Temporal Association Rule Mining: An Overview Considering the Time Variable as an Integral or Implied Component

Alberto Segura-Delgado[*], María José Gacto[†], Rafael Alcalá[‡], Jesús Alcalá-Fdez[§]

Article Type:

Overview

Abstract

Association rules are commonly used to provide decision-makers with knowledge that helps them to make good decisions. Most of the published proposals mine association rules without paying particular attention to temporal information. However, in real-life applications data usually change over time or presenting different temporal situations. Therefore, the extracted knowledge may not be useful, since we may not know whether the rules are currently applicable or whether they will be applicable in the future. For this reason, in recent years, many methods have been proposed in the literature for mining temporal association rules, which introduce a greater predictive and descriptive power providing an additional degree of interestingness. One of the main problems in this research field is the lack of visibility most works suffer since there is no standard terminology to refer to it, making it difficult to find and compare proposals and studies in the field. This contribution attempts to offer a well-defined framework that allows researchers both to easily locate the previous proposals and to propose well-grounded methods in the future. To accomplish both objectives, a two-level taxonomy is proposed according to whether the time variable is considered to provide order to the data collection and to locate some temporal constraints, or whether it is considered as an attribute within the learning process. Some recent applications, available software tools, and a

[*]Department of Computer Science and Artificial Intelligence, University of Granada, E-18071, Granada, Spain (e-mail: alberto.segura.delgado@gmail.com)

[†]Department of Computer Science, University of Jaén, 23071 Jaén, Spain (e-mail: mgacto@ujaen.es)

[‡]Department of Computer Science and Artificial Intelligence, University of Granada, E-18071, Granada, Spain (e-mail: alcala@decsai.ugr.es)

[§]Department of Computer Science and Artificial Intelligence, University of Granada, E-18071, Granada, Spain (e-mail: jalcala@decsai.ugr.es)

bibliographical analysis in accordance with the Web of Science are also shown. Finally, some critical considerations and potential further directions are discussed.

INTRODUCTION

In recent years, the amount of data collected from different application areas has led to a situation in which the extraction of interesting knowledge from large datasets is a very attractive and challenging task. Association discovery is one of the most common Data Mining (DM) techniques used to extract interesting knowledge from large datasets. The association rules are represented as $X \rightarrow Y$, where X and Y are sets of items (symbols) that satisfy $X \cap Y = \oslash$ (Agrawal, Imielinski, & Swami, 1993). These rules allow us to generate suitable descriptive or predictive models from collected data, which not only explain it better, but also enable us to predict new data (Han & Kamber, 2006; Zhang & Zhang, 2002). They have been successfully applied in a wide variety of areas, such as biomedicine (Buxton, Vohra, Guo, Fogleman, & Patel, 2019), traffic safety (A. Das, Ahmed, & Ghasemzadeh, 2019), software engineering (Miholca, Czibula, & Czibula, 2018), energy (Q. Xiao, Li, Tang, Li, & Li, 2019), demographic research (Degirmenci & ozbakir, 2018), and so on.

Many of the published methods mine association rules without paying particular attention to temporal information, considering that the extracted knowledge does not change over time, thus representing the same knowledge for all temporal situations. However, in real-life applications data usually change over time or represent different temporal situations, in such a way that this temporal information is usually included in the data collected. If such temporal components are not properly taken into account, the knowledge extracted may not be useful since we will not know whether the association rules are applicable at the present time or whether they will be applicable in the future (W. Lin, Orgun, & Williams, 2002; Roddick & Spiliopoulou, 2002). Moreover, interesting knowledge may be overlooked due to its association with a particular time period or with special events such as sports games (Saleh & Masseglia, 2011) among others. Designing DM algorithms capable of handling temporal information is therefore challenge for researchers.

Over the last few years, many methods have been proposed in the literature for mining

Temporal Association Rules (TARs) from datasets with temporal information, providing a greater predictive and descriptive power and an additional degree of interestingness (Ao, Luo, Wang, Zhuang, & He, 2018; Cariñena, 2014; Hong, Lan, Su, Wu, & Wang, 2016; Mannila, Toivonen, & Verkamo, 1997; Teng, Chen, & Lu, 1990; Y. Li, Ning, Wang, & Jajodia, 2003; Papapetrou, Kollios, Sclaroff, & Gunopulos, 2009). In addition to dependencies among items of the same transaction (intratransaction), these rules also represent dependencies among items of different transactions (intertransaction) that refer to different time moments in the antecedent and the consequent (Deogun & Jiang, 2005; Lu, Feng, & Han, 2000; Tung, Lu, Han, & Feng, 2003). Rules such as "after receiving a radiation treatment for 7 days, cancer patients suffer from both nausea and magnesium deficiency" are generally mined from sequential data in which the events are linearly ordered over time. Classical algorithms are usually extended to mine TARs establishing a time window in which they find the frequent itemsets (sequences, episodes, etc) that will be used to generate the rules (Mannila et al., 1997). In other cases, the time variable is also considered in order to represent some type of temporal constraint, such as temporal distance constraints among events, etc (Höppner, 2001; Harms & Deogun, 2004). Other approaches introduce the time variable in the learning process to analyze temporal areas in which rules occur. For instance, these approaches can induce rules representing situations that appear periodically over user-defined time intervals (Ozden, Ramaswamy, & Silberschatz, 1998; Y. Li et al., 2003).

While this research area has received significant attention in recent years, the field suffers from a lack of standard terminology, making it difficult to find and compare previous proposals and studies in the field. As an example, the following terms have been used in the literature to refer to this field: "sequential rules" (Teng et al., 1990), "episode rules" (Mannila et al., 1997), "inter-transaction association rules" (Lu et al., 2000), "prediction rules" (Deogun & Jiang, 2005), "cyclic association rules" (Ozden et al., 1998), "calendar association rules" (Ramaswamy, Mahajan, & Silberschatz, 1998), "periodical association rules" (D. Li & Deogun, 2005), "interval-based temporal association rules" (Höppner & Klawonn, 2001), and so on. Notice that some of these terms (such as "sequential rules" or "episode rules") could be even more general than TARs since the data can be sequentially ordered by different criteria other than time.

For all of the above reasons, this contribution aims to provide an overview of the current state of the TAR field by collecting, organizing and summarizing the previous proposals in order to help students and researchers who may be less familiar with the literature to easily locate the previous proposals and to propose well-grounded methods in the future. To that end, this contribution offers a well-defined framework proposing a new and easy to understand two-level taxonomy. Firstly, it establishes whether the time variable is considered to provide order to the data collection and to locate some temporal constraints, or whether it is considered as an attribute within the learning process. Secondly, it proposes different categories (identifying different use cases) based on the temporal information type or on how this temporal information is included in the learning process. It facilitates understanding of the state-of-the-art for researchers without a strong background in the field and enables them to search for articles related to the temporal type problems that they are trying to solve. Some of the most relevant proposals associated with each category are introduced as well as some contributions regarding previously recorded real-world applications, so that researchers facing problems with data collections that include temporal information might easily understand the previous proposals which contain potential solutions or inspiration relating to their specific problem.

In addition, this contribution includes a specific section comprising a list of recent real-world applications in which TARs have been successfully applied, highlighting the different application domains and linking them to the use cases in the proposed taxonomy. It also identifies some free and open source software (from several repositories: CRAN, Github, Pypi) and open-source datasets that support the use of TARs to address real-world applications. Notice that open source resources ease the application and adaptation of TARs to new problems and are essential to spreading TARs to industry and to new subjects.

A bibliographical analysis of the research field is also shown in accordance with the Web of Science (WoS) [1]. We also draw visual science maps (Moya-Anegón et al., 2004) based on the free software Science Mapping Analysis Tool (SciMAT) (M. J. Cobo, López-Herrera, Herrera-Viedma, & Herrera, 2012) and The Open Graph Viz Platform (Gephi). Finally, some critical considerations and potential research directions are discussed, in order to provide

[1] https://www.webofknowledge.com/

some recommendations that should be taken into account for further publications and some possible further directions, such as the need for new metrics and/or application of temporal information to IoT applications among others, in a similar way to that performed in temporal pattern mining (Radhakrishna, Kumar, & Janaki, 2015; Aljawarneh, Radhakrishna, Kumar, & Janaki, 2017; Radhakrishna, Aljawarneh, Kumar, & Janaki, 2019).

The paper is organized as follows. The second section presents some fundamental concepts regarding association rules and TARs. The third section presents the proposed taxonomy to classify the methods for mining TARs. The fourth and fifth sections introduce each of the categories of the taxonomy, presenting some of the most relevant proposals for each one. The sixth section focuses on available software tools and application papers of TARs. The seventh section shows a bibliographical analysis of the published proposals on TARs, revealing the main up-to-date research tendencies. The eighth section takes into account some critical reflexions and potential research directions. Finally, the ninth section makes some concluding remarks.

TEMPORAL ASSOCIATION RULES

Association rules enable the identification of dependences or correlations between the elements or values (items) of a dataset, which can be of different types (discrete, binary, quantitative, etc). These rules are defined as an expression of the type $X \rightarrow Y$ (with $X \cap Y = \oslash$), where X and Y are sets of items (itemsets) (Agrawal et al., 1993). For instance, the rule {high prices of airline 1}→{low prices of airline 2} could be extracted from travel agency data, indicating that the second airline offers low prices when the first airline raises the prices. They are commonly evaluated making use of the classical measures *support* and *confidence*. The support of an itemset I (which will be called $Sup(I)$) is defined as the frequency with which I appears in a dataset. Based on this definition, the measures support and confidence of the rule $X \rightarrow Y$ are defined as follows:

$$Support(X \rightarrow Y) = \frac{Sup(XY)}{\mid N \mid}, \quad Confidence(X \rightarrow Y) = \frac{Sup(XY)}{Sup(X)}, \tag{1}$$

where $\mid N \mid$ is the number of examples/transactions in the dataset, and $Sup(XY)$ and $Sup(Y)$

are the support of the itemsets XY and X, respectively. Many methods based on the classic support-confidence framework have been proposed to mine association rules. The user defines a minimum support (*minSup*) and a minimum confidence (*minConf*) to firstly extract all the itemsets with support greater than or equal to minSup (frequent itemsets) (Fournier-Viger, Lin, Vo, et al., 2017), and then, all the rules with confidence greater than or equal to minConf are obtained from them. However, it has been pointed out that they present several problems (Berzal, Blanco, Sánchez, & Vila, 2002). On the one hand, the confidence is unable to detect statistical independence or negative dependence between the antecedent and the consequent since it does not take into account the consequent support. On the other hand, itemsets with very high support can result in unhelpful rules since they are present in many examples of the datasets and any other itemset of the dataset (despite its meaning) may seem to be a good predictor of their presence. Therefore, several authors have proposed other measures to overcome these problems and to select and rank the rules on the basis of their potential interest to the user (Geng & Hamilton, 2006; Tan, Kumar, & Srivastava, 2002).

As we explained in the introduction, there is a need to explicitly consider the temporal information when it is available. For this reason, in recent years many methods have been proposed for mining TARs from temporal datasets. Temporal datasets can be defined in different ways depending on the time attribute: each transaction/sequence in the dataset is associated with a time-stamp; each item/event in a transaction is associated with a time-interval; and each item/event in a transaction/sequence is associated with a time-stamp in which it occurs. In contrast with the classical association rules, TARs usually add a time constraint (a point in time or a time range) to the rule to indicate when it holds. These rules can be defined from several temporal itemsets and frameworks, but two of the most popular frameworks to extract temporal itemsets are: sequential patterns and episodes.

The term of sequential pattern was introduced by Agrawal and SriKant (Agrawal & Srikant, 1995) from a customer dataset. We will consider the same customer dataset as an example in order to introduce this term, where a customer identifier (*id*), a time stamp, and a list of products (itemset) are included in each transaction. A new dataset can be generated from this one, in which the transactions with the same *id* are ordered by their timestamps and arranged in a single sequence, called $customer-sequence$. Notice that, each $customer-sequence$ is a sequence of itemsets. For instance, let us take a simple problem of a customer dataset with 4 customers and 5 items: a, b, c, d and e (see Table 1). Table 1a shows the values of the 9 transactions of the dataset ordered by *id* and time stamp. Table 1b shows the customer-sequences obtained from Table 1a.

