

Noise Sensitivity of an Information Granules Filtering Procedure by Genetic Optimization for Inexact Sequential Pattern Mining

Enrico Maiorino, Francesca Possemato, Valerio Modugno
and Antonello Rizzi

Abstract One of the most essential challenges in Data Mining and Knowledge Discovery is the development of effective tools able to find regularities in data. In order to highlight and to extract interesting knowledge from the data at hand, a key problem is frequent pattern mining, i.e. to discover frequent substructures hidden in the available data. In many interesting application fields, data are often represented and stored as sequences over time or space of generic objects. Due to the presence of noise and uncertainties in data, searching for frequent subsequences must employ approximate matching techniques, such as edit distances. A common procedure to identify recurrent patterns in noisy data is based on clustering algorithms relying on some edit distance between subsequences. However, this plain approach can produce many spurious patterns due to multiple pattern matchings on close positions in the same sequence excerpt. In this paper, we present a method to overcome this drawback by applying an optimization-based step lter that identifies the most descriptive patterns among those found by the clustering process, and allows to return more compact and easily interpretable clusters. We evaluate the mining systems performances on synthetic data in two separate cases, corresponding respectively to two different (simulated) sources of noise. In both cases, our method performs well in retrieving the original patterns with acceptable information loss.

Keywords Granular modeling · Sequence data mining · Inexact sequence matching · Frequent subsequences extraction · Evolutionary computation

E. Maiorino (✉) · F. Possemato · A. Rizzi
Department of Information Engineering, Electronics and Telecommunications (DIET),
SAPIENZA University of Rome, Via Eudossiana 18, 00184 Rome, Italy
e-mail: enrico.maiorino@uniroma1.it

V. Modugno
Dipartimento di Ingegneria Informatica, Automatica e Gestionale (DIAG), SAPIENZA
University of Rome, Via Ariosto 25, 00185 Rome, Italy
e-mail: valerio.modugno@uniroma1.it

1 Introduction

Nowadays, sequence data mining is a very interesting field of research that is going to be central in the next years due to the growth of the so called “Big Data” challenge. Moreover, available data in different application fields consist in sequences (for example over time or space) of generic objects. Generally speaking, given a set of sequences defined over a particular domain, a data mining problem consists in searching for possible frequent subsequences (patterns), relying on inexact matching procedures. In this work we propose a possible solution for the so called *approximate subsequence mining* problem, in which we admit some noise in the matching process. As an instance, in computational biology, searching for recurrent patterns is a critical task in the study of DNA, aiming to identify some genetic mutations or to classify proteins according to some structural properties. Sometimes the process of pattern extraction returns sequences that differ from the others in a few positions. Consequently, the choice of an adequate dissimilarity measure becomes a critical issue when we are designing an algorithm able to deal with this kind of problems. Handling sequences of objects is another challenging aspect, especially when the data mining task is defined over a *structured* domain of sequences [1, 2] Thinking data mining algorithms as a building block of a wider system facing a classification task, a reasonable way to treat complex sequential data is to map sequences to \mathbb{R}^d vectors by means of some *feature extraction* procedures in order to use classification techniques that deal with real valued vectors as input data [3–7]. The Granular Computing (GrC) approach [8] offers a valuable framework to fill the gap between the input sequence domain and the features space \mathbb{R}^d and relies on the so-called *information granules* that play the role of *indistinguishable* features at a particular level of abstraction adopted for system description. The main objective of Granular modeling consists in finding the correct level of information granulation that best describes the input data [9].

2 Frequent Substructures Mining and Matching Problem

The problem of *sequential patterns* mining was first introduced by Agrawal and Srikant [10] in a specific context: starting from a dataset of sequences of customer transactions, the objective consists in mining for *sequential patterns* in such dataset. In a dataset of sequences of customer transactions, the general object α_i of each sequence consists of the following fields: customer-id, transaction-time and the set of items purchased in the transaction. Agrawal et al. [10] introduce for the first time the notion of *itemset* as a non-empty set of items. This problem is often viewed as the discovery of “association rules”, that is strictly dependent on the task of mining frequent itemsets. In [11], the authors propose the very first algorithm able to generate significant association rules between items in databases. Manager of supermarkets as well as e-commerce websites have to make decisions about which products to put

on sale, how to design coupons and customize the offers in order to maximize their profits. This problem raises the need to analyze past transactions and predict future behaviors.

All the studies in this field are based on the notion of *market-basket* model of data [12]. It is used to describe relationship between *items* and *baskets*, also called “transactions”. Each basket consists in an itemset and it is assumed that the number of items in a basket is much smaller than the total number of items.

The *market-basket* model (also known as a priori-like) asserts that each itemset cannot be frequent if its items are not frequent or equivalently any super-pattern of infrequent patterns cannot be frequent. Using this principle Agrawal and Srikant proposed the AprioriAll algorithm in [10]. Their approach aims to extract frequent sequential patterns and is based on a candidate generation and test paradigm. Note that during the mining procedure, candidate frequent sequential patterns can be obtained only by joining shorter frequent sequential patterns. An example of a sequential pattern is “5 % of customers bought {Apple, Orange, Flour, Coffe} in one transaction, followed by {Coffee, Sugar} in a later transaction”. The weakness of the algorithm is that a huge set of candidate sequences are generated requiring an enormous amount of memory and many repeated database scans. This behavior gets worse with increasing size of sequences in the database.

In [13] a new algorithm named GSP (Generalized Sequential Patterns) is introduced. The authors propose a breadth-first search and bottom-up method to obtain the frequent sequential pattern. Moreover, they introduce a time constraint that fixes the minimum and maximum delay between adjacent elements in the candidate patterns and the possibility for items to be present in a set of transactions in a fixed time window. GSP overcomes the performances of the APrioriAll algorithm [10] reducing the number of candidate sequential patterns. However, all a priori-like sequential pattern mining methods tend to behave badly with large datasets, because they may generate a large set of candidate subsequences. Moreover, for such algorithms, multiple scans of the database are needed, one for each length of the candidate patterns and this becomes very time consuming for mining long patterns. Finally, another problem occurs with long sequential patterns: a combinatorial number of subsequences are generated and tested.

