



## Article The Semantic Adjacency Criterion in Time Intervals Mining

Alexander Shknevsky, Yuval Shahar \* D and Robert Moskovitch

Department of Software and Information Systems Engineering, Ben-Gurion University, Beer-Sheva 8410501, Israel; sheknabs@bgu.ac.il (A.S.); robertmo@bgu.ac.il (R.M.) \* Correspondence: yshahar@bgu.ac.il

**Abstract:** We propose a new pruning constraint when mining frequent temporal patterns to be used as classification and prediction features, the *Semantic Adjacency Criterion* [*SAC*], which filters out temporal patterns that contain potentially semantically contradictory components, exploiting each medical domain's knowledge. We have defined three SAC versions and tested them within three medical domains (oncology, hepatitis, diabetes) and a frequent-temporal-pattern discovery framework. Previously, we had shown that using SAC enhances the repeatability of discovering the same temporal patterns in similar proportions in different patient groups within the same clinical domain. Here, we focused on SAC's computational implications for pattern discovery, and for classification and prediction, using the discovered patterns as features, by four different machine-learning methods: Random Forests, Naïve Bayes, SVM, and Logistic Regression. Using SAC resulted in a significant reduction, across all medical domains and classification methods, of up to 97% in the *number* of discovered temporal patterns, and in the *runtime* of the discovery process, of up to 98%. Nevertheless, *the highly reduced set of only semantically transparent patterns*, when used as features, resulted in *classification and prediction models* whose performance was *at least as good* as the models resulting from using the complete temporal-pattern set.

**Keywords:** temporal data mining; machine learning; time intervals mining; semantics; frequent temporal pattern mining; classification; prediction; medicine

## 1. Introduction

This paper deals with the increasingly important topic of the discovery of frequent temporal patterns when given as input a set of *symbolic time intervals*, i.e., time periods over which hold one or more propositions, such as, in the medical domain, "The dose of the medication was High" or "The blood pressure was Low", and the temporal relationships among these periods. The discovered temporal patterns can then be exploited for clustering, classification, and prediction.

Analyzing time-oriented, multivariate clinical data enables researchers to discover new temporal knowledge and gain understanding regarding the temporal behavior and temporal associations of these data [1–9]. The main methods for the discovery of new knowledge in longitudinal multivariate data include multiple *Temporal Data Mining* (TDM) approaches, although an alternative, *Business Process Mining* (BPM) approach has also been used successfully when the data describe actual activities and processes; in such cases, temporal relations at a finer level of resolution are often less emphasized [10–13]. Unlike most TDM methods, which typically focus mainly on the analysis of the *raw timestamped* data, the use of *symbolic time intervals* can reduce inherent random noise in the data, avoid problems resulting from different sampling frequencies and at various temporal granularities, and often alleviate the problem of missing data [3,7,9,14,15]. Thus, to significantly enhance the capabilities for analysis of time-stamped data, a preprocessing step of meaningful summarization and interpretation of the time-stamped raw data (e.g., a series of hemoglobin values) into a set of interval-based abstractions or symbolic time intervals (e.g., periods of moderate anemia), known as *temporal abstractions*, can be



Citation: Shknevsky, A.; Shahar, Y.; Moskovitch, R. The Semantic Adjacency Criterion in Time Intervals Mining. *Big Data Cogn. Comput.* **2023**, 7, 173. https://doi.org/10.3390/ bdcc7040173

Academic Editor: Carson K. Leung

Received: 13 September 2023 Revised: 6 November 2023 Accepted: 7 November 2023 Published: 9 November 2023



**Copyright:** © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https:// creativecommons.org/licenses/by/ 4.0/). used [16–20]. The resulting interval-based summary can have multiple uses, such as to create natural-language free-text summaries (e.g., discharge letters) of large numbers of digital time-oriented clinical data [21], to visualize the data of individual patients [22], or to interactively explore associations among the time-oriented data and their abstractions [2].

Once a set of interval-based abstractions of the time-stamped raw data exists, a set of [sufficiently] frequent temporal patterns, incorporating these symbolic time intervals as components, can be discovered. We refer to these patterns as *time-interval-relation patterns* (*TIRPs*) [3]. Within TIRPs, all of Allen's seven basic temporal relations and their respective inverse relations [23] might hold, such as Before, During, Overlaps, Finishes, etc.

Recently, time-interval patterns have been increasingly used as features to classify multivariate temporal data [1,20,24–27]. Using that approach, the [sufficiently] frequent interval-based temporal patterns are used as the base features to induce a classifier. Furthermore, we have shown in an earlier study that frequent TIRPs can be consistently discovered, and in similar proportions, in different subsets of the same data set within three different medical domains, especially as the minimal threshold for frequency is raised, thus increasing their value for potential classification and prediction tasks [28]. This repeated discovery suggests that discovered TIRPs might indeed be good candidates for use as classification or prediction features.

However, our work in multiple clinical domains had suggested that many of the discovered frequent temporal patterns, although correct from the purely *syntactic* aspect, do not conform to the basic *semantics* of medical experts, who often assume a certain type of temporal adjacency among the temporal-pattern's components. Thus, such patterns are not *transparent* characterizations of the data. This semantic temporal adjacency, which domain experts seem to implicitly assume, means that no instance of the pattern includes any [additional] intermediate intervals within the scope of the pattern, which might contradict a potential interpretation of causality, or at least direct temporal association, among the pattern's components. Our new principle injects *semantics* into what are usually purely *syntactic* algorithms for discovering frequent temporal patterns in large data sets and demonstrates their effectiveness for classification and prediction. For example, this can be achieved by pruning away most of the potential machine-learning temporal features, without losing any classification or prediction accuracy.

An example of the idea at the core of our new semantic principle, which exploits also the semantics of the symbolic interval-based predicates and not just their temporal relations, is the discovery of the following frequent temporal pattern: <"A period during which a High dose of the medication" occurs *Before* "A period during which the Hemoglobin level was Low"> (instance No. 1 in Figure 1). A domain expert might assume that perhaps there is a causal association between the two symbolic intervals since they frequently seem to be found together in this specific order.

However, what if, after examining the symbolic intervals abstracted from the raw data, the expert finds that between these two symbolic intervals, there often exists, within the patient's original longitudinal record, an additional interval, during which a *High* level of Hemoglobin exists (Instance No. 2a in Figure 1). Alternatively, what if this expert finds one or more instances of the pattern in which, between the two components defining it, there exists in the patient's record an interval during which dose of the medication was actually *Low* (instance No. 2b in Figure 1)?

Although technically the original temporal relation still holds, its significance now, from a medical expert's point of view, might change considerably. That would be the case whether the discovered frequent pattern is used for human explanatory purposes, or for succinctly summarizing the data, or for machine-learning purposes.





**Figure 1.** An example of our new Semantic Adjacent Constraint. Three syntactically equal instances of the same interval-based temporal pattern, which includes the symbolic intervals, "<Medication-dose-level = High> Before <Hemoglobin [HGB]-level = Low>" are shown. Instance No. 1 describes a situation in which the two intervals are adjacent, and no contradicting value exists between them, and thus preserves semantic transparency. Instance No. 2a describes a situation in which the two intervals adjacent, since there is an unexpected (from the point of view of the domain expert) High hemoglobin-level value between them that contradicts the pattern's expected semantics. Instance No. 2b, similarly, contains an unexpected medication-dose level (Low) between the two symbolic intervals. Both of the instances of the non-SAC obeying patterns will be pruned out.

As we formally define in the Methods Section, we distinguish a TIRP (an abstract pattern with certain temporal qualitative constraints) from its TIRP *Instances*, which are found in the longitudinal records that are being analyzed. A TIRP is *frequent* if the proportion of the records within which its TIRP instances are discovered is higher than some threshold. However, within the temporal scope of some of the instances of frequent TIRPs, there might exist, in some of the patients' longitudinal records, additional symbolic intervals (which are not part of that instance), as shown in Figure 1, which seem to contradict the TIRP's intuitive semantics.

Note that experts, especially medical experts, often expect a meaningful frequent temporal pattern to convey some potential causal relationship, such as a High dose of a medication reducing the level of Hemoglobin, in the case of the temporal pattern depicted in Figure 1, Instance #1 (SAC-obeying TIRP). The fact that the true state of affairs is such that it rules that possibility out, as in the case of Instance #2, since in the patient's record there exists a High-Hemoglobin period after the administration of the medication, would not be expected by a clinician when hearing the description of the frequently discovered pattern as "<Medication-dose-level = High> Before <Hemoglobin [HGB]-level = Low>". Nor would this clinician expect that there might be an additional episode of medication administration, but with a Low dose, before the High Hemoglobin value. Given that description, from the point of view of a clinician, Instance #2 and Instance #3 lack *semantic coherence*.

In the current study, we are *not* exploring the purely *psychological* issue of the potential lack of transparency, to medical experts, of different temporal-pattern semantics. (Although such a lack might considerably reduce, for example, the patterns' explanatory value, their data-summarization value, and the efficacy of the experts in suggesting additional patterns to explore). What we *do* conjecture in this study concerns a purely quantitative *functional* issue: We believe that such "*semantically incoherent*" patterns, beyond being potentially less transparent to *medical experts*, might also be less useful, and perhaps even unnecessary, as classification and prediction features for a *machine-learning process*, precisely due to

their potential lack of semantic coherence. At the same time, discovering such redundant, "semantically incoherent" temporal patterns might require significant effort during the discovery time, as well as during classifier-induction time, without enhancing the accuracy of the resultant classifiers.

Of course, what precisely is and is not *semantically coherent* within a complete multipleinterval temporal pattern needs to be carefully and formally defined, as we do in Section 3. We shall see that in fact several options exist for exploiting the basic semantic intuition demonstrated in Figure 1, depending on which constraints exactly must hold in such a temporal interval triad, so as to comply with the notion of semantic coherence.

Although several of the earlier studies have noticed a potential redundancy during pattern discovery (in particular, the discovery of patterns containing repeating symbolic intervals as components, such as discovery of the pattern AAB in addition to the discovery of the pattern AB), or even considered patterns characterized by the *absence* of certain symbols [29], they have only considered the issue from a purely *computational* point of view (i.e., the complexity of the temporal-pattern discovery process) [1,3]. For example, no attention was paid to the relationship between pattern components denoted by different symbols, each of which represents a different proposition, which in a medical domain's ontology might in fact represent different values of the same concept. For example, both of the symbols "Low blood pressure" and "Hypertension" (i.e., High or Very-High values of the blood pressure) are propositions that assign different values to the same concept, namely, the concept that denotes the abstraction of the raw-data Blood Pressure measurement concept into a discrete symbolic value. In contrast to these studies, in the current study we shall refer to that potential problem from a *semantic* point of view (i.e., the potential *meaning* of the discovered pattern and of each of its components), as well as from a *functional* point of view (i.e., the implications for the *effectiveness* of the classification).

Thus, in the current study, we explored the application of *semantic considerations to symbolic time-intervals mining, and to classification and prediction tasks, in medical domains.* The current study complements and significantly extends into new grounds a previous study of ours in which, in addition to the main methods, we had very briefly introduced, as a secondary method, the SAC principle and had shown, among other related results, that pruning discovered temporal patterns in clinical data using the SAC constraint enhanced the consistency of discovering the same temporal patterns in similar rates in similar patient populations [28]. The current study examines, for the first time, through the use of several machine learning methods, the implications of the use of the SAC principle for a considerable reduction of the number of temporal-pattern features used by these machine-learning methods, without any loss of performance.

As we shall see, the use of domain-specific semantics (which is explained in detail in Section 3.3) can constrain the discovery of temporal patterns in symbolic time intervals data to only those patterns that include certain semantically meaningful relations amongst the symbolic time intervals of which they are composed, in the sense of not violating certain semantic constraints that we have formally defined. Note that *no new temporal patterns are discovered*; rather, a large number of candidate patterns are *pruned [filtered] out* during the discovery process. Thus, our main contribution in this study, beyond introducing, for the first time, a highly detailed and formal definition of the SAC principle and several of its variations, is the rigorous evaluation for classification and prediction purposes of a new pruning constraint for mining time intervals, the *Semantic Adjacency Criterion* (SAC). In fact, we have defined and explored three versions of the SAC criterion.

Consequentially, our core double-pronged hypothesis in this study is that:

- (a) It is more efficient, during the temporal data mining process, to discover only semantically coherent patterns [coherent in a sense that we shall formally define]. But nevertheless,
- (b) Imposing such semantic constraints, leading to the discovery of a significantly smaller set of patterns, will not cause any harm with respect to the performance of the discovered patterns as features for classification or prediction purposes, and might even enhance that performance.

As our current study demonstrates, using any of the SAC versions results in the discovery of temporal patterns whose overall *cardinality*, as well as *time* needed for discovery, are smaller by at least an order of magnitude than the respective resulting cardinality and required running time when not using the SAC constraint; but that whose *value* to classification and prediction tasks that use the discovered patterns as features is at least as good as the original full set of discoverable patterns.

The outline of this paper is the following: Section 2 provides the necessary background for the rest of the paper. Section 3 describes our computational framework, including the use of temporal abstractions, the discovery of TIRPs, and the formal definition of the three SAC versions. Section 3.9 describes the evaluation and the experiments we performed to assess the effect of using the three SAC versions to discover frequent temporal patterns and exploit them as features for classification and prediction purposes, using four different classifier-induction methods within each of three different clinical domains. Section 4 describes the results of our empirical evaluation. Section 5 discusses the results. Section 6 presents the main conclusions of this study.