ID	**Time Stamp**	**Purchased Items**
1	April 26, 2019	a
1	April 27, 2019	b
1	April 28, 2019	c,d
2	April 27, 2019	b,c
2	April 28, 2019	a
3	April 26, 2019	d,e
4	April 26, 2019	a
4	April 27, 2019	b
4	April 28, 2019	c

(a) A simple example of a customer dataset.

ID	**Customer-Sequence**
1	⟨(a), (b), (c,d)⟩
2	⟨(b,c), (a)⟩
3	⟨(d,e)⟩
4	⟨(a), (b), (c)⟩

(b) Dataset with customer-sequences.

Table 1: A simple problem with 4 customers and 5 items.

In general, any problem with sequential patterns where a sequence of itemsets can be obtained could be addressed in a similar way. In this way, a sequence of itemsets $\langle X_1, X_2, ..., X_n \rangle$ is considered a subsequence of another $\langle Y_1, Y_2, ..., Y_m \rangle$ with $n \leqslant m$ if $X_1 \subseteq Y_{i_1}, X_2 \subseteq Y_{i_2}, ..., X_n \subseteq Y_{i_n}$ and $i_1 < i_2 < ... < i_n$. Thus, the support of a sequence S ($Sup(S)$) is defined as the frequency with which S is a subsequence of the customer-sequences in the dataset. For instance, the support of the sequence $S = \langle (a), (c) \rangle$ is 0.5 because S is a subsequence of the customer-sequences 1 and 4 of Table 1b. Many methods have been proposed to mine sequential patterns from different types of datasets (time series, transactions, etc) and then to generate temporal rules from them (Harms & Deogun, 2004; Q. Li, Feng, & Wong, 2005).

Another approach is to discover TARs from a time-ordered sequence of events (such as alarm streams, web navigation logs, etc) in which events may occur at a point in time or at a time range (Mannila et al., 1997; Nguyen, Luo, Phung, & Venkatesh, 2018). This task consists of finding event collections (episodes) that appear frequently enough in a sequence in order to generate the corresponding rules that enable the prediction of a plausible subsequence continuation. These episodes can be classified into three classes (Mannila et al., 1997): Serial episode, when the order among events is total[2]; Parallel episode, when the order among events has no constraints; and Non-Serial and Non-Parallel episode, when the events have no constraints in the order but the order among some of them is total. Let us take a simple sequence with 4 different events (a, b, c and d), where each event has the associated time stamp in which it occurred:

$$S = \langle (a,1),(b,2),(d,3),(c,4),(a,5),(c,6),(a,7),(b,8),(c,9)\rangle$$

Figure 1 shows a simple example of each episode class together with a corresponding example subsequence S' in which it occurs. An episode A is considered a subepisode of another B ($A \preceq B$) if all events and the constraints in the order of A are included in B. Respectively, the episode B is considered a superepisode of the episode A. For instance, the serial episode (a, c) is a subepisode of the serial episode (a, c, d). An episode A occurs in an event sequence when A is a subepisode of the sequence. Thus, the support of an episode A is usually calculated establishing a time window w with a width within which the episode must occur. The support is defined as the fraction of windows w in which the episode occurs (Laxman, Sastry, & Unnikrishnan, 2005). For instance, let us take a window with width 5 for the event sequence S listed above. The serial episode (a, b, c) occurs in 5 of the 13 windows w on S (see Figure 2), its support being 0.38. Many methods have been proposed to mine frequent episodes and then to generate temporal rules from them (Ao et al., 2018).

Often, in many real applications interesting knowledge is associated with a particular time period or with special events. The problem is to find valid time intervals during which rules hold and the discovery of possible periodicities is possible. The periodicity constraint may even be relaxed in order to find interesting rules for which some of its occurrences will be misaligned due to noise events. Moreover, the periodicities of real-life events are not

[2]The order among events is total when all the events occur in this order in the sequence.

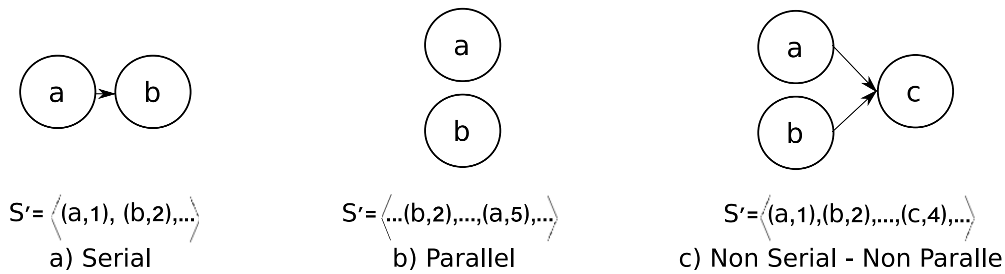

Figure 1: Example of the three classes of episodes.

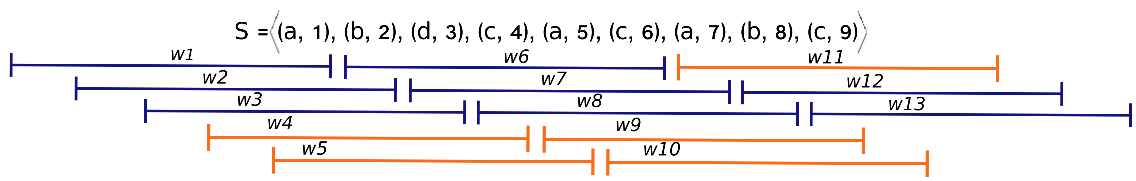

Figure 2: Windows in which the serial episode (a, b, c) occurs on S with a windows w with width 5.

always so regular and usually contain some disturbances. Discovering such rules may lead to the extraction of useful knowledge for the users.

The different frameworks and applications for which the temporal rules are used have resulted in a lack of standard terminology in the literature. For example, the following terms have been used to refer to this field: "sequential rules" (Teng et al., 1990), which represent temporal dependences from a set of sequences; "episode rules" (Mannila et al., 1997), which are defined as an expression $X \rightarrow Y$, where X and Y are episodes that occur in this order in a sequence of events; "inter-transaction association rules" (Lu et al., 2000), which are defined as "rules that express the association among items from different transaction records"; "prediction rules" (Deogun & Jiang, 2005), that are defined as "TARs that are useful for prediction, presenting a time lag between the antecedent and consequent and including only the most important factors in the antecedent"; "cyclic association rules" (Ozden et al., 1998), which are defined as "rules with the minimum confidence and support at regular time intervals"; "calendar association rules" (Ramaswamy et al., 1998), that are defined as "rules with the minimum confidence and support during every time unit contained in a calendar, modulo a mismatch threshold, which allows for a certain amount of error in the matching";

“periodical association rules” (D. Li & Deogun, 2005), that are defined as ”rules focused on finding the periodic events occurring at regular time intervals”; “interval-based temporal association rules” (Höppner & Klawonn, 2001), which are defined as ”rules that are mined from temporal data in which events/items are better treated as intervals rather than time points”; etc. Notice that some of these terms, such as “sequential rules”, “episode rules”, and so on, are more general than TARs since the data can be sequentially ordered by a criterion other than time.

This fact makes it difficult to find proposals and relevant studies in this research field. In what follows, we will introduce a taxonomy in order to offer a well-defined framework that allows researchers to easily locate the published proposals and to propose well-grounded methods in the future.

TAXONOMIC SCHEME

In order to analyze the state-of-the-art, we have looked at a broad variety of proposals to identify TARs from datasets with temporal information. To clear up and to classify them, we propose a taxonomy with two levels in Figure 3.

The first level categorizes the proposals depending on whether the time variable is considered as an implied or integral component. The following two main categories are thus obtained:

- Time as Implied Component. The time variable is considered to provide order to the data collection and/or to locate some temporal constraints that determine the relevance of an event/data/item with respect to another.

- Time as Integral Component. The time variable is considered as an attribute within the learning process and may be an integral component of the rule model.

In the second level, the proposals of the first category (*implied component*) are classified based on the type of temporal datasets from which the rules are extracted: sequential and inter-transaction. The category *sequential* gathers proposals that mine episode or sequential rules from a single sequence or multiple sequences in which the time is used to provide a linear

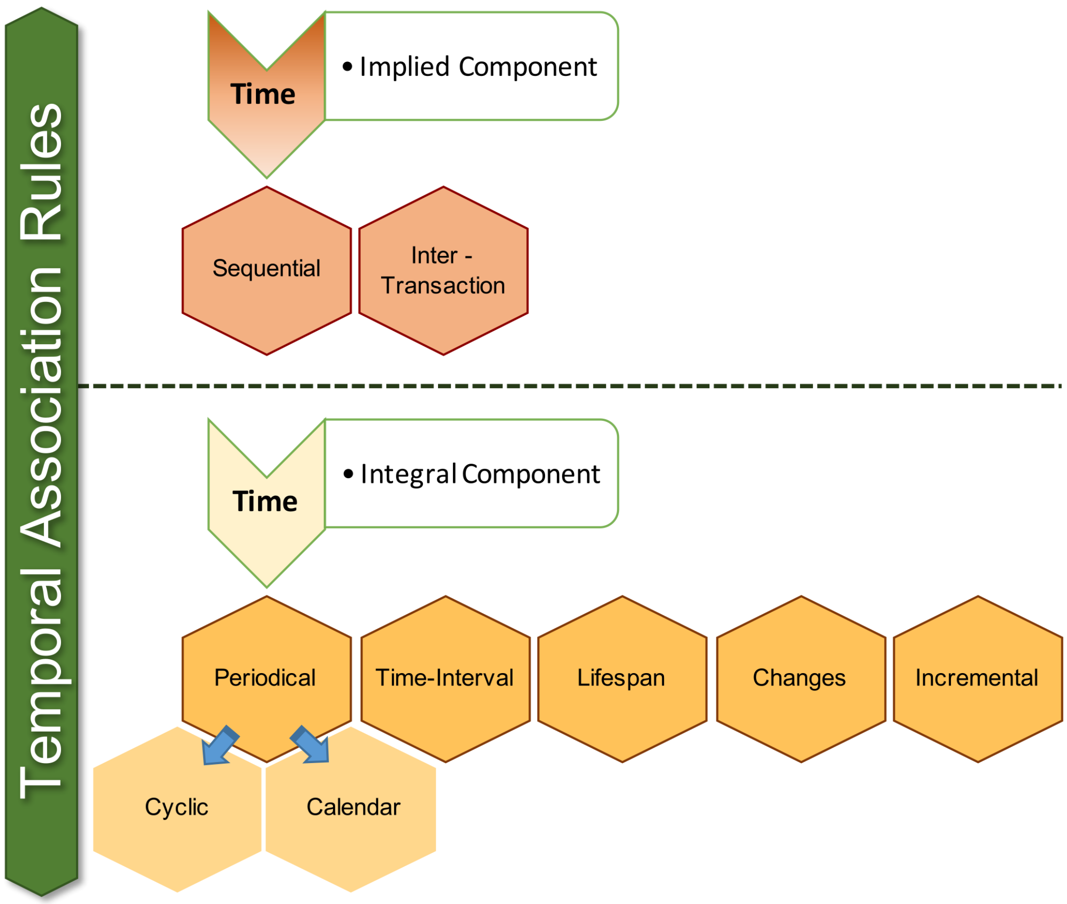

Figure 3: Proposed taxonomy for temporal association rules.

ordering of the occurrence of data/events. The category *inter-transactions* makes reference to groups of proposals that mine inter-transaction association rules from transaction datasets in which temporal components are usually included in each transaction as the date and/or time in which the data/events occurred.