To overcome these problems, in [14], a new algorithm named SPADE is introduced. The authors use a similar approach of GSP, however they use a vertical data format and divide the mining problem into smaller sub-problems reducing significantly the number of database scans required. In [15, 16] the authors introduce two algorithms FreeSpan and PrefixSpan. They are based on a completely different approach than APrioriAll and GSP: the pattern-growth approach for mining sequential patterns in large datasets. Each time new sequential patterns are generated, the whole dataset of sequences is projected into a set of smallest projected datasets using the extracted sequential patterns and bigger sequential patterns are grown in each projected dataset analyzing only locally frequent fragments. PrefixSpan introduces new techniques to reduce the size of the projected datasets.

All presented works describe search techniques for mining non-contiguous sequences of objects. However, these approaches are not ideal when the objective is to extract frequent sequential patterns, in which the contiguity of the component objects plays a fundamental role in the information extraction.

In particular, in computation biology, even though techniques for mining sequential *noncontiguous* patterns have many uses, they are not appropriate for many applications. Computational biology community has developed a lot of methods for detecting frequent patterns, that in this field are called *motifs*. Moreover, working with real-world data, the presence of some *noise* must be taken into account in the designing of the matching procedure [17–20]. In many fields and particularly in a biological context, patterns should have long lengths and high supports, but standard sequential pattern mining approaches tend to discover a large amount of “low quality” patterns, i.e. patterns having either short lengths or low supports. It is easy to observe that genome sequences contain errors, so it is unlikely that long subsequences generated from the same origin will be exactly identical. Moreover, the increase of the minimum number of occurrences of a subsequence in a database, in case of exact matching, obliges to accept shorter and shorter subsequences, with the possibility to obtain a massive quantity of data with a less specific meaning. In such cases, exact matches techniques can give only short and trivial patterns. So, by allowing some mismatches, it is possible to discover valuable sequential patterns, with longer length and higher approximate supports. Some works [17, 18] use Hamming distances to search for recurrent motifs in data. Other works employ suffix tree data structure [21], suffix array to store and organize the search space [22] or use a GrC framework for the extraction of frequent patterns in data [23].

The algorithm presented in [24] uses a suffix-tree data structure to mine frequent approximate contiguous subsequences (also called substrings). The procedure follows a “break-down-and-build-up” strategy. The “break-down” step aims at searching, by means of a suffix-tree based algorithm, for the longest subsequences which repeat, with an exact match, in the whole database. These subsequences represent the initial sequences (called *strands*), which will be iteratively assembled into longer strands by using a local search algorithm. The “build-up” step groups the obtained strands, forming the set from which all approximate subsequences will be identified.

The algorithm [25] uses a similar approach as [24], but taking into account the quality of sequential patterns. Good quality patterns can be obtained by balancing pattern length and pattern support. Short patterns are undesirable, particularly when sequences are long, since the meaning is less specific. Patterns with low supports are not desirable too, since they can be trivial and may not describe general phenomena. Thus, the algorithm is biased toward the search for longer subsequences, characterized by a sufficient frequency. It makes use of a suffix array to store and organize in a lexicographic order the search space (i.e., the set of subsequences). The search on such a suffix array follows a *prefix extension* approach, meaning that frequent subsequences are individuated analyzing the prefixes of the input sequences, tolerating inexactness during the evaluation.

Computational biology community has developed a lot of algorithms for mining frequent motifs using the Hamming distance as similarity measure. YMF [17] is based on the computation of the statistical significance of each motif, but its performances decrease as the complexity of motifs increases. Weeder [18] is a suffix-tree-based algorithm and is faster than YMF, because it considers only certain types of mismatches for the motifs, however it can not be used for different types of motifs. Another algorithm, MITRA [19], is a mismatch-tree-based approach and uses heuristics to prune the space of possible motifs.

Analysis and interpretation of time series is another challenging problem that many authors try to solve [26, 27]. Some works consider the problem of mining for motifs in time series databases in several applications: from the analysis of stock prices to the study of the ECG in medicine, to the analysis of measures from sensors. In particular in [28] it is showed how to discretize a time series, in order to obtain a sequence of symbols, defined over a fixed alphabet and use well known motif mining algorithms. However, in the discretization process, a lot of information is lost. Moreover, this algorithm uses exact matching procedures for mining patterns and is unusable in real cases with noisy data. In Chiu et al. [29] present another algorithm, based on [20], that considers the presence of noise in data. However, also in this case, a simple model of mismatches is considered. In [30] the algorithm FLAME is presented. It consists in a suffix-tree-based technique and can be used also with time series data sets, by converting such data into a sequence of symbols, discretizing the numeric data. All these approaches suffer from the loss of information during the discretization procedure.

In the following, we present a clustering-based subsequences mining algorithm that can be used with general sequence databases, choosing a suited similarity measure, depending on the particular application. Moreover, most methods focus only on the recurrence of patterns in data without taking into account the concept of “information redundancy”, or, in other words, the existence of overlapping among retrieved patterns [31]. Frequent pattern mining with approximate match is a challenging problem starting from the definition itself: even if one ignores small redundant patterns, there might be a huge number of large frequent redundant patterns. This problem should be taken in consideration, in a way that only some representatives of such patterns should survive after the mining process.

3 The Proposed Algorithm

In this work we present a new approximate subsequence mining algorithm called FRL-GRADIS (Filtered Reinforcement Learning-based GRanular Approach for DIScrete Sequences) [32] aiming to reduce the information redundancy of RL-GRADIS [33] by executing an optimization-based refinement process on the extracted patterns. In particular, this paper introduces the following contributions:

1. our approach finds the patterns that maximize the knowledge about the process that generates the sequences;
2. we employ a dissimilarity measure that can extract patterns despite the presence of noise and possible corruptions of the patterns themselves;
3. our method can be applied on every kind of sequence of objects, given a properly defined similarity or dissimilarity function defined in the objects domain;
4. the filtering operation produces results that can be interpreted more easily by application's field experts;
5. considering this procedure as an inner module of a more complex classification system, it allows to further reduce the dimension of the feature space, thus better addressing the *curse of dimensionality* problem.