## 2. Background and Related Work

In this section, we briefly present the background topics that are most relevant for later presenting our methodology in detail, including: Semantics of Symbolic Time Intervals, Time-Intervals Mining, and Classification using Temporal Patterns as features.

## 2.1. The Structure and Semantics of Symbolic Time Intervals

As the reader might have already gleaned from the examples mentioned in the previous section, the *symbols* that may hold on symbolic time intervals usually denote the combination of a *concept* and its *value*. A concept might represent *raw* data (e.g., a Blood-Pressure measurement) or an event (e.g., administration of the medication Insulin). The *values* of such concepts are often numeric, such as "90 mmHg" for a blood pressure measurement, or "2 units" for an Insulin administration. Other raw-data concepts, whose default value is "True", include events such as a total-hip-replacement surgery. A concept might also denote, however, a more abstract interpretation of the raw-data concept, such as the [discretized] level of the Blood Pressure raw-data concept, or the assessment of the dose of administered Insulin. In that case, the respective values might be 'High" or "Low dose" or "Very High".

Symbols that contain Abstract-concepts might also be the result of a *temporal-abstraction* process, in which a series of raw time points were transformed into one or more symbolic time intervals, such as the concept "The trend of the Blood-Pressure measurements" with the value "Decreasing".

Several types of abstract concepts, and in particular temporal abstractions, exist. To be clear and consistent throughout this paper, we use a simple well-known temporal-abstraction ontology, which has been deployed in multiple application domains, which is the ontology used by the *Knowledge-Based Temporal Abstraction* (KBTA) method [16]. *However, both the SAC principle and the results of using it are quite generic, and do not depend on the use of any particular temporal-abstraction ontology, nor on any particular methodology for generating the symbolic time intervals.* 

An abstract concept might denote a *State* abstraction, i.e., a classification of the value(s) of one or more raw-data concepts into a set of values of a single abstract concept, using a set of cutoff values. An example is the abstraction of the Hemoglobin-level value into several states, such as Normal\_Hemoglobin or Moderate\_Anemia. Similarly, the raw-data Height and Weight concepts might be jointly abstracted, using a simple arithmetical function (Weight/Height<sup>2</sup>), into the abstract concept of a *body mass index* (BMI). The BMI concept, in turn, might be further abstracted, using simple cutoff values, into the "BMI state" abstract concept, which might have values such as "underweight", "Normal\_weight", "overweight", or "obese". A state abstraction can be performed using cutoff values that

are provided by a domain expert [16]. Alternatively, the cutoff values might be derived directly from the raw data series [14,15,26].

Temporal abstractions are usually formed within a *context*, or a state of affairs, such as being a male or a female, an infant, or being under the influence of an Insulin injection [Shahar, 1997]. Thus, the knowledge necessary for correctly forming abstractions from raw data is context-sensitive, and the concepts, implicitly or explicitly, might include that context (e.g., "the State of the Hemoglobin-value of a young woman"). For our purposes in the current paper, we shall assume that the context is a part of the concept. Several approaches to the abstraction of time-oriented raw data into symbolic intervals exist. The KBTA methodology uses domain-specific and context-sensitive *classification knowledge* to generate the abstractions from raw concept values and applies *temporal interpolation knowledge* to bridge gaps between time points and time intervals and join them into longer [symbolic] time intervals. It was applied within multiple domains and to different tasks, such as within the domains of medicine, biology, information security, or traffic control, and to the tasks of summarization, visualization, exploration, classification, and prediction [15,22].

However, when no suitable domain knowledge exists, various data driven discretization methods exist [30], which typically focus on finding cut-off values (using various heuristics) for discretizing continuous data. Such methods include, for example, the *Equal Width Discretization* (EWD) method, which partitions the range of values into an equalwidth partitioning, and in practice, often suffices for the purposes of classification or prediction; the SAX method for discretization of time series, introduced by Lin et al. [14]; and the *Temporal Discretization for Classification* (TD4C) method, introduced by Moskovitch and Shahar [20], a discretization method that is specifically geared for the classification task. The TD4C method learns the state-abstraction cutoff values, which best separate the instances belonging to the predicted classes with respect to their differing distribution of abstraction values over time. The inter-distribution distance measure is either Cosine, Entropy, or the Kullback–Leibler measure. The TD4C method outperformed the EWD and SAX discretization methods, for the purpose of classification, in several different medical domains [20].

As we shall now see, one can mine symbolic time intervals to discover frequently occurring temporal patterns (combinations of the symbolic intervals with certain temporal relations among them) in the data of multiple subjects.

#### 2.2. Mining Symbolic Time Intervals

Typically, time interval mining methods use some subset or variation of Allen's temporal relations [23]. Allen defined thirteen temporal relations, based on seven relations (*before, meets, overlaps, finished-by, contains, start-by, equal*) and their inverses (note that the inverse of *equal* is *equal*). Another option, which was investigated also in the current study, is the one of using only three abstract temporal relations, two of which are defined by a disjunction of Allen's relations [3]: BEFORE, which is the disjunction of {*before, meets*}, OVERLAPS, which is the usual *overlaps*, and CONTAINS, which represents the disjunction of {*finished-by, contains, started-by, equal*}.

Höppner introduced a method to mine rules in symbolic time interval sequences using Allen's temporal relations, using a non-ambiguous representation through a conjunction of the pairwise temporal relations among the symbolic time intervals [31]. This time intervals patterns definition was later used to discover patterns more efficiently by several groups [3,24,32,33]. Other researchers used additional abstract relations [17] or other types of temporal relations, such as coinciding [34].

Further improvements to the specific problem of mining *frequent temporal patterns*, which appear above a particular percentage (i.e., support) threshold in the longitudinal records of a collection of entities, have been made. In particular, Lee, Lindgren, and Papapetrou [35] have introduced *Z*-*Miner*, an algorithm for solving this particular problem that employs two data structures: *Z*-*Table*, a hierarchical hash-based data structure for

time-efficient candidate generation and support count, and *Z*-*Arrangement*, a data structure for efficient memory consumption.

Recent progress has been made on the front of examining *time-interval relation patterns* (TIRPs) that include all of the information about TIRPs in the data set, referred to as *closed TIRPs* [36]. An efficient algorithm for discovering closed TIRPs, referred to as TRIPClo, was designed [9]; a recent evaluation has demonstrated the TIRPClo algorithm's considerable computational effectiveness on eleven real-world and four synthetic data sets, and has examined its implications [9].

Note that the task of *interval-based* temporal data mining is quite different from *sequential mining* (see Section 2.3) which, as its name would suggest, is purely sequential, and usually focuses on point-based episodes. The main task in the case of the multiple algorithms for mining patterns based on time intervals is to mine frequent patterns of repeating temporal relations among multiple time intervals. That task includes the determination of temporal relations such as *contains*, *overlaps*, and *finishes*, in addition to the standard *before* and *after* and *equals*.

However, it is important to note that unlike the SAC criterion that we define and exploit in this paper, *none* of the advanced frequent temporal-pattern discovery methods listed here, although exploiting several different temporal-pattern discovery principles [3,7,27,32], also considers the *internal semantics of the temporal patterns' components*. In other words, when discovering a TIRP such as  $T_1 = ((A \text{ before B}) \text{ and } (B \text{ before C}))$ , no attention is usually paid to the issue of the precise *meaning* of A, B, and C, and whether certain meanings make more sense, with respect to the temporal relation in question, than others.

Nevertheless, as shown in Figure 1, the *meaning* of these three symbolic intervals might well drastically change the *interpretation* of finding a frequent TIRP such as T<sub>1</sub>, the decision whether to even include it in the set of candidate TIRPs to be discovered while running an algorithm for frequent-temporal-pattern discovery, or the decision whether to use it as a feature for classification and prediction, as our evaluation of the SAC principle will demonstrate.

## 2.3. Classification and Prediction Based on Temporal Patterns

The field of sequential data mining and its use for TDM tasks has been explored in various ways, e.g., through the sequence classification task [37], or through sequence and motifs mining to extract features for classification [38–40]. However, its focus is on time-stamped events, and essentially only on the before relation, and so is not appropriate for the types of domains and tasks in which we are interested. We are interested in the more general nature of multivariate, time-stamped, and interval-based data. Quite simultaneously, several groups proposed using interval-based TIRPs as features for classifying multivariate time series [1,24,41], which were followed by more recent studies [3,7,26,27,35].

A somewhat different approach to the analysis of time series at a more abstract level is to partition the data into local groupings [42].

Interestingly, all of the above studies reported the use of temporal abstraction and the use of TIRPs for classification applied to biomedical data. Patel et al. proposed using IEClassifier to classify Hepatitis patients using TIRPs [1]. Batal et al. performed knowledge-based temporal abstraction, but used only two relations: before and co-occur, which is a specific case of an a priori sequential mining algorithm called STF-Mine [41]. Several studies had shown the advantages of using TIRPs over atemporal representations in classifying multivariate temporal data [1,24,26,41]. Other studies attempted to introduce several heuristics to decrease the number of discovered patterns that still maintain the same level of accuracy [25].

Moskovitch and Shahar presented KarmaLegoSification—a framework for classification of multivariate time series via temporal abstraction and time intervals mining [26]. Two new metrics for exploiting TIRPs as features for machine learning were defined: *horizontal support*, which represents the number of TIRP instances discovered for a specific entity, and *mean duration*, which measures the average time length of the supporting TIRP instances. Both new feature representation methods were shown to be superior to the default binary representation (i.e., simple existence or not of the TIRP) for classification. The use of the three abstract temporal relations (Section 2.2) was superior to the use of Allen's seven relations, and using knowledge-based State abstractions (see Section 2.1) when available performed better for classification purposes than using EWD or SAX [26]. However, in a later study, using the TD4C method [20] (see Section 2.1) to create State abstractions outperformed even the use of a knowledge-based abstraction method for that purpose. As we shall see when presenting our methods and results, we have used several insights from these preceding studies when assessing the value of our new semantic criteria for the purpose of significantly reducing the number of patterns, without reducing their classification performance.

Recent work has focused not only on the discovery of temporal patterns and their use for classification and prediction, but also on proposing *explanations* to the classifications and predictions [43–45].

Several research groups have been working recently on creating whole architectures for time series classification. Middlehurst et al. improved on their HIVE-COTE 1.0 suite to introduce the HIVE-COTE 2.0 meta-ensemble of time series classification [46]; Tan et al. have demonstrated the computational efficiency of their MultiRocket architecture, based on *Convolutional Neural Networks* (CNNs) for fast time-series classification, which has a performance comparable to the HIVE-2.0 meta-ensemble [47]; Lee, Lindgren, and Papapetrou have presented the Z-Time architecture [7]; and Sarafian and Moskovitch have introduced the *Saraswati* suite of tools, which can modify an algorithm for temporal-pattern discovery into an algorithm for the discovery of predictive temporal patterns [27].

However, unlike the case of other frameworks that discover and exploit temporal patterns as classification features [3,7,25–27,32,47], when using our new SAC principle, the decision whether to even include a candidate TIRP in the list of discovered TIRPs, and in particular, whether to exploit it as a feature for classification and prediction, can use also the *internal semantics of the components of the TIRP*, and as we shall show when defining the SAC principle, can lead to the decision to reject its inclusion to begin with. This [often very substantial] pruning can be made without any discernible decrease in the performance of the classifier that uses the TIRPs as features, as our evaluation of the SAC principle will demonstrate.

## 3. Materials and Methods

We start by first defining the basic interval-based TDM terminology. We then describe briefly the high-level overview of the general TDM algorithm we chose to deploy for this study, before formally introducing the SAC criterion and its semantics and several versions.

#### 3.1. The Time Intervals Mining Process

To formally define the problem of mining symbolic time intervals, and to better comprehend how the SAC constraint can be introduced into an interval-based frequentpattern discovery algorithm, we present several basic definitions used by common TDM algorithms, such as by the KarmaLego frequent-pattern discovery algorithm [3], which we are also using within the evaluation.

We define a *symbolic time interval*,  $I = \langle s, e, sym \rangle$ , as an ordered pair of time points, start time (*s*) and end time (*e*), and a symbol (*sym*) which represents one of the domain's symbols from the domain-specific set S. A *non-ambiguous lexicographic TIRP*  $P = \{\hat{I}, \hat{R}\}$  is defined as a set  $\hat{I}$  of *k* symbolic time intervals ( $I_1 \dots I_k$ ) ordered lexicographically by start time, then end time, then symbol, and a set  $\hat{R}$  of all of the pairwise temporal relations among each of the ( $k^2 - k$ )/2 pairs of symbolic time intervals in  $\hat{I}$ . Note that *each TIRP is an abstract temporal pattern* that represents a *class*, or a *set*, of specific *instances* of that TIRP (The input interval-based database is also ordered lexicographically).

The goal of discovering frequent TIRPs is to discover [sufficiently] frequently occurring abstract patterns within the instances of the given database; each pattern instance found in the database is a *TIRP instance*.

The fact that the database is ordered lexicographically enables us to use only seven of Allen's temporal relations (or even the three abstract relations) as defined in Section 2.2. Figure 2 presents a typical TIRP, represented as a half-matrix of temporal relations.



**Figure 2.** An example of a TIRP representation, containing five instances of symbolic time intervals of three types, A, B, and C, and all of their pair-wise temporal relations.

The *vertical support* for a TIRP is defined as follows: given an input database of |E| distinct entities (e.g., different patients), each represented as a set of symbolic time intervals (in which one or more symbols might repeat over different symbolic time intervals), the *vertical support* of a TIRP *P* that was discovered in the input data is denoted by the cardinality of the set  $E^P$  of distinct entities (within which at least one instance of *P* was discovered) divided by the cardinality of |E|.