The proposals of the second category (*integral component*) are organized depending on how the temporal information is included in the learning process: periodical, time-interval, lifespan, changes and incremental. The category *periodical* makes reference to groups of proposals in which periodicity constraints are considered in the extraction process. This category is divided into two subcategories (*cyclic* and *calendar*) depending on whether the temporal partitions are expressed by regular temporal periods or by a calendar scheme. The category *time-interval* includes proposals which represent temporal associations considering the duration of the events. The category *lifespan* gathers proposals that consider the temporal interval in which the items exist in the datasets. The category *changes* is related to proposals that mine temporal meta-rules to represent changes in association rules over time. Finally, the category *incremental* includes proposals that solve the problem of mining TARs from incremental datasets.

Moreover, we have considered other additional aspects such as whether the proposals are designed for a single sequence or for multiple sequences, or whether the proposals make use of particular techniques, such as the fuzzy logic, and so on, to sort the available proposals in some categories of the taxonomy. This taxonomy allows us to classify the search space in which we may find the existing proposals. A brief description of some related proposals for each category will be provided in the following sections.

CONSIDERING TIME AS AN IMPLIED COMPONENT IN THE MINING PROCESS

This section presents a brief introduction of the two subcategories in which the proposals consider the time variable to provide order to the data collection and/or to locate some temporal constraints that determine the relevance of an event/data/item with respect to another. They are: sequential and inter-transaction.

Sequential category

Many data collections are values/event sequences ordered linearly according to the time in which they occurred, such as sequences of alarms in a telecommunications network, human activities, states in the evolution of a disease, time-series microarray, and so on. Users may be interested in mining temporal rules from these data collections, in which antecedent and consequent refer to sequential moments in time, representing some kind of ordered dependence between data/events of the sequences with a given confidence that can be used for prediction (Lo, Khoo, & Wong, 2009; Fournier-Viger, Faghihi, Nkambou, & Nguifo, 2012). The sequential or episode rules are similar to the classical association rules but considering that the antecedent must occur before the consequent and that they are mined from sequence databases (Teng et al., 1990). In addition, the required confidence avoids the limitation of sequential patterns, which even while being frequent may have a low confidence, and which are therefore unhelpful for decision making or prediction.

These rules are usually generated from sequential patterns or episodes (Febrer-Hernández & Hernández-Palancar, 2012; Fournier-Viger, Lin, Kiran, Koh, & Thomas, 2017; Mooney & Roddick, 2013; Zimmermann, 2014) that appear in a single sequence (Mannila et al., 1997; Deogun & Jiang, 2005), across sequences (Lo, Ramalingam, Ranganath, & Vaswani, 2012; G. Das, Lin, Mannila, Renganathan, & Smyth, 1998), or common to multiple sequences (Lo et al., 2009; Fournier-Viger, Wu, Tseng, Cao, & Nkambou, 2015). Classical algorithms are usually extended making use of a sliding-window in which they find the frequent itemsets that will be used to generate the rules (Mannila et al., 1997; Fournier-Viger, Wu, Tseng, & Nkambou, 2012). Some algorithms also include other temporal constraints such as a maximum or minimum amount of time in which the rules must occur (Harms & Deogun, 2004; Nam, KiYoung, & Doheon, 2009) or a maximum elapsed time between any two consecutive events (Nam et al., 2009; Ao et al., 2018). Some variants try to relax the order constraint of the events for mining the rules (Fournier-Viger et al., 2015; Setiawan & Yahya, 2018) or define a rule template to mine relationships in which the user is interested (Bettini, Wang, & Jajodia, 1998; Sudkamp, 2005; Concaro, Sacchi, Cerra, Fratino, & Bellazzi, 2011). Most algorithms for mining sequential rules focus on generating all the possible rules (containing a high number of redundant rules), making mining inefficient. In order to solve this problem, some methods have specifically been designed to mine non-redundant sequential rules directly from sequence datasets (Tran, Le, Vo, & Hong, 2016). Other methods make use of specific data structures to improve the efficiency and interpretability of the rule mining (L. Wang, Meng, Xu, & Peng, 2018; Fournier-Viger, Gueniche, Zida, & Tseng, 2014).

The sequence or episode rules have been successfully applied to a wide variety of applications. For instance, they have been used to quickly identify unused spectrum bands that allow opportunistic access by radios seeking spectrum, alleviating the problem of constantly increasing demand for communication bandwidth (Heydari & Tajer, 2019), to improve system security through the analysis of event logs with historical information of potential security breaches (Khan & Parkinson, 2018), to predict the traffic congestion level (Wen, Zhang, Sun, Wang, & Xu, 2019), and so on.

Inter-Transaction Category

On the other hand, many real applications provide transaction datasets in which the temporal component is included in each transaction as the date and/or time in which they occurred. In such cases, users may be interested in mining dependences among items of different transactions instead of dependences among items within the same transaction. Although transactions occur in certain contexts, such as time, place, etc., such contexts have been ignored in classical association rule mining. When the items in the transaction database are organized by transaction time, the inter-transaction association rules enable the representation of associations through the dimension of time (Lu et al., 2000). These kinds of rules are usually generated making use of a sliding-window that allow a division of the dataset into equal length intervals, mining the interesting rules that span a certain number of intervals. To this end, some proposals present an extension of the classical method Apriori (Agrawal et al., 1993) for mining inter-transaction association rules (Lu et al., 2000; Feng, Dillon, & Liu, 2001; Q. Li et al., 2005; J.-W. Huang, Dai, & Chen, 2007). In order to improve the efficiency of the extraction process, some methods have specifically been designed for mining inter-transaction association rules, such as in Tung *et al.* (Tung et al., 2003) in which they first mine frequent intertransaction itemsets and they then form intertransaction rules from them, or in Wang C. S. (C.-S. Wang, 2015) in which the author characterizes non-redundant inter-transaction rules and presents a method for mining non-redundant inter-transaction rules efficiently. Some researchers have also presented proposals to select the size of the time-window to optimize the frequency of the extracted dependencies (Y. Xiao, Tian, & Zhao, 2014).

Fuzzy set theory has been used more and more frequently in intelligent systems because of its simplicity and similarity to human reasoning (Ishibuchi, Nakashima, & Nii, 2004). Some researchers have made use of fuzzy logic to propose new methods for mining these kinds of rules. For instance, in (Y.-P. Huang & Kao, 2005; Y.-P. Huang, Kao, & Sandnes, 2007) the authors propose an extension of some classical algorithms (Apriori and PrefixSpan) which consider a user-defined sliding window to generate only the fuzzy rules whose span is less than or equal to this window. Evolutionary algorithms have also been used to learn membership function contexts for mining the fuzzy rules (Matthews, Gongora, & Hopgood, 2011, 2013).

These rules have been successfully applied in many real applications. For instance, among other examples, they have been used to represent the dynamic characteristics of regional traffic congestion (Xie, Wang, & Zhao, 2020), and to generate rules from transaction databases in complex dynamical systems of financial markets making use of closed inter-transaction itemsets (Hsieh, Yang, Wu, & Chen, 2016).

CONSIDERING TIME AS AN INTEGRAL COMPONENT IN THE MINING PROCESS

In the following subsections, we present a brief description of different subcategories in which the proposals include time as an attribute within the learning process and in which it may be an integral component of the rule model. They are: periodical, time-interval, lifespan, changes and incremental.

Periodical Category

Sometimes users are interested in discovering recurring association rules that are found in many domains. For instance, rules that represent cyclical behaviours of consumers in E-Commerce in order to increase the sales of commercial organizations and make more profit (Thuan, Toan, & Tuan, 2012), rules that allow us to represent cyclical behaviors of people from the large amount of data available in the social networks (Tebourski, Abdessalem Kar, & Ghezala, 2015), and so on. These cyclic rules exist only in certain time intervals and they do not occur in the remainder. The cyclic association rules are similar to the classical association rules but represent dependences that occur periodically over time (Ozden et al., 1998). To generate these rules, the classical Apriori method is used to generate the association rules in each time unit (user-defined partitions of regular temporal periods). Then, a simple pattern matching algorithm is applied to detect the cycles (Ozden et al., 1998). In order to improve the efficiency of the extraction process, some methods also include user-specified constraints in the process (Tebourski et al., 2015) or are specifically designed for mining non-redundant cyclic association rules (Tebourski & Karaa, 2012). A more general approach is considered by Chen *et al.* (X. Chen & Petrounias, 1999) for finding by rule the longest set of contiguous time intervals in which each of them is periodical.

However, different cyclic rules may be discovered depending on the time intervals considered. To address this issue, the temporal partition may be defined with multiple granularities making use of a calendar scheme (for instance, week, day, and hour), reducing the need for prior knowledge to define the time intervals (Ramaswamy et al., 1998). The classical method for mining calendar association rules extends the second step of the classical cyclic TAR extraction method (Ozden et al., 1998), in order to consider a calendar scheme in the process (Ramaswamy et al., 1998). Some proposals introduce different techniques to improve the extraction process efficiency, such as new data structures to process each temporal partition (Verma & Vyas, 2005; Verma, Vyas, & Vyas, 2005). However, periodicities in the datasets are not usually very precise and have disturbances. In addition, users tend to define temporal constraints in a vague and imprecise way. For this reason, some authors make use of fuzzy logic to define the calendar schema in a way that replicates the natural language more closely (Ishibuchi et al., 2004). For instance, Lee *et al.* proposed two algorithms (W.-J. Lee & Lee, 2004; W.-J. Lee, Jiang, & Lee, 2008) that allow the calendar schema to be defined by means of fuzzy sets, as these are more flexible with the restrictions of the user in the rule mining process.

Many methods have been used to extract rules exhibiting periodic behavior, but the concept of periodic behavior is defined very strictly and many rules may be removed in the mining process. A periodic association rule always requires the same amount of time to have elapsed between two occurrences. Thus, these methods are fairly rigid, failing to capture a rule that would exhibit a pattern in most of the time intervals but not all the time. To address this issue, different proposals have been presented to make the concept of periodicity more flexible (Han, Gong, & Yin, 1998; Fournier-Viger et al., 2019). For instance, Li *et al.* (D. Li & Deogun, 2005) extend the approach of Han *et al.* (Han, Dong, & Yin, 1999) incorporating the concept of the slide window to extract rules whose periodicity is defined by the sliding window size.

Time-Interval Category

In many application domains, temporal datasets include not only time-stamped data, or time points, but also time intervals that are part of the input data (for instance, the administration of a vitamin for 1 month) or are abstractions/interpretations derived from them as part of the analysis process (for instance, a week of anaemia) (Bohlen, Busatto, & Jensen, 1998; Moskovitch, Elovici, & Rokach, 2008). This temporal information must be considered in the extraction process in order to obtain consistent interval-based temporal association rules. For instance, the duration of faults in an industrial process is essential information when we try to find correlations among them. Because of this, Laxman *et al.* (Laxman, Unnikrishnan, & Sastry, 2002; Laxman, Sastry, & Unnikrishnan, 2007) extended the definition of episode (Mannila et al., 1997) to incorporate event duration constraints.

Temporal rule discovery from these kinds of datasets is more complex and requires a different approach than those used to mine patterns from point-based datasets (such as sequential patterns or episodes). An interval involves a duration and, therefore, the extracted patterns may have a different semantic than simply before and after. Allen's temporal interval logic and its extensions (Allen, 1983; Roddick & Mooney, 2005) are typically used to represent the relationships among intervals (Höppner & Klawonn, 2001; Moskovitch & Shahar, 2015). For instance, Winarko *et al.* presented ARMADA (Winarko & Roddick, 2007), a method based on the algorithm for mining sequential patterns MEMISP (M.-Y. Lin & Lee, 2002). ARMADA generates richer temporal association rules from interval-based data considering a maximum gap time constraint to reduce the number of rules generated. Papapetrou *et al.* (Papapetrou et al., 2009) proposed three methods based on Apriori by including constraints to discover temporal relations from interval-based event sequences, so generating the top k rules that maximize a given interestingness measure. Lee *et al.* (Y. J. Lee, Lee, Chai, Hwang, & Ryu, 2009) presented a pre-processing algorithm to generate the temporal interval data and a method to extract temporal rules from them. Nazerfard *et al.* (Nazerfard, Rashidi, & Cook, 2011) presented TEREDA (Nazerfard et al., 2011), a method for mining temporal aspects (such as the usual start time and duration) of daily activities using TAR mining techniques. A method based on the algorithm Progressive-Partition-Miner (PPM)(C.-H. Lee, Chen, & Lin, 2003) is also proposed for mining academic relationships considering the duration of the academic activities (F. Huang, Zou, Liu, & He, 2012). Finally, a proposal is presented for mining TARs from point-like events and interval-like events considering user-defined rule templates to capture relationships in which the user is interested (Concaro et al., 2011).