This paper consists of three parts. In the first part we provide some useful definitions and a proper notation; in the second part we present FRL-GRADIS as a two-step procedure, consisting of a subsequences extraction step and a subsequences filtering step. Finally, in the third part, we report the results obtained by applying the algorithm to synthetic data, showing a good overall performance in most cases.

4 Problem Definition

Let $\mathcal{D} = \{\alpha_i\}$ be a *domain* of objects α_i . The objects represent the atomic units of information. A sequence S is an ordered list of n objects that can be represented by the set of pairs

$$S = \{(i \rightarrow \beta_i) \mid i = 1, \dots, n; \beta_i \in \mathcal{D}\},$$

where the integer i is the order index of the object β_i within the sequence S . S can also be expressed with the compact notation

$$S \equiv \langle \beta_1, \beta_2, \dots, \beta_n \rangle$$

A sequence database SDB is a set of sequences S_i of variable lengths n_i . For example, the DNA sequence $S = \langle G, T, C, A, A, T, G, T, C \rangle$ is defined over the domain of the four amino acids $\mathcal{D} = \{A, C, G, T\}$.

A sequence $S_1 = \langle \beta'_1, \beta'_2, \dots, \beta'_{n_1} \rangle$ is a *subsequence* of a sequence $S_2 = \langle \beta''_1, \beta''_2, \dots, \beta''_{n_2} \rangle$ if $n_1 \leq n_2$ and $S_1 \subseteq S_2$. The *position* $\pi_{S_2}(S_1)$ of the subsequence S_1 with respect to the sequence S_2 corresponds to the order index of its first element (in this case the order index of the object β'_1) within the sequence S_2 . The subsequence S_1 is also said to be *connected* if

$$\beta'_j = \beta''_{j+k} \quad \forall j = 1, \dots, n_1$$

where $k = \pi_{S_2}(S_1)$. Two subsequences S_1 and S_2 of a sequence S are *overlapping* if

$$S_1 \cap S_2 \neq \emptyset.$$

In the example described above, the complete notation for the sequence $S = \langle G, T, C, A, A, T, G, T, C \rangle$ is

$$S = \{(1 \rightarrow G), (2 \rightarrow T), (3 \rightarrow C), \dots\}$$

and a possible connected subsequence $S_1 = \langle A, T, G \rangle$ corresponds to the set

$$S_1 = \{(5 \rightarrow A), (6 \rightarrow T), (7 \rightarrow G)\}.$$

Notice that the objects of the subsequence S_1 inherit the order indices from the containing sequence S , so that they are univocally referred to their original positions in S . From now on we will focus only on connected subsequences, therefore the connection property will be implicitly assumed.

4.1 Pattern Coverage

The objective of this algorithm is to find a set of frequent subsequences of objects named as *patterns*. A pattern Ω is a subsequence of objects $\langle \omega_1, \omega_2, \dots, \omega_{|\Omega|} \rangle$, with $\omega_i \in \mathcal{D}$, that is more likely to occur within the dataset SDB. Patterns are unknown *a priori* and represent the underlying information of the dataset records. Moreover, each sequence is subject to noise whose effects include the addition, substitution and deletion of objects in a random uncorrelated fashion and this makes the recognition of recurrent subsequences more challenging.

Given a sequence $S \in SDB$ and a set of patterns $\Gamma = \{\Omega_1, \dots, \Omega_m\}$, we want to determine a quality criterion for the description of S in terms of the pattern set Γ . A connected subsequence $C \subseteq S$ is said to be **covered** by a pattern $\Omega \in \Gamma$ iff $d(C, \Omega) \leq \delta$, where $d(\cdot, \cdot)$ is a properly defined distance function and δ is a fixed tolerance (Fig. 1). The **coverage** $C_\Omega^{(\delta)}(S)$ of the pattern Ω over the sequence S is the union set of all non-overlapping connected subsequences covered by the pattern. We can write,

$$C_\Omega^{(\delta)}(S) = \bigcup_i \left[C_i \subseteq S \text{ s.t. } d(C_i, \Omega) \leq \delta \wedge C_i \cap C_j = \emptyset \forall i \neq j \right]. \quad (1)$$




Fig. 1 Coverage of the pattern Ω over the subsequence $C \subseteq S$ with tolerance δ . Black boxes and gray boxes represent respectively the covered and the uncovered objects of the sequence S . Notice that if $\delta > 0$ the sequences Ω and C need not to be of the same length

Formally, this set is still not well defined until we expand on the meaning of the property

$$C_i \cap C_j = \emptyset, \tag{2}$$

which is the requirement for the covered subsequences to be non-overlapping. Indeed, we need to include additional rules on how to deal with these overlappings when they occur. To understand better, let us recall the example of the DNA sequences presented above, where the dissimilarity measure between two sequences is the Levenshtein distance. The set of all covered subsequences C_i (in this context referred to as *candidates*) by the pattern Ω over the sequence S will consist only of sequences with values of length between $|\Omega| - \delta$ and $|\Omega| + \delta$. Indeed, these bounds correspond respectively to the extreme cases of deleting and adding δ objects to the subsequence. In case of two overlapping candidates C_i and C_j , in order to satisfy the property (2) of the coverage $C_\Omega^{(\delta)}(S)$, we have to define a rule to decide which subsequence belongs to the set $C_\Omega^{(\delta)}(S)$ and which does not. Candidates with smaller distances from the searched pattern Ω are chosen over overlapping candidates with higher distances. If the two overlapping candidates have the same distance the first starting from the left is chosen, but if also their starting position is the same the shorter one (i.e. smaller length value) has the precedence.

A coverage example in the context of the DNA sequences is shown in Fig. 2. The coverage of the pattern $\Omega = \langle A, G, G, T \rangle$ over the sequence S is $C_\Omega^{(\delta)}(S) = \langle A, C, G, T \rangle \cup \langle G, G, T \rangle \cup \langle A, C, G, G, T \rangle$.