When a TIRP has a vertical support above a given minimal predefined threshold *min\_ver\_sup*, it is referred to as *frequent*. Furthermore, the *horizontal support* of a TIRP *P* for an entity e (e.g., a single patient's record), hor\_sup(P, e), is the number of instances of P found in e. Accordingly, we also define the mean duration of the supporting instances of the same TIRP P within an entity e as the average time of the instances of P within a specific entity e, from its earliest start-time to its last end-time. Thus, given a minimum vertical support *min\_ver\_sup*, the goal of the mining task is to find the complete set of the frequent TIRPs, including all of their supporting instances vertically and horizontally [26].

The KarmaLego algorithm [3], which we chose to use in this study's evaluation to discover frequent temporal patterns due to its proven efficiency comparing it to several of the best-known alternatives (and it is provably complete in its discovery of all frequent TIRPs [26]), consists of two main phases. The first phase is called Karma, in which all of the frequent two-sized TIRPs, having two symbolic time intervals  $I^1$  and  $I^2$  and a temporal relation r among them, are discovered and indexed. In the second phase, called Lego, a recursive process extends the frequent two-sized TIRPs, referred to as  $T^2$ , through efficient candidate generation, into a tree of longer frequent TIRPs consisting of conjunctions of the two-sized TIRPs that were discovered in the Karma phase. The final output is an enumeration tree of all the frequent TIRPs discovered in the given database.

## 3.2. Adding Semantic Considerations to Time Intervals Mining

As explained in Section 1, there is a potential drawback inherent in the interval-based pattern mining task. Many of the discovered patterns are *syntactically* true, but *semantically* misleading. Addressing this problem requires the addition of semantic considerations to the time intervals mining task.

## 3.3. The Semantic Adjacency Criterion

Since frequent TIRPs mining algorithms generate *all* of the feasible TIRP candidates and search for them in the data, certain discovered TIRP instances (e.g., a period of High-dose medication of a certain type, is followed by a period of Low blood pressure), although *syntactically* accurate, i.e., corresponding to a TIRP formal definition, might not represent in a transparent fashion the common-sense semantics that a domain expert might assign to the real data (see Figure 1). Thus, many of the discovered TIRPs might not be sufficiently transparent to the expert. We shall now explore this observation in depth.

Recall that a symbol, and in particular a symbol that holds during the duration of a symbolic time interval, is composed of a [raw or abstract] concept and its value (Section 3.1). In Section 1, we presented a frequent TIRP that might be discovered (Figure 1), "<Medication-dose = High> occurs before <HGB-level = Low>". The TIRP seems to imply that administering a medication at a High dose is often [temporally] followed by a Low value of the Hemoglobin-value abstract concept (which is a State abstraction of the HGB-value raw-data concept; see Section 2.1). Such an association is not necessarily causal, of course, but it certainly might be, and justifies additional exploration.

To facilitate our discussion, for each symbol we will refer to its two components, the concept and its value, following purely for consistency and clarity reasons the KBTA theory's nomenclature [16] (see Section 2.1). In the example we just discussed, the *state* abstractions "Medication-dose-level" and "HGB-level" are the [abstract] concepts, and "High" or "Low" are their values. A concept can only have one value at any point in time; different values during the same time are considered mutually exclusive and therefore contradictory. However, the TIRP shown in abstract fashion in Figure 2 might in fact include, within its supporting instances group [for a given database] that defines its vertical (or *horizontal*) support (see Section 3.1), instances that include, somewhere within their overall temporal scope (although not at the same time), symbolic time intervals that represent values that are semantically *contradicting* (see explanation below) to those appearing in the formal TIRP definition. Two such cases were shown [as the grayed-out symbolic intervals of Instances #2a and #2b] in Figure 1. Contradictions are instances in which, between two of the TIRP's symbols, there is a symbol, composed of a concept and a value, such as, in this case, "HGB-level = High" (or "Medication-dose-level = Low"), in which the *concept* (which implicitly includes its abstraction type, such as State or Gradient, and its context, such as gender = Female) is *identical* to the concept of either of these two symbols, but its *value* is *different*. Either of these contradictory associations (with the first or the second symbols) might change, and in this case even reverse, the semantic meaning of the original temporal association, since it now seems that a *High* medication-dose level might be actually associated with a *High* Hemoglobin value (or conversely, that a *Low* medication-dose level is associated with Low hemoglobin level).

Note that even if the meaning of the intermediate symbol does not directly contradict the meaning of either of the temporal relation's symbolic time intervals, for example if both its concept name and abstraction type and its concept's value are the *same* as the concept name, type, and value of one of the two symbols, it might nevertheless change the overall pattern's semantics, or might simply be redundant. Thus, we might not wish to encounter even a *copy* of one of these two symbolic intervals between them.

It is important to note at this point that most time intervals mining algorithms, including KarmaLego, do not consider the semantics of the deeper structure of the symbols that hold over the symbolic time intervals, and thus view them as a single, non-decomposable symbol *sym*. However, as we have explained in Section 2.1, these symbols are in fact typically composed of a *concept* of some type and its *value* (e.g., the abstract concept "HGB-level" and its value "High"). Furthermore, using Shahar's KBTA ontology [16] (see Section 2.1), an abstract concept would include also the *abstraction type* (e.g., *State*) and usually also a *context* (e.g., gender = Female). For example, the following predicate is an example of an abstraction that might hold for some patient: "The *State* of the *HGB*-value, in the context of a *Female*, is *High*". (Note that the definition of the value *High* might vary in different contexts.). We often refer to the full representation of the concept embodied in each symbol (which typically also includes an abstraction type and a context) as its *semantic type*. In the case mentioned above, the semantic type would be "The *State* of the *HGB*-value, in the context of a *Female*". The value of that semantic type (i.e., of the symbol's concept), would be "*High*".

**Definition 1.** Symbolic time intervals  $I^i$  and  $I^j$  are of the same semantic type if each of the two symbols that hold over  $I^i$  and  $I^j$  represents some value for the same concept (which, as we recall, might include also a temporal abstraction type and a context). We denote that equivalence by the notation sem\_type( $I^i$ ) = sem\_type( $I^j$ ).

The semantic types (as defined above) of the symbols that might hold over all symbolic time intervals need to be pre-defined and used throughout the TIRPs discovery process. Such semantic types and the set of the allowed values for each abstraction of each concept in each context are a part of each domain's temporal-abstraction ontology [16] (see Section 2.1). Thus, there might be, for example, exactly five values for the state abstraction of blood-glucose values in the context of a patient who has diabetes, ranging from hypoglycemia to hyperglycemia. Our intent is to discover only TIRPs in which adjacent symbolic time intervals are semantically coherent, i.e., their symbols are composed of concepts and values that fulfil a new, semantically oriented criterion, the *Semantic Adjacency Criterion (SAC)*.

**A Semi-formal Definition:** The SAC guarantees that between two symbolic time intervals within a TIRP, there can exist no other symbolic time interval of the same semantic type as either of the two symbolic time intervals.

In particular, over such an intermediate "forbidden" symbolic interval there cannot hold a symbol whose *semantic type*, i.e., its conceptual aspect (e.g., Hemoglobin-State in a Female, or the State of the Medication-dose), is the *same* as the semantic type of one of the two symbolic time intervals, but with a *different value* (e.g., a LOW value of the Hemoglobin State abstraction, instead of a HIGH value; or a LOW value of the Medication-dose State abstraction). Our motivation in defining and using the SAC is that symbolic time intervals that appear between a pair of two other symbolic time intervals, but are of the same semantic type as that of one of the two symbolic intervals, and in particular, those that contradict the value of one of the symbols that hold over the two symbolic time intervals, are not easily understandable to a domain expert and thus, the discovered TIRP might not really represent the associations the expert expects to find in the data.

The expert will be notified that the temporal data mining algorithm discovered a frequent relation such as "A before B", as in the case of "a *high* medication-dose before a *low* value of the hemoglobin level", not realizing that there might be another symbolic time interval between them that contains a concept of the same *type* as that of A or B (i.e., a medication dose, or a hemoglobin level), but with a similar, or even a different *value*.

Instances of such a potentially contradictory intermediate symbolic time interval might also interfere with the learning (training) phase of an algorithm that induces a classifier, and reduce its classification power, which relies on features that are TIRPs that were discovered while using the (potentially deceptive) temporal relation between that pair of symbolic time intervals. The reason is that the real meaning of that temporal relation might change in a radical fashion, depending on what other symbolic intervals exist between the two members of the pair, for some of the TIRP's supporting instances. Thus, detecting in the patient's record an instance of Low blood pressure, and that two weeks ago she had taken a High dose of a medication that is similar to the medication that was taken yesterday at a Low dose, i.e., a temporal pattern that syntactically also complies with the temporal relation of being *before* a current instance of Low blood pressure, might potentially mislead the classifier looking for the temporal pattern "High medication-dose before Low blood pressure" as a feature. However, a medical domain expert analyzing the same data set might consider as meaningful only the *last* instance of the medication administration before the blood pressure measurement.

We conjecture that using the SAC constraint, in addition to the discovery of only semantically meaningful patterns, will also significantly reduce the number of potential frequent TIRPs to consider, thus reducing the computational requirements of TIRPs discovery.

The SAC constraint was inspired by the temporal interpolation mechanism from the KBTA methodology theory [16]. The *temporal interpolation mechanism* uses, in each domain, a domain-specific *interpolation function* (found as part of that domain's temporalabstraction ontology). The interpolation function is provided, as input, with two symbolic time intervals, both of which hold similar temporal abstraction types (see Section 2.1) (e.g., Gradient) of the same concept (e.g., HGB level), such as two Increasing HGB-level periods, each lasting for two weeks, with a gap of one week between them, and that returns an abstraction interpreted over an interval that joins the two intervals while bridging the gap between them, i.e., "five weeks of Increasing HGB level". The temporal interpolation mechanism [16] allows for a certain value-sensitive and context-sensitive maximal time gap to be bridged between the two symbolic time intervals. *It also ascertains that the values of the symbols that hold over the symbolic time intervals within the gap do not contradict in any way the values of the symbols that hold over the two symbolic intervals that are to be joined*.

Figure 3 closely examines the possible contradictions that might be hidden within a two-sized TIRP, thus revealing several possible versions of the SAC. For example, in the simple case of sequential data mining, a TIRP can be treated as a sequence of symbols (considering the *before* and *meets* relations only). The symbol that holds over the intermediate symbolic time interval cannot represent the same concept (i.e., have the same semantic type) as one of the two symbolic intervals, even if it has the same value.



**Figure 3.** Example of possible contradicting instances within a two-sized TIRP, when considering additional symbolic intervals that might exist in the same database, and the full range of temporal relations possible between two intervals. Cases 1, 2, 3, 6, 8, and 10 contradict the semantics of the TIRP defined above. Cases 4, 5, 7, and 9 appear outside the temporal relation gap within the TIRP and do not contradict it according to our current SAC definition.

## 3.4. The Sequential Semantic Adjacency Criterion: A Formal Definition

Using this simple version of the SAC constraint, any two *successive* symbolic time intervals of each temporal relation pair within a TIRP, when the symbolic intervals of the TIRP are ordered *lexicographically* in a canonical fashion, by start time, then end time,

then symbol (see Section 3.1), must be temporally adjacent in the sense of our semi-formal definition. That is, no symbolic time interval whose symbol has a semantic type equal to one of them can exist between them. The two symbols of the successive symbolic intervals themselves might include the same concept. This version of SAC is called the *Sequential Semantic Adjacency Criterion*.

**Definition 2.** The Sequential Semantic Adjacency Criterion (Sequential SAC or SSAC) holds over a TIRP  $P = \{\vec{l}, \vec{R}\}$ , where  $\vec{l} = \{I^1, I^2, ..., I^k\}$  iff:

 $\forall 1 < i < k-1 : (I^{i}.e < I^{i+1}.s) \\ \rightarrow \nexists t : ((I^{t}.e > I^{i}.e) \land (I^{t}.s < I^{i+1}.s)) \land (sem\_type(I^{i}) = sem\_type(I^{t}) \lor sem\_type(I^{i+1}) = sem\_type(I^{t}))$ 

(Note that  $I^i$  is before  $I^{i+1}$ ; relaxing this constraint and allowing for relations such as Contains or Overlaps leads to additional SAC versions that we introduce later.)

However, note that considering only successive symbolic time intervals within the TIRP definition when using the lexicographic ordering ignores the existence of all of Allen's temporal relations [23] and stays within the stricter limits of the sequential data mining approach. We shall demonstrate this observation with an example.

Figure 4 considers the TIRP "<Medication-dose = High> *before* <HGB = High> *before* <HGB = Normal>"; the relationship between the medication administration time interval and the third time interval is also *before*: "<Medication-dose> *before* <HGB = Normal>", but the two symbolic time intervals participating in this relation are not *successive* in the SSAC sense, since they do not follow each other in the lexicographic ordering.



**Figure 4.** An example of the difference between the sequential version of SAC and other possible versions that do not consider only successive symbolic time intervals.

Thus, in the Sequential SAC version, only the first two temporal relations of this particular TIRP's definition will be checked for the existence of any semantically contradictory (in the SAC sense) symbolic time interval instances between them. However, other semantically contradictory symbolic time interval instances might exist between the first and the third symbolic time intervals, which will *not* be checked using the SSAC. (For example, a Low medication dose within the dashed interval of Figure 4 would not contradict the relation "<Medication-dose = High> *before* <HGB = Normal>", since the components of that relation do not follow each other in the canonical lexicographic-ordering representation of this particular TIRP.)