Some researchers have made use of fuzzy logic to propose new methods for mining rules from interval-based data. Sudkamp (Sudkamp, 2005) proposes a method based on Apriori for mining fuzzy temporal relationships from multiple sources making use of fuzzy logic to consider temporal durations (for instance, a short period) and temporal constraints (for instance, shortly after). Wu (Wu, 2010) presents a method for mining fuzzy association rules from web user access patterns in web logs using fuzzy logic to represent the time duration that the users spent on the web pages.

These methods have been applied to a wide variety of applications. Further examples of application include but are not limited to: the recognition of complex activities of daily life in the area of mobile computing because the sensors included in such devices provide a large amount of interesting information for mobile applications (L. Liu et al., 2016); the prediction of medical diagnoses from hospital data stored during the different hospital stays of individual patients, considering as simultaneous the events that occur during the time that the patients are hospitalized in the same medical unit (Vandromme et al., 2017); the prediction of cardiorespiratory instability in step-down unit patients from continuous monitoring data of physiologic vital sign measurement, providing the predicted event's time of occurrence (Guillame-Bert et al., 2017).

Lifespan Category

An important aspect in the mining process is to consider temporal information about the existence of an item in transaction datasets. The lifespan of an item is defined as the time period between the first and the last time it appears in a transaction dataset. If the lifespan of the items is considered in order to calculate the support of the rules, a greater number of rules may achieve a support equal to or greater than minSup. For instance, let us take a simple dataset with 1000 transactions of an e-commerce store (such as Amazon, etc), which has been recorded during the last 10 years at 100 transactions per month, and a minSup of 0.5% (50 transactions). In this dataset, a product that has been bought 5 times per year is considered a frequent item while a new product included in the dataset in the last year that has been bought 40 times is an infrequent item. Although the new product has been very popular since its appearance (eight times more popular than the first), it appears in a small fraction of transactions with respect to the total number of transactions and will not be considered to generate rules. This temporary information (lifespan) is also used to prune datasets, removing those items that are not interesting for the user at present (Ale & Rossi, 2000).

Some researchers have made use of fuzzy logic to transform quantitative values in the transactions to fuzzy sets in order to generate fuzzy temporal association rules making use of the fuzzy count and lifespan of each item (C.-H. Chen, Lan, Hong, & Lin, 2016). However, the given membership functions may have a critical influence on the final mining results. For this reason, some proposals have also incorporated a learning or tuning of the functions (Chamazi & Motameni, 2019).

These methods have been mostly used on the transaction databases of different types of companies in order to detect consumer preferences (C.-H. Chen et al., 2016; Chamazi & Motameni, 2019), however they can be used in many other application domains. For instance, these methods could be used in Internet of Things applications to detect faults within a sensor network or to detect sensors whose activity has been considerably reduced, and they can be relocated or removed from the network.

Changes Category

The world is changing faster than ever due to technological development. This makes it essential to obtain useful knowledge which allows us to detect, evaluate and respond to changing conditions in a timely and intelligent manner (Boettcher, 2011). For this reason, a number of methods have been proposed to detect changes in rules or patterns extracted from data collected in different time periods. The main idea consist of introducing a level of abstraction that allows knowledge to be identified from the extracted elements by means of meta-rules. For instance, Huang *et al.* (T.-K. Huang, Huang, & Chuang, 2016) propose a method for mining association rules representing changes in students' learning performance and characteristics. Song *et al.* (Song, Kim, & Kim, 2001) present a method to detect changes in customer behavior by association rules generated from datasets at different time-periods. Au *et al.* (Au & Chan, 2002) generate fuzzy meta-rules to represent changes in the support and/or confidence of the association rules over time, which is useful to predict how association rules will be updated in the near future.

Incremental Category

Currently, temporal datasets are continuously updated and increased. Because of this, the rules that have previously been generated need to be updated, removing those rules that are no longer relevant and adding valid new rules. Re-running the temporal mining algorithm every time is ineffective since it neglects the previously discovered rules, and repeats the work done previously. Therefore, several methods have been proposed to solve the problem of mining incremental TARs. Gharib *et al.* (Gharib, Nassar, Taha, & Abraham, 2010) present an incremental algorithm based on the Sliding-Window Filtering algorithm (C.-H. Lee, Lin, & Chen, 2001), which maintains temporal frequent itemsets after the temporal transaction database has been updated in order to reduce the time needed for generating new candidates. Fouad *et al.* (Fouad & Mostafa, 2017) present a method for the efficient mining of incremental temporal association rules, making use of a new data structure and of previously discovered frequent temporal itemsets.

These methods have been applied successfully in a wide variety of applications. Teng *et al.* (Teng et al., 1990) present a method for mining sequential rules to detect anomalies in user behavior, which are adapted by means of a time-based inductive engine. Zhou *et al.* (Zhou & Hirasawa, 2017) propose a genetic algorithm for mining sequential rules to detect customer preferences, which are adapted on-line by making use of an ant colony optimization algorithm because customer preferences change over time.

REAL-WORLD APPLICATIONS AND AVAILABLE SOFTWARE TOOLS

The amount of data collected from different application areas in which temporal information is implicitly or explicitly included has led to a situation in which the extraction of interesting knowledge from temporal datasets is a very attractive and challenging task. TARs allow us to generate models with a greater predictive and descriptive power and an additional degree of interestingness, which can be successfully applied in a wide variety of applications. Table 2 shows a short list of recent application contributions that have been addressed making use of TARs. At a glance, we can see a large amount of quite different applications which are evidently related to very diverse general areas. Specifically, we can find applications in industry, security, medicine and healthcare, among many others, with the number of application proposals in the areas of medicine and healthcare particularly noteworthy. This demonstrates the great potentiality of TARs to be applied successfully to a wide variety of problems. In addition, the number of application contributions has increased in recent years, providing further evidence of the necessity of and growing interest in TARs for the resolution of real-world problems.

Although a high number of proposals on TARs have been published over the years, only a few researchers share the source codes associated with their proposals and the codes or descriptions available in books and journals often present inconsistencies (Thimbleby, 2003). This issue, along with the high complexity of some proposals, prevents the widespread use of TARs in real-world applications. Moreover, researchers need to reimplement methods to be able to compare their proposals, spending a great amount of time reimplementing methods that in many cases are not as efficient as the original versions. To address this issue, some free and open source software tools have been released with the aim of supporting the use of TARs to address real-world applications. Notice that the open source model facilitates the application and adaptation of TARs to new problems and is essential to spreading TARs to new subjects and industry (*Open Source Initative. 1998*, n.d.).

SPMF (Fournier-Viger, Gomariz, Gueniche, et al., 2014) is an open-source data mining library released under the GPL v3 license and developed in Java, so that, it can be run on all

Application Domain	**Description (References) Category**
Industry	Signal processing (Heydari & Tajer, 2019)•
Economy	Dynamical systems of financial markets (Hsieh et al., 2016) ◎
Security	Security event log analysis (Khan & Parkinson, 2018)•
Traffic Management	Intelligent transportation management (Wen et al., 2019)•, Traffic congestion (Xie et al., 2020) ◎
Human behaviors	Daily living activity recognition (L. Liu et al., 2016)◦, Changes in students' learning performance and characteristics (T.-K. Huang et al., 2016)↺, Detection of consumer preferences in supermarkets (Zhou & Hirasawa, 2017)↑ (Panchal & Prajapati, 2018) ◁ (Chamazi & Motameni, 2019) ⋄, Prediction of human behaviors in manufacturing companies (Setiawan & Yahya, 2018)•, Travel diaries analysis (Vu, Li, Law, & Zhang, 2018)•
Medicine and Healthcare	Screening potential drug interactions (Ji et al., 2016)•, Application to hospital data (Vandromme et al., 2017)◦, Drug safety signal detection (Arnaud et al., 2017), Forecast cardiorespiratory instability (Guillame-Bert et al., 2017)◦, Cancer treatment (Nguyen et al., 2018)•, Coronary heart disease diagnosis (Orphanou et al., 2018)◦, Prediction of developmental stages in drug abuse (Kilgore, Korneeva, Arnold, Trutschl, & Cvek, 2019)•, Health knowledge system (Giannoulis, Kondylakis, & Marakakis, 2019)◦, Heart diseases (Bou Rjeily, Badr, Hajjarm El Hassani, & Andres, 2019)•

• Sequential Category, ◎ Inter-Transaction Category, ◁ Periodical Category, ◦ Time-Interval Category, ⋄ Lifespan Category, ↺ Changes Category, ↑ Incremental Category.

Table 2: Recent application contributions on TARs by application domain.

major platforms. This library provides more than 110 algorithms for sequential rule mining, sequence prediction, association rule mining, sequential pattern mining, periodic pattern mining, episode mining, and so on, and several temporal datasets. In addition, SPMF can be embedded in other Java codes easily and can be run as a standalone software or by a graphical user interface. Several R packages can also be found in the CRAN repository for mining sequential rules, sequential pattern, and so on, and for handling temporal datasets [3]. For instance, the package arulesSequences provides interfaces to the C++ implementation of cSPADE, the package timetools provides several tools to manipulate sequential and seasonal time series, and the package eventInterval provides functions for the analysis of rate changes

[3] https://cran.r-project.org/web/packages/

in sequential events. The GitHub repository also provides several packages [4], such as the package CRFSuite, which is an implementation of Conditional Random Fields for sequential prediction problems, the package “Emotion Classification in Microblog Texts Using Class Sequential Rules”, which allows a classifier to be generated based on class sequential rules for emotion classification, and the package “Hierarchical Association Rules”, which allows association rules to be mined from sequential events. Several python packages can also be found in the Pypi repository [5]. For instance, the package “Maximal Sequential Patterns Mining” is a python wrapper between SPMF and Python, and the package prefixspan provides several methods for mining sequential patterns. We can also find some methods for mining sequential patterns as part of some open-source data mining suites, such as Weka (Witten, Frank, & Hall, 2016), Knime (Berthold et al., 2009) or Mahout (Lyubimov & Palumbo, 2016).

Several data collections are also available as open-source datasets to evaluate and compare the performance of algorithms. For instance, the webpage of SPMF provides a large number of datasets in the SPMF format, the Phyton package SecAlertSeqMining enables the generation of sequential databases in formats required by SPMF, and “IBM Generator” allows synthetic sequence databases to be generated.

Finally, notice that some of the most used algorithms are ARMADA (Winarko & Roddick, 2007), CMRules (Fournier-Viger, Faghihi, et al., 2012) and TRuleGrowth (Fournier-Viger et al., 2015). In addition, the algorithms CM-SPADE (Fournier-Viger, Gomariz, Campos, & Thomas, 2014) and SPAM (Ayres, Flannick, Gehrke, & Yiu, 2002) are also widely used to extract sequential patterns and then to generate temporal rules from them.

[4] https://github.com

[5] https://pypi.org

TEMPORAL ASSOCIATION RULES BIBLIOGRAPHICAL ANALYSIS

This section shows a bibliographical analysis of the research field in accordance with publications indexed at WoS [6]. This platform includes several large bibliographic databases with references to the main scientific publications of any science field and it is frequently used to track research for monitoring current developments (trends detection) and analyzing classical contributions (assessing the state-of-the-art). In this sense, we perform a bibliographical query as follows:

TS=((“temporal association rule” OR “temporal rule*”) OR ((“sequential rule*” OR “episode rule*” OR “periodic rule*” OR “meta-rule*”) AND (“temporal” OR “time”)) OR ((“prediction rule*” OR “change rule*”) AND “temporal”) OR (“cyclic association rule*” OR “cyclic rule*”) OR (“calendar association rule*” OR “calendar rule*” OR (“association rule*” AND “calendar”)) OR (“interval-based temporal association rule*” OR “interval-based temporal rule*” OR (“association rule*” AND (“time interval*” OR “temporal interval*”))))*

where TS means that these terms are searched in the fields Keywords Plus®, Author Keywords, Abstract, and Title, * means 0 or more characters, and $ means 0 or 1. The time stamp was fixed from 1998 to 2018 (the last complete year recorded at WoS). Note that the query has been made quite general in order to consider all terms that are usually included in published contributions, allowing for an overview of the research field.