Similarly, the compound coverage of the pattern set Γ is defined as

$$C_\Gamma^{(\delta)}(S) = \bigcup_{\Omega \in \Gamma} C_\Omega^{(\delta)}(S). \tag{3}$$

It is important to notice that, in this case, this set can include overlapping subsequences only if they belong to coverages of different patterns (i.e. it is assumed that different patterns can overlap). For example consider the case shown in Fig. 3. The coverage $C_{\{\Omega_1, \Omega_2\}}^{(\delta)}(S)$ for the patterns $\Omega_1 = \langle A, G, G, T \rangle$ and $\Omega_2 = \langle G, T, C \rangle$ is equal to $C_{\{\Omega_1, \Omega_2\}}^{(\delta)}(S) = \langle A, G, G, T, C \rangle$.




Fig. 2 Coverage examples in the case of DNA sequences. The searched pattern $\langle A, G, G, T \rangle$ is found 3 times with tolerance $\delta \leq 1$ using the Levenshtein distance. The three occurrences show all the edit operations allowed by the considered edit distance, respectively objects substitution, deletion and insertion




Fig. 3 Example of the compound coverage of multiple symbols, where the symbols $\langle G, T, C \rangle$ and $\langle A, G, G, T \rangle$ have Levenshtein distances from the corresponding subsequences equal to 0 and 1, respectively. Notice that different symbols can cover overlapping subsequences, while competing coverages of the same symbol are not allowed and only the most similar subsequence is chosen

5 The Mining Algorithm

In this section, we describe FRL-GRADIS, as a clustering-based sequence mining algorithm. It is able to discover clusters of connected subsequences of variable lengths that are frequent in a sequence dataset, using an inexact matching procedure. FRL-GRADIS consists in two main steps:

- the *symbols alphabet extraction*, which addresses the problem of finding the most frequent subsequences within a SDB. It is performed by means of the clustering algorithm RL-GRADIS [33] that identifies frequent subsequences as representatives of dense clusters of similar subsequences. These representatives are referred to as *symbols* and the pattern set as the *alphabet*. The clustering procedure relies on a properly defined edit distance between the subsequences (e.g. Levenshtein distance, DTW, etc.). However, this approach alone has the drawback of extracting many superfluous symbols which generally dilute the pattern set and deteriorate the *interpretability* of the produced pattern set.
- the *alphabet filtering* step deals with the problem stated above. The objective is to filter out all the spurious or redundant symbols contained in the alphabet produced by the symbols extraction step. To accomplish this goal we employ a heuristic approach based on evolutionary optimization over a validation SDB.

One of the distinctive features of this algorithm is its generality with respect to the kind of data contained in the input sequence database (e.g., sequences of real numbers or characters as well as sequences of complex data structures). Indeed, both steps outlined above take advantage of a dissimilarity-based approach, with the dissimilarity function being a whatever complex measure between two ordered sequences, not necessarily metric.

In the following, we first describe the main aspects of the symbols alphabet extraction procedure, then we present the new filtering method. For more details on the symbols alphabet construction we refer the reader to [33].

5.1 Frequent Subsequences Identification

Consider the input training dataset of sequences $\mathcal{T} = \{S_1, S_2, \dots, S_{|\mathcal{T}|}\}$ and a properly defined dissimilarity measure $d : \mathcal{T} \times \mathcal{T} \rightarrow \mathbb{R}$ between two objects of the training dataset (e.g., Levenshtein distance for strings of characters). The goal of the subsequences extraction step is the identification of a finite set of symbols $\mathcal{A}_e = \{\Omega_1, \Omega_2, \dots, \Omega_{|\mathcal{A}_e|}\}$,¹ computed using the distance $d(\cdot, \cdot)$ in a free clustering procedure. The algorithm we chose to accomplish this task is RL-GRADIS which is based on the well-known Basic Sequential Algorithmic Scheme (BSAS) clustering algorithm [33]. Symbols are found by analysing a suited set of variable-length subsequences of \mathcal{T} , also called n -grams, that are generated by expanding each input sequence $S \in \mathcal{T}$. The expansion is done by listing all n -grams with lengths varying between the values l_{\min} and l_{\max} . The parameters l_{\min} and l_{\max} are user-defined and are respectively the minimum and maximum admissible length for the mined patterns. The extracted n -grams are then collected into the SDB \mathcal{N} . At this point, the clustering procedure is executed on \mathcal{N} . For each cluster we compute its representative, defined by the Minimum Sum of Distances (MinSOD) technique [33, 34], as the element having the minimum total distance from the other elements of the cluster. This technique allows to represent the corresponding clusters by means of their most characteristic elements.

The quality of each cluster is measured by its *firing strength* f , where $f \in [0, 1]$. Firing strengths are used to track the dynamics describing the *updating rate* of the clusters when the input stream of subsequences \mathcal{N} is analyzed. A reinforcement learning procedure is used to dynamically update the list of candidate symbols based on their firing strength. Clusters with a low rate of update (low firing strength) are discarded in an on-line fashion, along with the processing of the input data stream \mathcal{N} . RL-GRADIS maintains a dynamic list of candidate symbols, named *receptors*, which are the representatives of the active clusters. Each receptor's firing strength (i.e. the firing strength of its corresponding cluster) is dynamically updated by means of two additional parameters, $\alpha, \beta \in [0, 1]$. The α parameter is used as a *reinforcement weight* factor each time a cluster \mathcal{R} is updated, i.e., each time a new input subsequence is added to \mathcal{R} . The firing strength update rule is defined as follows:

$$f(\mathcal{R}) \leftarrow f(\mathcal{R}) + \alpha(1 - f(\mathcal{R})). \quad (4)$$

The β parameter, instead, is used to model the speed of *forgetfulness* of receptors according to the following formula:

$$f(\mathcal{R}) \leftarrow (1 - \beta)f(\mathcal{R}). \quad (5)$$

¹The subscript “e” stands for “extraction” as in extraction step.

The firing strength updating rules shown in Eqs. (4) and (5) are performed for each currently identified receptor, after the analysis of each input subsequence. Therefore, receptors/clusters that are not updated frequently during the analysis of \mathcal{N} will likely have a low strength value and this will cause the system to remove the receptor from the list.

5.2 Subsequences Filtering

As introduced above, the output alphabet \mathcal{A}_e of the clustering procedure is generally redundant and includes many spurious symbols that make the recognition of the true alphabet quite difficult.