Note two interesting insights regarding the SAC principle: First, this version of SAC does allow for the discovery of what we refer to as "*Symbolic Gradient*" temporal patterns, e.g., Decreasing or Increasing values of the State abstractions of the same raw-data concept that hold over successive symbolic time intervals, such as a gradual decrease in the value of the Hemoglobin-State concept (see Figure 5). The reason is that the constraint will

only be checked between every two successive symbolic time intervals, and these can be of the same type. Second, SSAC also allows for the discovery of what we refer to as "Counting" temporal patterns. "*Counting*" temporal patterns are patterns that include some finite repetition of instances of a symbolic interval as part of the pattern. For example, the repeating of two or more successive Low Hemoglobin State abstraction values (see Figure 5). Such TIRPs might serve as useful features in certain domains.



**Figure 5.** A possible TIRP that might be discovered by using the SSAC; the first three symbols represent a "Symbolic Gradient" temporal pattern of decreasing values of the Hemoglobin State abstractions, while the last two symbols present a "Counting" temporal pattern of two successive low hemoglobin tests.

We have also defined two additional versions of SAC that capture slightly different semantics and enable different levels of expressivity of TIRPs that can be discovered.

## 3.5. The Conservative Semantic Adjacency Criterion (CSAC)

The first additional SAC version that we examine considers *every* pair of symbolic time intervals within the TIRP definition. While the previous version considered only successive symbolic time intervals, when the TIRP is represented in a canonical fashion using a lexicographic ordering, we might want to consider *all* of the temporal relations within the TIRP's definition (thus probably reducing even more the potential vertical support for such a TIRP and, correspondingly, the number of different TIRPs discovered in the output).

Allowing the SAC constraint to hold over *all* of the TIRP's pairwise relations enables us to discover, in the input interval-based data, only true SAC-obeying TIRPs, in the most restrictive interpretation. Thus, unlike the SSAC, this version of SAC will *not* allow for the discovery of any "Symbolic Gradient" temporal patterns or of any "Counting" temporal patterns at all. We refer to this version of SAC as the *Conservative Semantic Adjacency Criterion*.

**Definition 3.** The Conservative Semantic Adjacency Criterion (Conservative SAC or CSAC) holds over a TIRP  $P = \{\vec{I}, \vec{R}\}$ , where  $\vec{I} = \{I^1, I^2, ..., I^k\}$  iff:

# $\forall 1 < i < j < k : (I^{i}.e < I^{j}.s) \\ \rightarrow \nexists t : ((I^{t}.s < I^{j}.s) \land (I^{t}.e > I^{i}.e)) \land (sem\_type(I^{i}) = sem\_type(I^{t}) \lor sem\_type(I^{j}) = sem\_type(I^{t}))$

## 3.6. The Liberal Semantic Adjacency Criterion

The second additional version is a special variation of the conservative version, which considers the SAC for *all* of the temporal relations within the TIRP definition, but enforces it only between symbolic time intervals that represent *different* semantic types.

This new SAC version will allow for the discovery of "Symbolic Gradient" temporal patterns and of "Counting" temporal patterns. This is achieved since it allows for the discovery of patterns that involve successive symbolic time intervals whose symbol includes the same concept, such as  $A_3A_2A_1A_1A_1$ , where each  $A_i$  represents some value of the concept of type A. In some cases, this expressivity might be useful. This version of SAC is called the *Liberal Semantic Adjacency Criterion*.

**Definition 4.** The Liberal Semantic Adjacency Criterion (Liberal SAC or LSAC) holds over a TIRP  $P = \{\hat{I}, \hat{R}\}, where \hat{I} = \{I^1, I^2, ..., I^k\}$  iff:

$$\begin{array}{l} \forall 1 < i < j < k : (I^{i}.e < I^{j}.s) \land (sem\_type(I^{i}) \neq sem\_type(I^{j})) \\ \rightarrow \nexists t : ((I^{t}.s < I^{j}.s) \land (I^{t}.e > I^{i}.e)) \land (sem\_type(I^{i}) = sem\_type(I^{t}) \lor sem\_type(I^{j}) = sem\_type(I^{t})) \end{array}$$

However, unlike the Sequential SAC version, and certainly the Conservative SAC, the Liberal SAC constraint *does* allow for some potential semantics that might not be intuitive to domain experts. For instance, between the two LOW values of the Hemoglobin State abstractions in the TIRP displayed in Figure 5, the LSAC allows for the potential existence of a "hidden" HIGH Hemoglobin State value, because the constraint is not enforced between symbolic time intervals that represent the *same* semantic types.

## 3.7. The Computational Implications of Enforcing the SAC

Using the SAC, in addition to enabling the discovery of only patterns whose semantics enforce a stricter interpretation of the data relationships, is potentially also more functional and efficient, i.e., classification and prediction algorithms might benefit from features that represent more "reliable" temporal patterns and might also enable us to compute them within a briefer time span. Recently, there has been an increasing use of TIRPs as features for classification and prediction tasks [19,24,26,27,48], in which we would like to examine the potential contribution of the SAC.

The SAC is a stricter selection criterion for TIRPs discovery algorithms and many TIRP candidates might not be generated. Thus, our first hypothesis is that *using the SAC* (including the several versions we proposed) to discover TIRPs will generate fewer TIRPs than when not using the SAC for any given minimal vertical support, which is expected to lead to a shorter run time for the discovery phase of the TIRPs.

However, due to the semantic coherence of the discovered TIRPs, i.e., their more uniform meaning, our second hypothesis is that we expect that, *nevertheless, the resultant* [smaller] set of TIRPs, discovered using the SAC, will still induce a classifier that has the same or better classification and prediction performance, given the same minimal vertical support threshold for the TIRPs as features.

## 3.8. Adding the SAC Constraint to the KarmaLego Algorithm

The SAC is a highly general criterion, equally applicable to any time intervals mining algorithm, as well as sequential mining algorithms. However, we needed to assess it within a concrete framework. We decided to evaluate the SAC within the KarmaLego framework (see Section 3.1). Finally, the algorithm structure's natural modularity, composed of the Karma and Lego steps, greatly facilitated our task of integrating any SAC version.

To implement the SAC version within the KarmaLego algorithm, we first added the basic SAC pruning constraint to the Karma phase. That is, only two-sized TIRPs that obeyed the basic semi-formal SAC constraint (i.e., of not having any symbolic time interval between them, over which holds a symbol of the same semantic type as either of the two members of the potential two-sized TIRP) were added to the two-sized TIRP enumeration tree.

Then, during the Lego phase, given each SAC version to be applied, we decided which pairs of symbolic time intervals needed to be checked against the data when extending the TIRP from size k to size k + 1:

- 1. In the case of the SSAC version, we checked the constraint only between the (lexicographically ordered) *k*th and the new *k*th + 1 symbolic time intervals.
- 2. In the case of the CSAC version, we checked the constraint between the 1st, 2nd..., and the (lexicographically ordered) *k*th and the new *k*th + 1 symbolic time intervals.
- In the case of the LSAC version, we performed a similar procedure to that used to enforce the CSAC constraint, but only for pairs of symbolic time intervals over which hold symbols of different semantic types.

Appendix A contains a short pseudo-code of the SAC implementation with the particular frequent TIRP-discovery algorithm that we had selected, the KarmaLego algorithm.

## 3.9. Evaluation

To demonstrate our methods, we decided to use a highly efficient interval-mining algorithm that was recently introduced by the authors, called KarmaLego [Moskovitch and Shahar, 2015a], as our means for discovering TIRPs. However, the semantic enhancements that we introduced into KarmaLego are quite general. We measured the number of discovered TIRPs, the runtime, and the performance of the TIRPs when used as features for several classification and prediction tasks.

We evaluated the runtime of the KarmaLego algorithm and the number of discovered TIRPs with the different SAC versions. Given our informal hypotheses (see Section 3.3), which are based on reasonable arguments, but which need empirical verification, our specific research questions were:

- 1. Does using SAC indeed reduce the discovery *runtime*, compared to not using it? Which of the three SAC versions requires the shortest runtime?
- 2. Does using SAC indeed reduce the *number* of discovered TIRPs, compared to not using it? Which of the three SAC versions results in the smallest number of TIRPs?
- 3. Does using SAC maintain the classification and prediction *performance*, compared to not using it? Which of the SAC versions is best for classification and prediction?

This evaluation was performed across different state abstraction or discretization methods (KB, EWD, SAX, and TD4C-KL, which uses the Kullback–Leibler distance measure as explained in Section 2.1), each with three bins, different temporal relation sets (the three abstract temporal relations mentioned in Section 2.2, and the full set of Allen's seven temporal relations), and various minimal vertical support thresholds to measure the number of discovered TIRPs.

In addition, for the purpose of the classification and prediction tasks, we evaluated different TIRP feature-representation methods (Binary, Horizontal Support, and Mean Duration) (as discussed in Sections 2.3 and 3.1) and four different classification algorithms (Random Forest, Naïve Bayes, Support Vector Machines [SVM], and Logistic Regression). We expand on our classification-performance evaluation methods in Section 3.9.2.

## 3.9.1. The Data Sets

To evaluate the effect of using the SAC versions on the results of the TIRP discovery process and on the eventual performance of the models induced for classification and prediction, we used three clinical data sets.

The data sets included: (1) an oncology data set from the Rush Medical Center, Chicago, USA, including patients who had undergone either allogeneic or autologous bone-marrow transplantation; (2) a hepatitis data set describing patients who had either Hepatitis B or C, which is from a KDD conference challenge [49] and which is *publicly available* [50]; and (3) a diabetes data set from our local academic medical center [51], including Type II diabetes patients who had been followed (albeit sporadically) for at least five years, focusing on the future outcome of the level of albuminuria (protein in the urine, a measure of renal dysfunction) in the fifth year. (Only the Hepatitis data set is publicly available.)

Table 1 describes the characteristics of the three data sets used throughout all of the evaluations: the total number of data points, the number of patients, the number of

concepts, or the number of all semantic types (e.g., Hemoglobin State in a particular context, as explained in Section 3.3), and the average number of data points per patient.

Data Set	Data Points	Patients	Concepts	Total Number of Potential Values	Mean Data Points Per Patient
Oncology	76,468	207	12	41	369
Hepatitis	368,216	499	10	29	738
Diabetes	165,199	5178	4	12	32

**Table 1.** Descriptive statistics of the three data sets.

Note that the concepts were used as classification and prediction features. However, as shown in Appendix B, each concept (e.g., Hemoglobin State) might have from two to five different values (e.g., in the case of Hemoglobin State values, five values: Very-Low, Low, Moderately Low, Normal, High). Therefore, in the Oncology domain, the 12 concepts included a total of up to 41 potential values; in the Hepatitis domain, the 10 concepts included a total of up to 29 potential values; and in the Diabetes domain, the four concepts included a total of up to 12 potential different values.

The task in the case of the oncology data set was to classify patients who underwent bone-marrow transplantation into autologous bone-marrow transplantation versus an allogeneic bone-marrow transplantation; the task in the case of the hepatitis data set was to classify the patients into Hepatitis B patients versus Hepatitis C patients; and the task in the case of the diabetes data set was the prediction, within a variable period of up to 5 years, of the state abstraction of the albuminuria-value concept (a measure of the amount of protein secreted in the urine), and specifically, whether the patient will have a normal albuminuria level (denoting normal renal function) versus a micro-albuminuria or macro-albuminuria albuminuria level (indicating renal deterioration).

The full description of the three data sets, the definitions used within each domain in the case of the knowledge-based temporal state abstraction method, and additional details about the tasks within each domain, appear in Appendix B.

## 3.9.2. The Experimental Design and the Evaluation Measures

We based our evaluation on the KarmaLegoSification framework [26]. The input timestamped raw data were interpolated and abstracted into a set of symbolic time intervals, using either knowledge-based or automatic temporal abstraction methods. All of the frequent TIRPs that can be discovered were discovered from the symbolic time intervals output, with or without using any SAC version. In either case, we examined the effect of using either the abstract three temporal relations or the full seven temporal relations. The TIRPs were then used as features for the induction of classification and prediction models, by representing each TIRP using either a simple binary representation of the TIRPs, the mean horizontal support, or the mean duration of the TIRP within the entities.

For the evaluation of research questions 1 and 2, we performed a series of experiments recording the runtime in seconds and the number of discovered TIRPs using the KarmaLego algorithm with the three SAC versions, as well as without using any criterion at all. Note that we could use any other temporal data mining algorithm; we used KarmaLego because it is faster than several other approaches and it is complete [3]. The runtime and number of TIRPs were evaluated on the different temporal abstraction methods, different sets of relations, and various minimal vertical support thresholds.

Because these experiments measure runtime, each combination was executed separately and thus was isolated from other processes that might have influenced the CPU behavior. We used an AMD Opteron<sup>TM</sup> Processor 6128 2.00 GHz Machine with 32.00 GB RAM and Windows Server 2008 R2 Datacenter.

To answer research question 3, we evaluated the classification and prediction performance of the SAC using the *Area Under the Curve* (AUC). We compared the mean AUC with two statistical analysis methods: a one-way ANOVA and the post hoc Scheffé method, using IBM SPSS Statistics 20. The one-way ANOVA was applied to the general parameters, such as determining whether the different SAC versions performed differently and the Scheffé method was applied as post hoc examination to test differences within the SAC versions. Comparisons that were found to be significantly different ( $\alpha = 0.05$ ) are reported.