Figures 4a and 4b show the number of published proposals and citations per year according to the WoS, respectively. We can see how the number of citations increases quite linearly and the number of publications shows a growing trend over the years. These data allow us to say that this research field is currently a mature area with an active community working on challenging related issues and increasing the interest in this research field each year according to the number of citations.

In order to analyze the current research trends, a study has been carried out focusing on the distribution of the published proposals among the WoS categories from 2013 to 2018 (the last 6 years at the time of writing). Table 3 shows the top-ten WoS categories according

[6] https://www.webofknowledge.com/

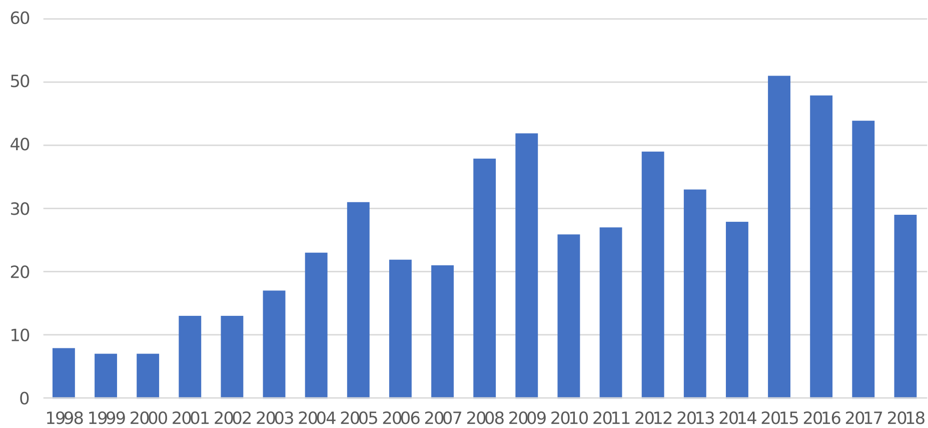

(a) Publications.

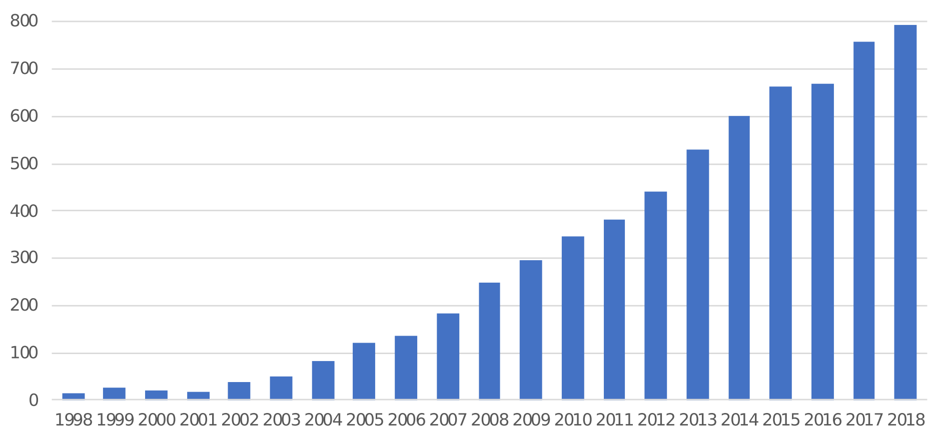

(b) Citations.

Figure 4: Number of publications and citations from 1998 to 2018 according to the WoS. Report produced in September 2, 2019; Total publications: 573; Total times cited: 6,482; Average citations per item: 11.31; h-index: 41.

to the percentage of publications that are associated with each of them. Notice that a publication can be associated with more than one category, so that the sum of the percentage column can be over 100%. This study presents a wide variety of categories in which TARs publications are associated, showing a wide applicability to different areas. Even though more than 80% of publications are associated with three subcategories of Computer Science (Artificial Intelligence, Theory Methods and Information Systems), notice that more than 50% of publications are also associated with categories for specific applications (such as Engineering Electrical Electronic, Software Engineering, Telecommunications, Neurosciences, and so on).

WoS Categories	**% Percentage**
Computer Science Artificial Intelligence	39.96
Computer Science Theory Methods	23.04
Computer Science Information Systems	21.12
Engineering Electrical Electronic	18.15
Computer Science Software Engineering	7.50
Computer Science Interdisciplinary Applications	7.33
Telecommunications	5.58
Neurosciences	4.19
Automation Control Systems	3.49
Computer Science Hardware Architecture	3.14

Table 3: The top-ten WoS Categories with the most associated publications.

For a deeper analysis, we have studied the evolution of the publications focusing on their co-concurrence in the main WoS categories from 2013 to 2015 and from 2016 to 2018. To that end, we have created visual science maps (Batagelj & Cerinšek, 2013; M. Cobo, López-Herrera, Herrera-Viedma, & Herrera, 2011; Moya-Anegón et al., 2004) (see Figure 5) in which each sphere represents a WoS category (the volume and darkness depend on the number of publications) and the link between two spheres represents the level of co-concurrence among them (the thickness and darkness depend on the level of co-concurrence). The level of co-concurrence between two categories (A and B) is calculated using the measure *equivalence index* (e_{AB}) (Callon, Courtial, & Laville, 1991), which is defined as $e_{AB} = c_{AB}^2/(c_A \cdot c_B)$ where c_{AB} is the number of publications simultaneously indexed in the categories A and B, c_A is the number of publications indexed in category A, and c_B is the number of publications indexed in category B. When the publications are always indexed in the two categories, the equivalence index is unity, when they are never indexed simultaneously in both categories, it is zero. These visual science maps have been created by means of the free software Science Mapping Analysis Tool (SciMAT) (M. J. Cobo et al., 2012) and the Open Graph Viz Platform (Gephi) [7].

[7] https://gephi.org/

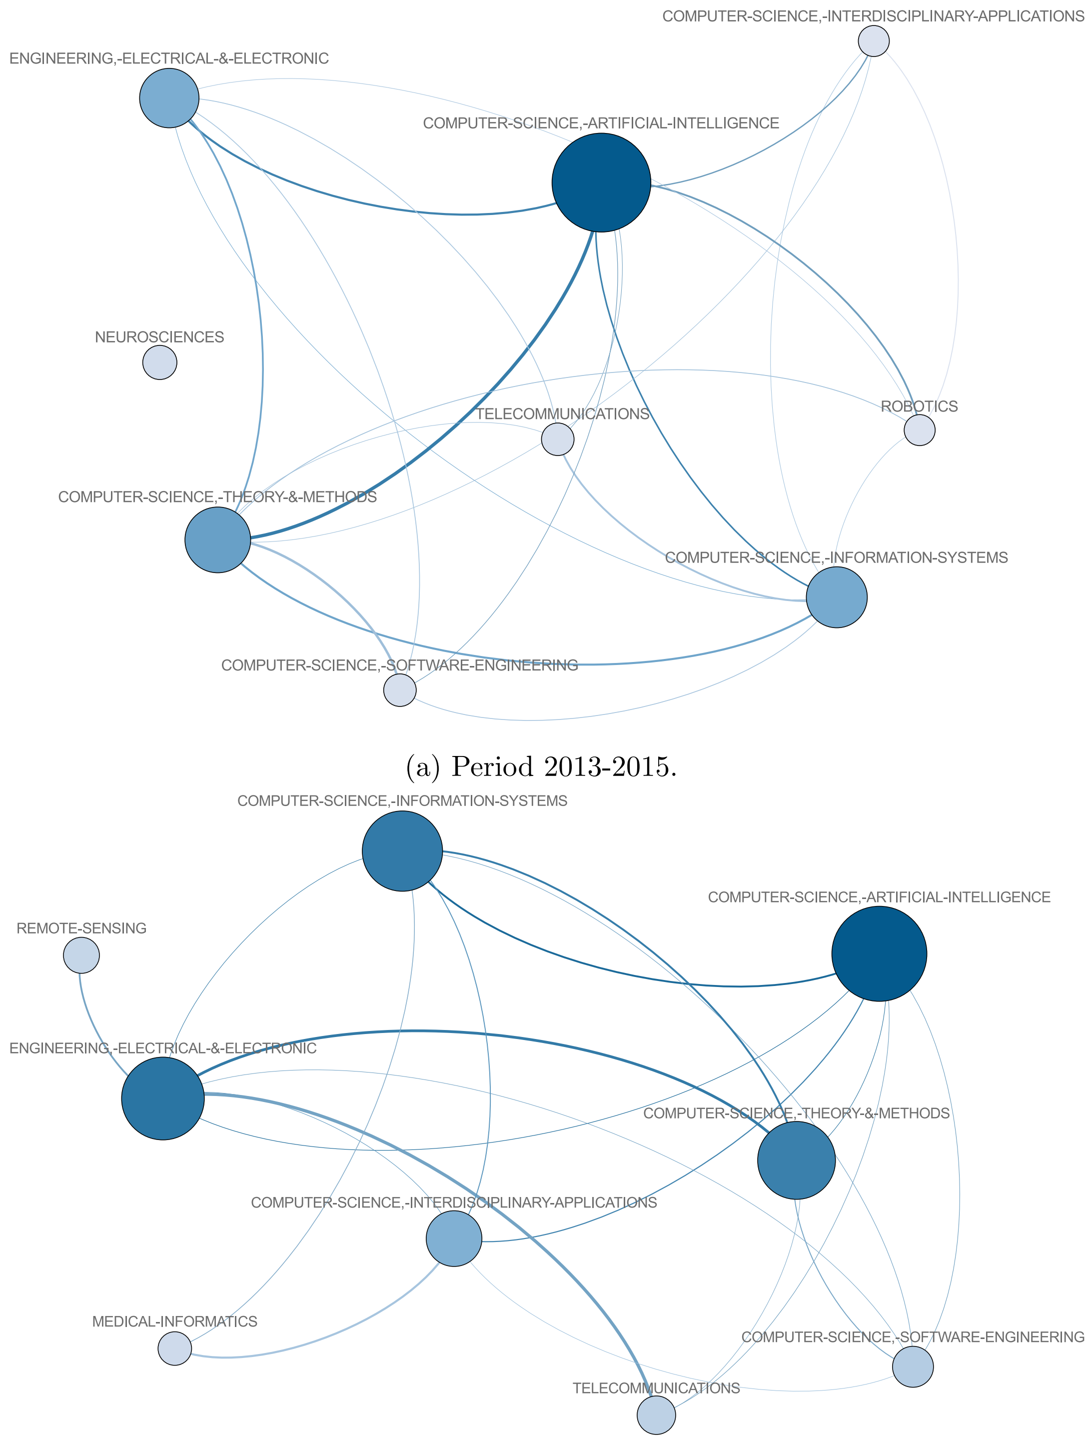

(a) Period 2013-2015.

(b) Period 2016-2018.