To deal with this problem, an optimization step is performed to reduce the alphabet size, aiming at retaining only the most significant symbols, i.e. only those that best resemble the original, unknown ones. Since this procedure works like a filter, we call the output of this optimization the *filtered alphabet* \mathcal{A}_f and, clearly, $\mathcal{A}_f \subset \mathcal{A}_e$ holds. Nevertheless, it is important for the filtered alphabet's size not to be smaller than the size of the true alphabet, since in this case useful information will be lost. Let $\Gamma \subset \mathcal{A}_e$ be a candidate subset of symbols of the alphabet \mathcal{A}_e and $S \in \mathcal{V}$ a sequence of a *validation SDB* \mathcal{V} . We assume the *descriptive power* of the symbols set Γ , with respect to the sequence S , to be proportional to the quantity $|C_\Gamma^{(\delta)}(S)|$ (cfr Eq. 3), i.e. the number of objects $\beta_i \in S$ covered by the symbols set Γ . In fact, intuitively, a lower number of uncovered objects in the whole SDB by Γ symbols can be considered as a clue that Γ itself will likely contain the true alphabet. The normalized number of uncovered objects in a sequence S by a pattern set Γ corresponds to the quantity

$$P = \frac{|S| - |C_\Gamma^{(\delta)}(S)|}{|S|}, \quad (6)$$

where the operator $|\cdot|$ stands for the cardinality of the set. The term P assumes the value 0 when the sequence S is completely covered by the pattern set Γ and the value 1 when none of the symbols in Γ are present in the sequence S . Notice that $C_\Gamma^{(\delta)}(S)$ depends on the parameter δ which represents the tolerance of the system towards the corruption of symbols' occurrences caused by noise.

On the other hand, a bigger pattern set is more likely to contain spurious patterns which tend to hinder the interpretability of the obtained results, so smaller set sizes are to be preferred. This property can be described with the normalized alphabet size

$$Q = \frac{|\Gamma|}{|\mathcal{A}_e|}, \quad (7)$$

where \mathcal{A}_e is the alphabet of symbols extracted by the clustering procedure described in the last section. Clearly, the cardinality of \mathcal{A}_e represents an upper bound for the

size of the filtered alphabet, so the term Q ranges from 0 to 1. The terms P and Q generally show opposite trends, since a bigger set of symbols is more likely to cover a bigger portion of the sequence and vice versa.

Finding a tradeoff between these two quantities corresponds to minimizing the convex objective function

$$G_S^{(\delta)}(\Gamma) = \lambda Q + (1 - \lambda) P \quad (8)$$

where $\lambda \in [0, 1]$ is a meta-parameter that weighs the relative importance between the two contributions. It is easy to verify that

$$0 \leq G_S^{(\delta)}(\Gamma) \leq 1. \quad (9)$$

More generally, for a validation SDB \mathcal{V} , the global objective function is the mean value of $G_S^{(\delta)}(\Gamma)$ over all sequences $S_i \in \mathcal{V}$, hence

$$G_{\mathcal{V}}^{(\delta)}(\Gamma) = \frac{\sum_{1 \leq i \leq |\mathcal{V}|} G_{S_i}^{(\delta)}(\Gamma)}{|\mathcal{V}|} \quad (10)$$

and the best symbols set after the optimization procedure is

$$\mathcal{A}_f = \operatorname{argmin}_{\Gamma \subset \mathcal{A}_e} G_S^{(\delta)}(\Gamma). \quad (11)$$

To solve the optimization problem described by Eq. (11) we employ a standard genetic algorithm, where each individual of the population is a subset Γ of the extracted alphabet $\mathcal{A}_e = \{\Omega_1, \dots, \Omega_{|\mathcal{A}_e|}\}$. The genetic code of the individual is encoded as a binary sequence E of length $|\mathcal{A}_e|$ of the form

$$E_{\Gamma} = \langle e_1, e_2, \dots, e_{|\mathcal{A}_e|} \rangle \quad (12)$$

with

$$e_i = \begin{cases} 1 & \text{iff } \Omega_i \in \Gamma \\ 0 & \text{otherwise} \end{cases}.$$

It is important not to mistake genetic codes with the SDB sequences described earlier, even if they are both formally defined as ordered sequences.

Given a validation dataset \mathcal{V} and a fixed tolerance δ , the fitness value $F(E_{\Gamma})$ of each individual E_{Γ} is computed as the following affine transformation of the objective function introduced in the last paragraph

$$F(E_{\Gamma}) = 1 - G_{\mathcal{V}}^{(\delta)}(\Gamma) \quad (13)$$

The computation is then performed with standard crossover and mutation operators between the binary sequences and the stop condition is met when the maximum fitness does not change for a fixed number N_{stall} of generations or after a given maximum number N_{max} of iterations. When the evolution stops, the filtered alphabet $\mathcal{A}_f = \tilde{\Gamma}$ is returned, where $\tilde{\Gamma}$ is the symbols subset corresponding to the fittest individual $E_{\tilde{f}}$.

6 Tests and Results

In this section, we present results from different experiments that we designed to test the effectiveness and performance of FRL-GRADIS in facing problems with varying complexity.

6.1 Data Generation

We tested the capabilities of FRL-GRADIS on synthetic sequence databases composed of textual strings. For this reason, the domain of the problem is the English alphabet

$$\mathcal{D} = \{A, B, C, \dots, Z\}.$$

Modeled noise consists in all cases of random characters insertions, deletions and substitutions to the original string. For this reason a natural choice of dissimilarity measure between sequences is the Levenshtein distance, that measures the minimum number of edit steps necessary to transform one string of characters into another. We conducted two different classes of tests, which accounted for two kinds of noise, respectively *symbols noise* and *channel noise*, presented in the following paragraphs.