Since mining TIRPs may result with different sets for each group of patients [26], we used a rigorous evaluation setup, including three-fold mining and ten-fold cross-validation classification. Thus, the data were split into three folds wherein TIRPs were discovered from one fold and then detected in the other two folds which were used for the classification experiment. This was repeated three times for the three-folds mining. We used four highly different types of induction algorithms: *Random Forest*, the best known application of the decision trees family (randomizing both the features and the data) [52], the classic pure probabilistic reasoning algorithm—*Naïve Bayes* [53], SVM—a very different family that uses a special type of linear optimization [54], and of course from the Linear classifiers, *Logistic Regression* [55], which is often used as the baseline statistical approach against which other methods are compared.

We did not perform any feature selection methods on the temporal patterns discovered, since previous studies [26] did not demonstrate their value, and also because we wanted to directly assess the value of using all features found with and without using any SAC variation. Once TIRPs were discovered, however, we exploited them for creating three different features from each TIRP: *Binary* (existence of the TIRP in the record), *Horizontal Support* (number of TIRP instances in the record), and *Mean Duration* of the TIRP in the record.

The KarmaLego method was implemented based on the original Moskovitch and Shahar study mentioned above [3]. We used the SAX algorithm, which we implemented based on Lin et al.'s description [14], and the TD4C-KL method, which was implemented based on Moskovitch and Shahar's original description [20], using the Kullback–Leibler divergence as the measure for deciding which value cut-off leads to the best separation between the outcome classes. We used the classification algorithm implementations available in WEKA 3.7.1 [56].

## 4. Results

In the following two subsections and their sub-subsections, we shall present our results in the following order: First, in Section 4.1, we shall present the runtime of the SAC-enhanced pattern-discovery algorithm and the number of TIRPs it discovered in each of the three evaluation domains using each of the three SAC variations, noting each time the output when not using any SAC version. Then, in Section 4.2, we shall present the classification and prediction performance of the TIRPs discovered in the same data sets when used as features in the same three domains, when using each of the three SAC variations, or when not using any of them.

## 4.1. The SAC Runtime and Number of Discovered TIRPs

For each data set, we ran the experiments with various minimal vertical support thresholds (for reasonable runtime and memory usage).

## 4.1.1. The Oncology Data Set

Figure 6 presents the runtime in seconds of the KarmaLego algorithm using Allen's seven temporal relations and Figure 7 does so using the abstract three relations.



**Figure 6.** The runtime of the KarmaLego algorithm for different minimal vertical support [MVS] thresholds on data mined using seven temporal relations in the oncology data set. Each graph represents one temporal abstraction method and displays all of the SAC versions (if any) used (the legend appears on the **upper right**).



**Figure 7.** The runtime of the KarmaLego algorithm for different minimal vertical support [MVS] thresholds on data mined using three temporal relations in the oncology data set. Each graph represents one temporal abstraction method and displays all of the SAC versions (if any) used (the legend appears on the **upper right**).

Figures 6 and 7 show that all of the SAC versions result in a faster runtime. Using SAC (and especially CSAC) allowed us to compute all of the TIRPs passing the minimal vertical support [MVS] threshold in a time that was almost an order of magnitude shorter than the time needed without using the SAC version.

Figure 8 presents the number of discovered TIRPs for when the data were abstracted using Allen's seven temporal relations and Figure 9 for when the data were abstracted using the abstract three relations. The same trends, as for runtime, hold for the number of discovered TIRPs.



**Figure 8.** The number of discovered TIRPs using seven temporal relations in the oncology data set for different minimal vertical support [MVS] thresholds. Each graph represents one temporal abstraction method and displays all of the SAC versions (if any) used (the legend appears on the **upper right**).

## 4.1.2. The Hepatitis Data Set

Figure 10 presents the runtime in seconds of the KarmaLego algorithm using Allen's seven temporal relations and Figure 11 does so using the abstract three relations. Using LSAC and CSAC allowed us to compute all of the TIRPs passing the minimal vertical support threshold in a time that was almost an order of magnitude shorter than the time needed using SSAC or without using any SAC. The most restrictive CSAC version is the fastest.

Figure 12 presents the number of discovered TIRPs for when the data were mined using Allen's seven temporal relations and Figure 13 does so for when the data were mined using the abstract three relations. Here too, the same trends are seen as in the runtime results. Using the SAC versions usually results in the discovery of a significantly smaller number of TIRPs and within a shorter runtime.



**Figure 9.** The number of discovered TIRPs for different minimal vertical support [MVS] thresholds using three temporal relations in the oncology data set. Each graph represents one temporal abstraction method and displays all of the SAC versions (if any) used (the legend appears on the **upper right**).



**Figure 10.** The runtime of the KarmaLego algorithm for different minimal vertical support [MVS] thresholds on data mined using seven temporal relations in the hepatitis data set. Each graph represents one temporal abstraction method and displays all of the SAC versions (if any) used (the legend appears on the **upper right**).



**Figure 11.** The runtime of the KarmaLego algorithm for different minimal vertical support [MVS] thresholds on data mined using three temporal relations in the hepatitis data set. Each graph represents one temporal abstraction method and displays all of the SAC versions (if any) used (the legend appears on the **upper right**).



**Figure 12.** The number of discovered TIRPs for different minimal vertical support [MVS] thresholds using seven temporal relations in the hepatitis data set. Each graph represents one temporal abstraction method and displays for each method all of the SAC versions (if any) used (the legend appears on the **upper right**).



**Figure 13.** The number of discovered TIRPs for different minimal vertical support [MVS] thresholds using three temporal relations in the hepatitis data set. Each graph represents one temporal abstraction method and displays all of the SAC versions (if any) used (the legend appears on the **upper right**).

## 4.1.3. The Diabetes Data Set

Figure 14 presents the runtime in seconds of the KarmaLego algorithm using Allen's seven temporal relations and Figure 15 does so using the abstract three relations. From both figures, we can see that using the SSAC and CSAC versions results in a faster extraction of the TIRPs. The most restrictive version, CSAC, was also the fastest, as would be expected. Using SSAC and CSAC allowed us to compute all of the TIRPs passing the minimal vertical support threshold within a much shorter time.

Figure 16 presents the number of discovered TIRPs of the KarmaLego algorithm when the data were mined using Allen's seven temporal relations and Figure 17 does so for when the data were mined using the abstract three relations. All SAC versions resulted in fewer TIRPs when compared to the number of non-SAC-obeying TIRPs.

We saw the best results in the diabetes data set, which is also the largest one. When running the experiment with 0.025 minimal vertical support, we discovered 7689 patterns (in about 731 s) when *not* using SAC, and only 253 patterns (in only about 15 s) when using CSAC. Thus, we obtained up to a 97% decrease in the number of discovered patterns in up to 98% less time. Overall, discovering SAC-obeying TIRPs is faster and the number of discovered TIRPs is much smaller when using the SAC. Moreover, the most restrictive CSAC version resulted in fewer TIRPs and the fastest runtime.



**Figure 14.** The runtime of the KarmaLego algorithm for different minimal vertical support [MVS] thresholds on data mined using seven temporal relations in the diabetes data set. Each graph represents one temporal abstraction method and displays all of the SAC versions (if any) used (the legend appears on the **upper right**).



**Figure 15.** The runtime of the KarmaLego algorithm for different minimal vertical support [MVS] thresholds on data mined using three temporal relations in the diabetes data set. Each graph represents one temporal abstraction method and displays all of the SAC versions (if any) used (the legend appears on the **upper right**).



**Figure 16.** The number of discovered TIRPs for different minimal vertical support [MVS] thresholds using seven temporal relations in the diabetes data set. Each graph represents one temporal abstraction method and displays all of the SAC versions (if any) used (the legend appears on the **upper right**).





## 4.2. Classification and Prediction Performance Using the SAC

For each data set, we calculated temporal abstractions based on the KB, EWD, SAX, and TD4C-KL temporal abstraction methods. We then discovered frequent TIRPs that are composed of the temporal abstractions and temporal relations among them, using the KarmaLego algorithm, with or without the SAC enhancement. To generate TIRPs, we examined both the use of Allen's full seven temporal relations as well as the use of only the three abstract temporal relations. The TIRPs were used as features to train a classifier using the various induction methods. Note that to produce the TIRP features, in each data set we used a different minimal vertical support, such that it produced features that characterize at least half of the patients, or at least produced a reasonable number of features (tens to hundreds of features).

We then trained a classifier using each of the four classifier-induction methods we had chosen (*Random Forest, Naïve Bayes, SVM,* and *Logistic Regression*) and evaluated the performance of the resultant classifiers using the methodology explained in Section 3.9.2.

Figure 18 displays the mean result of using the four resultant classifiers in the three domains when using any of the three SAC versions during the TIRP discovery process compared to the classification results when the TIRP discovery process did not use any SAC version (the results are averaged over all the SAC versions, representation methods, temporal abstractions, and temporal relations variations). As can be seen, using the greatly reduced set of discovered TIRPs achieved at least the same classification (or prediction) performance in all of the tested configurations as when using the original full set of TIRPs (without any SAC-based pruning). The classification performance, in each version of the experiment, was evaluated using the Binary (B), *Horizontal Support* (HS), or *Mean Duration* (MeanD) TIRP representation methods.



**Figure 18.** The mean AUC of using the four classifier-induction methods in all three domains when using any of the three SAC versions during the TIRP discovery process compared to not using any SAC version during that process. RF = Random Forest; NB = Naïve Bayes; SMO = Support Vector Machine; LR = Logistic Regression.

To present the results in more detail, we focus in the rest of this section only on the results of the *Random Forest* algorithm because (a) very similar results were achieved for all of the classifier-induction methods with respect to the effectiveness of either using or not using the SAC enhancement and we wanted to avoid a tedious repetition; and (b) its overall classification performance in the baseline case, when *not* using the SAC enhancement, was slightly better in a consistent fashion than that of the other induction methods.

## 4.2.1. The Oncology Data Set

For the oncology data set, we used a minimal vertical support of 0.5. All SAC versions performed the same, regarding classification accuracy, as when not using the SAC, in spite of using a much smaller number of TIRPs when using the SAC. We can see from our results of the empirical evaluation that using the SAC led to a slightly better performance, no matter which TIRP representation (see Figure 19) or abstraction method (see Figure 20) was used, although the differences are not significant. SAC also performed slightly better when using either three or seven temporal relations.



**Figure 19.** The classification performance results when using the different SAC versions in the oncology data set partitioned by TIRP representation methods.



**Figure 20.** The classification performance results when using the different SAC versions in the oncology data set partitioned by the temporal abstraction methods.

Note that each point in the figures represents the average AUC of multiple runs (see Evaluation Methods). For example, in Figure 19, each point represents the mean of 240 different experimental runs (3 pattern extractions from 1/3 of the data each time  $\times$  10 folds  $\times$  4 abstraction methods  $\times$  2 sets of temporal relations). Figure 20 points represent 180 different experimental runs (3 pattern extractions from 1/3 of the data each time  $\times$  10 folds  $\times$  3 feature representation methods  $\times$  2 sets of temporal relations).

The number of the TIRPs discovered without the use of the SAC was meaningfully larger. However, it did not result in a superior classification performance, compared to the use of the reduced sets of TIRPs that resulted when using the SAC. The full set of TIRPs did not have additional classification or prediction power; rather, it even slightly reduced the performance. (See the Discussion section for several possible implications.)

## 4.2.2. The Hepatitis Data Set

For the hepatitis data set, we used a minimal vertical support of 0.7. With and without, all SAC versions performed the same (see Figures 21 and 22), in spite of the use of a much

smaller set of discovered TIRPs. No significant difference between using three or seven temporal relations with respect to classification performance was found.



**Figure 21.** The classification performance results when using the different SAC versions in the hepatitis data set partitioned by TIRP representation methods.



**Figure 22.** The classification performance results when using the different SAC versions in the hepatitis data set partitioned by the temporal abstraction methods.

Although the data set is dense and there are multiple instances of the same TIRP for each patient (as opposed to the sparser oncology data set), using the reduced set of TIRPS as features led to a classification performance that was as good as that of using the large set of TIRPs discovered without using the SAC. Thus, using the full set of TIRPs was not superior to the reduced set of TIRPs discovered using the SAC.

### 4.2.3. The Diabetes Data Set

For the diabetes data set, we used a minimal vertical support of 0.1. Using the features discovered by using all three SAC versions and the original TIRP discovery process without using SAC led to a similar level of prediction performance (see Figures 23 and 24), in spite of the use of a much smaller set of discovered TIRPs. There was no significant difference in performance when using three versus seven temporal relations.



**Figure 23.** The prediction performance results when using the different SAC versions in the diabetes data set partitioned by TIRP representation methods.



**Figure 24.** The prediction performance results when using the different SAC versions in the diabetes data set partitioned by the temporal abstraction methods.

Using the full set of TIRPs did not seem to have any additional benefit regarding prediction performance; it was sufficient to use the reduced set of TIRPs discovered using the SAC. This data set is sparser in comparison to the two others; thus, there are fewer SACobeying TIRPs, but the performance stays the same. In this domain, the maximal gap for the before relation was the largest, and thus, a larger number of potential semantic contradictions may have been avoided, compared to not using any semantic considerations; thus, using the SAC led to a good performance regardless of the small number of discovered TIRPs, which was an order of magnitude less than when not using any semantic considerations. From a practical point of view, it seems that CSAC produced the smallest set of TIRPs while maintaining the classification and prediction performance, and even less when using the seven temporal relations. In parallel, the Mean Duration TIRP representation method and the Knowledge-based discretization method were the best for classification and prediction tasks. In summary, the reduced set of discovered TIRPs with a much faster discovery runtime using the various versions of the SAC maintained the same classification (or prediction) performance in all of the tested configurations. These results complement the result of discovering a much smaller number of TIRPs when using the SAC. Note also the

clear trend towards higher performance when using the various SAC versions, in spite of the much reduced feature set.

