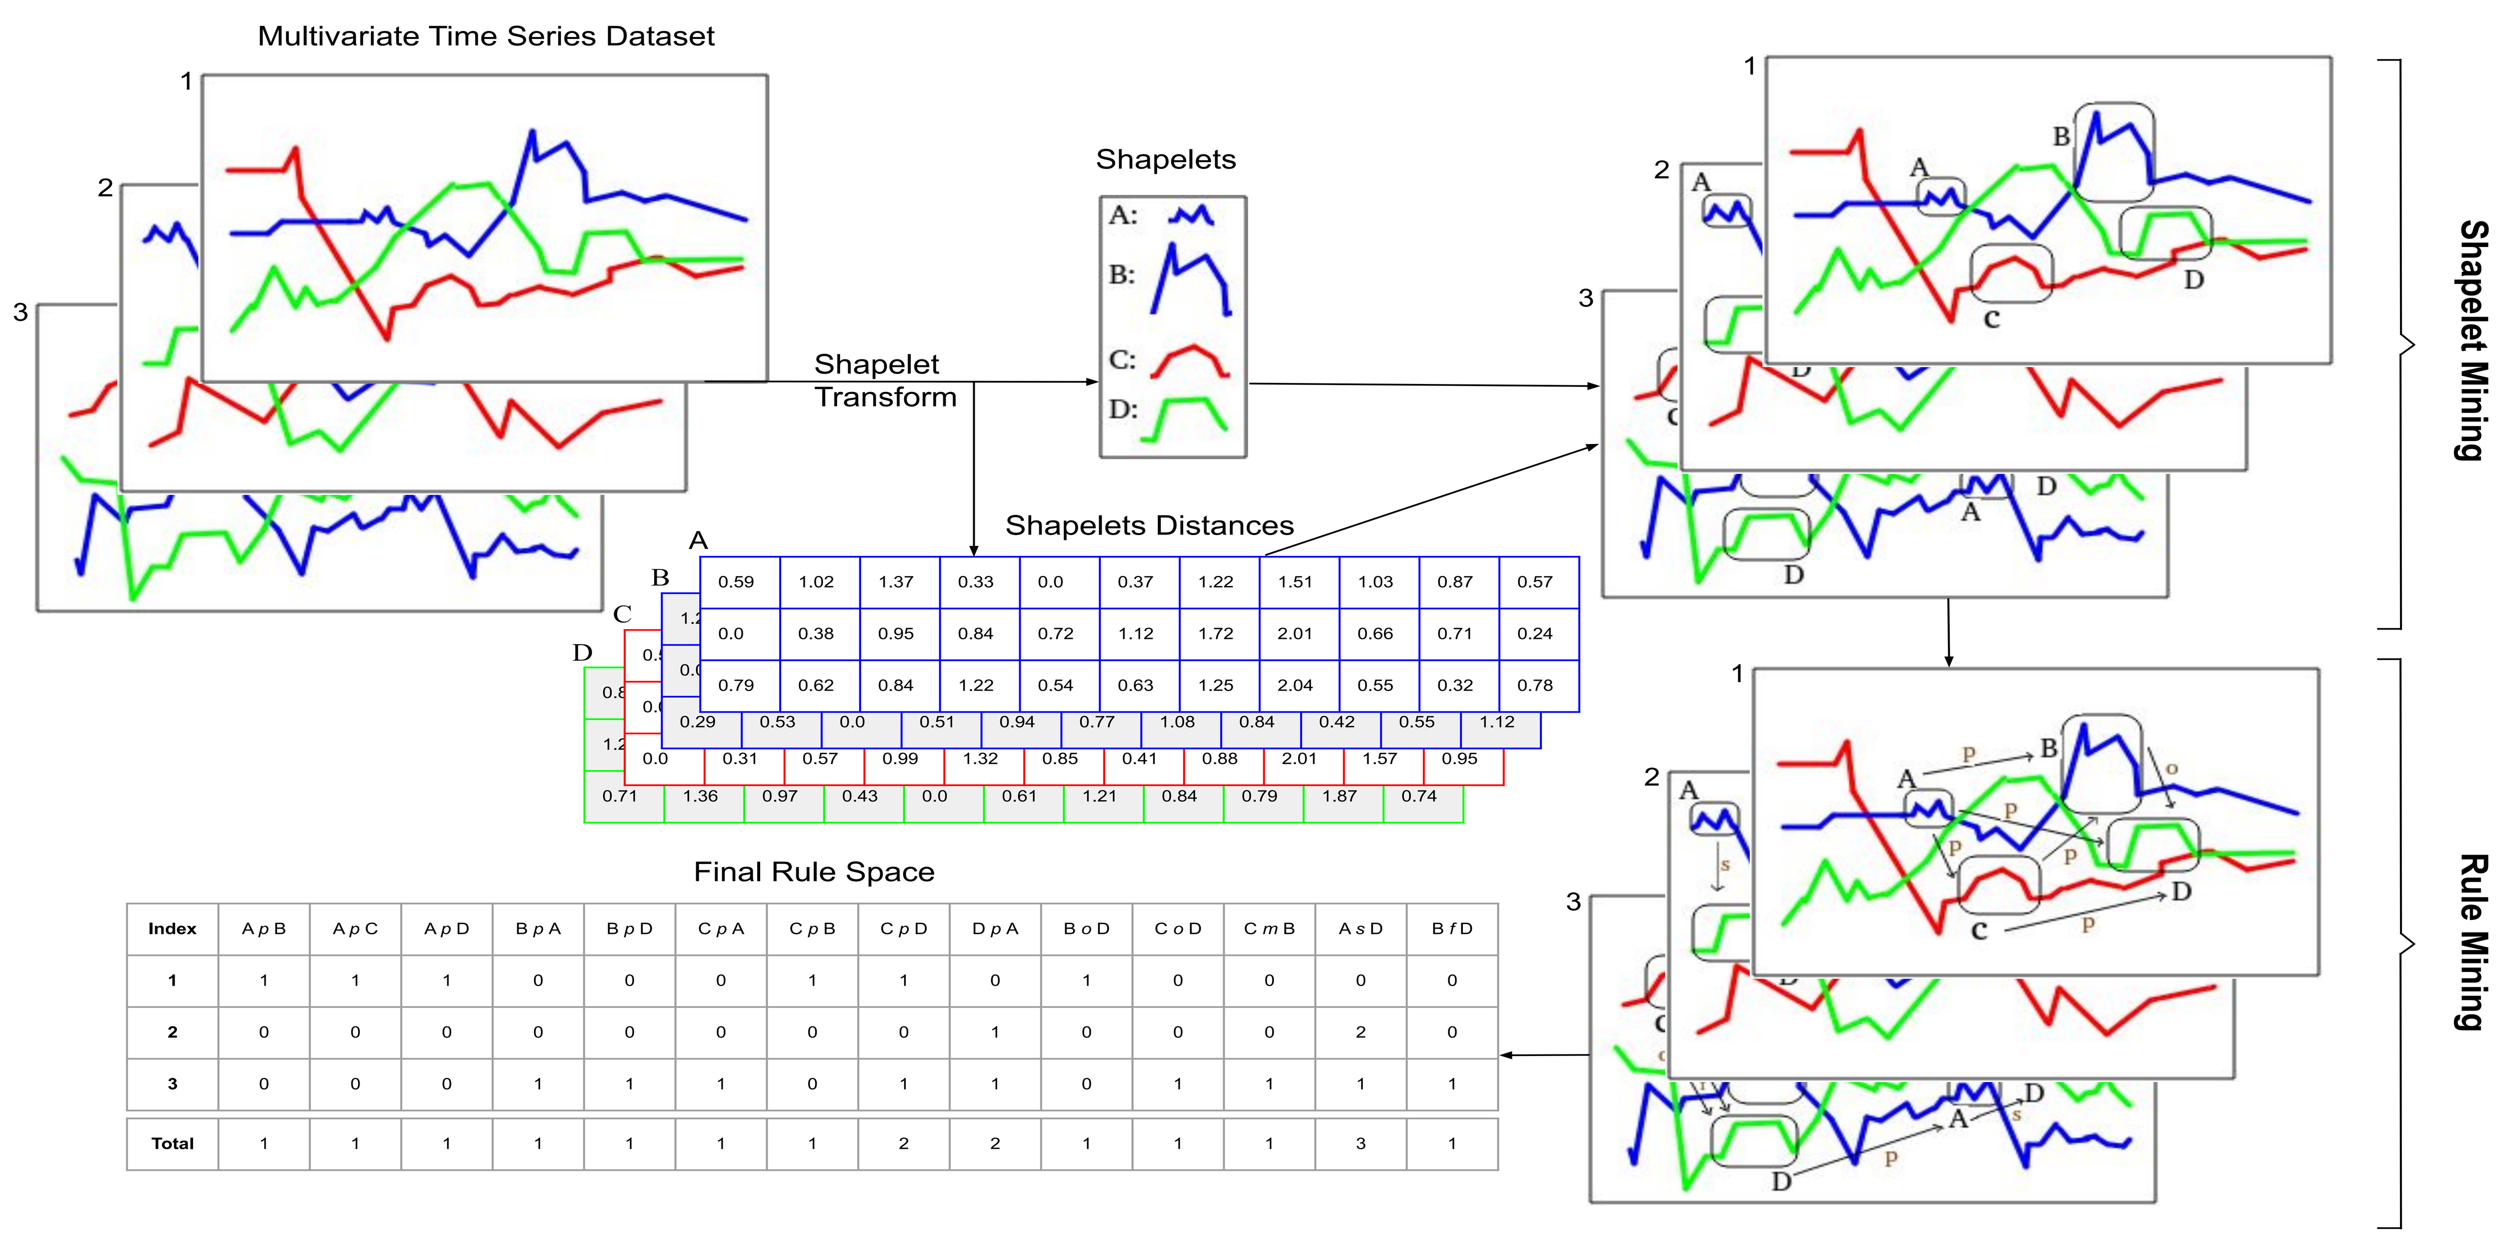

Figure 1: Rule Transform on a Toy Dataset

1. The most discriminative shapelets are extracted from each dimension using ST.
2. Rules are created based on Allen's interval relationships and using the shapelets as time intervals.
3. The support of each rule is computed as the total number of its occurrences.
4. The dataset is transformed to the final rule space. Each rule represents a feature variable, with its support as value.
5. Feature selection is performed to select the most discriminative rule set.

Experiments

Benchmarking:
We evaluate the performance of RT and compare it to Fast Shapelets (FS), Learning Times series Shapelets (LTS) Efficient learning Interpretable Shapelet (ELIS), and ST on the UEA multivariate time series classification archive. After performing feature selection on the rule space using Fisher scores, the results in Table 2.a show that RT performs better than the compared baselines.