Figure 5: Visual science maps on TAR publications

Figures 5a and 5b show the map for each period, presenting the categories on which the researchers have mainly focussed their research over the last 6 years. For the period 2013-2015, we can see how three subcategories of Computer Science (Artificial Intelligence, Theory Methods and Information Systems) and the category Engineering Electrical Electronic are the 4 WoS categories with the most associated publications, with the subcategory Artificial Intelligence presenting a higher number of publications than the remaining categories and sharing a strong link to the category Theory Methods. These main categories are also connected to 5 other WoS categories for specific applications such as Telecommunications, and so on. For the second period (2016-2018) we can appreciate how the main categories are the same, but the categories Theory Methods, Information Systems and Engineering Electrical Electronic have increased their number of publications with respect to the category Artificial Intelligence. Now, the category Theory Methods shares a strong link to Engineering Electrical Electronic and Information Systems, and the category Engineering Electrical Electronic also shares a strong link to the category Telecommunications. In addition, the subcategory Interdisciplinary-Applications of Computer Science increases its number of publications in relation to the period 2013-2015 and several new categories arise as emergent categories, showing a growing interest in fields such as Medical-Informatics and Remote-Sensing. This demonstrates the great capacity of TARs to be applied successfully to a wide variety of problems and the interest that TARs have generated with regard to solving real-world problems.

CRITICAL CONSIDERATIONS AND FURTHER DIRECTIONS

In recent years, a lot of proposals have been published relating to TARs. Each new publication has aroused great interest in the novelties presented with respect to the existing proposals in the literature. But, which are the critical aspects of the current publications? In what follows, we provide some recommendations that should be taken into account in further publications:

- Many of the published proposals on TARs make use or extend some of the well-known classical methods for mining classical association rules. However, these methods usually present problems of complexity (due to the large number of rules generated and the number of items included in each rule) and scalability (in terms of memory and time consumed). In the literature, we can find a number of recent proposals that perform an efficient mining process and optimize several interesting measures to obtain more useful and concise knowledge (Radhakrishna et al., 2015, 2019). The use of these proposals would allow more interesting TARs to be obtained, avoiding some of the problems of scalability and complexity.

- The publication of new proposals requires an experimental analysis comparing the performance of new algorithms to the previous best ones from the specialized literature. This analysis would allow researchers to provide good criteria to show the effectiveness of their proposal. We would recommend an experimental study comparing the obtained results to the widely and well-known classical proposals and to the most recent ones in the topic, making use of several benchmark datasets and statistical tests where possible. However, researchers in this field do not always include this analysis in their proposals or only compare them to classical proposals which were surpassed many years ago. This fact makes it more difficult for researchers to determine the needs in the area and to properly focus their work on significant further proposals.

- A few free and open source software tools are available for mining sequential rules, sequential patterns, and so on, and for the handling of temporal datasets, but few researchers share the source codes associated with their proposals. This issue, along with the high complexity of some proposals, prevents the widespread use of TARs in real applications and makes it difficult for researchers to compare their proposals with other algorithms. We recommend that authors make the source code of their proposals public since this would have a very positive impact both on the development of better algorithms and their applicability to very diverse areas.

From our point of view, there are still many thing to do. We present some possible further directions:

- The technological revolution of the last few years has allowed a huge amount of data to be stored in databases. For instance, one of the main sources of these temporal data collections is the Internet of Things (IoT), which gathers information from collections of sensors, smartphones, among many others, which are capable of operating within the existing Internet infrastructure (Aljawarneh et al., 2017). These temporal data collections, in some cases known as Big Data, present a challenge to the mining process since the processing power of most of the existing techniques is not enought to handle them. Analytics solutions based on MapReduce paradigm (Hadook-based and Spark-based), Deep Learning and Fog Computing, among others, allow researchers to deal with this huge amount of data and to harness their temporal dimension, though the use of these solutions also introduces new challenges such as maintaining data security and privacy. Some proposals for mining TARs from Big Data scenarios have appeared recently, but further developments for big data mining are still expected.

- Some recent proposals also incorporate the spatial information into the mining process to obtain Spatio-Temporal Association Rules (STARs) since the occurrence of some events is often associated with a particular temporal period and location (C. Liu et al., 2019). However, several authors have noticed problematic issues arising from this spatial association rule mining. Firstly, these data collections also consist of spatial attributes that determine the scope of objects and their spatial location, and the dependencies among them are not explicitly represented in the database. Because of this, a pre-processing stage is needed to transform the spatial data into spatial predicates and connectors (such as close_to, etc.) which allows us to represent topological relationships, etc. This may exceed the memory capacity of traditional techniques when the database consists of a large amount of spatial data, whereby it is necessary to find a good trade-off between the pre-processing and mining process. Secondly, some authors also find several problems of interpretation when STARs are generated since the interpretation of some spatial predicates like close_to() usually depend on the user or the information stored in the dataset. For this reason, further developments for STARs mining are still expected.

- Discovery of High-Utility TARs seems another prolific area (Zida, Fournier-Viger, Wu, Lin, & Tseng, 2015). Most of the published methods assume that all items in a dataset have the same relevance, however in many real applications each item has an associated unit profit that determines its relevance (for instance, the price of each item). The aim is to extract high-utility rules that consider items with a high profit value and with high confidence. Another recent approach is to consider that each event in a sequence has a cost associated. The aim is to extract low-cost rules with high values for different quality measures (confidence, length, etc.) (Dalmas, Fournier-Viger, & Norre, 2017).

ACKNOWLEDGEMENTS

This work was supported in part by the Spanish Ministry of Science, Innovation and Universities under Grant TIN2017-89517-P, and in part by the Andalusian Government under Grant P18-RT-2248.

CONCLUSIONS

In real-life applications, data usually change over time and this temporal information is generally included in the data collected. As a result, the discovery of TARs allows us to obtain models with a greater predictive and descriptive power providing an additional degree of interestingness, which may be successfully applied in a wide variety of applications. One of the main problems in this research field is the lack of visibility most works suffer from, since there is no standard terminology to refer to it, making it difficult to find and compare proposals and studies in the field. In this review work, we have attempted to offer a well-defined framework that allows researchers to easily find existing proposals and to present well-grounded future proposals.

To that end, we have presented a two-level taxonomy to organize the proposals depending on whether the time variable is considered to provide order to the data collection and to locate some temporal constraints (category *implied component*), or whether time is considered as an attribute within the learning process (category *integral component*). In the second level, the proposals of the category *implied component* are also classified based on the type of temporal datasets, and the proposals of the category *integral component* are organized depending on how the temporal information is included in the learning process.

Some recent applications with TARs and available software tools have also been shown. Further, a bibliographical analysis of the research field has been carried out in accordance with publications indexed at WoS. From this analysis, we can conclude that this field seems to be mature with a high number of published proposals and an increasing number of citations per year. Nevertheless, many of these proposals are extensions of classical algorithms (such as Apriori, etc.) and they do not include an experimental analysis regarding the existing proposals in the field. In view of the proposals studied, we recommend the use of more advanced algorithms from the specialized literature and the carrying out of experimental studies comparing the results obtained to the classical algorithms and the most recent ones in the field by making use of several benchmark databases and statistical tests whenever possible. Moreover, we recommend the authors to share the source code of their proposals since this would have a very positive impact, both on the development of better algorithms and on their applicability to new subjects and the industry.

Aside from the classical WoS categories Computer Science Artificial Intelligence, Computer Science Theory Methods, Computer Science Information Systems and the category Engineering Electrical Electronic, we can also observe a recent increase in the number of publications associated with several categories for specific applications, such as Telecommunications, Medical-Informatics, Remote-Sensing, etc. This evidences the great capacity of TARs to be applied successfully to a wide variety of problems and the interest that TARs has awakened in recent years with regard to their capacity to solve real-world problems. Finally, we have identified some interesting further directions, such as the development of new proposals for solving temporal big data problems making use of different approaches, the use of other dimensions in the process of knowledge extraction (such as space) and the extraction of High-Utility TARs that maximize some unit profit.

References

Agrawal, R., Imielinski, T., & Swami, A. (1993). Mining association rules between sets of items in large databases. In *Sigmod* (p. 207-216). Washington D.C., USA.

Agrawal, R., & Srikant, R. (1995). Mining sequential patterns. In *11th international conference on data engineering (ICDE 1995)* (p. 3-14). Washington, DC, USA.

Ale, J., & Rossi, G. (2000). An approach to discovering temporal association rules. In *ACM symposium on applied computing (SAC 2000)* (Vol. 1, p. 294-300). Como, Italy.

Aljawarneh, S. A., Radhakrishna, V., Kumar, P. V., & Janaki, V. (2017). G-SPAMINE: an approach to discover temporal association patterns and trends in internet of things. *Future Generation Computer Systems*, *74*, 430-443.

Allen, J. F. (1983). Maintaining knowledge about temporal intervals. *Communications of the ACM*, *26*(11), 832-843.

Ao, X., Luo, P., Wang, J., Zhuang, F., & He, Q. (2018). Mining precise-positioning episode rules from event sequences. *IEEE Transactions on Knowledge and Data Engineering*, *30*(3), 530-543.

Arnaud, M., Begaud, B., Thurin, N., Moore, N., Pariente, A., & Salvo, F. (2017). Methods for safety signal detection in healthcare databases: a literature review. *Expert Opinion on Drug Safety*, *16*(6), 721-732.

Au, W.-H., & Chan, K. C. (2002). Fuzzy data mining for discovering changes in association rules over time. In *IEEE international conference on fuzzy systems (FUZZ-IEEE 2002)* (p. 890-895). Honolulu, HI, USA.

Ayres, J., Flannick, J., Gehrke, J., & Yiu, T. (2002). Sequential pattern mining using a bitmap representation. In *International conference on knowledge discovery and data mining (ACM SIGKDD 2002)* (p. 429-435). Edmonton, AB, Canada.

Batagelj, V., & Cerinšek, M. (2013). On bibliographic networks. *Scientometrics*, *96*(3), 845-864.

Berthold, M. R., Cebron, N., Dill, F., Gabriel, T. R., Kötter, T., Meinl, T., ... Wiswedel, B. (2009). KNIME–the konstanz information miner: Version 2.0 and beyond. *ACM SIGKDD Explorations Newsletter*, *11*(1), 26-31. Retrieved from http://www.knime

.com/

Berzal, F., Blanco, I., Sánchez, D., & Vila, M. A. (2002). Measuring the accuracy and interest of association rules: A new framework. *Intelligent Data Analysis*, *6*(3), 221-235.

Bettini, C., Wang, X. S., & Jajodia, S. (1998). Mining temporal relationships with multiple granularities in time sequences. *IEEE Data Engineering Bulletin*, *21*, 32-38.

Boettcher, M. (2011). Contrast and change mining. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, *1*(3), 215-230.

Bohlen, M., Busatto, R., & Jensen, C. (1998). Point-versus interval-based temporal data models. In *Proceedings 14th international conference on data engineering* (p. 192-200). Orlando, USA.

Bou Rjeily, C., Badr, G., Hajjarm El Hassani, A., & Andres, E. (2019). Medical data mining for heart diseases and the future of sequential mining in medical field. In G. A. Tsihrintzis, D. N. Sotiropoulos, & L. C. Jain (Eds.), *Machine learning paradigms: Advances in data analytics* (p. 71-99). Cham: Springer International Publishing.

Buxton, E., Vohra, S., Guo, Y., Fogleman, A., & Patel, R. (2019). Pediatric population health analysis of southern and central illinois region: A cross sectional retrospective study using association rule mining and multiple logistic regression. *Computer Methods and Programs in Biomedicine*, *178*, 145-153.

Callon, M., Courtial, J., & Laville, F. (1991). Co-word analysis as a tool for describing the network of interactions between basic and technological research: The case of polymer chemsitry. *Scientometrics*, *22*, 155-205.

Cariñena, P. (2014). Fuzzy temporal association rules: combining temporal and quantitative data to increase rule expressiveness. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, *4*(1), 64-70.

Chamazi, M., & Motameni, H. (2019). Finding suitable membership functions for fuzzy temporal mining problems using fuzzy temporal bees method. *Soft Computing*, *23*(10), 3501-3518.

Chen, C.-H., Lan, G.-C., Hong, T.-P., & Lin, S.-B. (2016). Mining fuzzy temporal association rules by item lifespans. *Applied Soft Computing*, *41*, 265-274.

Chen, X., & Petrounias, I. (1999). Mining temporal features in association rules. In *European* *conference on principles of data mining and knowledge discovery (PKDD 1999)* (p. 295-300). Prague, Czech Republic.