#### 5. Discussion

We defined, formalized, and assessed in detail the *Semantic Adjacency Criterion*, a filtering principle that in the past we had only briefly introduced. Note that *no new temporal patterns were discovered*, of course, during the filtering-out process using the SAC principle, but the number of temporal patterns found, and the time needed to discover them, was greatly reduced by the various SAC pruning versions that we had used, without hurting classification and prediction performance. Three versions of the SAC criterion were suggested: The *Sequential Semantic Adjacency Criterion* (SSAC), which enforces the constraint only over pairs of temporally successive symbolic time intervals within the TIRP's canonical lexicographical order definition; the *Conservative Semantic Adjacency Criterion* (CSAC), which enforces the constraint over *every* pair of symbolic time intervals within the TIRP's definition (recall that in a TIRP, a temporal relation is defined in an unambiguous fashion between *each* pair of SSAC, which enforces the constraint only over pairs of symbolic time intervals); and the *Liberal Semantic Adjacency Criterion* (LSAC), a variation of CSAC, which enforces the constraint only over pairs of symbolic time intervals that have *different* semantic types.

It is important to note that each type of SAC can serve a different purpose and has a different expressivity. For instance, the LSAC version implicitly enables a *counting* of symbolic time intervals of the same type (e.g., patterns such as A1A1A1 denoting, for example, three overall administrations of the same range of the dose of the medication, although different ranges of doses of the medication might have been administered between them), while constraining intervals of different semantic types. Another example is the SSAC version, which enables a more restricted version of counting but, like the LSAC version, enables the discovery of a TIRP that implicitly contains a "Symbolic Gradient" temporal pattern (e.g., patterns such as A1A2A3B denoting, for example, a Low level of the medication-dose, followed by a Medium level, and then a High level, followed by some side effect). The CSAC version seems the most useful to maintain compactness in the number of discovered TIRPs, while preserving the strictest semantics of the TIRPs, although it prevents the discovery of repetitions of symbolic time intervals within the same TIRP or the discovery of Symbolic Gradients.

Note also that as the minimal vertical support increases, and thus the number of potentially discoverable patterns decreases, the returns for using the SAC principle are diminishing; however, this is exactly the phenomenon that might enable researchers and physicians to extract useful patterns, using the SAC principle, from bigger, and even very big, data sets.

We evaluated the classification and prediction performance of the features discovered using the three SAC versions on three different medical domains: oncology, infectious hepatitis, and type II diabetes. *Note that we did not choose any simulated or artificial data set*, since the main point of the evaluation was to test the SAC principle within *real* clinical domains that incorporate *real semantics*, and in particular, potential *causality*. We believe that the true value of the SAC principle can only be apparent within *real-world* data, since it is precisely the lack of coherence of most real temporal patterns that is being filtered out by that principle.

To further bolster the assessment process and its conclusions, we performed the evaluation using classification algorithms from *four different classifier-induction families*: Random Forest, Naïve Bayes, SVM, and Logistic Regression.

It is important to emphasize that the main objective of this study was to demonstrate the possibility of reducing the number of pattern features that need to be discovered in a large time-oriented data set by at least an order of magnitude (and enhance their intuitive meaning to a domain expert, due to their increased transparency), without losing any classification performance. Our goal was not to enhance any classification performance. The different SAC versions behaved slightly differently in each domain and for each classifier version. However, overall, using all of them required much less *time*, up to 98% less than when not using any SAC version, depending on the minimal vertical support specified, to discover all of their respectively relevant TIRPs, compared to not using any SAC at all. Using the various SAC versions also resulted in a significantly reduced *number* of TIRPs, up to 97% less, depending on the minimal vertical support threshold. This reduced set of TIRPs, however, did not lead to any reduced *performance* in any of the three medical domains, i.e., the resulting classifiers performed when using this reduced set of features as well as when using the full original set of TIRPs discovered in the standard KarmaLego methodology.

We infer that, at least in the medical domains in which we assessed our methodology, SAC-obeying TIRPs seem to contain most of the information important for classification and prediction.

Most of the data sets we used were relatively small (at least compared to current big data sets, although the data sets we experimented with contained about 70,000, 160,000, and 360,000 data points). However, the clear trend noted above towards a *higher* performance when using the *reduced* set of TIRPs filtered using various SAC versions, in spite of the much smaller feature set, suggests that repeating our studies with much larger data sets might, in fact, not only show that the much smaller set of TIRP-based features is sufficient, but might even demonstrate a significant *improvement* in the classification performance, and future studies might elucidate that aspect.

However, in any case, a reduction of the number of temporal pattern features to be discovered in a big data set has significant computational implications. On a similar note, it is interesting to consider that Fradkin and Mörchen's conclusions in their study [25] were that the main advantage of their new proposed sequential mining algorithm, BIDE-DC, lies in generating a *smaller number of patterns*, while preserving the same classification performance.

As we noted in the Introduction, we have recently shown that frequent TIRPs can be consistently discovered and in similar proportions in different subsets of the same data set within three different medical domains, thus increasing their value for potential patient trajectory clustering, classification, and prediction tasks [28]. (In fact, we used the same three domains mentioned in the current study.) This study has also shown that consistent discovery can be increased by increasing the minimal support threshold for frequency and, interestingly, by using the SAC principle to prune in each subset the patterns that are candidates for discovery.

Note also that the SAC principle is quite general and is not specific to the KarmaLego algorithm on which we demonstrated it or even to the family of multivariate interval-based TDM algorithms that KarmaLego is a part of. For example, sequential mining algorithms such as SPADE [57] start with a set of time-stamped events, each containing several items, and discover qualitative associations that involve the Before temporal relation. Using SPADE to generate a set of temporal sequences that will be used as classification features might well benefit similarly through the addition of semantic considerations similar to the SAC variations, while enhancing its semantic transparency to domain experts.

Several limitations of our study can be noted. We did not measure the runtime of the classification and prediction phases, but obviously, representing a larger number of features, especially when using various functional methods (e.g., computing the mean duration of each TIRP) requires more time. That might be an additional advantage of using the SAC principle which we did not assess. Other factors that might also require more time are selecting and using various feature selection algorithms and inducing a classifier from a larger set of features.

Note that a trivial case for semantic equivalence is the one in which all concepts are different (e.g., different events, each with its own symbol); semantic type equivalence between two symbolic intervals will then consist of having the same symbol hold over both intervals.

Our main intent in the current study, however, was to explore the non-trivial case in which most concepts might have more than one value or, at least, in which there is some domain knowledge that assigns types to the various concepts. However, exploring the potential implications of the SAC principle for the simple case in which all symbols are different and no domain knowledge exists can certainly be explored in a future study.

Another potential limitation might be noted. We did not assess the actual transparency of the SAC-obeying TIRPs, as opposed to the non-SAC-obeying TIRPs, *in the eyes of medical domain experts* in the three domains. That was not an objective of the current study, which focused on the pure objective computational aspects of using the SAC principle, but it might be interesting to assess this subjective cognitive aspect explicitly in future studies.

The use of the four different abstraction and discretization methods led *qualitatively* to the same results, with respect to the number of TIRPs discovered, the time needed to discover them, and the performance of the TIRPs as features for classification and prediction purposes, in all three domains using four different classification algorithms. Nevertheless, when using the EWD abstraction method, we noted in the specific case of the hepatitis data set that the SSAC's runtime (and the number of discovered TIRPs) was close to the runtime achieved without the use of any SAC (see Figures 10 and 11). SSAC is a sequential version of SAC; thus, the most reasonable explanation for this phenomenon is probably that the hepatitis data were not sequential and most of the concepts appeared at the same time. Still, the use of SSAC did not significantly reduce the performance compared to not using any semantic criterion.

Not all of the SAC versions performed equally well in the case of the diabetes data set (see Figures 14 and 15). Using LSAC, which is the liberal version of SAC, meaning that it restricts the criterion to hold only over pairs of semantically *different* concepts, led to a worse performance when compared to the other SAC versions. The reason might be that the diabetes data set includes a small number of concepts measured repeatedly over a long time and is pretty sparse, but there are several laboratory tests that are very common and, in the case of the liberal version of the SAC, relations among pairs of intervals of the *same* concept were considered, just as in the case of not using *any* semantic criterion; the result is a runtime that is pretty close to that of not using *any* semantic considerations, at least in some of the configurations.

The last two examples, i.e., the exceptions in our results of the empirical evaluation, demonstrate that one must learn the data and select the most appropriate SAC version, as well as the other parameters, e.g., discretization and representation methods. However, overall, the CSAC version performed best, no matter which configuration was chosen.

## 6. Conclusions and Future Work

We defined and formalized in detail a new Semantic Adjacency Criterion [SAC] for pruning temporal interval relation patterns [TIRPs] during their discovery, which increases the transparency of the discovered TIRPs for domain experts, and which can exploit even very basic domain knowledge. We have demonstrated a significant reduction, up to an order of magnitude, in the number of TIRPs discovered when using the SAC, as well as in the runtime needed to extract these TIRPs. Nevertheless, this reduced set of TIRPs, when serving as features for classification and prediction using any of four different families of classifier-induction algorithms in three different clinical domains, proved to be as good as the whole set with respect to classification and prediction performance. Overall, the CSAC version, the most restrictive of the SAC versions, seemed to be the most promising for inducing the smallest set of TIRPs while maintaining the same classification and prediction performance. We have examined three variations of the SAC principle; future studies can examine the implications of using additional variations on our basic concept of exploiting domain semantics to prune temporal relations within temporal patterns. Furthermore, the subjective implications of interpreting discovered frequent TIRPs by domain experts, when using or not using the SAC principle for pruning, can be examined in future studies as well. **Author Contributions:** Conceptualization: Y.S. and A.S.; Methodology: Y.S. and A.S.; Software: A.S.; Validation: A.S., Y.S. and R.M.; Formal Analysis: A.S., Y.S. and R.M.; Investigation: Y.S. and A.S.; Resources: Y.S.; Data Curation: A.S.; Writing—Original Draft Preparation: A.S. and Y.S.; Writing—Review and Editing: Y.S. and R.M.; Visualization: A.S.; Supervision: Y.S.; Project Administration: Y.S.; Funding Acquisition: Y.S. All authors have read and agreed to the published version of the manuscript.

**Funding:** A. Shknevsky and Y. Shahar were partially supported by the European Union (EU) MobiGuide project, partially funded by the European Commission 7th Framework Programme grant No. 287811. Y. Shahar was also partially supported by the USA Office of Naval Research (ONR) award No. N629091912124.

**Institutional Review Board Statement:** This study was conducted in accordance with the Declaration of Helsinki, approved by the Institutional Review Board of the Palo Alto Veterans Association Health Care Center, as an additional analysis of the anonymous retrospective data used within the Martins et al. study [Martins, 2008], in the case of the Oncology data set; by the Ethics Committee of Ben Gurion University's Academic Medical Center, as part of an additional analysis of the anonymous retrospective data used within M. Gordon's Ph.D. thesis [Gordon. 2012], in the case of the Diabetes data set; and by the Ethics Committee of Chiba University Hospital, approved for use as a publicly available anonymous retrospective data set, as part of the KDD 2003 Conference challenge [Ho and Nguen, 2003], in the case of the Hepatitis data set.

**Data Availability Statement:** The data sets included: (1) an oncology data set from the Rush Medical Center, Chicago, USA, including patients who had undergone either allogeneic or autologous bonemarrow transplantation; (2) a hepatitis data set describing patients who had either Hepatitis B or C, from a KDD conference challenge [Ho and Nguyen 2003], which is *publicly available* [Berka et al. 2002]; and (3) a diabetes data set from our local academic medical center [Gordon 2012; Klimov et al. 2015], including patients who had been followed (albeit sporadically) for at least five years, focusing on the future outcome of the level of albuminuria (protein in the urine, a measure of renal dysfunction) in the fifth year. Only the Hepatitis data set is publicly available.

Acknowledgments: The authors wish to thank all their clinical collaborators for assisting in developing the clinical knowledge bases.

Conflicts of Interest: The authors declare no conflict of interest.

## Appendix A. The SAC Pseudo-Code within the KarmaLego Algorithm

The KarmaLego algorithm (Algorithm A1) consists of two main phases [Moskovitch and Shahar, 2015a]. The first phase is called Karma and it enumerates all two-sized TIRPs (Algorithm A2). The second phase is called Lego and it is a recursive method that extends each k-sized TIRP into the possible (k+1)-sized TIRPs (Algorithm A3).

Note that the supplied algorithms are coherent with Moskovitch and Shahar's original paper, except for the <u>underlined modifications</u>, added to support the SAC principle. Additional details can be found in Alex Shknevsky's Ben Gurion University M.Sc. thesis.