6.1.1 Symbols Noise

This kind of noise simulates those situations in which symbols are altered during the composition of the sequence. In fact, each instance of a symbol being added to the data sequence has a fixed probability of being mutated by one addition, deletion or modification of its objects. Moreover, a variable number of uncorrelated objects are added between contiguous instances of symbols in the sequence, to simulate the presence of irrelevant data separating actual symbols. The detailed process of data generation is described below:

1. the true symbols alphabet \mathcal{A}_1 is generated. This alphabet consists of N_{sym} symbols with lengths normally distributed around the mean value L_{sym} . Each character is chosen in \mathcal{D} with uniform probability and repeated characters are allowed;

2. a training SDB \mathcal{T} and a validation SDB \mathcal{V} respectively composed of N_{tr} and N_{val} sequences are generated. Each of these sequences is built by concatenating N_{symseq} symbols chosen randomly from \mathcal{A}_t . Notice that generally $N_{\text{symseq}} > N_{\text{sym}}$ so there will be repeated symbols;
3. in each sequence, every symbol will be subject to noise with probability μ . The application of noise to a symbol in a sequence corresponds to the deletion, substitution or insertion of one character to that single instance of the symbol. This kind of noise is referred to as *intra-pattern* noise;
4. a user-defined quantity of random characters is added between instances of symbols in each sequence. This noise is called *inter-pattern* noise. Such quantity depends on the parameter η that corresponds to the ratio between the number of characters belonging to actual symbols and the total number of character of the sequence after the application of inter-pattern noise, that is,

$$\eta = \frac{(\# \text{ symbol characters})}{(\# \text{ total characters})}.$$

Notice that the amount of inter-pattern noise is inversely proportional to the value of η .

6.1.2 Channel Noise

In this case we simulate a noise affecting an hypothetical channel through which the sequence is transmitted. This kind of noise alters the objects in a uncorrelated manner, without keeping track of the separation between symbols.

In this case we generate the original SDB \mathcal{T} and \mathcal{V} in the same manner as described in steps 1 and 2 of Sect. 6.1.1. The noise is then added to these datasets by iterating through the objects of the sequence and altering each object with a fixed probability p . The alteration consists with equal probability in either:

- the substitution of the object with another randomly chosen object;
- the deletion of the object;
- the addition of another randomly chosen object to the right of the current object position in the sequence.

The generated datasets \mathcal{T} and \mathcal{V} are then ready to be used as input of the FRL-GRADIS procedure. Notice that the true alphabet \mathcal{A}_t is unknown in real-world applications and here is used only to quantify the performance of the algorithm.

6.2 Quality Measures

We now introduce the quality measures used in the following tests to evaluate the mining capabilities of the FRL-GRADIS algorithm. These measures are computed

for the resulting alphabets obtained from both the extraction and the filtering steps presented in Sect. 5, in order to highlight the improvement made by the filtering procedure (i.e. the improvement of FRL-GRADIS over RL-GRADIS).

The *redundance* R corresponds to the ratio between the cardinality of the alphabet \mathcal{A} and the true alphabet \mathcal{A}_t , that is,

$$R = \frac{|\mathcal{A}|}{|\mathcal{A}_t|} \quad (14)$$

Clearly, since the filtering step selects a subset \mathcal{A}_f (filtered alphabet) of the extracted alphabet \mathcal{A}_e , we always have that

$$R_f < R_e.$$

The redundancy measures the amount of unnecessary symbols that are found by a frequent pattern mining procedure and it ranges from zero to infinite. When $R > 1$ some redundant symbols have been erroneously included in the alphabet, while when $R < 1$ some have been missed, the ideal value being $R = 1$.

It is important to notice that the redundancy depends only on the number of symbols reconstructed, but not on their similarity with respect to the original alphabet. For this purpose we also introduce the *mining error* E , defined as the mean distance between each symbol Ω_i of the true alphabet \mathcal{A}_t and its best match within the alphabet \mathcal{A} , where the best match means the symbol with the least distance from Ω_i . In other words, considering $\mathcal{A}_t = \{\Omega_1, \dots, \Omega_{|\mathcal{A}_t|}\}$ and $\mathcal{A} = \{\tilde{\Omega}_1, \dots, \tilde{\Omega}_{|\mathcal{A}|}\}$, the mining error corresponds to

$$E = \frac{\sum_i d(\Omega_i, \tilde{\Omega}_{(i)})}{|\mathcal{A}_t|} \quad (15)$$

where

$$\tilde{\Omega}_{(i)} = \operatorname{argmin}_{\tilde{\Omega} \in \mathcal{A}} d(\Omega_i, \tilde{\Omega}).$$

This quantity has the opposite role of the redundancy, in fact it keeps track of the general accuracy of reconstruction of the true symbols regardless of the generated alphabet size. It assumes non-negative values and the ideal value is 0. For the same reasons stated above the inequality

$$E_f \geq E_e$$

holds, so the extraction procedure's mining error constitutes a lower bound for the mining error obtainable with the filtering step.

6.3 Results

We executed the algorithm multiple times for different values of the noise parameters, to assess the different response of FRL-GRADIS to increasing amounts of noise. Most parameters have been held fixed for all the tests and they are listed in Table 1.

As a first result, we present the synthetic tests performed by adding varying quantities of symbols noise to the data sequences, performed with $\mu = 0.5$ and variable amounts of inter-pattern noise η . It means that about half of the symbols in a sequence are subject to the alteration of one character and increasing amounts of random characters are added between symbols in each sequence. The results obtained with this configuration are shown in Figs. 4 and 5.

The redundancy plot in Fig. 4 shows an apparently paradoxical trend of the extraction procedure's redundancy: with decreasing amounts of inter-pattern noise (i.e. increasing values of η) the extraction algorithm performs more poorly, leading to higher redundancies. That can be easily explainable by recalling how the clustering procedure works.

Higher amounts of inter-pattern noise mean that the frequent symbols are more likely to be separated by random strings of characters. These strings of uncorrelated characters generate very sparse clusters with negligible cardinality that are very likely to be deleted during the clustering's reinforcement step. Clusters corresponding to actual symbols, instead, are more active and compact, their bounds being clearly defined by the noise characters, and so they are more likely to survive the reinforcement step.