Algorithm	Wins	Average Rank
FS	6	2.50
LTS	4	2.92
ELIS	0	3.96
ST	5	2.85
RT	3	1.69

(a) UEA Archive

Algorithm	Accuracy (%)	STD
FS	46.40	1.60
LTS	43.50	2.21
ST	44.39	2.79
RT	72.82	1.77

(b) Solar Flare Dataset

Table 2: Performance Comparison

Case Study: Solar Flare Classification
We apply RT to a real-life solar flare-based multivariate time series dataset. The original data is collected in the form of solar vector magnetograms, captured by the Helioseismic Magnetic Imager (HMI) onboard NASA's Solar Dynamics Observatory (SDO). Each sample in the dataset is made of 33 dimensions representing active region magnetic field parameters recorded at 12 minutes intervals for a duration of 12 hours (60 time steps in total) and labeled according to the largest solar flare that occurred in the next 12 hours (5 different classes: X, M, C, B, and Q). Table 2.b shows that RT was superior in this case too. **In addition, visualizing the temporal rules extracted by RT can provide significant insights to domain experts.** In particular, it is highly interesting to examine rules that happen uniquely under one output class, such as the *A overlaps B* (Figure 2.a) and an *A precedes B* (Figure 2.c) rule which happen uniquely under the X class. Furthermore, shapelet *A* in the *A precedes B* (Figure 2.d) happens also under class Q, which highlights the importance of the rule in discriminating the two classes, and in this case, in preventing high-risk solar flare events.

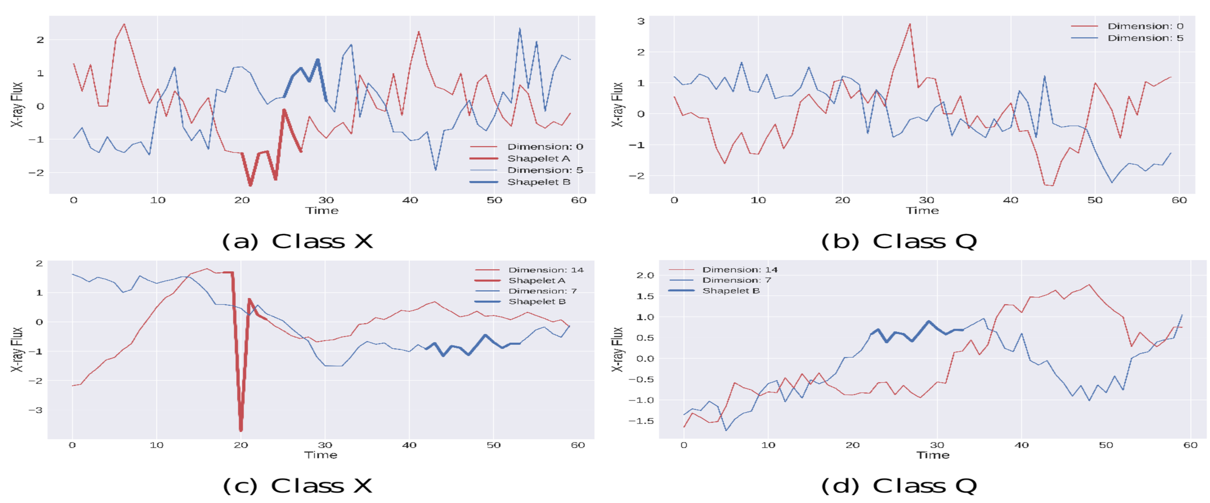

Figure 2: Two Discriminative Rules Extracted