Algorithm A1. The KarmaLego main loop.
Input:
db—A database of  E  entities.
min_ver_sup—The minimal vertical support threshold.
Output: T—An enumerated tree of all frequent TIRPs.
Algorithm:
1. T = Karma(db, min_ver_support)
2. Foreach $t \in T^2 / / T^2$ is T at the 2 <sup>nd</sup> level
2.1. Lego(T, t, min_ver_support)
3. Return T

Algorithm A2. The Karma algorithm with the SAC modifications.
Input:
db—A database of $ E $ entities, representing for each entity $e \in E$ the lexicographically
sorted vector of its symbolic time intervals, e.I
min_ver_sup—The minimal vertical support threshold.
Output: T—An enumerated tree of up to two-sized frequent TIRPs.
Algorithm:
$\overline{1. \underline{T} \leftarrow \varnothing}$
2. Foreach $e \in E / / T^2$ is T at the $2^{nd}$ level
2.1. Foreach $I^i$ , $I^j \in e.I \land i < j$
2.1.1. <b>r</b> $\leftarrow$ the temporal relation among $I^i$ , $I^j$
2.1.2. if SSAC $\lor$ (LSAC $\land e.I_{sym}^{i} \neq e.I_{sym}^{j} \lor$ CSAC
2.1.2.1. Foreach $I^k \in e.I \land k \neq i \land k < j$
2.1.2.1.1. If $I^k \cdot e > I^i \cdot e \land$
$\left(sem\_type\left(e.I^{i} ight) = sem\_type\left(e.I^{k} ight) \lor sem\_type\left(e.I^{k} ight) = sem\_type\left(e.I^{j} ight) ight)$
2.1.2.1.1.1. <u>break</u>
2.1.3. Index( $T^2$ , $< e.I_{sym}^i, r, e.I_{sym}^j >$
3. Foreach $t \in T^2$
3.1. If ver_sup(t) < min_ver_sup
3.1.1. Prune(t)
4. Return T

Note that regarding the use of the sem\_type method, we saved for each symbolic time interval (STI) its symbol (sym) as a pair of concept and value. In this case, checking the semantic equivalence of symbolic intervals is interpreted as a comparison of the concepts (and not the values), corresponding to Definition 1.

Algorithm A3. The Lego algorithm.
Input:
T—The enumeration tree created by Karma.
t—A TIRP that has to be extended.
min_ver_sup—The minimal vertical support threshold.
Output: void.
Algorithm:
1. Foreach $sym \in T^1$
1.1. Foreach $r \in R$
1.1.1. Create new $t^c$ of size (t.size +1)
1.1.2. $t^c.s[t^c.size-1] \leftarrow sym$
1.1.3. $t^c.tr[t^c.tr\_size - 1] \leftarrow r$
1.1.4. $C \leftarrow 0$
1.1.5. $C \leftarrow$ generate cantidate TIRPs from $t^c$
1.1.6. Foreach $c \in C$ // candidates
1.1.6.1. Search_Supporting_Instances(c, $T^2$ )
1.1.6.1.1. If $ver_sup(c) > min_ver_sup$
1.1.6.1.1.1. $T \leftarrow T \cup c //c$ is frequest
1.1.6.1.1.2. Lego(T, c, min_ver_sup)
2. Return T

The Search\_Supporting\_Instances method (Algorithm A4) receives as input the candidate TIRP c and the set of the two-sized TIRPs in T2. In line 1, the next (new) symbol that was added (in Algorithm A3) is set to next\_sym; then, for each instance in the extended TIRP's supporting instances, the search is made. First, the temporal relation rel between the next time interval and the latest (in the extended TIRP) is set (line 2.1), and then the latest symbol of the extended TIRP sym is set (line 2.2). GetNextSTIs (Line 2.3) searches the symbolic time interval (sti) two-dimensional square array  $T^2$  using indices defined by the symbols sym and last\_sym and the temporal relation rel, for the instances starting with the latest time interval in the instance ins.sti. GetNextSTIs might return several new symbolic time intervals next\_stis.

The method *NoInst*? searches for an instance of a pair of symbolic time intervals having the given temporal relation; thus, it obtains sym, rel, and next\_sym as indices to its fourth argument, the appropriate  $T^2$  array entry, in which it queries the HashMap for the pair based on the entity id of the new instance (*inst<sup>new</sup>.e\_id*) and the first symbolic time interval instance (*inst<sup>new</sup>.sti*[*c.size* -2 - i]). It returns True if no instance of the relation was found, or else it returns False.

The method *SAC*? (e\_id, sti\_1, sti\_2,  $T^2$ [sti\_1.sym, rel, *sti*\_2.sym]) checks the relation between a given pair of symbolic time intervals, sti\_1 and sti\_2, for the entity e\_id in the  $T^2$  SAC compatibility. If we are checking without SAC, then the result is True. If we are checking using LSAC, but sem\_type(sti\_1) = sem\_type(sti\_2), then the result is True. Otherwise (i.e., when using SSAC, CSAC, or LSAC but sem\_type(sti\_1)  $\neq$  sem\_type(sti\_2)), and then if there is a gap between sti\_1 and sti\_2, then *SAC*? will return True if  $T^2$ [sti\_1.sym, rel, *sti*\_2.sym] contains the pair (sti\_1, sti\_2), which means that it was previously discovered as obeying the SAC criterion.

Algorithm A4. Search_Supporting_Instances.
Input:
c-TIRP to extend by searching supporting symbolic intervals instances.
$T^2$ – A the 2-dimential array of two-sized TIRPs instances.
min_ver_sup—The minimal vertical support threshold.
<b>Output:</b> c-extended by the supporting instances.
Algorithm:
1. $next\_sym \leftarrow c.sym[c.size - 1]$
2. Foreach inst $\in c.insts$
2.1. rel $\leftarrow$ c.tr[c.tr_size-1]
2.2. $sym \leftarrow c.sym[c.size-2]$
2.3. next_stis $\leftarrow$ GetNextSTIs(inst.sti[c.size-2], $T^2$ [sym, rel, next_sym])
2.4. Foreach next_sti ∈ next_stis
2.4.1. $inst^{new} \leftarrow inst$
2.4.2. <u>if <math>\neg NoSAC \land \neg SAC</math>? (inst<sup>new</sup> e id, inst<sup>new</sup> sti[c.size - 2].</u>
<u>next_sti, T<sup>2</sup>[sym, rel, next_sym])</u>
2.4.2.1. <u>break</u>
2.4.3. For (i=1; i <c.size-1; i++)<="" td=""></c.size-1;>
2.4.3.1. rel $\leftarrow$ c.tr[c.tr_size-1-i]
2.4.3.2. $sym \leftarrow c.sym[c.size-2-i]$
2.4.3.3. <u>if NoInst? (inst<sup>new</sup>. e_id, inst<sup>new</sup>. sti[c.size <math>-2 - i</math>]</u> ,
<u>next sti,T<sup>2</sup>[sym,rel,next sym])</u>
2.4.3.3.1. <u>break</u>
2.4.3.4. $if \neg NoSAC \land \neg SSAC \land$
$\neg SAC?$ (inst <sup>new</sup> . e id, inst <sup>new</sup> . sti[c.size $-2 - i$ ].
<u>next sti, T<sup>2</sup>[sym, rel, next sym]</u>
2.4.3.4.1. <u>Dreak</u>
2.4.4. c.insts $\leftarrow$ c.insts $\cup$ <i>insts</i> <sup>now</sup>
2.5. remove inst from c.insts
3. Return c

## Appendix B. Data Sets and Knowledge Base Definitions

We describe here the data sets used in our experiments, and the definitions we used in the case of the knowledge-based temporal abstraction method.

## Appendix B.1. The Oncology Data Set

A medical oncology domain knowledge source that was used for the evaluation (of the KB temporal abstraction method) was an oncology knowledge database specific to the bone-marrow transplantation domain. It includes in total more than 350 concepts from 1991 to 1997 including more than 200 raw concepts (e.g., laboratory tests—White Blood Cell count, Hemoglobin), internal events (e.g., bone marrow transplantations—BMT), and more. The data source used for the evaluation was of bone-marrow transplantation patients who were followed for 2–4 years at the Rush Medical Center, Chicago, USA. The knowledge base and databases were previously elicited from our clinicians' colleagues. Table A1 presents the knowledge base that was used for the purposes of this study. Note that in case of an overlap, the maximum value will be taken.

We used 207 patients who had a bone marrow transplantation and data for the following 12 laboratory tests: Platelet count, Hemoglobin, White Blood Cell count (WBC), Glucose levels, Total Bilirubin, Alkaline Phosphatase, Hematocrit, Monocytes, Lymphocytes, Eosinophil granulocyte count (EOS), Neutrophilic band forms (Bands), Basophil granulocyte count (Basos).

Platelet count		Hem	oglobin	V	WBC	
High	$\geq 400$	High	$\geq 16$	Very_High	$\geq 20$	
Normal	100-400	Normal	11–16	High	12-20	
Moderately_Low	50-100	Moderately_Low	9–11	Normal	2.5–12	
Low	20-50	Low	7–9	Moderately_Low	0.5–2.5	
Very_Low	<20	Very_Low	<7	Low	0.1-0.5	
-		-		Very_Low	<0.1	
Glucose level		Total I	Total Bilirubin Alkaline Phosph		Phosphatase	
Very_High	≥250	Very_High	$\geq 10$	Very_High	≥225	
High	151-250	High	3–10	High	110-225	
Normal	75–151	Normal	1.5–3	Normal	35–110	
Low	<75	Low	<1.5	Low	<35	
Hematocrit		Mor	Monocytes		Lymphocytes	
High	$\geq 46.9$	High	$\geq 10$	High	≥52	
Normal	34.9-46.9	Normal	3–10	Normal	18–52	
Low	<34.9	Low	<3	Low	<18	
EOS		Bands		Basos		
Very_High	≥12	High	>=6	High	>=3	
High	6–12	Normal	<6	Normal	<3	
Normal	<6					

Table A1. The knowledge base for the oncology data set.

For the interpolation, for producing intervals out of the time-stamped raw data, we treated each time-stamped point as good for one day after and one day before. The task was to classify patients who went through autologous bone-marrow transplantation (137 patients) versus allogeneic bone-marrow transplantation (70 patients) based on TIRPs discovered from the mentioned laboratory tests.

## Appendix B.2. The Hepatitis Data Set

The hepatitis data set contains the results of laboratory examinations on hepatitis B and C patients who were admitted to Chiba University Hospital in Japan. Hepatitis A, B, and C are viral infections that affect the liver of the patient. Hepatitis B and C chronically inflame the hepatocytes, whereas hepatitis A acutely inflames them. Hepatitis B and C are especially important because they have a potential risk for developing liver cirrhosis or hepatocarcinoma. The data set contains long time-series data of laboratory examinations. The subjects are 771 patients with hepatitis B and C who were examined between 1982 and

2001. Table A2 presents the relevant knowledge that was extracted from a public KDD conference challenge [Ho and Nguyen, 2003] and was used for our evaluation. In case of an overlap, the maximum value was taken.

The data set is publicly available [Berka et al., 2002].

We used 499 patients who had a biopsy result of hepatitis B (204 patients) or C (295 patients) and the ten most frequent tests (occurring in most of the patients), including the following: Glutamic-Oxaloacetic Transaminase (GOT), Glutamic-Pyruvic Transaminase (GPT), Lactate DeHydrogenase (LDH), Total Protein (TP), ALkaline Phosphatase (ALP), Albumin (ALB), Uric Acid (UA), Total BILirubin (T-BIL), Indirect BILirubin (I-BIL), and Direct BILirubin (D-BIL). For the interpolation, we treated each time-stamped point as good for 15 days before and after each point. The task was to classify the patients as Hepatitis B versus Hepatitis C, based on the TIRPs discovered from the mentioned 10 most frequent tests.

GOT			GPT		LDH		
High	≥40 	High	$\geq 40$	High	≥450		
Normal	7–40	Normal	7–40	Normal	216-450		
Low	<7	Low	<7	Low	<216		
ТР			ALP		ALB		
High	$\geq 8.2$	High	≥206	High	≥5.1		
Normal	6.5-8.2	Normal	72-206	Normal	3.9-5.1		
Low	<6.5	Low	<72	Low	<3.9		
UA			T-BIL		I-BIL		
High	$\geq 8$	High	≥1.2	High	$\geq 0.9$		
Normal	2.5-8	Normal	0.2-1.2	Normal	0.2-0.9		
Low	<2.5	Low	<0.2	Low	<0.2		
D-BIL							
High	$\geq 3$						
Normal	<3						

Table A2. The knowledge base for the hepatitis data set.

## Appendix B.3. The Diabetes Data Set

The diabetes data set was provided by the National Institute for Biotechnology in the Negev (NIBN) in a joint study with Soroka University Medical Center [Gordon, 2012]. The subjects are 26k anonymous patients (and about 12 million raw data records) who had diabetes and various laboratory tests between 2004 and 2008. The data include static information (e.g., gender) and temporal records (e.g., High-density lipoprotein, Lowdensity lipoprotein, Triglycerides, Glucose, Hemoglobin A1c, Creatinine, Total cholesterol, and Albuminuria). The main interest in this data was on the investigation of factors associated with changes in renal function (mostly focusing on the level of albuminuria, or secretion of protein in the urine), exploring its predictive risk factors.

We used 5178 patients who had Albumin-to-creatinine ratio or Albumin-24 h from urine in the last fifth year of the data set, and who also had these tests and also Glycosylated hemoglobin (HbA1c) and Creatinine (CREATININE) in the first four years of the data set. For the interpolation, we treated each time-stamped point of Albuminuria ACR and Albuminuria U24h as good for 3 months before and after each point, for Creatinine as good for 2 months before and after each point, and for HBA1C as good for 4 months before and after each point. The task was to predict Albuminuria-normo (3231 patients) versus micro-or macro-albuminuria (1947 patients) in the fifth year based on TIRPs discovered in the first four years. Table A3 presents the relevant knowledge that was supplied by Gordon [Gordon, 2012] and other clinicians who worked on other projects as well.