In case of negligible (or non-existent) inter-pattern noise, instead, different symbols are more likely to occur in frequent successions that cause the generation of many clusters corresponding to spurious symbols, obtained from the concatenation

Table 1 Fixed parameters adopted for the tests

Parameter	Value	Parameter	Value
N_{tr}	50	N_{val}	25
N_{sym}	5	N_{symseq}	10
l_{min}	4	l_{max}	12
δ	1	λ	0.5
N_{pop}	100	N_{elite}	0.1
p_{cross}	0.8	p_{mut}	0.3
N_{max}	100	N_{stall}	50

The parameter δ corresponds to the tolerance of the Levenshtein distance considered when calculating the coverage as in Eq. (1) while λ weighs the two terms of the objective function of Eq. (8). The values shown in the second part of the table refer to the genetic algorithm's parameters. N_{pop} corresponds to the population size, N_{elite} is the fraction of individuals who are guaranteed to survive and be copied to the new population in each iteration, p_{cross} and p_{mut} are respectively the crossover and mutation probabilities. The evolution terminates if N_{evol} iterations have been performed or if for a number N_{stall} of iterations the maximum fitness has not changed

Fig. 4 Plot of the redundance R of the extraction (RL-GRADIS) and filtering (FRL-GRADIS) steps with variable inter-pattern noise and $\mu = 0.5$

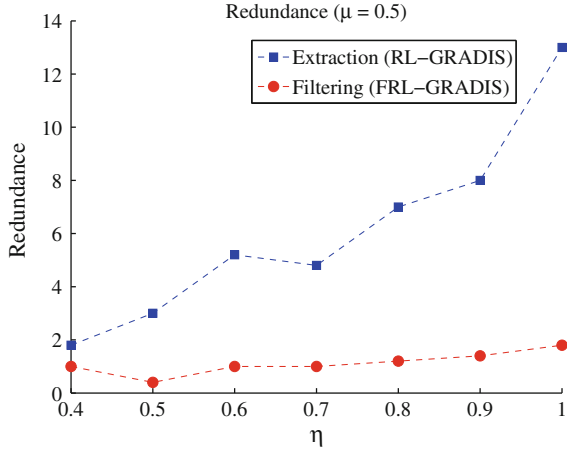



Fig. 5 Plot of the mining error E of the extraction (RL-GRADIS) and filtering (FRL-GRADIS) steps with variable inter-pattern noise and $\mu = 0.5$



of parts of different symbols. The filtering procedure overcomes this inconvenience, as it can be seen from Fig. 4 that it is nearly not affected by the amount of inter-pattern noise. As it is evident, the filtering procedure becomes fundamental for higher values of the parameter η , where the clustering produces highly redundant alphabets that would be infeasible to handle in a real-world application. Figure 5 shows that the mining error after the filtering procedure remains mostly the same for all values of η , which means that the system is robust to the moderate alteration of the input signal.

In the second pool of tests we show the response of the system to increasing quantities of channel noise p . In Fig. 6 and 7 are shown the redundance and the mining error measured for different values of p . While FRL-GRADIS shows slightly higher mining error levels than RL-GRADIS, its redundancy is still significantly lower and,




Fig. 6 Plot of the redundance R of the extraction (RL-GRADIS) and filtering (FRL-GRADIS) steps with variable channel noise. Error bars represent the standard deviation over 3 runs of the algorithm with the same parameters

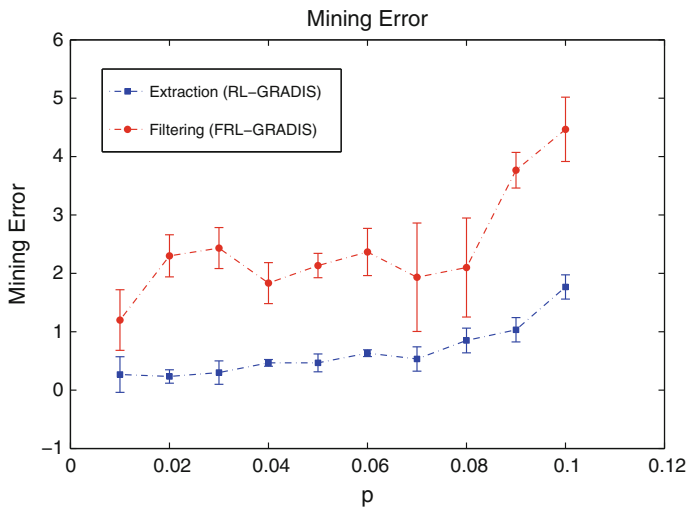


Fig. 7 Plot of the mining error E of the extraction (RL-GRADIS) and filtering (FRL-GRADIS) steps with variable channel noise. Error bars represent the standard deviation over 3 runs of the algorithm with the same parameters

to a large extent, insensitive to the quantity of noise in the system. Clearly, lower mining error levels are obtainable by setting suitable values of the parameter λ (at the expense of resulting redundancy) or stricter convergence criteria of the genetic algorithm (at the expense of convergence time).

In general, we can conclude that the system allows for a remarkable synthesis of the extracted alphabet despite of a modest additional mining error.

7 Conclusions

In this work we have presented a new approach to sequence data mining, focused on improving the interpretability of the frequent patterns found in the data. For this reason, we employed a two-steps procedure composed of a clustering algorithm, that extracts the frequent subsequences in a sequence database, and a genetic algorithm that filters the returned set to retrieve a smaller set of patterns that best describes the input data. For this purpose we introduced the concept of coverage, that helps in recognizing the true presence of a pattern within a sequence affected by noise. The experiments were performed on two cases of synthetic data affected by two different sources of noise. The results have shown a good overall performance and lay the foundations for improvements and further experiments on real data.