Cobo, M., López-Herrera, A., Herrera-Viedma, E., & Herrera, F. (2011). An approach for detecting, quantifying, and visualizing the evolution of a research field: A practical application to the fuzzy sets theory field. *Journal of Informetrics*, *5*(1), 146-166.

Cobo, M. J., López-Herrera, A. G., Herrera-Viedma, E., & Herrera, F. (2012). SciMAT: A new science mapping analysis software tool. *Journal of the American Society for Information Science and Technology*, *63*(8), 1609-1630.

Concaro, S., Sacchi, L., Cerra, C., Fratino, P., & Bellazzi, R. (2011). Mining health care administrative data with temporal association rules on hybrid events. *Methods of Information in Medicine*, *50*(2), 166-179.

Dalmas, B., Fournier-Viger, P., & Norre, S. (2017). Twincle : A constrained sequential rule mining algorithm for event logs. *Procedia Computer Science*, *112*, 205-214.

Das, A., Ahmed, M., & Ghasemzadeh, A. (2019). Using trajectory-level shrp2 naturalistic driving data for investigating driver lane-keeping ability in fog: An association rules mining approach. *Accident Analysis and Prevention*, *129*, 250-262.

Das, G., Lin, K.-L., Mannila, H., Renganathan, G., & Smyth, P. (1998). Rule discovery from time series. In *International conference on knowledge discovery and data mining (ACM SIGKDD 1998)* (p. 16-22). New York, USA: ACM Pres.

Degirmenci, T., & ozbakir, L. (2018). Differentiating households to analyze consumption patterns: A data mining study on official household budget data. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, *8*(1), 1-15.

Deogun, J., & Jiang, L. (2005). Prediction mining – an approach to mining association rules for prediction. In D. Slezak, J. Yao, J. F. Peters, W. Ziarko, & X. Hu (Eds.), *Rough sets, fuzzy sets, data mining, and granular computing* (p. 98-108). Berlin, Heidelberg: Springer Berlin Heidelberg.

Febrer-Hernández, J. K., & Hernández-Palancar, J. (2012). Sequential pattern mining algorithms review. *Intelligent Data Analysis*, *16*(3), 451-466.

Feng, L., Dillon, T., & Liu, J. (2001). Inter-transactional association rules for multidimensional contexts for prediction and their applications to studying meteological data.

Data and Knowledge Engineering, *37*, 85-115.

Fouad, M., & Mostafa, M. (2017). IndxTAR: An efficient algorithm for indexed mining of incremental temporal association rules. *International Journal of Computer Information Systems and Industrial Management Applications*, *9*, 103-113.

Fournier-Viger, P., Wu, C., Tseng, V. S., Cao, L., & Nkambou, R. (2015). Mining partially-ordered sequential rules common to multiple sequences. *IEEE Transactions on Knowledge and Data Engineering*, *27*(8), 2203-2216.

Fournier-Viger, P., Faghihi, U., Nkambou, R., & Nguifo, E. M. (2012). Cmrules: Mining sequential rules common to several sequences. *Knowledge-Based Systems*, *25*(1), 63-76.

Fournier-Viger, P., Gomariz, A., Campos, M., & Thomas, R. (2014). Fast vertical mining of sequential patterns using co-occurrence information. In V. S. Tseng, T. B. Ho, Z.-H. Zhou, A. L. P. Chen, & H.-Y. Kao (Eds.), *Advances in knowledge discovery and data mining* (pp. 40–52). Springer International Publishing.

Fournier-Viger, P., Gomariz, A., Gueniche, T., Soltani, A., Wu, C.-W., & Tseng, V. S. (2014). SPMF: A java open-source pattern mining library. *Journal of Machine Learning Research*, *15*, 3569-3573. Retrieved from http://jmlr.org/papers/v15/fournierviger14a.html

Fournier-Viger, P., Gueniche, T., Zida, S., & Tseng, V. S. (2014). ERMiner: Sequential rule mining using equivalence classes. In H. Blockeel, M. van Leeuwen, & V. Vinciotti (Eds.), *Advances in intelligent data analysis XIII* (p. 108-119). Cham: Springer International Publishing.

Fournier-Viger, P., Lin, J.-W., Kiran, R., Koh, Y., & Thomas, R. (2017). A survey of sequential pattern mining. *Data Science and Pattern Recognition*, *1*(1), 54-77.

Fournier-Viger, P., Lin, J.-W., Vo, B., Chi, T., Zhang, J., & Le, H. (2017). A survey of itemset mining. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, *7*(4), 1-18.

Fournier-Viger, P., Wu, C.-W., Tseng, V. S., & Nkambou, R. (2012). Mining sequential rules common to several sequences with the window size constraint. In L. Kosseim & D. Inkpen (Eds.), *Advances in artificial intelligence* (p. 299-304). Berlin, Heidelberg:

Springer Berlin Heidelberg.

Fournier-Viger, P., Yang, P., Lin, J., Duong, Q., Dam, T.-L., Frnda, J., … Voznak, M. (2019). Discovering periodic itemsets using novel periodicity measures. *Advances in Electrical and Electronic Engineering*, *17*(1), 33-44.

Geng, L., & Hamilton, H. J. (2006). Interestingness measures for data mining: A survey. *ACM Computing Surveys*, *38*(3).

Gharib, T. F., Nassar, H., Taha, M., & Abraham, A. (2010). An efficient algorithm for incremental mining of temporal association rules. *Data and Knowledge Engineering*, *69*(8), 800-815.

Giannoulis, M., Kondylakis, H., & Marakakis, E. (2019). Designing and implementing a collaborative health knowledge system. *Expert Systems with Applications*, *126*, 277-294.

Guillame-Bert, M., Dubrawski, A., Wang, D., Hravnak, M., Clermont, G., & Pinsky, M. (2017). Learning temporal rules to forecast instability in continuously monitored patients. *Journal of the American Medical Informatics Association*, *24*(1), 47-53.

Han, J., Dong, G., & Yin, Y. (1999). Efficient mining of partial periodic patterns in time series database. In *International conference on data engineering (ICDE 1999)* (p. 106-115). Sydney, Australia.

Han, J., Gong, W., & Yin, Y. (1998). Mining segment-wise periodic patterns in time-related databases. In *International conference on knowledge discovery and data mining (KDDM 1998)* (p. 214-218). New York, USA.

Han, J., & Kamber, M. (2006). *Data mining: Concepts and techniques* (2nd. ed.). USA: Morgan Kaufmann.

Harms, S., & Deogun, J. (2004). Sequential association rule mining with time lags. *Journal of Intelligent Information Systems*, *22*(1), 7-22.

Heydari, J., & Tajer, A. (2019). Quickest search and learning over correlated sequences: Theory and application. *IEEE Transactions on Signal Processing*, *67*(3), 638-651.

Hong, T.-P., Lan, G.-C., Su, J.-H., Wu, P.-S., & Wang, S.-L. (2016). Discovery of temporal association rules with hierarchical granular framework. *Applied Computing and Informatics*, *12*(2), 134-141.

Höppner, F. (2001). Discovery of temporal patterns. In *European conference on principles of data mining and knowledge discovery (PKDD 2001). Lecture notes in computer science* (Vol. 2168, p. 192-203). Freiburg, Germany: Springer Berlin Heidelberg.

Höppner, F., & Klawonn, F. (2001). Finding informative rules in interval sequences. In *Advances in intelligent data analysis. Lecture notes in computer science* (Vol. 2189, p. 125-134). Cascais, Portugal.

Hsieh, Y.-L., Yang, D.-L., Wu, J., & Chen, Y.-C. (2016). Efficient mining of profit rules from closed inter-transaction itemsets. *Journal of Information Science and Engineering*, *32*(3), 575-595.

Huang, F., Zou, Z., Liu, X., & He, J. (2012). Association rules mining for academic cooperation based on time extension and duration accumulation. In *International conference on computer science and service system* (p. 2007-2012). Nanjing, China.

Huang, J.-W., Dai, B.-R., & Chen, M.-S. (2007). Twain: Two-end association miner with precise frequent exhibition periods. *ACM Transactions on Knowledge Discovery from Data*, *1*(2), 1-33.

Huang, T.-K., Huang, C.-H., & Chuang, Y.-T. (2016). Change discovery of learning performance in dynamic educational environments. *Telematics and Informatics*, *33*(3), 773-792.

Huang, Y.-P., & Kao, L.-J. (2005). A novel approach to mining inter-transaction fuzzy association rules from stock price variation data. In *14th IEEE international conference on fuzzy systems* (p. 791-796). Reno, USA.

Huang, Y.-P., Kao, L.-J., & Sandnes, F. (2007). Data mining and fuzzy inference based salinity and temperature variation prediction. In *IEEE international conference on systems, man and cybernetics* (p. 2074-2079). Montreal, Canada.

Ishibuchi, H., Nakashima, T., & Nii, M. (2004). *Classification and modeling with linguistic information granules: Advanced approaches to linguistic data mining*. Berlin: Springer-Verlag.

Ji, Y., Ying, H., Tran, J., Dews, P., Lau, S.-Y., & Massanari, R. (2016). A functional temporal association mining approach for screening potential drug-drug interactions from electronic patient databases. *Informatics for Health and Social Care*, *41*(4), 387-

404.

Khan, S., & Parkinson, S. (2018). Eliciting and utilising knowledge for security event log analysis: An association rule mining and automated planning approach. *Expert Systems with Applications*, *113*, 116-127.

Kilgore, P. C. S. R., Korneeva, N., Arnold, T. C., Trutschl, M., & Cvek, U. (2019). Gatewaynet: a form of sequential rule mining. *BMC Medical Informatics and Decision Making*, *19*(1), 87.

Laxman, S., Sastry, P., & Unnikrishnan, K. (2007). Discovering frequent generalized episodes when events persist for different durations. *IEEE Transactions on Knowledge and Data Engineering*, *19*(9), 1188-1201.

Laxman, S., Sastry, P. S., & Unnikrishnan, K. P. (2005). Discovering frequent episodes and learning hidden markov models: a formal connection. *IEEE Transactions on Knowledge and Data Engineering*, *17*(11), 1505-1517.

Laxman, S., Unnikrishnan, K. P., & Sastry, P. S. (2002). Generalized frequent episodes in event sequences. In *International conference on knowledge discovery and data mining, workshop on temporal data mining (ACM SIGKDD 2002)* (p. 1-7). Edmonton, Alberta, Canada.

Lee, C.-H., Chen, M.-S., & Lin, C.-R. (2003). Progressive partition miner: an efficient algorithm for mining general temporal association rules. *IEEE Transactions on Knowledge and Data Engineering*, *15*(4), 1004-1017.

Lee, C.-H., Lin, C.-R., & Chen, M.-S. (2001). Sliding-window filtering: An efficient algorithm for incremental mining. In *Proceedings of the tenth international conference on information and knowledge management (CIKM 2001)* (p. 263-270). Atlanta, Georgia, USA.

Lee, W.-J., Jiang, J.-Y., & Lee, S.-J. (2008). Mining fuzzy periodic association rules. *Data and Knowledge Engineering*, *65*(3), 442-462.

Lee, W.-J., & Lee, S.-J. (2004). Discovery of fuzzy temporal association rules. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, *34*(6), 2330-2342.

Lee, Y. J., Lee, J. W., Chai, D. J., Hwang, B. H., & Ryu, K. H. (2009). Mining temporal interval relational rules from temporal data. *Journal of Systems and Software*, *82*(1), 155-167.

Li, D., & Deogun, J. S. (2005). Discovering partial periodic sequential association rules with time lag in multiple sequences for prediction. In *International symposium on methodologies for intelligent systems* (p. 332-341). NY, USA.

Li, Q., Feng, L., & Wong, A. (2005). From intra-transaction to generalized inter-transaction: Landscaping multidimensional contexts in association rule mining. *Information Sciences*, *172*(3), 361-395.

Li, Y., Ning, P., Wang, X. S., & Jajodia, S. (2003). Discovering calendar-based temporal association rules. *Data and Knowledge Engineering*, *44*(2), 193-218.