Albuminuria ACR/Albuminuria U24h							
	Macro	>300		Macro	>300		
F 1	Micro	30-300	1.1	Micro	30-300		
Female	Normo-High	15-30	Male	Normo-High	13-30		
	Normo-Low	0–15		Normo-Low	0–13		
CREATININE							
	Very_High	>4		Very_High	>4		
<b>F</b>	High	2–4	1.1	High	2–4		
Female	Moderately_High-2		Male	Moderately_High.2-2			
	Normal	<1		Normal	<1.2		
HbA1c							
Very_High	>10.5						
High	9–10.5						
Moderately_High	7–9						
Normal	<7						

Table A3. The knowledge base for the diabetes data set.

## References

- 1. Batal, I.; Valizadegan, H.; Cooper, G.F.; Hauskrecht, M. A Temporal Pattern Mining Approach for Classifying Electronic Health Record Data. *ACM Trans. Intell. Syst. Technol. (ACM TIST)* **2013**, *4*, 1–22. [CrossRef]
- Klimov, D.; Shknevsky, A.; Shahar, Y. Exploration of patterns predicting renal damage in diabetes type II patients using a visual temporal analysis laboratory. J. Amer Med. Inform. Assoc. 2015, 22, 275–289. [CrossRef] [PubMed]
- Moskovitch, R.; Shahar, Y. Fast time intervals mining using the transitivity of temporal relations. *Knowl. Inform. Syst.* 2015, 42, 21–48. [CrossRef]
- Sacchi, L.; Dagliati, A.; Bellazzi, R. Analyzing Complex Patients' Temporal Histories: New Frontiers in Temporal Data Mining. Data Min. Clin. Med. 2015, 1246, 89–105.
- Che, Z.; Purushotham, S.; Cho, K.; Sontag, D.; Liu, Y. Recurrent neural networks for multivariate time series with missing values. *Sci. Rep.* 2018, *8*, 6085. [CrossRef] [PubMed]
- Yu, K.; Zhang, M.; Cui, T.; Hauskrecht, M. Monitoring ICU Mortality Risk with A Long Short-Term Memory Recurrent Neural Network. *Biocomp* 2020, 25, 103–114.
- 7. Lee, Z.; Lindgren, T.; Papapetrou, P. Z-Time: Efficient and effective interpretable multivariate time series classification. *Data Min. Knowl. Disc.* **2023**. [CrossRef]
- 8. Lee, J.M.; Hauskrecht, M. Modeling multivariate clinical event time-series with recurrent temporal mechanisms. *Artif. Intell. Med.* **2021**, *112*, 102021. [CrossRef]
- 9. Harel, O.; Moskovitch, R. TIRPClo: Efficient and complete mining of time intervals-related patterns. *Data Min. Knowl. Disc.* 2023, 37, 1806–1857. [CrossRef]
- 10. Aalst, W.V.; Weijters, T.; Maruster, L. Workflow mining: Discovering process models from event logs. *IEEE Trans. Knowl. Data Eng.* **2004**, *16*, 1128–1142. [CrossRef]
- 11. Aalst, W.V. Business process simulation revisited. In *Workshop on Enterprise and Organizational Modeling and Simulation*; Springer: Berlin/Heidelberg, Germany, 2010; pp. 1–14.
- 12. Aalst, W.V.; Schonenberg, H.M.; Song, M. Time prediction based on process mining. Inform. Syst. 2011, 36, 450–475. [CrossRef]
- 13. Aalst, W.V.; Adriansyah, A.; van Dongen, B. Replaying history on process models for conformance checking and performance analysis. *Wiley Interdisc. Rev. Data Min. Knowl. Discov.* **2012**, *2*, 182–192. [CrossRef]
- Lin, J.; Keogh, E.; Lonardi, S.; Chiu, B. A symbolic representation of time series, with implications for streaming algorithms. In Proceedings of the 8th ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery-DMKD '03, San Diego, CA, USA, 13 June 2003; ACM Press: New York, NY, USA, 2003; pp. 2–11.
- Mörchen, F.; Ultsch, A. Optimizing time series discretization for knowledge discovery. In Proceedings of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery in Data Mining, Chicago, IL, USA, 21–24 August 2005; ACM Press: New York, NY, USA, 2005; pp. 660–665.
- 16. Shahar, Y. A framework for knowledge-based temporal abstraction. Artif. Intell. 1997, 90, 79–133. [CrossRef]
- 17. Sacchi, L.; Larizza, C.; Combi, C.; Bellazzi, R. Data mining with Temporal Abstractions: Learning rules from time series. *Data Min. Knowl. Disc* 2007, *15*, 217–247. [CrossRef]
- 18. Verduijn, M.; Sacchi, L.; Peek, N.; Bellazzi, R.; de Jonge, E.; de Mol, B.a.J.M. Temporal abstraction for feature extraction: A comparative case study in prediction from intensive care monitoring data. *Artif. Intell. Med.* 2007, 41, 1–12. [CrossRef] [PubMed]

- Batal, I.; Sacchi, L.; Bellazzi, R. Multivariate Time Series Classification with Temporal Abstractions. In Proceedings of the Twenty-Second. International FLAIRS Conference; American Association for Artificial Intelligence: Washington, DC, USA, 2009; pp. 344–349.
- 20. Moskovitch, R.; Shahar, Y. Classification-driven temporal discretization of multivariate time series. *Data Min. Knowl. Disc.* **2015**, 29, 871–913. [CrossRef]
- 21. Goldstein, A.; Shahar, Y. An automated knowledge-based textual summarization system for longitudinal, multivariate clinical data. *J. Biomed. Inform.* **2016**, *61*, 159–175. [CrossRef]
- 22. Martins, S.B.; Shahar, Y.; Goren-Bar, D.; Galperin, M.; Kaizer, H.; Basso, L.V.; McNaughton, D.; Goldstein, M.K. Evaluation of an architecture for intelligent query and exploration of time-oriented clinical data. *Artif. Intell. Med.* **2008**, *43*, 17–34. [CrossRef]
- 23. Allen, J.F. Maintaining Knowledge about Temporal Intervals. *Comm. ACM* **1983**, *26*, 832–843. [CrossRef]
- 24. Patel, D.; Hsu, W.; Lee, M.L. Mining Relationships Among Interval-based Events for Classification. In Proceedings of the ACM SIGMOD International Conference on Management of Data, Philadelphia, PA, USA, 12–17 June 2008; pp. 393–404.
- 25. Fradkin, D.; Mörchen, F. Mining sequential patterns for classification. Knowl. Inform. Syst. 2015, 45, 731–749. [CrossRef]
- 26. Moskovitch, R.; Shahar, Y. Classification of multivariate time series via temporal abstraction and time intervals mining. *Knowl. Inform. Syst.* **2015**, 45, 35–74. [CrossRef]
- 27. Sarafian, N.; Moskovitch, R. Predictive temporal patterns discovery. Exp. Syst. Appl. 2023, 226, 119974. [CrossRef]
- Shknevsky, A.; Shahar, Y.; Moskovitch, R. Consistent discovery of frequent interval-based temporal patterns in chronic patients' data. J. Biomed. Inform. 2017, 75, 83–95. [CrossRef]
- Höppner, F.; Peter, S. Temporal interval pattern languages to characterize time flow. Wiley Interdiscip. Rev. Data Min. Knowl. Disc. 2014, 4, 196–212. [CrossRef]
- Garcia, S.; Luengo, J.; Saez, J.A.; Lopez, V.; Herrera, F. A Survey of Discretization Techniques: Taxonomy and Empirical Analysis in Supervised Learning. *IEEE Trans. Knowl. Data Eng.* 2013, 25, 734–750. [CrossRef]
- Höppner, F. Learning Temporal Rules from State Sequences. In Proceedings of the IJCAI Workshop on Learning from Temporal and Spatial Data, Seattle, WA, USA, 25–31 August 2001.
- Papapetrou, P.; Kollios, G.; Sclaroff, S. Discovering Frequent Arrangements of Temporal Intervals. In Proceedings of the Fifth IEEE International Conference on Data Mining (ICDM'05), Houston, TX, USA, 27–30 November 2005.
- Winarko, E.; Roddick, J.F. Discovering Richer Temporal Association Rules from Interval-Based Data. In *Data Warehousing and Knowledge Discovery*; Springer: Berlin, Heidelberg, 2005; pp. 315–325.
- Moerchen, F. Algorithms for time series knowledge mining. In Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining-KDD '06, Philadelphia, PA, USA, 20–23 August 2006; pp. 668–673.
- Lee, Z.; Lindgren, T.; Papapetrou, P. Z-miner: An efficient method for mining frequent arrangements of event intervals. In Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '20, Virtual Event, 6–10 July 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 524–534.
- Chen, Y.C.; Weng, J.T.Y.; Hui, L. A novel algorithm for mining closed temporal patterns from interval-based data. *Knowl. Inform.* Syst. 2016, 46, 151–183. [CrossRef]
- 37. Xing, Z.; Pei, J.; Keogh, E. A brief survey on sequence classification. ACM SIGKDD Explor. Newslett. 2010, 12, 40-48. [CrossRef]
- 38. Buza, K.; Schmidt-Thieme, L. Motif-based classification of time series with Bayesian Networks and SVMs. In *Advances in Data Analysis, Data Handling and Business Intelligence;* Springer: Berlin/Heidelberg, Germany, 2010; pp. 105–114.
- 39. Ferreira, P.; Azevedo, P. Protein sequence classification through relevant sequence mining and bayes classifiers. In *Progress in Artificial Intelligence*; Springer: Berlin/Heidelberg, Germany, 2005; pp. 236–247.
- Lesh, N.; Zaki, M.J.; Ogihara, M. Mining features for sequence classification. In Proceedings of the Fifth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining-KDD '99, San Diego, CA, USA, 15–18 August 1999; pp. 42–346.
- Batal, I.; Fradkin, D.; Harrison, J.; Moerchen, F.; Hauskrecht, M. Mining recent temporal patterns for event detection in multivariate time series data. In Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining-KDD '12, Beijing, China, 12–16 August 2012; pp. 280–288.
- Lee, Z.; Trincavelli, M.; Papapetrou, P. Finding Local Groupings of Time Series. In *Machine Learning and Knowledge Discovery in Databases*; ECML PKDD 2022; Lecture Notes in Computer Science; Amini, M.R., Canu, S., Fischer, A., Guns, T., Kralj Novak, P., Tsoumakas, G., Eds.; Springer: Cham, Switzerland, 2023; Volume 13718.
- 43. Mochaourab, R.; Venkitaraman, A.; Samsten, I.; Papapetrou, P.; Rojas, C.R. Post Hoc Explainability for Time Series Classification: Toward a Signal Processing Perspective. *IEEE Signal Process. Mag.* **2022**, *3*, 119–129. [CrossRef]
- Fauvel, K.; Fromont, É.; Masson, V.; Faverdin, P.; Termier, A. XEM: An explainable-by-design ensemble method for multivariate time series classification. *Data Min. Knowl. Disc.* 2022, *36*, 917–957. [CrossRef]
- Cabello, N.; Naghizade, E.; Qi, J.; Kulik, L. Fast, accurate and explainable time series classification through randomization. *Data Min. Knowl. Disc.* 2023, 1–23. [CrossRef]
- Middlehurst, M.; Large, J.; Flynn, M.; Lines, J.; Bostrom, A.; Bagnall, A. HIVE-COTE 2.0: A new meta ensemble for time series classification. *Mach. Learn.* 2021, 110, 3211–3243.
- 47. Tan, C.W.; Dempster, A.; Bergmeir, C.; Webb, G.I. MultiRocket: Multiple pooling operators and transformations for fast and effective time series classification. *Data Min. Knowl. Disc.* **2022**, *36*, 1623–1646.

- Höppner, F.; Peter, S.; Berthold, M.R. Enriching Multivariate Temporal Patterns with Context Information to Support Classification. In *Computational Intelligence in Intelligent Data Analysis*; Springer: Berlin, Heidelberg, 2013; pp. 195–206.
- 49. Ho, T.B.; Nguyen, T.D. Mining Hepatitis Data with Temporal Abstraction. In Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, 23–27 August 2003; pp. 369–377.
- Berka, P.; Rauch, J.; Tsumoto, S. ECML/PKDD 2002 Discovery Challenge. 2002. Available online: https://sorry.vse.cz/~berka/ challenge/PAST/ (accessed on 1 August 2020).
- 51. Gordon, M. Development and Implementation of Computational Methodologies for a Systems Level Analysis of Bio-Medical Data. Ph.D. Dissertation, Ben Gurion University, Beer Sheva, Israel, 2012.
- 52. Breiman, L. Random Forests. Mach. Learn. 2001, 45, 5–32. [CrossRef]
- John, G.H.; Langley, P. Estimating Continuous Distributions in Bayesian Classifiers. In Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence, Montréal Qué, Canada, 18–20 August 1995; Morgan Kaufmann Publishers Inc.: San Francisco, CA, USA, 1995; pp. 338–345.
- Keerthi, S.S.; Shevade, S.K.; Bhattacharyya, C.; Murthy, K.R.K. Improvements to Platt's SMO Algorithm for SVM Classifier Design. Neural Comput. 2001, 13, 637–649. [CrossRef]
- 55. Landwehr, N. Logistic Model Trees. Mach. Learn. 2005, 59, 161–205.
- Frank, E.; Hall, M.; Holmes, G.; Kirkby, R.; Pfahringer, B.; Witten, I.H. WEKA A Machine Learning Workbench for Data Mining. In Data Mining and Knowledge Discovery Handbook; Springer: New York, NY, USA, 2005; pp. 1305–1314.
- 57. Zaki, M.J. SPADE: An efficient algorithm for mining frequent sequences. Mach. Learn. 2001, 42, 31-60. [CrossRef]

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.