References

1. Possemato, F., Rizzi, A.: Automatic text categorization by a granular computing approach: facing unbalanced data sets. In: The 2013 International Joint Conference on Neural Networks (IJCNN), pp. 1–8 (2013)
2. Modugno, V., Possemato, F., Rizzi, A.: Combining piecewise linear regression and a granular computing framework for financial time series classification (2014)
3. Bianchi, F., Livi, L., Rizzi, A., Sadeghian, A.: A granular computing approach to the design of optimized graph classification systems. *Soft Comput.* **18**, 393–412 (2014)
4. Bianchi, F.M., Scardapane, S., Livi, L., Uncini, A., Rizzi, A.: An interpretable graph-based image classifier. In: 2014 International Joint Conference on Neural Networks (IJCNN), pp. 2339–2346. IEEE (2014)
5. Rizzi, A., Del Vescovo, G.: Automatic image classification by a granular computing approach. In: Proceedings of the 2006 16th IEEE Signal Processing Society Workshop on Machine Learning for Signal Processing, pp. 33–38 (2006)
6. Del Vescovo, G., Rizzi, A.: Automatic classification of graphs by symbolic histograms. In: IEEE International Conference on Granular Computing. GRC 2007, pp. 410–410 (2007)
7. Del Vescovo, G., Rizzi, A.: Online handwriting recognition by the symbolic histograms approach. In: IEEE International Conference on Granular Computing. GRC 2007, pp. 686–686 (2007)
8. Bargiela, A., Pedrycz, W.: *Granular Computing: An Introduction*. Springer (2003)
9. Livi, L., Rizzi, A., Sadeghian, A.: Granular modeling and computing approaches for intelligent analysis of non-geometric data. *Appl. Soft Comput.* **27**, 567–574 (2015)
10. Agrawal, R., Srikant, R.: Mining sequential patterns. In: Proceedings of the Eleventh International Conference on Data Engineering, pp. 3–14. IEEE (1995)
11. Agrawal, R., Imieliński, T., Swami, A.: Mining association rules between sets of items in large databases. In: ACM SIGMOD Record, vol. 22, pp. 207–216. ACM (1993)
12. Rajaraman, A., Ullman, J.D.: *Mining of Massive Datasets*. Cambridge University Press, New York (2011)
13. Srikant, R., Agrawal, R.: *Mining Sequential Patterns: Generalizations and Performance Improvements*. Springer (1996)
14. Zaki, M.J.: Spade: an efficient algorithm for mining frequent sequences. *Mach. Learn.* **42**, 31–60 (2001)
15. Han, J., Pei, J., Mortazavi-Asl, B., Chen, Q., Dayal, U., Hsu, M.C.: Freespan: frequent pattern-projected sequential pattern mining. In: Proceedings of the Sixth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, pp. 355–359 (2000)

16. Pei, J., Han, J., Mortazavi-Asl, B., Pinto, H., Chen, Q., Dayal, U., Hsu, M.C.: Prefixspan: mining sequential patterns efficiently by prefix-projected pattern growth. In: 2013 IEEE 29th International Conference on Data Engineering (ICDE), IEEE Computer Society, pp. 0215–0215 (2001)
17. Sinha, S., Tompa, M.: YMF: a program for discovery of novel transcription factor binding sites by statistical overrepresentation. *Nucleic Acids Res.* **31**, 3586–3588 (2003)
18. Pavesi, G., Mereghetti, P., Mauri, G., Pesole, G.: Weeder web: discovery of transcription factor binding sites in a set of sequences from co-regulated genes. *Nucleic Acids Res.* **32**, W199–W203 (2004)
19. Eskin, E., Pevzner, P.A.: Finding composite regulatory patterns in dna sequences. *Bioinformatics* **18**, S354–S363 (2002)
20. Buhler, J., Tompa, M.: Finding motifs using random projections. *J. Comput. Biol.* **9**, 225–242 (2002)
21. Zhu, F., Yan, X., Han, J., Yu, P.S.: Efficient discovery of frequent approximate sequential patterns. In: Seventh IEEE International Conference on Data Mining. ICDM 2007, pp. 751–756. IEEE (2007)
22. Ji, X., Bailey, J.: An efficient technique for mining approximately frequent substring patterns. In: Seventh IEEE International Conference on Data Mining Workshops. ICDM Workshops 2007, pp. 325–330. IEEE (2007)
23. Rizzi, A., Possemato, F., Livi, L., Sebastiani, A., Giuliani, A., Mascioli, F.M.F.: A dissimilarity-based classifier for generalized sequences by a granular computing approach. In: IJCNN, IEEE, pp. 1–8 (2013)
24. Zhu, F., Yan, X., Han, J., Yu, P.S.: Efficient discovery of frequent approximate sequential patterns. In: Proceedings of the 2007 Seventh IEEE International Conference on Data Mining, Washington, DC, USA, IEEE Computer Society, pp. 751–756 (2007)
25. Ji, X., Bailey, J.: An efficient technique for mining approximately frequent substring patterns. In: Proceedings of the Seventh IEEE International Conference on Data Mining Workshops. ICDMW '07, Washington, DC, USA, IEEE Computer Society, pp. 325–330 (2007)
26. Fu, A.W.C., Keogh, E., Lau, L.Y., Ratanamahatana, C.A., Wong, R.C.W.: Scaling and time warping in time series querying. *VLDB J. Int. J. Very Large Data Bases* **17**, 899–921 (2008)
27. Vlachos, M., Kollios, G., Gunopulos, D.: Discovering similar multidimensional trajectories. In: 18th International Conference on Data Engineering. Proceedings. IEEE, pp. 673–684 (2002)
28. Patel, P., Keogh, E., Lin, J., Lonardi, S.: Mining motifs in massive time series databases. In: 2002 IEEE International Conference on Data Mining. ICDM 2003. Proceedings. IEEE, pp. 370–377 (2002)
29. Chiu, B., Keogh, E., Lonardi, S.: Probabilistic discovery of time series motifs. In: Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, pp. 493–498 (2003)
30. Floratou, A., Tata, S., Patel, J.M.: Efficient and accurate discovery of patterns in sequence data sets. *IEEE Trans. Knowl. Data Eng.* **23**, 1154–1168 (2011)
31. Matsui, T., Uno, T., Umemori, J., Koide, T.: A new approach to string pattern mining with approximate match. In: *Discovery Science*, pp. 110–125. Springer (2013)
32. Maiorino, E., Possemato, F., Modugno, V., Rizzi, A.: Information granules filtering for inexact sequential pattern mining by evolutionary computation (2014)
33. Rizzi, A., Del Vescovo, G., Livi, L., Frattale Mascioli, F.M.: A new granular computing approach for sequences representation and classification. In: Proceedings of the 2012 International Joint Conference on Neural Networks, pp. 2268–2275 (2012)
34. Del Vescovo, G., Livi, L., Frattale Mascioli, M., Rizzi, A.: On the problem of modeling structured data with the minsod representative. *Int. J. Comput. Theory Eng.* **6**, 9–14 (2014)