Lin, M.-Y., & Lee, S.-Y. (2002). Fast discovery of sequential patterns by memory indexing. In Y. Kambayashi, W. Winiwarter, & M. Arikawa (Eds.), *Data warehousing and knowledge discovery* (p. 150-160). Berlin, Heidelberg: Springer Berlin Heidelberg.

Lin, W., Orgun, M., & Williams, G. (2002). An overview of temporal data mining. In *Australian data mining workshop (ADM 2002)* (p. 83-90). Canberra, Australia.

Liu, C., Zhang, Q., Luo, H., Qi, S., Tao, S., Xu, H., & Yao, Y. (2019). An efficient approach to capture continuous impervious surface dynamics using spatial-temporal rules and dense landsat time series stacks. *Remote Sensing of Environment*, *229*, 114-132.

Liu, L., Wang, S., Peng, Y., Huang, Z., Liu, M., & Hu, B. (2016). Mining intricate temporal rules for recognizing complex activities of daily living under uncertainty. *Pattern Recognition*, *60*, 1015-1028.

Lo, D., Khoo, S.-C., & Wong, L. (2009). Non-redundant sequential rules - theory and algorithm. *Information Systems*, *34*(4-5), 438-453.

Lo, D., Ramalingam, G., Ranganath, V.-P., & Vaswani, K. (2012). Mining quantified temporal rules: Formalism, algorithms, and evaluation. *Science of Computer Programming*, *77*(6), 743-759.

Lu, H., Feng, L., & Han, J. (2000). Beyond intratransaction association analysis: Mining multidimensional intertransaction association rules. *ACM Transactions on Information Systems*, *18*(4), 423-454.

Lyubimov, D., & Palumbo, A. (2016). *Apache mahout: Beyond mapreduce* (1st ed.).

USA: CreateSpace Independent Publishing Platform. Retrieved from https://mahout.apache.org

Mannila, H., Toivonen, H., & Verkamo, A. I. (1997). Discovery of frequent episodes in event sequences. *Data Mining and Knowledge Discovery*, *1*(3), 259-289.

Matthews, S. G., Gongora, M. A., & Hopgood, A. A. (2011). Evolving temporal association rules with genetic algorithms. In M. Bramer, M. Petridis, & A. Hopgood (Eds.), *Research and development in intelligent systems XXVII* (p. 107-120). London: Springer.

Matthews, S. G., Gongora, M. A., & Hopgood, A. A. (2013). Evolutionary algorithms and fuzzy sets for discovering temporal rules. *International Journal of Applied Mathematics and Computer Science*, *23*(4), 855-868.

Miholca, D.-L., Czibula, G., & Czibula, I. (2018). A novel approach for software defect prediction through hybridizing gradual relational association rules with artificial neural networks. *Information Sciences*, *441*, 152-170.

Mooney, C. H., & Roddick, J. F. (2013). Sequential pattern mining – approaches and algorithms. *ACM Computing Surveys*, *45*(2), 1-39.

Moskovitch, R., Elovici, Y., & Rokach, L. (2008). Detection of unknown computer worms based on behavioral classification of the host. *Computational Statistics & Data Analysis*, *52*(9), 4544-4566.

Moskovitch, R., & Shahar, Y. (2015, 01). Fast time intervals mining using the transitivity of temporal relations. *Knowledge and Information Systems*, *42*(1), 21-48.

Moya-Anegón, F., Vargas-Quesada, B., Herrero-Solana, V., Chinchilla-Rodríguez, Z., Corera-Álvarez, E., & Muñoz-Fernández, F. J. (2004). A new technique for building maps of large scientific domains based on the cocitation of classes and categories. *Scientometrics*, *61*(1), 129-145.

Nam, H., KiYoung, L., & Doheon, L. (2009). Identification of temporal association rules from time-series microarray data sets. *BMC Bioinformatics*, *10*(3), 1-9.

Nazerfard, E., Rashidi, P., & Cook, D. (2011). Using association rule mining to discover temporal relations of daily activities. In *International conference on smart homes and health telematics (ICOST 2011)* (p. 49-56). Montreal, Canada.

Nguyen, D., Luo, W., Phung, D., & Venkatesh, S. (2018). LTARM: A novel temporal association rule mining method to understand toxicities in a routine cancer treatment. *Knowledge-Based Systems*, *161*, 313-328.

Open source initative. 1998. (n.d.). http://www.opensource.org/docs/osd.

Orphanou, K., Dagliati, A., Sacchi, L., Stassopoulou, A., Keravnou, E., & Bellazzi, R. (2018). Incorporating repeating temporal association rules in naive bayes classifiers for coronary heart disease diagnosis. *Journal of Biomedical Informatics*, *81*, 74-82.

Ozden, B., Ramaswamy, S., & Silberschatz, A. (1998). Cyclic association rules. In *International conference on data engineering (ICDE 1998)* (p. 412-421). Orlando, Florida, USA.

Panchal, M. C., & Prajapati, G. I. (2018). Calendric association rule mining from time series database. In K. Saeed, N. Chaki, B. Pati, S. Bakshi, & D. P. Mohapatra (Eds.), *Progress in advanced computing and intelligent engineering* (pp. 283–293). Springer Singapore.

Papapetrou, P., Kollios, G., Sclaroff, S., & Gunopulos, D. (2009). Mining frequent arrangements of temporal intervals. *Knowledge and Information Systems*, *21*(2), 133-171.

Radhakrishna, V., Aljawarneh, S. A., Kumar, P. V., & Janaki, V. (2019). ASTRA - A novel interest measure for unearthing latent temporal associations and trends through extending basic gaussian membership function. *Multimedia Tools and Applications*, *78*(4), 4217-4265.

Radhakrishna, V., Kumar, P., & Janaki, V. (2015). An approach for mining similarity profiled temporal association patterns using gaussian based dissimilarity measure. In *Proceedings of the the international conference on engineering & mis (icemis 2015)* (p. 1-6). Istanbul, Turkey.

Ramaswamy, S., Mahajan, S., & Silberschatz, A. (1998). On the discovery of interesting patterns in association rules. In *International conference on very large data bases (VLDB 1998)* (p. 368-379). San Francisco, CA, USA.

Roddick, J., & Mooney, C. (2005). Linear temporal sequences and their interpretation using midpoint relationships. *IEEE Transactions on Knowledge and Data Engineering*, *17*(1), 133-135.

Roddick, J., & Spiliopoulou, M. (2002). A survey of temporal knowledge discovery paradigms and methods. *IEEE Transactions on Knowledge and Data Engineering*, *14*(4), 750-767.

Saleh, B., & Masseglia, F. (2011). Discovering frequent behaviors: time is an essential element of the context. *Knowledge and Information Systems*, *28*(2), 311-331.

Setiawan, F., & Yahya, B. N. (2018). Improved behavior model based on sequential rule mining. *Applied Soft Computing*, *68*, 944-960.

Song, H., Kim, J., & Kim, S. (2001). Mining the change of customer behavior in an internet shopping mall. *Expert Systems with Applications*, *21*(3), 157-168.

Sudkamp, T. (2005). Discovery of fuzzy temporal associations in multiple data streams. In *Soft computing: Methodologies and applications* (p. 3-13). Springer.

Tan, P., Kumar, V., & Srivastava, J. (2002). Selecting the right interestingness measure for association patterns. In *8th international conference on knowledge discovery and data mining (KDD 2002)* (p. 32-41). Edmonton, Canada.

Tebourski, W., Abdessalem Kar, W., & Ghezala, H. (2015). From concrete to inferred knowledge: Enhanced mining constraint-based cyclic association rules from medical social network. *International Journal of Knowledge-Based and Intelligent Engineering Systems*, *19*(2), 109-116.

Tebourski, W., & Karaa, W. B. A. (2012). Cyclic association rules mining under constraints. *International Journal of Computer Applications*, *49*(20), 30-37.

Teng, H. S., Chen, K., & Lu, S. C. (1990). Adaptive real-time anomaly detection using inductively generated sequential patterns. In *IEEE computer society symposium on research in security and privacy* (p. 278-284). CA, USA.

Thimbleby, H. (2003). Explaining code for publication. *Software - Practice and Experience*, *33*, 975-1001.

Thuan, N. D., Toan, N. G., & Tuan, N. L. V. (2012). An approach mining cyclic association rules in e-commerce. In *International conference on network-based information systems* (p. 408-411). Melbourne, Australia.

Tran, M.-T., Le, B., Vo, B., & Hong, T.-P. (2016). Mining non-redundant sequential rules with dynamic bit vectors and pruning techniques. *Applied Intelligence*, *45*(2), 333-342.

Tung, A. K. H., Lu, H., Han, J., & Feng, L. (2003). Efficient mining of intertransaction association rules. *IEEE Transactions on Knowledge and Data Engineering*, *15*(1), 43-56.

Vandromme, M., Jacques, J., Taillard, J., Hansske, A., Jourdan, L., & Dhaenens, C. (2017). Extraction and optimization of classification rules for temporal sequences: Application to hospital data. *Knowledge-Based Systems*, *122*, 148-158.

Verma, K., & Vyas, O. P. (2005). Efficient calendar based temporal association rule. *ACM SIGMOD Record*, *34*(3), 63-70.

Verma, K., Vyas, O. P., & Vyas, R. (2005). Temporal approach to association rule mining using t-tree and p-tree. In *International workshop on machine learning and data mining in pattern recognition* (p. 651-659). Leipzig, Germany.

Vu, H., Li, G., Law, R., & Zhang, Y. (2018). Travel diaries analysis by sequential rule mining. *Journal of Travel Research*, *57*(3), 399-413.

Wang, C.-S. (2015). Mining non-redundant inter-transaction rules. *Journal of Information Science and Engineering*, *31*(6), 1849-1865.

Wang, L., Meng, J., Xu, P., & Peng, K. (2018). Mining temporal association rules with frequent itemsets tree. *Applied Soft Computing Journal*, *62*, 817-829.

Wen, F., Zhang, G., Sun, L., Wang, X., & Xu, X. (2019). A hybrid temporal association rules mining method for traffic congestion prediction. *Computers & Industrial Engineering*, *130*, 779-787.

Winarko, E., & Roddick, J. (2007). Armada - an algorithm for discovering richer relative temporal association rules from interval-based data. *Data and Knowledge Engineering*, *63*(1), 76-90.

Witten, I., Frank, E., & Hall, M. (2016). *The WEKA workbench. online appendix for data mining: Practical machine learning tools and techniques.* Morgan Kaufmann *fourthedition*. Retrieved from http://www.cs.waikato.ac.nz/ml/weka/

Wu, R. (2010). Mining generalized fuzzy association rules from web logs. In *International conference on fuzzy systems and knowledge discovery (FSKD 2010)* (p. 2474-2477). Yantai, China.

Xiao, Q., Li, C., Tang, Y., Li, L., & Li, L. (2019). A knowledge-driven method of adaptively optimizing process parameters for energy efficient turning. *Energy*, 142-156.

Xiao, Y., Tian, Y., & Zhao, Q. (2014). Optimizing frequent time-window selection for association rules mining in a temporal database using a variable neighbourhood search. *Computers & Operations Research*, *52*, 241-250.

Xie, D., Wang, M.-H., & Zhao, X.-M. (2020). A spatiotemporal apriori approach to capture dynamic associations of regional traffic congestion. *IEEE Access*, *8*, 3695-3709.

Zhang, C., & Zhang, S. (2002). *Association rule mining: Models and algorithms.* Springer.

Zhou, H., & Hirasawa, K. (2017). Evolving temporal association rules in recommender system. *Neural Computing and Applications*, 1-15.

Zida, S., Fournier-Viger, P., Wu, C.-W., Lin, J. C.-W., & Tseng, V. S. (2015). Efficient mining of high-utility sequential rules. In P. Perner (Ed.), *Machine learning and data mining in pattern recognition* (p. 157-171). Cham: Springer International Publishing.

Zimmermann, A. (2014). Understanding episode mining techniques: Benchmarking on diverse, realistic, artificial data. *Intelligent Data Analysis*, *18*(5), 761-791.